UNIVERSITY OF CALGARY

Contextualism and the Reference Class Problem

by

Masashi Kasaki

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES

IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE

DEGREE OF DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHILOSOPHY

CALGARY, ALBERTA

September, 2010

© Masashi Kasaki 2010

UNIVERSITY OF
CALGARY

The author of this thesis has granted the University of Calgary a non-exclusive license to reproduce and distribute copies of this thesis to users of the University of Calgary Archives.

Copyright remains with the author.

Theses and dissertations available in the University of Calgary Institutional Repository are solely for the purpose of private study and research. They may not be copied or reproduced, except as permitted by copyright laws, without written authority of the copyright owner. Any commercial use or re-publication is strictly prohibited.

The original Partial Copyright License attesting to these terms and signed by the author of this thesis may be found in the original print version of the thesis, held by the University of Calgary Archives.

Please contact the University of Calgary Archives for further information:
E-mail: uarc@ucalgary.ca
Telephone: (403) 220-7271
Website: http://archives.ucalgary.ca

Abstract

The reference class problem is typically associated only with a particular theory of knowledge – type reliabilism. This thesis aims to show that the reference class problem is generalizable to virtually all other theories of knowledge, and to offer a version of contextualism as a general solution to it.

To fully articulate the force of the reference class problem, Chapters 1 through 3 are devoted to elucidating the formal-substantive distinction within both probability theory and the theory of knowledge. All substantive theories of knowledge share the probabilistic apparatus with one or another substantive theory of probability. It is because any substantive theory of probability faces one or more variants of the reference class problem that the corresponding substantive theory of knowledge cannot evade the reference class problem. Chapter 4 examines alleged solutions to the reference class problem and points out their drawbacks.

In Chapters 5 through 7, I argue that what I call 'task-sensitive contextualism' militates against the reference class problem in the theory of knowledge. My argument consists of five independent sub-arguments. In addition, Chapter 7 reveals that task-sensitive contextualism has significant ramifications for both the substantive and the formal theory of knowledge.

Acknowledgements

During the time of writing, I received tangible and intangible support from many people. In particular, my deepest gratitude goes to my supervisor, Professor Jeremy Fantl, for his generous support, constant encouragement, precise understanding, and most importantly, patient instruction and helpful criticism at every step of writing. Without his help, this thesis could never have gotten off the ground. I am also deeply grateful to the Philosophy Department of the University of Calgary that provided a good environment for my study and research. Its warm atmosphere also eased the despair unavoidably arising from writing a doctoral thesis. The University of Calgary helped me by offering scholarships for my research and travel. Last, but not least, I thank my parents and sister who were supportive and understanding, even from overseas.

The thesis was read and corrected by many, but I am the only person responsible for any remaining errors and shortcomings.

Table of Contents

List of Abbreviations

BKS: Basic Knowledge Structure

I-S: Inductive-Statistical

KR: The Requirement of Knowledge of Reliability

MCK: Multi-Premise Closure of Knowledge

MCW: Multi-Premise Closure of Warrant

MTK: Multiple-Premise Transmission of Knowledge

MTW: Multiple-Premise Transmission of Warrant

NSS: No Self-Support Principle

RBH: The Requirement of Broadest Homogeneity

RMS: The Requirement of Maximal Specificity

SCK: Single-Premise Closure of Knowledge

SCW: Single-Premise Closure of Warrant

SES: The Shift-in-Epistemic-Standard Account

S-R: Statistical-Relevance

SRC: The Shift-in-Reference-Class Account

SSI: Subject-Sensitive Invariantism

STK: Single-Premise Transmission of Knowledge

STW: Single-Premise Transmission of Warrant

UP: Underdetermination Principle

WAM: Warranted Assertability Maneuver

Chapter 1: The Notion of Epistemic Position

Introduction

The purpose of Chapters 1 through 3 is to establish that a problem akin to the reference class problem in probability theory arises in the theory of knowledge. For this purpose, I start with an elucidation of the discipline of the theory of knowledge, and propose that the theory of knowledge is to be best understood on the model of probability theory.

To begin with, what is commonly called 'probability theory' refers to two different disciplines: in one sense, it refers to a formal system consisting of a set of axioms; in the other sense, it refers to a substantive theory of what probability, as ordinarily conceived, is, or alternatively, a substantive theory of what the ordinary concept of probability is.[1] A probability theory in the second sense is taken to give a certain interpretation of the probability formally defined in the probability theory in the first sense. While Kolmogorov's formal theory establishes a privileged status of orthodoxy, there are many different substantive theories or interpretations of probability. The relationship between the two disciplines is normative: other things being equal, a substantive theory of probability must conform to the axioms in the formal theory of probability. The formal theory thus places constraints on substantive theories of probability, and thereby each of its axioms works as a criterion for the adequacy of any substantive theory.

[1] For probability, as well as for knowledge, I do not distinguish between the concept and its referent unless the distinction is relevant.

This chapter aims to draw a distinction within the theory of knowledge that is similar to the distinction between the two probability theories, and elucidate the formal, non-substantive side of the distinction (the substantive side is the topic of Chapter 3). In §1, I introduce two different kinds of approaches to epistemological problems: the partisan and the non-partisan approach. The partisan approach is driven by the goal of the substantive theory of knowledge, i.e., to identify, in a non-circular way, one or a set of *epistemic* factors – the components of knowledge which turn true belief into a case of knowledge. Various substantive theories of knowledge differ as to what factors are epistemic. The traditional, partisan approach to epistemological problems explanatorily relies on the epistemic factors thus identified. However, it is possible to pursue epistemological problems from a non-partisan perspective, i.e., without assuming any substantive theory of knowledge. The non-partisan approach to the problem of what factors are epistemic is the formal theory of knowledge, and it studies formal properties of epistemic factors in general.

In §2, an important formal notion, *epistemic position*, is introduced on the basis of Keith DeRose's writings. This notion seems to capture most, if not all, relevant features that any substantive theory of knowledge must take account of in identifying the relevant epistemic factors. Then, in §3, I summarize five formal theses about epistemic position, and propose that each of them has an axiomatic status in the sense that it offer a criterion for the adequacy of any substantive theory of knowledge. Each sub-section of §3 articulates one of the five theses: **Degree**, **Limitation**, **Threshold**, **Normativity**, and **Supervenience**.

§1. Two Kinds of Approaches to Knowledge

Epistemology is concerned with a variety of problems, prominent among which are general problems regarding knowledge. They include, for example, the definitional problem of *what knowledge is*, the explanatory problem of *why the subject has knowledge*, the normative problem of *why knowledge is valuable*, and the relational problem of *how knowledge is related to other epistemic, mental, and practical activities, states, and properties*. In the attempt to answer these questions, one way to proceed is to give a pride of place to the definitional question. The branch of epistemology that is primarily concerned with the definitional question is most properly called the 'theory of knowledge.' Traditionally, the theory of knowledge aims to analyze or explicate knowledge in terms of other conditions that are conceptually prior to the concept of knowledge, and to derive a set of *a priori* conditions that are individually necessary and jointly sufficient for knowledge; in other words, it purports to provide reductive analysis of knowledge in terms of its components.

So far, a great majority of practitioners of the substantive theory of knowledge agree that knowledge involves at least two components, truth and belief, and accordingly propose the following necessary conditions for knowledge: necessarily, for all subject *S* and proposition *p*, *S* knows *p*, only if

(Factivity) *p* is true, and

(Belief) *S* believes *p*.[2]

Factivity follows from the linguistic fact that 'know' is factive, i.e., 'knowing *p*' entails '*p*.' **Belief** seems fairly obvious; in order for *S* to know a proposition, it needs to be psychologically attained by *S*, and, for propositions, believing is the primary mode of attainment. At least on the standard account, neither **Factivity** nor **Belief** is a purely epistemic condition; **Factivity** is a semantic condition, since whether it is satisfied simply depends on what truth value *p* has; and **Belief** is a psychological condition, since whether it is satisfied depends on psychological facts about *S*.

The focus of the theory of knowledge is *epistemic* rather than semantic or psychological. As Stanley (2005, p. 2) points out, there are at least two distinct senses of 'epistemic.' In a broad sense of 'epistemic,' it denotes the factors in virtue of which true belief turns into a case of knowledge.[3] What I call a 'substantive theory of knowledge' is a theory that aims to specify what factors or components are epistemic in this sense, and different substantive theories of knowledge differ as to what factors are epistemic.

---

[2] Instead of the notion of belief, some, e.g., Lehrer (2000), employ the notion of acceptance: on Lehrer's definition, to accept *p* is to take *p* to be true on the basis of reflective evaluation purely for the purpose of attaining truth and avoiding falsity. I prefer **Belief** to Lehrer's acceptance condition, mainly because it has greater generality and suits my purpose for describing the general framework of the theory of knowledge.

[3] This sense of 'epistemic' is what Plantinga calls 'warrant,' that which "together with truth makes the difference between knowledge and mere true belief" (Plantinga, 1993b, p. 3). Just for unity of my description, I describe various philosophers' views using the concept of epistemic factor rather than that of warrant, even though they are originally formulated in terms of the latter.

When proponents of a substantive theory of knowledge attempt to answer epistemological problems, they are prone to appeal to the epistemic factor(s) identified by the theory. This way of approaching epistemological problems has two distinct features: first, it is *atomistic*, since explanatory priority is given to the epistemic component(s) of knowledge, and thus the obtaining, properties, values, and relations of knowledge are ultimately explained in terms of those of the epistemic component(s); second, it is *partisan*, since it presupposes the commitment to a particular substantive theory of knowledge.[4]

The relationship between being atomistic and being partisan needs clarification. 'Atomistic' here means both conceptually and explanatory atomistic. This may seem strange at first glance, since conceptual priority is, in principle, not to be conflated with explanatory priority. On the one hand, when a theory involves reductive analysis, it presupposes that the analysans is conceptually prior to the analysandum. On the other hand, explanatory priority is relative to a theory's explanatory purpose or significance, and it is generally independent of conceptual priority. So it is not necessarily the case that conceptual and explanatory orders go hand-in-hand. In epistemology, however, the explanandum of epistemological inquiry, knowledge, is the very analysandum of the theory of knowledge. Then, not in general, but in particular in epistemology, does it seem to be the case that conceptual and explanatory orders coincide.

That an epistemological account is partisan entails that it is atomistic, if it is true that the same order of analysis and explanation is in place in epistemology; one cannot be committed to a substantive theory of knowledge while denying that the epistemic factors it

[4] The term 'partisan' is taken from D. Howard-Synder, F. Howard-Synder, & Feit (2003), although they only use the term 'nonpartisan.'

identifies have conceptual and explanatory priority over knowledge. Thus, one way for an epistemological account to be non-partisan is to be non-atomistic. Blome-Tillmann (2007), Sutton (2007), and Williamson (2000a) deny that an atomistic approach is even possible for knowledge. According to Williamson, knowledge has no proper components – those components into which knowledge is reductively analyzable. Williamson calls his approach to epistemological problems 'knowledge first,' since, on his account, both conceptual and explanatory priority are given to knowledge rather than its components. Thus, adopting the knowledge-first approach is one way to be non-partisan. Note that Williamson does not deny that **Factivity** and **Belief** capture necessary conditions for knowledge, in the sense that knowledge entails truth and belief. What he does deny is that the latter are conceptually prior to the former.[5]

That an epistemological account is atomistic does not entail that it is partisan. One can approach epistemological problems without being committed to a substantive theory of knowledge. To affirm the explanatory priority of epistemic factors does not require endorsing any particular theory of knowledge. In addition, it is possible to pursue epistemological problems from non-partisan perspectives without endorsing the knowledge-first approach, i.e., without giving up on the explanatory priority of epistemic factors over knowledge. This may sound contradictory; to figure out what factors are epistemic is the main goal of the substantive theory of knowledge. Then, how is it possible

[5] Sutton endorses the classical analysis of knowledge that knowledge is justified true belief. But his endorsement is nothing substantial; following Williamson's lead that only knowledge justifies, Sutton adds that only knowledge is justified. Then, justification is not a proper component of knowledge, since *p is justified* entails *p is known*.

to pursue epistemological problems without relying on epistemic factors while remaining atomistic? The answer is that it is possible because *formal* characteristics of epistemic factors can be delineated without presupposing any substantive theory of knowledge.

This means that a non-partisan approach is possible even for addressing the definitional problem. For example, we might ask whether instantiating an epistemic factor, whatever it is, entails truth.[6] Such a non-partisan study of epistemic factors falls under what I call a 'formal theory of knowledge' (henceforth, I differentiate between the substantive theory of knowledge and the formal theory of knowledge' as parts of the theory of knowledge, and refer to a particular theory in each discipline with an indefinite article). A formal theory of knowledge contributes to the aim of the substantive theory of knowledge in the following way: it offers criteria for the adequacy of any substantive theory of knowledge. Insofar as one accepts the atomistic approach, the obtaining, properties, values, and relations of knowledge need to be explained in terms of those of its components (unless knowledge has some emergent property).[7] Thus, nonpartisan studies, as it were, issue promissory notes to be cashed out by a substantive theory of knowledge. More precisely, a substantive theory of knowledge must be such that the components it identifies can explain the facts about knowledge studied from non-partisan perspectives.

[6] See (Coffman, 2008), (D. Howard-Synder et al., 2003), (Huemer, 2005), and (Merricks, 1995, 1997).

[7] Especially for the value of knowledge, this might seem to involve the flat-out denial of the thesis, advocated by Franz Brentano and G. E. Moore, that knowledge is an organic unity: if *x* is organic unity, the value of *x* as a whole exceeds the values of its parts. Zagzebski (2003, 2004) interprets Sosa (2003) as being committed to the thesis of organic unity. It seems that Zagzebski's interpretation relies on the assumption that relations between parts are not parts. I am using the term 'components' to include relations between components.

Of course, it may turn out in the final analysis that the atomistic approach is hopeless, i.e., there is no epistemic factor conceptually independent of knowledge. But even if this possibility cannot be ruled out at the outset – after all, we do not yet have a firm grasp on what epistemic factors knowledge involves –, the formal theory of knowledge is still important; it is one thing to claim that a concept has no conceptual priority over another, but it is another to figure out what it is like. Thus, the formal theory is important for epistemology, whether the atomistic approach is correct or not. For it offers a way to pursue the epistemological problems independently of the issue concerning conceptual priority.

In what follows, I expound, in more detail, how the formal theory of knowledge is to proceed, and what ramifications it has for the substantive theory of knowledge.

§2. DeRose's Characterization of 'Epistemic Position'

Keith DeRose coined the term 'epistemic position' and used it in a series of writings (1992, 1995, 1996b, 1999, 2009) to illustrate a version of contextualism in epistemology. Although he defined this term in part by reference to his favored theory of knowledge, what I call the 'counterfactual robustness theory', it shortly gained wider, non-partisan usage in the literature as the non-partisan approach became more and more popular. Versions of contextualism and its rivals are all non-partisan, in that they do not necessarily presuppose a particular substantive theory of knowledge.

Contextualism in general is a linguistic position that a certain term is context-sensitive, i.e., depends for its semantic content on context of use. More particularly, contextualism in epistemology is the position that ‘know’ and its cognates are context-sensitive (I henceforth refer to contextualism in epistemology simply as ‘contextualism’). This position is typically derived from data taken to show that knowledge ascriptions or denials, sentences of the form ‘*S* knows *p*’ or ‘*S* doesn’t know *p*,’ have different semantic contents or truth conditions in different contexts of use. Although virtually all contextualists are committed to some specific substantive theory of knowledge and formulate their contextualist position in terms of it, the linguistic thread of contextualism is independent of any substantive theory of knowledge.

One of the major rivals to contextualism, subject-sensitive invariantism (henceforth, SSI), proposed by Fantl & McGrath (2002, 2007, 2009b), Hawthorne (2004a), and Stanley (2005), is also a non-partisan position, although, unlike contextualism, it is not linguistic.[8] As mentioned above, Stanley distinguishes between the two different senses of ‘epistemic.’ In the broad sense, a factor is epistemic iff it makes a difference between mere true belief and knowledge. Or, if there are multiple such factors, they are epistemic iff they contribute to such a difference. In the narrow sense, a set of factors are epistemic iff they are truth-conducive, i.e., “their existence makes the belief more likely to be true, either objectively or from the point of view of the subject” (Stanley, 2005, p.1). He characterizes SSI as the denial of what he calls ‘intellectualism’, viz., that these two senses of ‘epistemic’ are co-extensional.

[8] The term ‘subject-sensitive invariantism’ comes from DeRose (2004b). Although DeRose is reluctant to attribute it to the account proposed by Hawthorne (2004a), the term, as commonly used in the literature, refers to the accounts of these philosophers. See also footnote 22.

The core idea of SSI, then, is that some non-epistemic factor in the narrow sense counts as epistemic in the broad sense. Thus, there is a sense in which SSI is a substantive theory of what determines the obtaining of the knowledge-relation between *S* and *p*. Nevertheless, this sense is so weak that SSI remains non-partisan, since the core idea of SSI is silent on what substantive theory of knowledge it is to be combined with, though, in practice, the SSIists are committed to some specific theory of knowledge.[9]

As interest in these two non-partisan positions has increased, epistemologists tend to employ DeRose's notion of epistemic position. DeRose defines it as follows, though with the proviso 'roughly': *S* satisfies the predicate 'know *p*' iff

> *S* has a true belief that *p* and is in a *good enough* epistemic position with respect to *p*. (DeRose, 1992, p. 922)

[9] They champion some sorts of evidentialism, because of the decision-theoretic framework they adopt. Fantl & McGrath are clear that their favored theory of knowledge is internalist evidentialism; whereas, at some points, Stanley and Hawthorne, under the influence of Williamson (2000a), are more inclined toward the idiosyncratic version of evidentialism that entails that the strength of *S*'s epistemic position varies with non-epistemic factors. Hawthorne (2004a, p. 178) suggests that high stakes lower epistemic probability. As he is aware in his (2004b, p. 517), the consequence of this is that stakes affect the truth of knowledge ascriptions by shifting how high *S*'s epistemic position is; Stanley (2005, pp. 181-2) implies that stakes affect what evidence *S* has, and this has the same consequence about the strength of epistemic position. These points are only suggested by Hawthorne and Stanley, and so I focus on the 'official' discipline of SSI, according to which stakes or other non-epistemic factors shift the epistemic standard, rather than the strength of epistemic position.

The notion of epistemic position seems to stem from our everyday locution '*S* is in a position to know.' When '*S* knows *p*' is true, '*S* is in a position to know *p*' is trivially true. On the other hand, when '*S* is in a position to know *p*' is true, '*S* knows *p*' is not necessarily true. *S* might be precluded from knowing *p* for non-epistemic reasons: *S* may be semantically precluded in case that *p* is false, due to the failure of **Factivity**[10]; *S* may be psychologically precluded in case that *S* does not believe *p*, due to the failure of **Belief**. But when it is true that *S* is in an *epistemic* position to know *p*, it is not the case that *S* is precluded from knowing *p* for epistemic reasons: "you are in a position to know that *p* iff no epistemic weaknesses with respect to *p* stand in the way of your knowing that *p*" (Fantl & McGrath, 2009b, p. 84). Thus, *S* is not precluded from knowing for epistemic reasons iff *S*'s epistemic position is strong enough for knowledge.

If DeRose is right that being in a strong enough position to know is a component of knowledge, it provides the third necessary condition for knowledge. Necessarily, *S* knows *p* only if:

(Epistemic Position) *S* is in a strong enough epistemic position with regard to *p*.

Epistemic Position is open to any substantive theory of knowledge, precisely because it is a *formal* condition: it characterizes the term 'epistemic position' only *negatively*, i.e.,

[10] This presupposes that the truth or falsity of *p* does not affect the strength of epistemic position with regard to *p*. For example, if Merricks (1995, 1997) is right that instantiating the relevant epistemic factor (or being warranted) entails truth, the falsity of *p* is an epistemic factor to preclude *S* from knowing *p*. For the discussions on Merricks' arguments, see the literature cited in footnote 6.

merely as the gradable component or the sum of the gradable components of knowledge gained by subtracting non-epistemic components, such as semantic and psychological ones, from the total components of knowledge. Thus, within the formal theory of knowledge, 'epistemic position' remains primitive.[11] Furthermore, **Epistemic Position** does not specify what factors determine the strength of *S*' epistemic position. For these reasons, **Epistemic Position** needs to be supplemented by some substantive theory of knowledge. The task of any substantive theory of knowledge, then, is to give an account of what it is for *S* to be in a strong enough epistemic position with regard to *p*.

Proponents of the knowledge-first approach would contend that **Epistemic Position** does not capture a proper component of knowledge: it is only in virtue of knowledge that *S* is in a strong enough position. But yet, **Epistemic Position** is completely compatible with the knowledge-first approach. Even if it is true that nothing conceptually prior to knowledge is available for answering the question as to what the concept of epistemic position really is, it does not follow that **Epistemic Position** is false.[12] Indeed, even the

---

[11] I am not alone in calling the theory of this kind 'formal.' Merricks (1995, p. 841) describes the negative characterization of warrant (sufficiently high epistemic position in our terminology) as 'purely formal characterization.' Humberstone (2000) and Williamson (2000a, p. 32-3), however, point out that the subtraction procedure is not feasible for certain concepts. For example, *being red* may be reckoned *being colored plus something more*. But it is difficult to see what remains once *being colored* is subtracted from *being red*.

[12] If the truth of *p* is an epistemic factor, it may be possible that **Epistemic Position** is satisfied merely because *p* is true. This may seem absurd on the ground that the truth of *p* only increases the strength of epistemic position regarding *p* to a slight degree, and it cannot be high enough for knowledge. How truth

non-partisan approach to the definitional problem, if it aims to state a necessary condition for knowledge besides **Factivity** and **Belief**, must identify what factors are epistemic. (The substantive theory of knowledge is, by definition, atomistic, but for lack of better term, I henceforth use ‘the substantive theory of knowledge’ loosely so as to include the non-partisan approach to the definitional problem with this aim in its extension, unless the difference is of importance.)

For illustrative purposes, let us see how DeRose fleshes out the formal notion of epistemic position with his counterfactual robustness theory:

> An important component of being in a strong epistemic position with respect to *p* is to have a belief as to whether *p* is true match the fact of the matter as to whether *p* is true, not only in the actual world, but also at the worlds sufficiently close to the actual world. That is, one’s belief should not only be true, but also should be non-accidentally true, where this requires one’s belief as to whether *p* is true to match the fact of the matter at nearby worlds. The further away one gets from the actual world, while still having it be the case that one’s belief matches the fact at worlds that far away and closer, the stronger a position one is in with respect to *p*. (DeRose, 1995, p. 34)

The counterfactual robustness theory is a substantive theory of knowledge, since it predicates that counterfactual robustness is one, if not the only, epistemic factor in both the

---

contributes to strengths of epistemic position, however, is a complex matter. I will discuss this question in §2.3 of Chapter 7.

broad and the narrow sense; and that the strength of *S*'s epistemic position with regard to *p* is (at least in part) determined by how counterfactually robust *S*'s belief *p* is. Thus, DeRose uses counterfactual robustness to explain why one knows or does not know in any given case. As we will see in Chapter 3, construing the notion of epistemic position in terms of counterfactual robustness is just one way to give interpretation of it.

§3. The Formal Theory of Knowledge

The notion of epistemic position is widely used in recent epistemological literature, although, oddly enough, little dedicated research has been done about it. In this section, I engage in such research, and articulate the essential characteristics of the notion of epistemic position. In the standard usage by DeRose and others, the concept of epistemic position has the following characteristics[13]:

(Degree) Epistemic position comes in degrees.

(Limitation) There are upper- and lower-bounds on the strength of epistemic position.

(Threshold) There is a threshold of epistemic position for knowledge.

(Normativity) The stronger the epistemic position is, the better it is.

---

[13] The five characteristics are intended to be at least exclusive but not comprehensive. For example, if it turns out to be true that instantiating an epistemic factor entails truth, it could counts as an addition to the five characteristics.

(Supervenience) The strength of *S*'s epistemic position with regard to *p* supervenes on *S*'s epistemic factors regarding *p*.

These five theses constitute a formal theory of knowledge. And, it is because virtually all epistemologists are committed to it that it may be called the 'orthodox' theory. Each of the five theses has a status similar to an axiom in the formal probability theory, or so shall I argue. Since this point is discussed in more detail in §2 of Chapter 2, I only briefly mention the relevant respects of similarity.

First, just as axioms of a formal probability theory are abstracted from our ordinary practice of estimating, calculating, and assigning probability, the five theses are abstracted from our epistemic practice of evaluating, attributing, and denying one's epistemic status. Second, the five theses, by their own nature, are non-partisan, and hence open to a wide variety of interpretations given by substantive theories of knowledge. An interpretation must specify (i) what type(s) of factor satisfies these theses, and (ii) how the strength of epistemic position is determined by such a type(s) of factor in any given case. Third, the five theses capture, at least *prime facie*, the conditions all epistemologists need to take account of in developing a substantive theory of knowledge. So, it is *prime facie* mandatory for any substantive theory of knowledge to conform to them. They place constraints on epistemological theorizing. In these regards, the relation between the five theses and the substantive theory of knowledge is analogous to that between formal and substantive theories of probability. Of course, there are differences; for example, the axioms of any formal probability theory determine how probabilities are distributed over the sample space, and the five theses do nothing like that. The three respects in which they are analogous, however, are enough for my purpose of illuminating the relationship between the partisan and the non-partisan approach to the epistemological problems.

§3. 1. Degree

Degree says that one's epistemic position with regard to a proposition admits of degrees. DeRose expresses this by several different locutions: *S*'s epistemic position can get better or worse, it can get higher or lower, or it can get stronger or weaker. Two rationales for **Degree** are found in DeRose's writings. First, epistemic position can be strengthened or improved. On his counterfactual robustness theory, the strength of *S*'s epistemic position with regard to *p* is determined by how robust *S*'s belief *p* counterfactually is. *S* can strengthen her epistemic position by the means of acquiring more accurate discriminatory ability about *p*; by this means, *S*'s belief *p* would become more counterfactually robust, i.e., match a wider range of possible worlds. In addition, DeRose suggests that *S* can strengthen her epistemic position with regard to *p* by the means of gathering more evidence for *p* (1995, p. 33; 2000, p. 135). This implies that DeRose is a pluralist about epistemic factors, in that he admits that at least two factors, counterfactual robustness and evidence, contribute to the overall strength of epistemic position.

Second, we can compare how well one is epistemically positioned inter- or intra-subjectively. We seem to have an intuitive grasp on how one is epistemically better or worse off than oneself or others with respect to one or different propositions. Such a grasp is reflected in our ordinary language, when we say, e.g., "S_1 has better discriminatory ability to perceive *p* than S_2 does." or "S_1 has better evidence for *p* than S_2 does.", as well as

"S_1 has better evidence for *p* than S_1 does for *q*." or "S_1 has better discriminatory ability to perceive *p* than S_1 does to perceive *q*." When we assert these or similar sentences, certainly, we credit S_1 with a good epistemic score on *p*, compared with S_2's score on *p* or S_1's score on *q*.

Degree entails that there is a scale for epistemic positions on which they take a certain value; the value may be either a number-value or an interval-value.[14] Either way, the interpretation of **Degree** cannot be settled independently of the interpretation of **Supervenience**. For **Supervenience** is relevant for what the strength of epistemic position is a measure of.

§3. 2. Limitation

Limitation is divided into two parts: that there is a maximal epistemic position, a point above which epistemic position cannot get any higher; and that there is a minimal epistemic position, a point below which epistemic position cannot get any lower. On most substantive theories of knowledge, the existence of one of the upper- or the lower-bound entails the existence of the other, since the epistemic position regarding *p* is inversely proportional to the epistemic position regarding ~*p* (this follows if the strength of epistemic position conforms to Kolmogorov's axioms).

---

[14] It seems that most epistemologists conceptualize epistemic positions as being of a number-value. Wunderlich (2009) argues that degrees of justification are interval-valued rather than point-valued. If so, Wunderlich argues, degrees of justification cannot be compared with each other in some cases.

DeRose's remarks about epistemic position do not explicitly touch on **Limitation**. Nevertheless, his counterfactual robustness theory contains not merely **Limitation** but also a concrete interpretation of it, viz., that *S*'s epistemic position with regard to *p* is maximal when *S*'s belief *p* matches the fact as to whether *p* is true across all possible worlds, and *S*'s epistemic position with regard to *p* is minimal when *S*'s belief *p* matches the fact concerning *p* across no possible worlds.[15] In this regard, the counterfactual robust theory satisfies the formal requirement imposed by **Limitation**. Likewise, other substantive theories of knowledge involve their own ways to interpret **Limitation**.

In general, the highest epistemic position is typically identified with some sort of certainty or infallibility.[16] Infallibilism is the position that knowledge requires infallibility, and fallibilism is its denial. Combined with **Limitation**, fallibilism entails the following thesis:

---

[15] As noted in §2, DeRose is also a proponent of evidentialism, but it is unclear on his pluralist account how evidence contributes to the overall strength of epistemic position. Depending on how it does, different interpretations of **Limitation** would emerge (or, **Limitation** might be abandoned if his account entails that there is no upper limit on the strength of epistemic position). Note that the counterfactual robustness theory entails that the truth or falsity of *p* at the world of evaluation, the world in which *S* is located, is an epistemic factor. For the minimal epistemic position is such that *S*'s belief fails to match the fact concerning *p* at any world including *S*'s own. *S*'s epistemic position with regard to *p*, then, can be improved to a slight degree if *S*'s belief matches the fact as to whether *p* is true at *S*'s world.

[16] Leite (2004) denies this: on his account, knowing fallibly has the highest status. His account faces a challenge from **Normativity**, since it, at the very least, intuitively plausible that infallible knowledge is epistemically better than fallible knowledge. Leite tries to handle this problem by appeal to our everyday practice, though I do not think that it is successful. Either way, Leite's account is consistent with **Limitation**.

> (Fallible Knowledge) *S* knows *p* fallibly iff (i) *S* knows *p*, and (ii) *S*'s epistemic position with regard to *p* is not maximal.[17]

This captures at least one central strand of fallibilism, but it is not all there is to fallibilism. Fallibilism is often formulated in terms of possibilities of error: knowing fallibly is knowing despite a possibility of error. This entails a different conception of what knowing fallibly amounts to:

> (Fallible Knowledge') *S* knows *p* fallibly iff (i) *S* knows *p*, and (ii) *S*'s epistemic position with regard to *p* is compatible with the falsity of *p*.

Some gloss is in order. What is meant by (ii) is that it is possible for *S* to have the same strength of epistemic position with regard to *p* as *S* actually does, while *p* is false. More

[17] The definition of fallibilism that is closest to this is found in Fantl & McGrath (2007, p. 559): "*S* has fallible knowledge that *p* iff *S* knows that *p* but *S*'s strength of epistemic position regarding *p* is not maximal with respect to justification (i.e., there are [higher or better] epistemic positions regarding *p* with respect to justification)" (the terms inside the square brackets are suggested by one of the authors, Jeremy Fantl). They claim that this definition has greater generality than other ones since it applies to various theories of knowledge, such as evidentialism, type reliabilism, and token reliabilism – though token reliabilism is not usually taken to be about justification; the traditional definition only focuses on internalist evidentialism. For various definitions of fallibilism, see (Fantl & McGrath, 2009b, ch. 1), (Reed, 2002), and (Vahid, 2008). Nearly all definitions found there can be derived from **Fallible Knowledge** by interpreting it with certain theories of knowledge.

concretely, (ii) holds, relative to a set of epistemic factors determined by a substantive theory of knowledge, iff instantiating such a set of factors does not entail the falsity of *p*.

Fallible Knowledge and **Fallible Knowledge'** are not equivalent as they stand, though they are commonly taken to be so. The alleged equivalence is motivated by a particular conception of epistemic factor we have already encountered, viz., that epistemic factors are truth-conducive. This conception of epistemic factor is probabilistic; truth-conducive factors are such that instantiating them make *p* more probable than otherwise. Thus, it is natural to construe the strength of epistemic position in probabilistic terms: the strength of *S*'s epistemic position with regard to *p* is assimilated to the probability or likelihood of *p* conditional on the obtaining of a set of epistemic factors. With the probabilistic conception of epistemic factor in place, **Fallible Knowledge** allows *S* to know *p* even when *S*'s instantiating a set of epistemic factors makes *p* probable to some non-maximal extent, i.e., the probability of *p* conditional on *S*'s instantiating such a set of epistemic factors is less than 1. With the assumption that the probability of *p* is less than 1 iff there is some possible world in which *p* is false,[18] **Fallible Knowledge** is equivalent to **Fallible Knowledge'**.

[18] Since this assumption is a biconditional, it includes two conditionals: if the probability of *p* is less than 1, then there is some possible world in which *p* is false, and if there is some possible world in which *p* is false, then its probability is less than 1. Hájek (2007a) notes that the former conditional is a received one, whereas the latter conditional is more controversial. Even if the latter conditional is denied, insofar as the former conditional is true, with the probabilistic conception of epistemic factor, it is still the case that **Fallible Knowledge** entails **Fallible Knowledge'**. Regardless of this point, **Fallible Knowledge** has greater generality than **Fallible Knowledge'**.

Note that **Limitation**, even when combined with the probabilistic conception of epistemic factor, remains a formal thesis; all it means is simply that the lowest epistemic position is the position whose probability is 0, and the highest epistemic position is the position whose probability is 1. That is, epistemic positions are standardized (just as probabilities are in Kolmogorov's formal theory). The probabilistic conception does not determine what factor(s) is relevant for determining the strength of epistemic position. Of course, it is possible to interpret 'probable' in the probabilistic conception in one way or another, by adopting some substantive theory of probability. Indeed, as I will argue in Chapters 3, each substantive theory of knowledge is to be understood on the model of one or another substantive theory of probability. Thus, to determine what interpretation of probability is relevant for the probabilistic conception would fix what interpretation of the epistemic position is true.

Limitation is intimately related to **Degree** for an obvious reason: the upper- and the lower-bound are certain points on the epistemic scale. Thus, to set the upper- and the lower-bound as 1 and 0 respectively, as the probabilistic conception of epistemic factor does, entails that the epistemic scale is probabilistic: an epistemic position takes a point- or interval-value within the interval [0, 1]. The probabilistic measure can be complex; for example, Carrier (1993) and Vahid (2008), while accepting the probabilistic conception of epistemic factor, hold that infallibility requires second-order knowledge, i.e., knowledge of one's own knowledge.[19] On their view, the probabilistic scale is enlarged so as to include a dimension along which second-order knowledge is evaluated.

§3. 3. Threshold

If fallibilism is true, *S* can know even though the strength of *S*'s epistemic position is not maximal. But, at the same time, it is not plausible that *S* can know even though the strength of *S*'s epistemic position is minimal. It is quite reasonable for the fallibilist to assume that there is a point on the epistemic scale set by one or another interpretation of **Degree** such that *S* is in a position to know only if the strength of *S*'s epistemic position equals or exceeds it. In other words, there is a threshold for the strength of epistemic position required for knowledge. **Threshold** captures this idea. DeRose designs the notion of epistemic position with an eye to **Threshold**. He claims that what he calls the 'epistemic

[19] As far as **Limitation** is concerned, this position is not entirely absurd. Obviously, due to **Factivity**, (a) that *S's knowing that S knows p* entails *S's knowing p* is true, but its converse (b) that *S's knowing p* entails *S's knowing that S knows p* is false. Then, insofar as strengths of epistemic position conform to Kolmogorov's axioms, it follows from the truth of (a) that *S* is in at least as strong an epistemic position with regard to *p* as with regard to *S knows p*; whereas, it follows from the falsity of (b) that *S* is not in at least as strong an epistemic position with regard to *S knows p* as with regard to *p*. Putting together, *S* is a stronger position with regard to *p* than with regard to *S knows p*. This entails that second-order knowledge always puts *S* in a stronger epistemic position than does first-order knowledge, whether first-order knowledge is fallible or not. Reed (2002, p. 148) argues against Carrier that he is committed to a level-confusion in Alston's (1980) sense.

standard' determines how strong an epistemic position *S* must be in, in order for *S* to be attributed knowledge truly.[20]

It is a highly controversial matter whether the epistemic standard, and accordingly, the epistemic threshold are always constant or not. Furthermore, it is also a controversial matter what it means that the threshold is constant or variable. Let us call the position that the threshold is constant 'constantism.' Infallibilism entails constantism, since it requires *S*'s epistemic position to be maximal in order for *S* to be in a position to know *p*, for any *S* and *p* (insofar as what position is maximal is constant). That is, the threshold is held fixed at the highest value, regardless of any parameter such as proposition, time, or place. On the other hand, obviously, the converse of the entailment does not hold, and hence constantism is compatible with fallibilism. There are three main fallibilist positions about **Threshold**.

First, fallibilist constantism is possible. The epistemic threshold is always held fixed at some point on the epistemic scale between the maximal and the minimal value. This seems like the position most traditional invariantists are committed to. Invariantism is the denial of contextualism, and the position that the semantic content of 'know' or what relation 'know' expresses does not vary with context of use. If combined with constantism, the semantic content of 'know' is *being in an epistemic position whose strength is equal to or higher than d*, where *d* is a constant value on the epistemic scale. Of course, traditional invariantists differ as to what the epistemic scale is, having different interpretations of **Degree**.

---

[20] Although most substantive theories entail its denial, it may be possible that the threshold set by the epistemic standard is minimal. If this is the case, **Epistemic Position** is satisfied even when *S* is in no strong epistemic position.

Second, a prominent version of contextualism (henceforth, standard contextualism), championed by DeRose, Cohen (1988, 1999), and Lewis (1996), is in a subtle position with regard to constantism. Standard contextualism is the position that the semantic context of 'know' uttered by an attributor in a context is associated with an epistemic standard, which determines whether 'know' is satisfied by *S* (with regard to *p*) in that context. To use DeRose's terminology, 'attributor factors' – factors belonging to the attributor of putative knowledge in context – set epistemic standards:

> Attributor factors set a certain standard the putative subject of knowledge must live up to in order to make the knowledge attribution true: They affect *how good an epistemic position the putative knower must be in to count as knowing*. They thereby affect the truth conditions and the content or meaning of the attribution. Subject factors, on the other hand, determine whether or not the putative subject lives up to the standards that have been set, and thereby can affect the truth value of the attribution *without* affecting its content: They affect *how good an epistemic position the putative knower actually is in*. (DeRose, 1992, pp. 921-2)

Thus, with standard contextualism in sight, attributor factors are to be distinguished from subject factors – factors belonging to the subject of putative knowledge. The latter determine the strength of *S*'s epistemic position, as made explicit in **Supervenience**.[21]

[21] The proponents of standard contextualism differ as to what substantive theory of knowledge they adopt in interpreting **Supervenience**. Cohen favors internalist evidentialism; Lewis is committed to an interpretation

Standard contextualists claim that attributor factors include a pragmatic factor, how much is at stake on *p* for the attributor, and a psychological factor, what alternatives to *p* are salient to her.

Standard contextualism only entails that 'know' expresses different relations in different contexts of use, depending on what epistemic standard is in play in the contexts. Each of the relations 'know' is capable of expressing is associated with one epistemic standard or threshold. So, standard contextualism, in effect, entails constantism, insofar as constantism is about knowledge rather than 'know.' What standard contextualism denies is constantism': the epistemic threshold for the satisfiability of 'know' is always constant. For, even when the triplet parameters of subject *S*, proposition *p*, and time *t* are held fixed, the truth condition of '*S* knows *p* at *t*' varies across contexts where different epistemic standards are in play.

Third, SSI, while being a version of invariantism, denies constantism. DeRose (2004b), in which the term 'subject-sensitive invariantism' originates, defines SSI as "any view on which the standards a subject must meet for an attribution of knowledge to her to be true vary, but on which they are set by the subject's context in such a way that they comprise the key truth condition for any speaker's attribution of knowledge to that subject" (p. 346).[22] Subject factors, or more precisely, truth-conducive subject factors, have been

---

of the relevant alternatives theory in terms of counterfactuals. The relevant alternatives theory is a formal theory of knowledge, not a substantive one. For more on this point, see footnote 43 of Chapter 7.

[22] DeRose carefully avoids attributing this position to Hawthorne (2004a), since he finds no talk about epistemic standards in Hawthorne's discussions of his position (although DeRose only refers to a draft of

Hawthorne's book, the published version contains no change in this regard). Hawthorne's silence on the topic may be explained by his liberal conception of epistemic standard:

> I do suspect, however, that the notion of 'standards' is sometimes taken too seriously in these discussions. To say that you have higher standards than I for 'know' is just to say that there are subject-time-proposition triples that fall under my extension of 'know' but not under yours – and not vice versa. We needn't reify *standards* into some kind of index around which to build a semantics for 'know'. We should also be careful not to take for granted that 'one's standards' for knowledge are a priori available, as if we had ready access to how one's dispositions – in combination with the world – determine an extension of 'know' in one's own mouth. (Hawthorne, 2004a, fn. 23, ch. 2, p. 57)

Here Hawthorne claims that the epistemic standard shifts iff the extension of 'know' varies while the triplet parameters of *S*, *p*, *t*, are held fixed. SSI does not involve the variability of epistemic standard in this sense, as I mention in a moment. But the epistemic standard in Hawthorne's sense is different from that in our sense derived from the usage of the term by DeRose and others. If the latter sense is adopted, it is appropriate to construe Hawthorne's position as a version of SSI.

Hawthorne's conception of epistemic standard is both too broad and too narrow. First, it is too broad, since it fails to distinguish between standard contextualism and non-standard contextualism I defend in this thesis; according to non-standard contextualism, the extension of 'know' varies across contexts, because contexts affects the strength of epistemic position rather than the epistemic standard. Second, it is too narrow, since it fails to include, for example, the position that each proposition is associated, in one way or another, with a certain threshold for it to be known, though I know no one who is committed to such a position. Such a position is not entirely absurd: for example, suppose that a proposition intrinsically has a measurable information value, and a proposition with higher value requires one to be in a stronger epistemic position to know than does a proposition with lower value. Then, it is possible that the extension of 'know' does not vary between propositions and so it does not shift in Hawthorne's sense, while a proposition requires meeting a higher standard than another in our sense. Fantl & McGrath (2009b, pp. 194-200) points out that a similar traditionally taken to be determinants of the strength of epistemic position. SSI enlarges the range of such determinants by including non-truth-conducive factors. SSIists appeal to the subject factors akin to the attributor factors employed by standard contextualists: Fantl & McGrath (2002, 2007, 2009b) and Stanley (2005) focus on stakes on the side of the subject, and Hawthorne (2004a) adds salience of error possibilities to the subject. The function of these factors is to shift the epistemic standard, just as the function of the attributor factors is on contextualism. Unlike contextualism, SSI does not entail that the epistemic standard varies while *S* and *t* are held fixed, because they are *subject* factors – I will discuss this point in more detail in §3. 5.[23]

§3. 4. Normativity

Normativity is implicitly assumed when DeRose uses 'being in a stronger epistemic position' and 'being in a better epistemic position' interchangeably, or when he reckons being in a stronger epistemic position as 'improvement' in some relevant sense. That is, the dimension along which one's epistemic positions go up and down is also an evaluative dimension, or equivalently, the epistemic scale on which epistemic positions

---

view, which concerns the cognitive value for *S* of believing a proposition, has the same consequences as SSI in DeRose's sense, and attributes it to Lehrer (2000).

[23] As Fantl & McGrath (2009a) notes, contextualism and SSI are not logically contradictory. It is a perfectly consistent position that both subject and attributor factors determine epistemic standards while only the latter affect the semantic content of 'know.' A version of such 'subject-sensitive' contextualism is espoused by John Greco (2003, 2004a, 2008, 2010).

takes a value is also an evaluative scale. Then, the question is what type of value **Normativity** is about. This is a topic much discussed in recent epistemology; many regard epistemology to currently be in the process of a 'value turn,' a movement of reconstructing and reconsidering epistemological issues in value terms. Since problems concerning epistemic value will take us too far afield, I only discuss why **Normativity** is important for the theory of knowledge.

Normativity and **Supervenience** are intimately interrelated. **Supervenience** says that the strength of epistemic position is determined by a set of epistemic factors. **Normativity** requires that an epistemic factor be that which puts *S* in a normatively better epistemic position than otherwise.[24] This point is highly important for the substantive theory of knowledge. For ignoring this would end up counting extraneous factors as epistemic.

Huemer (2005) argues that there is no uniquely correct theory or analysis of knowledge. Suppose that a correct theory predicates that the epistemic condition $\Phi(p)$, with **Factivity** and **Belief**, is necessary and sufficient for knowledge. But then, Huemer adds, $p \supset \Phi(p)$ should count as an epistemic condition, since the conjunction of **Factivity**, **Belief**, and '$\Phi(p)$' is logically equivalent to the conjunction of **Factivity**, **Belief**, and '$p \supset \Phi(p)$.' One way to avoid this result is to appeal to **Normativity**; there is no sense in which satisfying the condition $p \supset \Phi(p)$ makes *S* epistemically better off than otherwise, since it is

[24] I intend epistemic factors to include anti-Gettiering conditions, which are obviously epistemic in both of the two senses articulated by Stanley. Kvanvig (2003) holds that no anti-Gettiering condition adds further value to justified true belief. For critical discussions of Kvanvig, see (DePaul & Grimm, 2007) and (Pritchard, 2008b).

automatically satisfied if *p* is false. Huemer is aware of the normative aspect of epistemic factors, but refuses it without reason.[25] His argument, then, can be taken as a *reductio* for **Normativity**: neglecting **Normativity** would result in a deficient analysis. This is the constraint **Normativity** places on the substantive theory of knowledge.

Still, this does not answer the question of what value is involved in **Normativity**. It is commonplace in epistemology to reckon the epistemic dimension of evaluation to be teleological or instrumental: instantiating an epistemic factor is valuable relative to our epistemic goal.[26] Despite its widespread acceptance, however, it is controversial what this conception of epistemic value amounts to. There are mainly three respects in which philosophers disagree regarding the teleological or instrumental conception of the epistemic goal: (a) what is it the goal of? (b) what is the goal in question? (c) what does 'being relative' mean?

As regard to (a), there are mainly three positions: one position, influenced by Peirce, focuses on the goal of intellectual inquiry; another position takes the goal in question to be of our cognition in general; and a more narrow attention is often paid to the goal of belief.

[25] Huemer indeed is ignoring a point implied by his target, put forth in Merricks (1995), to the effect that instantiating an epistemic factor is 'improvement' in the epistemic sense. Another way to block Huemer's argument is to require that an epistemic factor be epistemic in both the broad and the narrow sense. Satisfying the condition $p \supset \Phi(p)$ does not increase the strength of epistemic position regarding *p*, and hence is not epistemic in the narrow sense; again because it is satisfied when *p* is false.

[26] Proponents of the teleological conception of epistemic value are virtually all epistemologists. They, however, disagree as to whether there is only one epistemic goal. Kvanvig (2003, 2005) espouses pluralism by emphasizing that understanding is, among others, our important epistemic goal (see also (Zagzebski, 1996)). Stich (1990) proposes a radically pluralistic view on cognitive goals on the basis of his pragmatism.

The differences among these positions, however, may not be substantive. For one, to follow Peirce, belief is the product of inquiry; more precisely, intellectual inquiry ends with either producing, reconfirming, or suspending belief. For another, believing is a primary mode of cognition. Thus, putting the ing-ed ambiguity aside, all the three positions may be upheld without inconsistency.

Differences regarding (b) are more substantive. Pragmatists, such as Rorty (1995) and Rosenberg (2002), typically regard the goal in question to be justification; whereas, others, such as Conee & Feldman (1992), Sartwell (1991, 1992), Williamson (2000a), and Zagzebski (1996), take the goal to be knowledge. But by far the majority of epistemologists endorse the view that our epistemic goal is truth. Although Alston (2005) and David (2001) propose a charitable interpretation of this view on which the truth goal involves both attainment of truth and avoidance of error, and this is indeed what many epistemologists have in mind, one should not be identified with the other, as William James emphasized. Furthermore, certain theories are indeed specifically built around one rather than the other. For example, Levi (1980) reckons the goal of inquiry at a certain stage to consist solely in avoidance of error; sensitivity theory is focused on avoidance of error, as opposed to safety theory, which is focused on attainment of truth – sensitivity and safety theories are discussed in detail in §3. 3 of Chapter 3.

(c) amounts to the question of what it is for instantiating an epistemic factor to be instrumentally valuable for the epistemic goal. For simplicity, I only discuss the truth-goal thesis, that the epistemic goal is to attain every truth. David (2001, 2005) notes that the truth-goal thesis should be construed to be that $\forall p\ [G\ (Tp \supset Bp)]$, where p is a variable ranging over propositions, G is a sentential operator abbreviating ‘it is my goal that’ or ‘I want it to be the case that,' and T and B are predicates for being true and believing, respectively. This is the *de re* reading of the truth-goal thesis, on which one wants every proposition to be such that one believes it if it is true.[27] [28] The instrumental value of something can be measured in terms of how effective it is as a means to the goal in question. However, it is not as obvious how this consequentialist measure is to be applied to the truth-goal thesis, once it is construed in David's fashion. Many, e.g., Alston (1985, 2005), BonJour (1985), and Kvanvig (2003), articulate the consequentialist measure probabilistically.[29] That is, an epistemic factor regarding p is more epistemically valuable iff instantiating it raises the likelihood or probability of the truth of p. This is again the

---

[27] David contrasts this construal of the truth-goal thesis with the one that $G\ [\forall p\ (Tp \supset Bp)]$. The latter is the *de dicto* reading of the truth-goal thesis, on which one wants to believe all true propositions. For an attempt to formulate the goal thesis in the *de dicto* fashion more precisely, see (Latus, 2000). Horwich (2006) and Sosa (2003) provide formulations of the goal thesis in the *de re* fashion. David (2001) points out that the *de dicto* reading, with the consequentialist measure, has the absurd result that one is never in a good epistemic position with regard to a proposition which *de facto* leads to many false beliefs.

[28] I put aside the question as to what the domain of the universal quantifier is. Chisholm (1977, p. 14) and Sosa (2003) restrict the domain to those proposition that are of interest for us, on the ground that trivial truths are of no value (though Sosa's view is subtle, since he, in his (2004b), distinguishes true propositions from true beliefs, and suggests that the latter may be intrinsically valuable). Fantl & McGrath (2009b, pp. 163-170) defend the view that any true belief is *prima facie* valuable to a slight degree (see also (David, 2005, p. 298-9)).

[29] They talk as if the teleological conception of epistemic value is identical with the probabilistic conception of epistemic factor, though Bonjour equivocates between the two construals of the truth-goal thesis noted in footnote 27.

probabilistic conception of epistemic factor. Thus, the probabilistic conception has a place in a natural answer to (c).

The value turn, indeed, puts into doubt that the instrumental value relative to the truth goal is all there is to epistemic value. For the instrumental value fails to explain why instantiating an epistemic factor adds extra value to true belief. When *p* is true, the truth goal is trivially achieved. Why is it more valuable to instantiate an epistemic factor regarding *p*? Insofar as the value of instantiating an epistemic factor is not distinctive, i.e., it is merely parasitic on the value of truth, there is no way to answer this question.[30]

This is the leading question in the value turn, and its solution is being sought by many. For example, Kvanvig (2003) and Zagzebski (2003) respond that the consequentialist measure of instrumental value misses a delicate but important distinction between effective means and intentional means to a goal; it only focuses on the quantitative efficacy of a means to bring about the goal, and ignores the qualitative aspects of exercising a means for the purpose of bringing about the goal. They argue that if a factor is able to play a role in intentionally bringing about the epistemic goal (Kvanvig) or being motivated to achieve it (Zagzebski), it is distinctively valuable; Kvanvig implies that this point may speak in favor of internalism, and Zagzebski deploys it on behalf of her virtue theory.

A lesson to be drawn from the value turn is that instrumentally valuable factors, as they are picked up by the probabilistic-consequentialist measure, may not be epistemic in

[30] This problem is called the '*Meno* problem' or the 'swamping problem' – it dates back to Plato's *Meno*. The revival of the *Meno* problem is induced, especially, by Zagzebski (1996, 2003). Although she reckons it to be the problem for type reliabilism, any theory of knowledge cannot evade it, at least easily. For more on the *Meno* Problem, see (Kvanvig, 2003) and (Pritchard, 2007).

the strong sense in which it is of distinctive epistemic value. To settle the question of what factors are of such value is beyond the scope of the formal theory of knowledge, since it requires substantive inquiry into epistemic value.[31]

§3. 5. Supervenience

Supervenience says that epistemic factors regarding *p* are supervenience bases for epistemic positions with regard to *p*.[32] [33] No two subjects, *S* and *S*', can differ with respect to their epistemic positions without also differing with respect to epistemic factors. Among various notions of supervenience, **Supervenience** is to be construed in terms of wide supervenience:

[31] The *Meno* problem does not require that epistemic value be intrinsic. Brogaard (2006) and Pritchard (2007, 2008b) argue that instantiating an epistemic factor may add up to final value: final value of a thing depends on its extrinsic properties but is still valuable for its own sake. The concept of final value comes from Rabinowicz & Rønnow-Rasmussen (2000, 2003).

[32] If truth or falsity is to count as an epistemic factor, **Supervenience** can be reformulated in terms of 'regarding true proposition *p*.'

[33] **Supervenience**, as far as I know, is first explicitly formulated by D. Howard-Synder et al. (2003) in reconstructing one of the arguments in Merricks (1995). They criticize Merricks' internalist interpretation of **Supervenience**, according to which the strength of epistemic position only supervenes on internalist factors. Theses similar to **Supervenience** are found in Feldman & Conee (2001), Silins (2005), and Sosa (1980), although they are formulated narrowly in terms of internalist justification.

> Property *A* strongly supervenes on property *B* iff for all possible world *w* and *w*' and individuals *x* in *w* and *x*' in *w*', if there is no difference between *x* and *x*' with regard to *B*, there is no difference between them with regard to *A*.

Supervenience is the axiom that is most deeply entrenched in the enterprise of the theory of knowledge, for several reasons: first, it is important for the formal theory of knowledge, since, as mentioned above, interpreting it in one way or another partially determines how the other axiomatic features are interpreted; second, in order to give a certain interpretation of **Supervenience**, a substantive theory of knowledge must specify the epistemic factor(s) that determines the strength of *S*'s epistemic position. Substantive theories of knowledge compete for the best candidate factor(s) to give interpretation of **Supervenience** (in this regard, no substantive theory of knowledge violates it); third, **Supervenience** follows from the reductionist tenet of substantive theories; Jackson (1998) points out that when a theory explains a property *A* by another property *B* in a reductive manner, it is not enough that the occurrences of *A* always correlate with those of *B*. There must be a necessary connection between *A* and *B*, and such a necessary connection is accounted for by a supervenience relation: for all possible world *w*, *x* and *y* are *A*-alike in *w*, only if they are *B*-alike in *w*. **Supervenience** is simply such a supervenience relation with regard to epistemic position.

Practitioners of the substantive theory of knowledge, thus, deploy **Supervenience** when they engage in the definitional question of what knowledge is. They are prone to use the method of counterexample in settling the question as to which substantive theory is on the right track; a counterexample to a theory is constructed by finding a hypothetical case or a possible world in which a counterpart of the subject instantiates the factor specified by that theory but fails to know.

The ramification of **Supervenience** for the substantive theory of knowledge may be put this way: a substantive theory of knowledge specifies the function that maps a proposition or belief *p* onto a value on the scale pertaining to **Degree**, in such a way that the value corresponds to the strength of epistemic position with regard to *p*. If the theory of knowledge is pluralist, the value-assigning function may be a high-order function including a weighting function of each of the epistemic factors, which specifies how much each of them contributes to the overall strength of epistemic position.[34]

Standard contextualism respects **Supervenience**. Epistemic factors are subject factors in DeRose's sense. Attributor factors set the relevant epistemic standard, and accordingly fix the semantic content or truth condition of '*S* knows *p*' in a given context. Subject factors, as determinants of the strength of *S*'s epistemic position, on the other hand, determine whether *S* satisfies the truth condition thus fixed. What subject factors *S* instantiate at a given time is not context-dependent.

Supervenience concerns the strength of epistemic position *simpliciter*, and therefore is not to be conflated with what Fantl & McGrath call 'purism,' a similar but significantly different supervenience thesis, which concerns the strength of epistemic position *enough to know* rather than the strength of epistemic position *simpliciter*:

---

[34] Most pluralist theories of knowledge, however, fail to satisfy this requirement. For instance, DeRose is utterly silent on how the two epistemic factors he reckons relevant, counterfactual robustness and evidence, are weighted.

> (Purism) For any two possible subjects *S* and *S*', if *S* and *S*' are alike with respect to the strength of their epistemic position regarding a true proposition *p*, then *S* and *S*' are alike with respect to being in a position to know that *p*. (Fantl & McGrath, 2007, p. 558)

Fantl and McGrath characterize SSI as the denial of purism. Purism is equivalent to intellectualism, as Stanley formulates it, that the two senses of 'epistemic' are not co-extensional. Stanley takes epistemic factors in the broad sense as those which make a difference between true belief and knowledge. If **Epistemic Position**, with **Factivity** and **Belief**, is sufficient for knowledge, then being in a strong enough position with regard to *p* is the only condition that makes a difference between true belief that *p* and knowledge that *p*. However, *S*'s instantiating epistemic (truth-conducive) factors regarding *p* in the narrow sense, according to SSI, is not sufficient to put *S* in a strong enough epistemic position to know *p*. For non-epistemic (non-truth-conducive) subject factors, such as stakes and salience, may induce shift in epistemic standard while *S*'s truth-conducive factors are being held fixed. Suppose that *S*, with instantiating the relevant truth-conducive factors, meets a lax standard in place and believes truly *p* at t_1. Then, *S* knows *p* at t_1. If the epistemic standard has risen at t_2 while everything else is held fixed, then *S* may fail to meet the elevated standard at t_2. Then, *S* does not know *p* at t_2.

The result is the denial of purism. Being in a strong enough position is not determined solely by truth-conducive subject factors. Rather, it is determined by both truth-conducive subject factors and non-truth-conducive subject factors. This is tantamount to the denial of intellectualism; non-truth-conducive factors are relevant for the difference between *S*'s knowing at t_1 and *S*'s failure to know at t_2, and so they must count as epistemic in the broad sense.[35] The argument for the denial of purism or intellectualism presupposes **Supervenience**, and hence SSI does not violate **Supervenience**.

[35] As noted above, there is a sense in which SSI is a substantive theory, in a yet non-partisan way. More precisely, it is substantive in that it offers a necessary condition or component for knowledge; and yet, it is non-partisan for at least two reasons: first, the core idea of SSI assumes no commitment to a particular theory of knowledge; second, it is open to various theories of what subject factor determines the epistemic standard.

Chapter 2: The Reference Class Problem in Probability Theory

Introduction

This chapter sets up the following one, which aims to establish that there is a family of similar problems for substantive theories of knowledge. These problems are lumped together under the label of 'reference class problem'; they are lumped together, since they are all structurally alike; and they are thus labeled, since they arise in the same ways as (variants of) the reference class problem in probability theory does. Therefore, the next chapter will draw a parallel between the theory of knowledge and probability theory. Before proceeding, then, an exposition of the reference class problem in probability theory is in order.

What is commonly called the reference class problem in probability theory is not a single problem, but rather a host of structurally similar problems. Although the reference class problem is typically reckoned a problem fatal to frequentism, other substantive theories of probability face problems of similar kinds. The main purpose of this chapter, then, is to give a summary of the specific forms in which the reference class problem generate for substantive theories of probability.[1] Substantive theories of knowledge face the reference class problem in similar forms, because they employ probabilistic apparatus

[1] Hájek (2003, 2007b) offers fully detailed discussions on the variants of the reference class problems for virtually all theories of probability, and I owe much of this chapter to his papers. I critically examine Hájek's argument in §4. 1.

in one way or another, and thus inherit the same difficulties from substantive theories of probability.

The distinction between formal and substantive theories of knowledge is modeled on a similar distinction within probability theory. §1 introduces the formal theory of probability and depicts the relationship between formal and substantive theories of probability. Then, §2 describes the reference class problem in its general form, and draws the distinction between a metaphysical problem and an epistemological problem involved in the reference class problem. There are a variety of substantive theories of probability, but the discussion is selective; of interest here are only theories of the probabilistic apparatus which are relevant for theories of knowledge: frequentism, the propensity theory, and the epistemic probability theory. In §3, I present each theory in turn, and see what form the reference class problem takes for it. §4 concerns the scope of the reference class problem. I examine the relationship between the metaphysical and the epistemological reference class problem at length, and also articulate the three types of cases in which the reference class problem becomes acute.

§1. The Formal Theory of Probability

The distinction between formal and substantive theories within the theory of knowledge is drawn by modeling it on a similar distinction in probability theory. Even though the formal theory of knowledge is not perfectly assimilated to the formal theory of probability, the modeling serves the purposes of elucidating the relationship between formal and substantive theories of knowledge.[2]

Just as in any other mathematical discipline, axiomatization in probability theory is carried out by taking certain terms or concepts as primitive. In probability theory, of course, one such primitive term is 'probable' or 'probability.' Although there are possible to be, and in fact are different formal systems of probability, Kolmogorov's axiomatic system is in the position of orthodoxy, and it is what philosophers usually refer to as 'the formal theory of probability.' In Kolmogorov's system, probabilities are attributed to events, though the term 'event' remains primitive, being only defined in set-theoretic terms; or alternatively, probabilities are attributed to propositions – here, I choose the latter option.

Probability is a function from one or a set of propositions to a real number or an interval in the closed unit interval [0, 1]. The domain of the probability function is called the 'sample space.' A sample space Ω is a non-empty set of propositions A, B, C, … closed under truth-functional operations. Then, Kolmogorov's axioms are formulated as follows:

(Non-Negativity) $P(A) \geq 0$ for all $A \in \Omega$.

---

[2] One prominent difference may be that axioms of a formal theory of probability are amenable to mathematical treatment. It is possible to build such axioms into a formal theory of knowledge so as to allow for mathematical treatment. Those epistemologists who countenance the epistemic probability theory would be willing to do so. Whether such augmentation is permissible in epistemology is controversial. For example, **Nominalization** below requires any logically true proposition to have probability 1, and any logically false proposition probability 0. In order to apply these points to human subjects, logical omniscience must be assumed. Mathematical stringency and clarity is gained only by such idealization.

(Normalization) If T is a tautology, then $P(T) = 1$.

(Finite Additivity) $P(A \vee B) = P(A) + P(B)$ for all $A \in \Omega$ and $B \in \Omega$ such that A and B are logically inconsistent.[3]

Kolmogorov, in his axiomatization, commits himself to a certain conceptual order, and many follow him here as well as elsewhere: unconditional probability is conceptually prior to conditional probability. For a conditional probability is defined as a ratio of two unconditional probabilities:

(Ratio) $P(A \mid B) = \dfrac{P(A \wedge B)}{P(B)}$ (if $P(B) > 0$, otherwise undefined)

The formal theory of probability is, after all, 'formal' and it leaves important, substantive questions regarding the probability function unanswered: among others, the most important ones are (i) what type of entity is the bearer of probability, and (ii) how the value of the probability function for any given argument is to be determined. Answering (i) and (ii) in one way or another amounts to giving a certain interpretation of the primitive term 'probability.' As we will see at length in §3, each substantive theory of probability differs with regard to (i) and (ii).

---

[3] **Finite Additivity** is often extended to countable additivity, so as to be applied to a countably infinite sequence of (pairwise) disjoint sets, each of which is an element of Ω. Many substantive theories of probability fail to satisfy countable additivity.

The distinction between formal and substantive theories within probability theory is analogous to that within the theory of knowledge. First, the task of substantive theory of knowledge corresponds to (i) and (ii), as mentioned in Chapter 1 and seen in more detail in Chapter 3. Second, in the pursuit of (i) and (ii) (and their counterparts), substantive theories in probability theory and the theory of knowledge are subject to similar constraints. Except for trivial constraints, like logical consistency, simplicity, and so on, there are two important constraints, neither of which, however, is absolute but only *prima facie*: (a) there, of course, is the formal constraint imposed by the formal theory: other thing being equal, any substantive theory must conform to it; and (b) there is a constraint imposed by the ordinary concept or conception of probability.[4] A substantive theory must do justice to what we reckon to be instances of probability or knowledge. Satisfying (a) does not make for a theory of *probability* or *knowledge*. Such a theory must reflect how we use the corresponding concept in our practice. We may have no explicit grip on the ordinary concept or conception of probability or knowledge, and thereby satisfying (b) requires analysis of how the concept in question is used and what role it plays in our ordinary practice.[5]

---

[4] What I mean by 'conception' is quite broad: it consists of our theoretical and practical grips on the concept in question. Moreover, it may as well include our intuitions concerning application of the concept. For discussions on the constraint from the ordinary concept of probability, see (Eagle, 2004) and (Hájek, 1996).

[5] Other constrains may be needed. For example, Salmon (1966, pp. 63-4) proposes three criteria for the adequacy of any interpretation of probability, which are, roughly, the following:

One thing to note is that the conceptual constraint is in place even if the ordinary concept of probability is ambiguous. Indeed, it is taken for granted that probabilistic expressions – and their close relatives, modal expressions – have at least two different senses: one is objective, and the other is more subjective or epistemic. Generally, objective theories purport to capture the first sense, and subjective or epistemic theories the second sense. But this does not change the formal constraint that the concept or conception concerning the probability in one sense must be taken account of in theorizing the probability in that sense.[6]

Neither constraint, however, is absolute. First of all, the formal constraint is relative to the choice of a formal system. But which system is to be privileged? The fact that Kolmogorov's system is orthodox does not guarantee that it is theoretically superior, without further arguments. Indeed, there are many unorthodox theories not in accordance

---

(Admissibility) The meaning assigned to the primitive terms transforms the formal axioms, and consequently all the theorems, into true statements. That is, the probability concepts interpreted satisfy the mathematical relations specified in the axiomatic system.

(Ascertainability) Interpretation provides some method by which, in principle at least, we can ascertain values of probabilities.

(Applicability) Interpretation has practical predicative importance. As Bishop Butler's famous aphorism expresses, probability must be the very guide to life.

[6] In the substantive theory of knowledge, some, e.g., Goldman (1988), claim that an epistemic term, 'justification,' is ambiguous between the internalist sense and the externalist sense. But externalism and internalism disagree as to which sense captures the genuine epistemic factor relevant for the strength of epistemic position. A radically meaning-pluralist approach to the epistemic position is proposed in Alston (2005).

with some of Kolmogorov's axioms, and such theories offer complementary axioms.[7] Moreover, the conceptual constraint may be ignored on theoretical grounds. It may turn out in the final analysis that the ordinary concept and conception are itself inconsistent or contain some deficiency. And the formal constraint may be reckoned to outweigh the conceptual constraint, though, in general, it is not easy to determine how much weight each is to carry.[8] At any rate, in what follows, as is standard, Kolmogorov's formal system is assumed and referred to as 'the probability calculus.'

§2. Probability Theory and the Reference Class Problem

As mentioned in the Introduction, 'the reference class problem' refers to a host of problems rather than a single problem. The core of the problem consists in what Hájek (2007b) calls the 'essential relativity of probability assignment.' He characterizes this relativity in terms of conditional probabilities, and formulates the reference class problem in its most general form as follows:

---

[7] See footnote 3. In addition, some theories even fail to satisfy **Finite Additivity** or **Normalization**.

[8] The conceptual constraint may be in conflict with the formal constraint; normal human beings may not follow formal theories of probability, due to their computational and memory limits. Furthermore, there is ample evidence that our probabilistic judgments are not always congenial to Kolmogorov's formal theory. For these reasons, Kahneman & Tversky (1979) differentiate a *normative* theory from a *descriptive* theory of probability and decision-making, and reckon the formal theory to be the former.

> Let X be a proposition. It seems that there is *one unconditionalized* probability of X; but all we find are *many conditional* probabilities of the form $P(X$, given $A)$, $P(X$, given $B)$, $P(X$, given $C)$, etc. that differ from each other. Moreover, we cannot recover $P(X)$ from these conditional probabilities by the law of total probability, since we likewise lack unconditional probabilities for A, B, C, etc. (and in any case A, B, C, etc. need not form a partition). Relativized to the condition A, X has one probability; relativized to the condition B, it has another; and so on. Yet none of the conditions stands out as being the *right* one (ibid., p. 565).

Some gloss is in order here. First, this is merely a general format of the problem, and each substantive theory of probability faces the reference class problem in a particular form. As we will see in the next section, different substantive theories spell out the probability function P differently; most relevantly, they differ with regard to what type of factor conditional probability is dependent on, i.e., what type of factor is plugged into the denominator of conditional probability $P(X \mid Y)$. Thus, substituting a different type of factor for Y, a different form or version of the reference class problem is generated from the general format. The scope of the problem is not restricted to the theories on which the arguments of the probability function are propositions. Hájek picks a proposition as the denotation of 'X' only for the illustrative purpose.

Second, as Hájek implies, even if what type of factor is to be Y is held fixed by a theory of probability, there are many candidates for the value of Y. The probability of the one and the same argument X may radically vary with a choice of Y. The reference class problem is perplexing, for there seems no way to determine the unique Y relevant for $P(X)$.

Third, Hájek, at least at some points, expounds that '*X*' here refers to a single event or a singular proposition, a proposition about a single individual or event. In general, two kinds of probability, *definite probability* and *indefinite probability*, are to be distinguished.[9] Definite probability is that which is attributed to a single event (event token) or a singular proposition, by making use of a singular term, as in 'the probability of *a* being *G* is …' or 'the probability of the *F* being *G* is ...'; indefinite probability is that which is attributed to an attribute (a class or type of event) or a general proposition, typically in 'the probability of a *F* being *G* is …'. For some theories, such as frequentism and the propensity theory, the reference class problem arises in the context of making sense of definite probability. Although Hájek is not explicit about this, for the epistemic theory of probability, the reference problem is generalized to indefinite probability as well.

In what follows, I will use the class-member relationship interchangeably with the type-token relationship. Moreover, individuating a single event or event token in one way or another is to assign it to some class, insofar as the description relevant for the individuation of the event token includes an extensional predicate to which an event class corresponds.

Hájek divides the reference-class problem into two sub-problems: the *metaphysical* and the *epistemological* reference class problem. On the one hand, the metaphysical problem concerns the very nature of probability itself. It is quite natural to assume that there is a fact of the matter about what probability *X* has. But, given the essential relativity

[9] The distinction between the two kinds of probability comes from Pollock (1990). Definite probabilities are also called 'single-case' or 'singular' probabilities; indefinite probabilities are also called 'general' or 'generic' probabilities.

of probability assignment, there is no unconditional probability of *X*, rather many conditional probabilities of it. Then, the problem is which, if any, conditional probability is *the* probability of *X*. On the other hand, the epistemological reference class problem concerns how one can rationally assign a probability to *X*. Probability is, in fact, our guide to life; it, in one way or another, directs and rationally constrains our practice, typically in connection with inductive inference and decision-making. Barring skeptical worries about the rationality of our probabilistic judgments, it seems that our probability assignments are generally rational, and hence we manage to pick up one rational reference class out of many. However, the problem remains as to what grounds or reasons there are to justify our selection of a particular reference class for the probability assignment to *X*.[10]

From the ubiquity of the reference class problem in substantive probability theories, Hájek argues that the reference class problem has an important ramification for Kolmogorov's formal theory, in which unconditional probability is conceptually prior to conditional probability. Hájek uses the metaphysical reference class problem as a *reductio* against this conceptual order; probability is essentially a two-place notion, and so unconditional probability is to be defined in terms of conditional probability, and not *vice*

---

[10] Hájek formulates the epistemological reference class problem in several different ways. Most of them are in terms of the notion of the guide to life, and then, the problem concerns what (conditional) probability among many should guide us; the other is in terms of the rationality of the choice of the right reference class, and then the problem concerns how one can rationally pick out the right reference class problem out of many. Although it is not quite clear how he thinks of the relationship between the two formulations, a hint may be found in his statement: "[i]f probability is to serve as a guide to life, it should in principle be possible to designate one of these conditional probabilities as the right one" (Hájek, 2007b, p. 584).

versa.[11] Thus, he resolves the metaphysical problem rather than solves it by denying the existence of the fact of the matter regarding unconditional probabilities.

This resolution, as Hájek recognizes, only militates against the metaphysical reference class problem, and leaves the epistemological reference class problem untouched. Even if it is granted that there is no fact of the matter with regard to the probability of *X*, one conditional probability may be privileged as being rational to choose in our practice. Unless some justification or reason for our choice of the relevant reference class is given, our practice of using probabilities remain utterly mysterious.[12]

---

[11] This type of analysis, while it remains unorthodox, has some proponents, e.g., Popper (1959a).

[12] Hájek seems to hold that if we find a solution to the metaphysical reference class problem, we will also find a solution to the epistemological problem as well. The idea behind this may be that if we can have a principled way to specify the correct reference class for the probability of *X*, we have reason to justify why it is rational to choose it for the probability assessment of *X*. There is a subtle problem for this idea, since 'we' here is ambiguous; suppose that a theoretician manages to solve the metaphysical reference class problem by offering a complicated, highly technical procedure for determining the correct reference class in any given case. Then, is it rational if one happens to choose the reference class for a probability assignment that the procedure prescribes, while one has no access to it? If the epistemological problem concerns the question of whether there is reason to rationally prefer a reference class over others, the question asks for impersonal justification; but if it concerns the question of whether a subject has reason for rationally preferring a reference class over others, it asks for personal justification. The distinction between personal and impersonal justification may be a distinction without difference, if each rational subject is implicitly sensitive to some general procedure for choosing the right reference class, and such implicit sensitivity is enough for personal justification. Something like this is commonly assumed in the discussions on the epistemological reference class problem, and efforts are made to find such a procedure.

Hájek is not alone in proposing the resolution of the metaphysical problem in the way just described; Gillies (2000) and Howson & Urbach (1993) may join the camp, for they, in the face of the reference class problem, deny that there are objective definite probabilities, all definite probabilities being subjective.

§3. Substantive Theories of Probability

Now, I turn to each individual form of the reference class problem that arises for a substantive theory of probability. Hájek seems to think that each form of the problem is cross-classified with the distinction between the metaphysical and the epistemological problems. For simplicity, then, I will omit that distinction for a while, and come back to it in §4. 1. The substantive theories of probability that are of interest here are frequentism, the propensity theory, and the epistemic probability theory.

§3. 1. Frequentism

The reference class problem is always associated with, and taken to be devastating for frequentism. The core features of frequentism are twofold: (i) the bearer of probability is primarily a class or sequence of events in the objective sense.[13] This makes frequentism

[13] A sequence is the bearer of probability on limiting frequentism. On actual frequentism, the bearer of probability is also identified with an attribute or a event type. With the simple assumption that an attribute, class, or type, of event corresponds to a set of event tokens to which they belong, differences in formulating (i) make no essential difference. In what follows, I shall use ‘attribute,’ ‘class,’ and ‘type’ interchangeably.

a theory of objective indefinite probability; (ii) probability consists in relative frequency. Although frequentism comes in two different versions, actual frequentism and limiting frequentism, I mainly focus on actual frequentism.

On actual frequentism, the probability of an event class F is the actual relative frequency of F within an event class G; more precisely, the actual relative frequency of F in G is the ratio of the cardinality of actual event tokens of class $\{F \wedge G\}$ to the cardinality of actual event tokens of class G. Event classes F and G are usually referred to as the 'target class' or the 'attribute class' and the 'reference class,' respectively.

Both target and reference classes are essential for relative frequency. Suppose that you want to know the frequency of a home run being hit. This sets the target class F as {home runs}. F alone, however, is not enough to estimate the frequency in question. You may be interested in a batter's home run rate, e.g., Alex Rodriguez's, or in the home run rate for a stadium, e.g., the new Yankee Stadium. In the former case, the reference class G is set as {at bats of Alex Rodriguez}, and in the latter case, as {at bats in the new Yankee Stadium}. The frequency of F within G is calculated by dividing the number of every occurrence (tokens) of $\{F \wedge G\}$ by the number of every occurrence (tokens) of $\{G\}$. Actual frequentism restricts the relevant range of occurrences of $\{F \wedge G\}$ and $\{G\}$ to those that have actually happened. On frequentism in general, then, probability is inevitably relativized to a reference class. Different reference classes often, if not always, lead to different probability assignments to one and the same target class.

Actual frequentism faces numerous problems, prominent among which is the problem of the single case. To use a well-known example by Venn (the inventor of Venn diagrams), suppose that we are concerned with the probability that John Smith, a consumptive Englishman aged 50, will live to 61. The event of John's living to 61 is a unique, unrepeatable event, and it either obtains or not. Then, the actual frequency of this single event is either 0 or 1. The problem for actual frequentism, then, is how to assign a non-trivial value to definite probabilities.

A similar problem arises for limiting frequentism, on which the indefinite probability of *F* is the relative frequency that *F* approaches at the limit within the infinite sequence of tokens of *G* (*G* is called the 'reference sequence'). As a matter of fact, some frequentists refuse to apply frequentism to single events or cases. For example, a limiting frequentist, von Mises (1957), remarks that definite probability is not the proper subject matter of probability theory, since, on limiting frequentism, probabilities are primarily indefinite; in a similar vein, Reichenbach (1949), another limiting frequentist, argues that single-case probability statements are meaningless, on accounts of their failure to satisfy the verificationist criterion for meaning – though he allows us to talk in such a way that an analog of probability, what he calls 'weight,' obtains for single cases.

Other frequentists attempt to deal with the problem of the single case by identifying definite probability with indefinite probability: typically, a single event *A* is denoted by a definite description 'the *F*,' when it is a member of event class *F*.[14] Then, the definite probability of *A* is identified with the relative frequency of *F* within some reference class or sequence *G*. Given this, actual frequentism, for definite probabilities, can be formulated as follows:

[14] I take it that this is the reason that there is no controversy on what is to be the relevant target class in assigning a probability to a single event. When a question is raised about the definite probability of the *F*, the target class is naturally set as {*F*}.

(Actual Frequentism) $P(A) = P(A \mid A \in G)$ = x iff the actual frequency of F within reference class $G = x$, where $A \in F$[15]

§3. 2. The Reference Class Problem for Frequentism

Whether the identification of definite probability with indefinite probability works against the problem of the single case or not, the identification leads to the reference class problem. A single event belongs to many (even infinitely many) event classes. What is the relevant reference class G? John Smith, for example, can be attributed to reference classes {consumptive}, {Englishman}, {at the age of 50}, {male}, {in the 19th century}, and so forth – and every union and intersection of these classes. The frequency of *living to 61* varies considerably in value with the choice of the reference class. Thus, on frequentism, the single event of *John Smith's living to 61* has different degrees of probability, depending on how it is classified. Unless there is some principled way to single out a reference class as relevant, there is no determinate probability pertaining to a single event. This is the form the reference class problem takes for actual frequentism.

§3. 3. The Propensity Theory

[15] To follow Hájek, I represent how unconditional probability is relativized to a reference class in the form of a conditional probability, and it must be accompanied by the clause that its denominator has non-zero probability. I omit this complication here, and the same goes for the formulations below.

What is called the 'propensity theory' or the 'propensity interpretation' involves two different views, the long-run propensity theory and the single-case propensity theory, dubbed by Kyburg (1974a).[16] As with frequentism, propensity theories are concerned with objective probability. The long-run and the single case propensity theory are both proposed and developed, at least in part, to overcome the shortcomings of frequentism. Thus, both have the following circumventing features over frequentism: (i) probability distribution is dependent on a chance set-up, i.e., a set of repeatable arrangements of experimental apparatus that generate admissible sequences of events; (ii) it is (purports to be) capable of dealing with definite probability; (iii) probabilities are 'counterfactually robust.'[17]

(i) through (iii) are interrelated. First, Popper (1957, 1959b), the founder of propensity theory, emphasizes the importance of (i) by considering a case of definite probability: take two dice, one 1/4 biased in favor of 6 and the other fair. Then, consider a sequence, consisting almost entirely of throws of the biased die, interspersed with one or two throws of the fair die. What is the probability of the fair die getting a 6 on one of the interspersed throws? On frequentism, it may approximate 1/4 because the throw belongs to the interspersed sequence; but, on an intuitive level, the probability should be 1/6. The problem of frequentism, Popper argues, is that it cannot exclude ontologically heterogeneous reference classes or sequences. In the face of this problem, Popper claims, admissible reference sequences must be "either virtual or actual sequences which are

[16] Eagle (2004) and Gillies (2000) make a similar distinction, and further divides them into sub-categories, though their sub-categories are different. Gillies' book contains a good exposition of the reference class problem for propensity theories.

[17] This phrase is taken from (Eagle, 2004, p. 374)

characterized by a set of generating conditions—by a set of conditions whose repeated realisation produces the elements of the sequences" (Popper, 1959b, p. 34).

A chance set-up is 'a set of generating conditions' as just cited. It consists of not only experimental apparatus but also (part of) the surrounding environment. Popper (1957) illustrates the importance of environmental factors by two interesting examples: first, consider the probability of getting heads on a toss of a coin biased toward heads. If this coin is thrown in a lower gravitational field, say, on the Moon, the bias would have less effect than it does on the earth, and consequently the probability would have a lower value; second, suppose that we throw a fair coin not on a flat surface, but on a surface with many slots cut. The finite sequence produced by the coin, then, would include a class of outcomes, {edges}, in addition to {tails} and {edges}, where {edges} comprise events of the coin sticking to one of the slots. Relative to this sequence, the probability of heads is lower than otherwise. Therefore, the probability of heads depends not only on the experimental apparatus, the coin, but also on the local environment.

Second, probability is closely connected with possibility. A problem of actual frequentism is that it severs this connection. A coin, even if never thrown, clearly has a probability of landing heads; and, another coin, even if thrown only once and landing heads, has some probability of landing heads between 0 and 1. The problem of the single case arises prominently for actual frequentism, since probability, on that account, is too tightly hooked up to actuality. Thus, a way to solve the problem is to associate probability with modal or counterfactual profiles of the chance set-up. That is, (ii) propensity theory purports to evade the problem of the single case by (iii) rendering probability counterfactually robust.

The long-run and the single-case propensity theory differ as to how (iii) is implemented. On the one hand, the long-run propensity theory maintains the frequentist idea that probability is identified with relative frequency, but unlike actual frequentism, the frequency pertinent to probability is hypothetical. The chance set-up of a coin has the disposition or propensity to produce the outcomes of target class {heads} within reference class {tosses} with a certain frequency, even if it has never been thrown, i.e., the disposition has never been manifested (see (Popper, 1957, p. 67)). Probability, then, is identified with hypothetical frequency with which a chance set-up would produce outcomes, were its disposition manifested repeatedly.

The long-run propensity theory requires that a chance set-up be repeatable, and hence that it constitute a class or type. More precisely, a chance set-up class is defined as the union of a class of experimental apparatus and a class of surrounding environment each of which is causally relevant for producing a class of events. The long-run propensity theory is primarily a theory of indefinite probabilities relativized to a chance set-up class. Hence, on this theory, the ultimate bearer of probability is a class of chance set-up. Indefinite probabilities are further relativized to a pair of target class and reference class, but they are, in part, causally constrained by the class of chance set-up, since the candidates for the relevant pair are restricted to those which the class of chance set-up in question is disposed to produce. Just like frequentism, the long-run propensity theory identifies the definite probability of a single event A with the indefinite probability relativized to a set of chance set-up, target, and reference classes:

(Long-Run Propensity) $P(A) = P(A \mid A$ is produced by a token of chance set-up class K and $A \in G) = x$ iff tokens of K has the propensity to produce events of target class F on a long sequence of events of class G with relative frequency x, where $A \in F$

On the other hand, the single-case propensity theory primarily concerns definite probabilities. It identifies the definite probability of A with the strength or degree of the propensity or disposition of a chance set-up token to produce A (and its counterparts). That is, the bearer of probability is a chance set-up token rather than a class. Although there exist several different versions of the single-case propensity theory, a rough characterization will do here.

(Single-Case Propensity) $P(A) = P(A \mid A$ is produced by chance set-up token $K) = x$ iff K has the propensity to produce A to degree x[18]

§3. 4. The Reference Class Problem for the Propensity Theory

Neither version of the propensity theory can circumvent the reference class problem. First of all, the long-run propensity theory immediately inherits the reference class problem

[18] The propensity theory in general has a notorious difficulty with dealing with conditional probabilities. Thus, it is problematic to formulate both theories in terms of conditional probability. As Eagle (2004) submits, this raises a doubt about Hájek's general format of the reference class problem. For lack of better way to state the general format, I ignore this point here.

from frequentism. The reference class *G* needs to be specified in order to estimate hypothetical frequency. But again, candidates for *G* are numerous, and with the choice of *G* may hypothetical frequency vary. Second, the long-run propensity theory also faces another form of the reference class problem. Probability is relativized to a chance set-up class, and there are numerous different classes K_1, K_2, . . . to which a chance set-up token belongs. Hypothetical frequency varies with the choice of the relevant chance set-up class. The problem is again that there is no principled way to specify the relevant class for the chance set-up token in question.

For example, consider again Popper's case of the throw of a biased coin on the moon. The experimental apparatus causing the die to land on 4 on the next toss are of many classes, such as {dice}, {biased dice}, {dice biased in favor of 4}, {dice biased in favor of 4 made in the 20th century}, and so forth. Furthermore, the surrounding environment belongs to many classes, such as {gravitational field}, {low gravitational field}, {universe}, and so forth. The classes of which a chance set-up is a token are unions of each of these experimental apparatus and environment classes. A chance set-up token results in different hypothetical frequencies, depending on what class is assigned to it.

The same type of problem arises for the single-case propensity theory. The strength of the propensity to produce *A* varies with how the chance set-up token is individuated, in much the same way just seen; insofar as individuating a chance set-up token involves describing it with a certain predicate, since the extension of a predicate constitutes a class. Virtually all proponents of the single-case propensity theory are aware of this problem, and design their accounts to solve this variant of the reference class problem. On the account by (Giere, 1973) and (Miller, 1994), propensities at a given time supervene on "the complete situation of the universe (or the light-cone) at the time." (Miller, ibid., p. 185)[19] As they make explicit, this way there is no room for the chance set-up to be further relativized, attaining the maximal specificity for its specification. On the other hand, Fetzer (1982) relativizes propensities at a given time in a world to "a complete set of (nomically and/or causally) relevant conditions … which happens to be instantiated in that world at that time" (p. 195).

Giere and Miller's accounts are subject to many objections; they deprive probability of scientific usage, since the propensity determined by the entire state of the universe at a given time is untestable (Gillies, 2000); they makes it impossible to extrapolate probability of a single case to similar cases (Howson, 1984); they are no-theory theories in that they cannot give a guide to life (Hájek, 2007b) (though as I discuss in §4. 1, I doubt that Hájek's objection is well-founded). Fetzer's account cannot evade the reference class problem. What is a complete set of relevant conditions relative to which propensities are determined? Fetzer gives no satisfactory answer to this question.

§3. 5. The Epistemic Probability Theory

Lastly, consider epistemic probability theories. As the name indicates, these theories belong to both probability theory and epistemology. The core idea of the epistemic probability theory is that probability is a measure of an epistemic relation of one sort or another, as opposed to the objective relation elucidated in frequentism or the propensity

[19] Popper (1983) applies the same idea to his long-run propensity theory.

theory. Thus, a common approach to the strength of epistemic position, adopted by such figures as Fantl & McGrath (2002, 2009b), Hawthorne (2004a), Stanley (2005), and Williamson (2000a), is to identify it with epistemic probability of some sort, though most of them do not detail the version of epistemic probability they countenance.[20] A variety of theories fall under the heading of 'epistemic probability theory,' but of interest here are only three: Bayesianism, the logical theory, and Kyburgian theory.

On the logical theory, probability is a measure of the logical relationship between sentences or propositions, and thereby it is the bearer of probability. More precisely, Carnap (1950, 1962), the founder of the logical theory, reckons probability as the degree of confirmation or support evidence E gives for A, where A and E are sentences in a formal language. The confirmation relation is a generalized form of logical entailment: in the extreme cases where E logically entails A, the degree of confirmation of A conditional on E, $C(A \mid E)$, is 1, and where E logically entails $\sim A$, $C(A \mid E)$ is 0; in other intermediate cases where, to use Carnap's terminology, E partially entails A, $C(A \mid E)$ takes some value in the interval [0, 1], depending on how inductively strong support E gives for A. Probability as the degree of confirmation, thus, is always relative to some evidence. For Carnap, logic is an *a priori* discipline, and this means that values of the confirmation function C are determined *a priori*.

[20] Williamson proposes an idiosyncratic view of epistemic probability. He takes the epistemically interesting kind of probability to be conditional on evidence and calls it the 'evidential probability,' but he does not identify it with actual or counterfactual credence. Rather, the initial evidential probability function "measures something like the intrinsic plausibility of hypotheses prior to investigation" (Williamson, 2000a, p. 211), and updates on new evidence.

Bayesianism takes probability as a measure of rational degrees of belief or degrees of rational credence. Thus, probabilities are assigned to contents of beliefs, i.e., propositions one believes. The bearer of probability is inferential relations between believed propositions. For this reason, the credence function is personal, being always relative to a subject. The Bayesian typically holds that one's rational credence toward a proposition *A*, *Cr*(*A*), is estimated as one's betting quotient on *A*; *Cr*(*A*) is the amount of money one would take to be a fair bet on a ticket that will pay 1.00$ if *A* is true and nothing if *A* is false; likewise for conditional probabilities, *Cr*(*A* | *E*) is the betting quotient one takes to be fair, if *S* knows or rationally believes that *E* is true. Ramsey (1926) develops this type of theory in terms of the notion of utility.

The Bayesian typically imposes two constraints on *Cr*: (a) the probability calculus as a synchronic constraint, and (b) conditionalization as a diachronic constraint. As for (a), unless one's credence is distributed over propositions one believes at time *t* by obeying the probability calculus, it is in principle possible to construct a Dutch book, a series of gambles which inevitably results in one's loss. On pain of irrational betting behaviors, credence must be in accordance with the probability calculus[21]; (b) prescribes that one's credence must be updated by new information by conditionalizing on it, i.e., $Cr_{\mathrm{new}}(A) = Cr_{\mathrm{old}}(A \mid E)$, where *E* is new evidence.

Kyburgian theory shares some features with Bayesianism and the logical theory: on the one hand, like Bayesianism, what Kyburg calls 'epistemic' or 'evidential' probability is at least in part personal, since it is relative to *S*'s rational corpus, i.e., the body of

[21] The Dutchbook argument is not without problems. For attempts to show (a) without relying on the Dutchbook argument, see (Kaplan, 1996) and (Howson & Urbach, 2005).

knowledge or evidence *S* possesses; on the other hand, Kyburg himself takes his theory as a successor of Carnap's logical probability theory, and regards probability as a measure of strength of logical support that holds among beliefs construed as statements in a formal language. However, his theory is idiosyncratic in several respects: first, epistemic probability is interval-valued rather than point-valued, and takes as a value some subset of the unit interval [0, 1]; second, as a result of this, epistemic probability may fail to accord with general principles such as **Ratio** and conditionalization; third, most importantly, the value of epistemic probability is determined by reference to statistical data possessed by *S* *via* direct inference, and thereby it reflects objective, empirical features of the world.

Direct inference is inference from indefinite probability to definite probability. Kyburg's idea is that definite epistemic probability tracks objective indefinite probability known for *S*. While this idea is consistent with any substantive theory of objective indefinite probability, Kyburg countenances actual frequentism as the best candidate for use in direct inference. Suppose that *S* believes of a single object or event *a* that *a* is *F*, and *S*'s body of knowledge Γ contains the statement that *a* is *G*, and one or a set of statistical (actual frequentist) statements Δ concerning the target class *F* and the reference class *G* to which *a* is assigned (all of these are statements in Kyburg's formal language). Δ may contain no point-valued statistical data, and the relevant data may be interval-valued at [*p*, *q*]. When these conditions are met, *Ep*(*Fa*), the epistemic probability of *S*'s belief that *a* is *F*, or to use Kyburg's terms, the degree of *prima facie* support Δ confers on the belief that *a* is F, relative to the body of knowledge Γ, equals to $[p, q]$.[22] Kyburg formulates it as the 'principle of support':

(Support) $Ep(Fa \mid Ga \wedge \mathrm{Freq}(F \mid G \wedge \Gamma) = [p, q])$ for S is $[p, q]$.[23]

Support, as Kyburg articulates it, requires epistemic probability to be dependent on *S*'s totality of knowledge, but it can be formulated in terms of total evidence rather than total knowledge (Kyburg does not explicitly distinguish between them). The notion of total evidence stems from Carnap's (1950, 1962) requirement of total evidence: "In application of inductive logic to a given knowledge situation, the total evidence available must be taken as a basis for determining the degree of confirmation" (Carnap, 1962, p. 211). This requirement is much respected and accommodated across different theories of probability; needless to say in the logical theory, but also in Bayesianism, rational credence for *S* is relative to *S*'s total evidence (thus, $Cr_{\text{old}}(A)$, which conditionalizes on E, is indeed conditional on the total evidence acquired prior to E. Appeal to the total evidence or knowledge in **Support** is required to handle cases where Δ contains more than one statistical data for the same F but based on different Gs. As we will see in §4. 1, the requirement of total evidence is at the heart of the reference class problem for epistemic probability theories in general. In addition, it is controversial what is to count as evidence.

---

22 The notion of *prima facie* support comes from Pollock (1990). Pollock's view itself is influenced by Kyburg's writings.

23 Kyburg uses his own notation, and my formulation of **Support** is a bit different from Kyburg's original, which is found in Kyburg & Teng (2001, p. 211).

Here, I assume that sentence or proposition E is part of S's total evidence or in S's evidence set only if it is known or rationally believed by S.

§3. 6. The Reference Class Problem for the Epistemic Theory of Probability

Now, we turn to the variants of the reference class problem that arise for epistemic theories. First, as for the logical theory, Ayer (1963) argues against Carnap that logic cannot dictate which evidence $C(A)$ is conditional on. Carnap's project was to define C for any A and E. Thus, in Carnap's light, all the sentences of the form '$C(A \mid E) = x$' are on a par, all being *a priori* and necessarily true, while x may take a radically different value for each E. Then, there is no logical ground on which one instance of $C(A \mid E)$, one with E being total evidence available, is preferable over others. Ayer correctly points out that the difficulty here is similar to the reference class problem for frequentism; the logical probability of a sentence may vary with evidence E on which it is conditional, but there are many candidates for E.[24]

Ayer, however, misses the fact that Carnap does not take the total evidence requirement to be a logical principle.[25] As the citation above makes clear, it is a principle

---

[24] Hájek (2003, 2007a, 2007b) further remarks that Carnap's logical probability is relative to a choice of formal language, and a variant of the reference class problem arises on the level of language, given that there can be many different formal language in which probability is elucidated or analyzed. Carnap's project to analyze logical probability syntactically has little proponent in the recent literature, mainly because of Goodman's new riddle of induction, which is discussed in §4. 2.

[25] This point is emphasized against Ayer by Hempel (1960).

about the application of logical probability to human rational behavior, and this is why Bayesianism and Kyburgian theory do justice to it; *S* may have multiple bits of evidence E_1, …, E_n for *A*, which give individually weak but jointly strong support for *A*. In such a case, it is intuitive that *S* is rational to believe *A* only on the ground of all E_1, …, E_n, but not on the ground of any individual *E*. The underlying thought is that ignoring the relevant evidence results in irrational behavior.[26]

Hájek (2003, 2007b) offers an excellent analysis of how the epistemic theory of probability in general flies in the face of the reference class problem, and I will follow his discussions in what follows; though, in the next section, I critically examine Hájek's argument. First, he claims that there must be external constraints on any epistemic probability function other than the probability calculus (and conditionalization); otherwise, epistemic probability remains too subjective in that it licenses probability assignments that do not reflect the way the world is. A merely subjective theory, Hájek claims, is able to evade the reference class problem at the cost of being a 'no-theory theory.'

Different objective constraints on credence are already proposed. To treat them in a unified manner, Hájek adopts the notions of expert function and expert assignment from Gaifman (1988). The general idea is that *S*'s rational credence in a singular proposition *A* conforms to the value assigned to *A* by an expert function *Exp*:

(Expert) $Cr(A \mid Exp(A) = x) = x$

[26] More thorough arguments for the requirement of total evidence are given in Good (1967) and Horwich (1982). Good's argument shows that when further evidence is available, to get it maximizes expected utility.

This is the canonical form of expert function, and a particular expert function needs to be substituted for *Exp*. Different kinds of function can be substituted. First, any expert in the usual sense of the term can be an expert function. In Hájek's example, if *S* regards a local weather forecaster as an expert on meteorological matters, and knows that she assigns .1 to *it is raining tomorrow*, then, for *S*, $Cr(\text{rain} \mid Exp(\text{rain}) = .1) = .1$. Second, Hájek follows van Fraassen (1989) in taking an expert function to be a probabilistic theory on the relevant subject, e.g., quantum mechanics, and use **Expert** to elucidate acceptance of such a theory. A scientific theory may attribute probability .1 to *a radium atom decaying in 1000 years*. If *S* accepts the theory, *S*'s rational credence in the proposition that a radium atom decays in 1000 years is .1. Third, most importantly, the probability function, as it is specified by a certain substantive probability theory, can be an expert function. For example, Carnap (1962), indeed, envisages application of his logical probability theory in such a way that the confirmation function *C* is an expert function:

(Logical Expert) $Cr(A \mid C(A \mid E) = x \wedge E) = x$

Moreover, not only the logical theory but also theories of objective probability can play the role of an expert function. As Hájek points out, this idea is incorporated into the principal principle, proposed in Lewis (1980):

(Principal Principle) $Cr(A \mid Ch(A) = x \wedge AE) = x$

'*Ch*' denotes the chance function, '*A*' a singular proposition, and '*AE*' what Lewis calls 'admissible evidence.'[27] Lewis identifies chance with definite objective probability, and construes it in terms of some sort of the single-case propensity theory.[28] This principle, then, states that *S*'s rational credence about a particular event conforms to the known or rationally believed objective chance of that event, if *S* has no inadmissible evidence. Lewis does not give a precise definition of admissibility, but defends it by appeal to examples: information on the chance set-up, statistical record of the chance set-up, historical information on the past, and information on dependence of chance on history are admissible. These types of information are admissible, since, though they carry evidential weights for *A*, once the chance in question is specified, they are evidentially unloaded, i.e., become evidentially irrelevant. For this reason, having such evidence does not undermine the epistemic status of *S*'s belief that $Ch(A) = x$. On the other hand, information undermining $Ch(A)$ is inadmissible. Note that this definition of the admissibility of evidence is hardly explanatory, since it amounts to a sufficient condition: *E* is admissible if it does not affect the true value of $Ch(A)$. Unless the factors to determine the value of

[27] Lewis regards **Principal Principle** to be a diachronic constraint on the initial credence function at a certain time. However, it is commonly construed as a synchronic constraint on the credence function at any time. I am following this construal. Howson & Urbach (1993) defends **Principal Principle** by a Dutchbook argument (though they no longer appeal to Dutch book anywhere in their (2005)). Similar arguments may well be constructed for other expert probability functions.

[28] For unity of exposition, then, **Principal Principle** should be named **Single-Case Propensity Expert**. On Lewis's account, the single-case propensity for *A*, $Ch(A)$, at time *t*, is determined by a set of true 'history-to-chance conditionals' that describe how $Ch(A)$ at t depends on history up to *t*.

$Ch(A)$ is specified, this definition give no hint as to what evidence is admissible or inadmissible.

Hájek thus enlists many expert functions as objective constraints on rational credence.[29] With the plurality of expert functions, Hájek claims, the reference class problem necessarily arises for the epistemic probability theory in two different ways.

First, rational credence is relativized to an expert function. But expert functions may conflict with one another, assigning different values to A, while they are all rational to choose. As conditional on an expert assignment EXP_1, A takes a certain value x_1; as conditional on an expert assignment EXP_2, A takes a certain value x_2 , and so on.

Second, Hájek points out that when the expert probability function in play is an objective probability function, as in **Principal Principle**, it inherits the reference class problem from theories of objective probability. Hájek does not rehearse this point, but what he is pointing at would be the following: on **Principal Principle**, rational credence is relativized to a single-case chance in *S*'s evidence set. Hitchcock (2001) and Levi (1983) claim that it is in practice very difficult, if not impossible, to apply **Principal Principle** as it stands. The reason is that $Ch(A)$ is an objective probability function about definite probability. Objective definite probabilities are not easily accessible to us (they are even humanly impossible if Giere and Miller's single-case propensity theory is true). Furthermore, I add, insofar as *S* rationally believes that determinism is true, it follows by **Principal Principle** that rational credence about a single event is only 0 or 1. This is the

[29] Hájek also cites van Frrassen's reflection principle (van Fraassen, 1984, 1995). It states that *S*'s future probability assignment constrains *S*'s current probability assignment, i.e., *S*'s future self is an expert for *S*'s current credence assignment: (Reflection) $Cr_t(A \mid Cr_{t+\Delta}(A) = x) = x$

problem of the single-case for the epistemic theory of probability. As in cases of objective probability theory, the solution to this problem may be to identify definite probabilities with indefinite probabilities relativized to some reference class. Then, $Ch(A)$ would become a more complex function. Replacing $Ch(A)$ with each definition of $P(A)$ given above, a specific version of expert probability function is generated. For example, assuming that A is a singular proposition of the form '*a F*s' or '*a* is *F*,' actual frequentism would yield the following expert function:

(Frequentist Expert) $Cr(a \text{ is } F \mid AFreq(F \mid G) = x) = x$, where $a \in F$ and $a \in G$[30]

---

[30] Hacking (1965, p. 193) offers an expert probability function, called the 'frequency principle.' It concerns indefinite rather than definite probabilities. Hacking is a proponent of the long-run propensity theory, and regards probability as hypothetical frequency with which a chance set-up of class K is disposed to produce outcomes of class F within a reference class G (henceforth, $HFreq\ (F \mid G \wedge K)$). Hacking holds that when S knows that $HFreq(F \mid G \wedge K) = x$, and nothing else about G and K,

(Frequency Principle) $Cr(\text{a } F \text{ is } G) \mid HFreq(F \mid G \wedge K) = x \wedge AE) = x$, where $a \in F$, $a \in G$, and the chance set-up producing the single-case event '*a* is *F*' $\in K$

The frequency principle inherits the reference class problem from the long-run propensity theory, when S has in her evidence set multiple data about G or K. The expert function Hitchcock and Levi suggests is obtained from the frequency principle, by rendering it be about definite probabilities:

(Long-Run Propensity Expert) $Cr(a \text{ is } F \mid HFreq(F \mid G \wedge K)) = x) = x$, where $a \in F$ and $a \in G$

Frequency Expert is indeed no more and no less than **Support**. Hitchcock and Levi suggest another expert function on the basis of the long-run propensity theory. It is not clear how the expert function constructed from the long-run propensity theory can circumvent the objection they raise against Lewis's **Principal Principle**; it is in practice difficult to access to the modal profile of a class of chance set-up. As a matter of fact, whereas propensity theorists develop their theories as metaphysical accounts of objective probability, they tend to agree that statistical data on actual relative frequency play an evidential role in determining long-run or single-case propensity.

> … observed frequencies often provide the best evidence we have concerning long-run frequencies and the strengths of probabilistic causes. (Salmon, 2005, p. 151)

> … in the absence of a well-developed theoretical backgrounds, observed relative frequencies may provide the only evidence for propensity statements. (Giere, 1973, p. 478)

Salmon is a proponent of the long-run propensity theory, and Giere the single-case propensity theory. They share the idea of the propensity theory as throughout metaphysical, and it is not expected to play an epistemic role in guiding our lives. If such an epistemic

---

As Hitchcock notes, the reference class problem arises when *S* has in her evidence set multiple data that are of the form '$HFreq(F \mid G \wedge K)$ but about different *G*s and *K*s. Beebee & Papineau (1997), too, discuss the reference class problem in defending a principle very similar to **Long-Run Propensity Expert**.

role is only assigned to observed frequency, **Principal Principle** may be reducible to **Support** or **Frequency Expert**.

Hájek contemplates that the two versions of the reference class problem for the epistemic theory of probability are one at the root. The first problem arises when multiple expert functions assign conflicting values to the same proposition, and the second problem arises when an expert probability function assigns different values to the same proposition relative to different reference classes. The source of both problems is that epistemic probabilities are relativized to an expert assignment.

Hájek treats the reference class problem for the epistemic probability theory as parallel to that for the objective probability theory. However, there is a difference. The general format of the reference class problem says that the problem consists in the essential relativity of definite probabilities to a reference class. On frequentism and the long-run propensity theory, the relativity is inevitable because the bearer of objective probabilities is ultimately classes of event or chance set-up. By contrast, on the epistemic probability theory, the bearer of probabilities is logical or inferential relations among sentences or propositions. Hájek's point is that epistemic probabilities are always relativized to a set of evidence. Thus, the relativity of epistemic probability to a reference class is indirect: it is relativized to a reference class only if it is included in the probabilistic data or evidence pertaining to an expert assignment.

This makes for two more differences between the reference class problem for the epistemic theory and the one for the objective theory. First, unlike in the objective theory, the reference class problem for the epistemic theory is equally applied to indefinite probabilities. Suppose that *S* is familiar with different theories of objective probabilities, and how they differ in indefinite probability assignment for a certain reference class *G*. Then, *S*'s belief that a *F* is *G* may be assigned different values by different expert functions.[31] Second, the reference class problem does not arise for epistemic probabilities across the board. The reference class problem is only generated if *S* has multiple expert assignments in her evidence set; and *S* may have only one, leaving no room for further relativization.

§4. The Scope of the Reference Class Problem

We have seen that each substantive theory of probability faces the reference class in one or more specific forms. In this section, I will consider ramifications of the reference class problem. First, in §4. 1, I argue that Hájek's argument fails to establish the intended conclusion, viz., that there is no fact of the matter about definite probabilities. Second, in §4. 2, I distinguish among the three cases in which the reference class problem manifest itself, and hence, among the three sub-problems involved in the reference class problem.

§4. 1. The Metaphysical-Epistemological Distinction of the Reference Class Problem

[31] Nevertheless, the reference class problem concerning indefinite probabilities is not as serious a problem for the epistemic theory as the reference class problem concerning definite probabilities is. To begin with, the theory is little concerned with indefinite probabilities, for the reason that indefinite probabilities barely rationally guide our lives.

Hájek's argument for the reference class problem can be reconstructed as follows:

(1). A substantive theory of probability *T* is either a no-theory theory or not.

(2). If it is a no-theory theory, then *T* is not a genuine theory of probability.

(3). If it is not a no-theory theory, then, on *T*, unconditional probabilities are of a value only relative to a reference class.

(4). *T* provides no principled way to specify the relevant reference class in any given case.

(5). On *T*, there is no fact of the matter about what reference class is relevant.

(6). (5) is true for any genuine theory of probability.

Conclusion (6) is meant to establish that there is no fact of the matter about definite probabilities, i.e., the metaphysical reference class problem for the substantive theory of probability in general is unsolvable. This argument has two related problems, both being based on a conflation between the metaphysical and the epistemological level.

First, he appeals to (2) in refusing the single-case propensity theory and the subjective theory, though on different grounds. The subjective theory is reckoned a no-theory theory, since subjective probability is not an *effective* guide to life; the single-case propensity theory is not a guide to life *at all*. Granted, Hájek is right about the subjective theory, but it seems that the move he makes against the single-case propensity theory is quite problematic.

When Hájek accuses the single-case propensity theory of being a no-theory theory, one of his targets is Giere's account; as cited above, however, Giere assigns no epistemic role to his single-case propensity theory. Hence, Hájek's ground for (2) is off the mark, and the same may hold for other theories. Given this, Hájek's objection may boil down to a weak one that Giere's account fails to solve the epistemological reference class problem. But again, given that the single-case propensity theory is metaphysical, it is not clear why it is to be expected to solve the epistemological problem. A possible defense of (2), at best, would be that the supervenience of definite probabilities on the entire universe is proposed only for the purpose of evading the problem of single-case, and has an *ad hoc* flavor.

Second, a similar line of defense of epistemic theories is possible. The reference class problem for them is deeply related to the requirement of total evidence, which is built into the expert probability functions as a clause about the admissibility of evidence. Even though a probability expert function is formulated in such a way that involves relativization to a reference class, as in **Frequentist Expert** or **Support**, it is put to work only if there is no inadmissible evidence. Hitchcock, in offering the expert probability function based on the long-run propensity theory, suggests the following sufficient condition for the admissibility of evidence: a piece of information is admissible if it does not undermine $Cr(A)$ by changing the relevant reference class.[32] This is precisely why Kyburg constrains

[32] The function he is concerned with is **Long-Run Propensity Expert**, mentioned in footnote 30. Suppose that *S* knows that a coin *d* is disposed to land heads with frequency 1/2. Then, if *S* has no inadmissible evidence, it follows by **Long-Run Propensity Expert** that *S*'s rational credence in *d lands heads on the next toss* is 1/2. If more information on the constitution of the coin, the exact way in which the coin is tossed, and the past record of the tosses of the coin are taken account of, *S*'s rational credence may result in a different value. These are all about specific features of the chance-set up, which determine what class it is a token of. Note that what evidence counts as admissible or inadmissible differs across expert probability functions; for

Support by the total evidence requirement; *S*'s evidence set may include evidence to the effect that *a* is attributed to different reference classes G_1, …, G_n, and also that different frequency is associated with each *G*. Then, *Ep*(*Fa*) varies with the choice of the relevant reference class. To prohibit such *G*s from affecting *Ep*(*Fa*), application of **Support** needs to be restricted to cases where only admissible evidence counts: a piece of evidence is admissible if it does not undermine *Ep*(*Fa*) by changing the relevant reference class *a* is assigned to.

The need for an admissibility clause amounts to the total evidence requirement. The standard interpretation of the total evidence requirement is in terms of Carnap's notion of inductive relevance. The principle is introduced to ascertain that there is no information with which the value of the confirmation function varies. Thus, *C*(*A*) for *S* does not necessarily require it to be conditional on the totality of evidence available to *S*; it is sufficient for it to be conditional on the portion of the total evidence such that adding further evidence to that portion does not change the value of *C*(*A*); added information is inductively irrelevant. In other words, the total evidence requirement only demands that *C*(*A*) is conditional on all the relevant evidence available to *S*. (When evidence *E* entails *A*, the requirement is trivially satisfied, since the entailment relation is monotonic.)

In general, when two or more expert assignments to a proposition conflict with each other, the total evidence requirement guarantees that there is a unique value of epistemic

---

example, information on the chance set-up is admissible for **Frequentist Expert**, since the chance set-up does not affect frequency as it is defined on actual frequentism.

probability.[33] However, neither the admissibility clause nor the total evidence requirement provides for an explanatory condition for evidence to be admissible or irrelevant.[34] Hitchcock is surely right that the difficulty of the reference class problem for the epistemic theory consists in the perplexity of explaining why some evidence is admissible or

---

[33] A solution to the reference class problem for epistemic theories might be to posit some high-order weighting function. But Hájek argues that such a weighting function, too, must be externally constrained; otherwise, whatever weight *S* prefers might admissibly be put to any expert assignment. This amounts to relativization of expert first-order assignment to a high-order expert assignment. There might be multiple high-order expert functions which weigh first-order expert assignments differently. Then, a variant of the reference class problem is generated at the second-order level. The role of high-order expert function is to be played by the total evidence requirement. For the total evidence requirement is supposed to that which adjudicate conflicts of different expert assignments (regardless of their orders).

[34] Kyburg has struggled with the reference class problem for more than four decades, and developed his solution in a series of writings (Kyburg, 1974b, 1977, 1983, 2003; Kyburg & Teng, 2001). His latest (2001, 2003) position is to impose three conditions for statistical evidence relevant for application of **Support**: roughly, when Δ contains more than one piece of statistical evidence, the relevant statistical evidence is gained by taking any two pieces of evidence out of Δ, and compare them by the following selection procedure: where E_1 and E_2 are to the effect that $Freq(F \mid G) = [p, q]$ and that $Freq(F \mid G') = [p', q']$, respectively,

> Step (1): If E_1 contains a full joint distribution and E_2 a marginal distribution, disregard E_2.
>
> Step (2): If *S* knows or rationally believes that $\{G\}$ is narrower than $\{G'\}$, disregard E_2.
>
> Step (3): If $[p, q] \subseteq [p', q']$, disregard E_2.

By repeating steps (1) through (3) for any remaining pair of statistical evidence, the relevant one is picked out. Kyburg, however, concedes that this procedure does not always work.

irrelevant: first, it must not be, and indeed is not the case that any evidence is irrelevant or admissible; second, it must not be the case that any evidence is relevant or inadmissible; otherwise, the problem of single-case threatens.

These considerations show that it is not easy to make a move from (4) to (5). For the move may be blocked for reasons independent of the reference class problem, to the effect that (5) is false. The requirement of total evidence simply reflects the intuitive norm that epistemic probability is assigned only against the background of the totality of relevant evidence available to *S*. This is tantamount to the thesis that epistemic probability supervenes on the fact about the total evidence. Then, in order to show that there is no fact of the matter about definite epistemic probabilities, it needs to be the case that there are two different probability assignments to one proposition based on the same total evidence. This is, by definition, impossible. The move from (4) to (5) is fatal only when the metaphysical thesis underlying a substantive theory of probability *T* is *ad hoc*, i.e., merely designed for the purpose of blocking the very move (though, even if it is not, the reference class problem may mitigate the force of the metaphysical thesis). In this regard, the epistemic theory is on a better footing than the single-case propensity theory. The epistemic theory, on the other hand, still faces the epistemological reference class problem. The difficulty of explaining why a certain reference class is relevant and others are not is precisely at the heart of the problem.[35]

[35] It is not clear whether Hájek accepts the total evidence requirement. In his (2007b), he notes, in passing, that Carnap's concept of total evidence, such that the total evidence is the strongest proposition available to *S* as evidence, is not well-defined. Hájek's ground is that evidence may be infinitely many, and hence there is no the strongest evidence. This is not enough to show that the relevant pieces of evidence are infinitely many.

Therefore, it requires more efforts to draw out a metaphysical conclusion for the epistemic theory of probability from the reference class problem. Even worse, this puts Hájek in the state of inconsistency. Hájek accepts, in formulating the epistemological reference class problem, that "a rational agent apparently can assign only one (unconditional) probability to *X*" (Hájek, 2007b, p. 565). The epistemological problem, thus, presupposes that there is a uniquely rational probability assignment to *X*, and hence there is a unique reference class relevant for such a probability assignment. The problem consists in our lack of explanation or justification of the choice of the reference class. But then, unless he accepts skepticism about our probability assignments, the very presupposition of the problem seems to be inconsistent with his claim that there is no fact of the matter with regard to epistemic probabilities.

§4. 2. Three Specific Reference Class Problems

Even though (5) is not directly derivable from (4), if (4) holds for a substantive theory of probability, it still shows that the theory is incomplete, in that it fails to deliver a unique value of definite probability, on accounts of the multiplicity of reference classes. The multiplicity, however, manifests itself in different ways from case to case. Thus, it is worthwhile articulating the ways in which multiple reference classes become acute.

One way is similar to Goodman's new riddle of induction. Goodman (1955) formulates his riddle by making use of the predicate 'grue': *x* is grue iff *x* is examined before a given future time *t* and green, or it is not examined before *t* and blue. Suppose, as we do normally, that we have a large number of samples of emeralds observed before *t*, and they are all green. The hypothesis that all emeralds are green, then, is confirmed or justified, by enumerative induction, to a sufficiently high degree *r*. However, whatever evidence for this green-hypothesis is also evidence for the grue-hypothesis that all emeralds are grue, and it is likewise confirmed to the degree *r*. Since the two hypotheses are inconsistent, by the probability calculus, *r* must be less than .5. *r* becomes even lower if more hypotheses are introduced with other *gruesome* predicates. The difficulty to arise from allowing such gruesome predicates in enumerative induction is that the epistemic probability of a hypothesis, supported *via* enumerative induction, may take any arbitrary value.

Gruesome predicates give rise to the same kind of difficulties for direct inference. To point this out, Kyburg (1970b, p. 172) presents the following case: suppose that a fair coin has been tossed a sufficiently large number of times, and the frequency of a toss of this coin getting heads is known to be .5. Now, one wants to know the definite probability *r* of the coin landing heads on the next toss. By direct inference, it seems to follow that *r* is .5, but this presupposes that the relevant reference class for this direct inference is {tosses of the coin}. Suppose that a disjunctive predicate is introduced and the reference class is defined by that predicate as $\{H \vee (\{\text{next toss}\} \wedge T)\}$, where *H* denotes tosses yielding heads, and *T* tosses yielding tails. Relative to this gruesome reference class, the probability of the next toss landing heads is arbitrarily close to 1. Again, the difficulty to arise from allowing gruesome predicates in direct inference is that the epistemic probability of a singular proposition, derived *via* direct inference, may take any arbitrary value.

The exact nature of gruesome predicates is difficult to define. Goodman construes it in terms of projectibility, but this is of no help. For a predicate *G* being projectible with regard to a predicate *F* just means that *G* is admissible in induction leading to a hypothesis of the form 'all *F*s are *G*.'[36] At the very least, it is trivially true that a gruesome predicate contains a disjunction. Nonetheless, this non-trivially entails that what counts as a gruesome predicate is relative to language, since what counts as a disjunction in a language depends on what primitive terms it contains. However, not every disjunctive predicate is gruesome; some are clearly projectible.

Goodman suggests that projectible predicates are entrenched in our language system. His solution to the new riddle of induction, if any, is similar to Hájek's resolution of the metaphysical reference class problem: Goodman does not specify the conditions in metaphysical terms under which indefinite probability assignments involving gruesome predicates are false. Instead, he appeals to a contingent feature of our linguistic usage, in order to explain why our choice of particular predicates in inductive inference is rational or justified.[37] The problem his solution is designed to solve is epistemological rather than metaphysical.

---

[36] To follow (Pollock, 1990), I take 'projectible' as a two-place predicate. In this regard, my usage of 'projectibility' may diverge from Goodman's.

[37] Wittgenstein's (1953) problem of rule-following, as Kripke (1984) construes it, may well be taken as a special case of the new riddle of induction, and hence of the reference class problem. The problem, put in my own terms, is the following: a finite set or sequence of outputs does not determine what function it encodes. For illustrative purposes, Kripke appeals to 'quaddition' (+*): $x +^* y = x + y$, if $x, y < 57$, and = 5, if otherwise. Each of addition and quaddition can be represented as a function:

{Addition} = {<<1, 1>, 2>, <<1, 2>, 3>, …, <<1, 57>, 58>, …}

{Quaddition} = {<<1, 1>, 2>, <<1, 2>, 3>, …, <<1, 57>, 5>, …}

Now, we have a finite sequence of ordered-triples of the form $<<x, y>, z>$, where ‘z’ is the answer to ‘$x + y$’ and $x, y < 57$. What function does this sequence encode, addition or quaddition? Note that a function is defined as a set of ordered n-tuples, and hence the problem is structurally similar to the reference class problem: A sequence of individual applications of ‘+,’ when $x, y < 57$, can be assigned to one of the two different classes, and there is no principled way to exclude one rather than another as irrelevant. The reference class problem may appear as a problem concerning functional identity, and I discuss this aspect of the reference class problem in §3. 1 of Chapter 4 and §2. 2 of Chapter 7.

Kripke, in passing, mentions the analogy between the problem of rule-following and Goodman’s new riddle of induction (ibid., p. 58). It is controversial to what degree these problems are analogous to one another. Allen (1989) regards them to be not only structurally similar but also essentially the same. Hacking (1993) objects to Allen, and examines the alleged analogy. Hacking’s careful examination, however, does not undermine the structural similarity between the two problems.

Stern (1996) points out that Kripke’s skepticism contains two skeptical claims: (a) “there is no fact about whether plus or quus is meant” and (b) “there is no justification for one response rather than another” (p. 153). (a) corresponds to the metaphysical reference class problem, in that answering it requires identifying some metaphysical fact to determine which function is realized; (b) corresponds to the epistemological reference class problem, in that answering it requires explaining why it is rational to choose one function rather than another. Kripke maintains that if one manages to give a satisfactory response to (a), then it is enough to be a satisfactory answer to (b). This line of response is called the ‘straight’ solution to the skepticism about rule-following. However, Kripke disregards any attempt at straight solution as hopeless. Contrasted with the straight solution is the ‘skeptical’ solution, and it is to give an answer to (b) without assuming an answer to (a). Kripke argues that it is because the linguistic community uses ‘+’ as addition that it is rational to choose addition as the correct function. The skeptical solution resembles Hájek’s resolution of the metaphysical reference class problem. Kripke and Hájek both concede that there is no fact of the matter with regard to the relevant reference class. While Hájek has no solution to the epistemological problem, Kripke tries to explain why it is rational to pick one reference class rather than another in non-metaphysical

Non-projectible predicates or reference classes generate a special case of the reference class problem for probability theories, generalized from Goodman's new riddle of induction.[38] With this, it is possible to distinguish three types of cases where the reference class problem manifests itself, and correspondingly, three specific problems involved in the reference class problem:

(1) The problem of conflicting reference class: how is the conflict of reference classes to be adjudicated?

(2) The problem of trivializing reference class: how are trivializing reference classes, the ones which result in probability 0 or 1, to be excluded?

(3) The problem of non-projectible reference class: how are gruesome reference classes to be excluded?[39]

The distinction among the three problems is made across the border between the metaphysical and the epistemological reference class problem, and over substantive theories of probability. The three problems partly overlap: (1) is most general, and solving it must involve the solutions to (2) and (3); (2) is simply the problem of the single case reformulated in terms of the reference class problem. As we have seen, the reference class

---

terms (Unlike in the case of the reference class problem, the subtle issue regarding the distinction between personal and impersonal justification, mentioned in footnote 12, does not arise for Kripke. For the type of justification he has in mind is clearly impersonal).

[38] (Pollock, 1990) contains useful discussions on the problems caused by non-projectible reference classes.

[39] These three problems are implicit in Kyburg's writings, and explicitly distinguished in Thorn (2007).

problem arises for frequentism in the attempts to circumvent (2), and the epistemic theory succumbs to the same problem. Indeed, most alleged solutions to (1) have notorious difficulties with dealing with (2). I will turn to the attempts to solve the reference class problem in Chapter 4. Before proceeding, the next chapter is devoted to an exposition of the reference class problem for the theory of knowledge.

Chapter 3: The Reference Class Problem in the Theory of Knowledge

Introduction

In Chapter 2, I delineated the reference class problem in probability theory, and discussed what form(s) it takes for each substantive theory of probability. In this chapter, I will argue that virtually all substantive theories of knowledge face one or more variants of the reference class problem. It is not a coincidence that the two disciplines equally face the reference class problem; rather, it is rooted in the fact that each theory of knowledge employs a probabilistic apparatus. It is because each substantive theory of knowledge has the theoretical apparatus in common with some substantive theory of probability that the former inevitably faces one or more versions of the reference class problem given risen to for the latter. In other words, a reference class problem in a substantive probability theory carries over to the corresponding substantive theory of knowledge.

In §1, I introduce the reference class problem for the theory of knowledge at the general level. §2 provides a general landscape of the substantive theory of knowledge, with the focus on the distinction between internalism and externalism. Then, in §3, I move on to each substantive theory of knowledge and the version(s) of the reference class problem to which it is susceptible: first, in §3. 1 and 3. 2, I argue that type reliabilism is structurally similar to actual frequentism or its relative, and for that very reason, cannot evade the same type of the reference class problem as arises for actual frequentism. §3. 3 is concerned with token reliabilism, which is divided into sensitivity and safety theory. Sensitivity theory shares the probabilistic apparatus with the single-case propensity theory, and they are prone to the same types of the reference class problem; safety theory is constructed on the model of sensitivity theory, and equally falls prey to the same types of the problem. In §3. 4 and 3. 5, I explain why internalist evidentialism is no less susceptible to the reference class problem than externalist theories are.

§1. The Theory of Knowledge and the Reference Class Problem

Just as in Chapter 2, it is better to introduce a general format of the reference class problem in the theory of knowledge at the outset, and move on to the individual form(s) it takes for each substantive theory of knowledge. By way of coordination, the general format of the problem is formulated in terms of conditional strength of epistemic position, where '*EP*' refers to a function mapping a proposition or belief onto a value within the unit value [0, 1], and a value represents a certain strength of epistemic position – in the case of the theory of knowledge, however, talk of conditional strength is more metaphorical, only indicating relativization to a reference class.[1] To be parallel to Hájek's general format of the reference class problem in probability theory, its relative in the theory of knowledge can be formulated as follows:

> Let p be a proposition or belief. It seems that there is *one unconditionalized* strength of epistemic position with regard to p; but all we find are *many conditional* strengths of epistemic position of the form $EP(p$, given $A)$, $EP(p$, given $B)$, $EP(p$,

[1] A substantive theory of knowledge may or may not accord with the axioms of probability theory. For this reason, $EP(p \mid q)$ only means $EP(p)$ being relative to q.

given *C*), etc. that differ from each other. Moreover, we cannot recover *EP*(*p*) from these conditional strengths of epistemic position by the law of total probability, since we likewise lack unconditional strengths of epistemic position for *A*, *B*, *C*, etc. (and in any case *A*, *B*, *C*, etc. need not form a partition). Relativized to the condition *A*, *p* has one strength; relativized to the condition *B*, it has another; and so on. Yet none of the conditions stands out as being the *right* one. (compare with (Hájek, 2007b, p. 565))

Substantive theories of knowledge specify the function *EP* in different ways, and thereby differ as to what kind(s) of epistemic factor contributes to the determination of the value of *EP*. As a result, they plug different kind(s) of factor into the reference class place *q* of *EP*(*p* | *q*). Whatever theory is adopted, the candidates for *q* in any given case are many, and may result in different value assignments to *p*. Thus, the strength of epistemic position is always relativized to a reference class. If a principled way to specify the relevant reference class in a given case is absent, *EP* has no determinate value in that case.

As with probability theory, a metaphysical and an epistemological reference class problem can be distinguished in the theory of knowledge: the metaphysical problem concerns the very nature of epistemic position. It is natural to assume that there is *a fact of the matter* about what strength of epistemic position *S* has for *p*. But, given the essential relativity of epistemic position assignment, there is no unconditional (unrelative) strength of epistemic position assigned to a pair of *S* and *p*, rather many conditional (relative) strengths of epistemic position. Then, the problem is which, if any, conditional strength of epistemic position is the strength of *S*'s epistemic position with regard to *p*; on the other hand, the epistemological reference class problem concerns how one can rationally assign a certain strength of epistemic position to a pair of *S* and *p*. We overtly or covertly evaluate *S*'s epistemic position in our epistemic practice without much ado, and barring skeptical worries, this practice indicates that in most cases, we manage to pick up the relevant reference class for a given pair of *S* and *p*. The epistemological problem is to show what grounds or reasons there are to justify our selection of a particular reference class for the epistemic position assignment to a pair of *S* and *p*.

The epistemological reference class problem arises equally for the knowledge-first approach and the partisan approach. On the other hand, the metaphysical reference class problem may not be as pressing for the knowledge-first approach. It has an easy solution: *q* may be operationally specified in such a way that it is the reference class relative to which *S*'s epistemic position with regard to *p* is sufficiently high if *S* knows *p*, and not if *S* does not. This type of maneuver, however, is not available to the partisan approach, on pain of its reductionist tenet that the determinants of the strength of epistemic position must be more fundamental than knowledge. A reference class is one of the determinants of the strength of epistemic position. So, if the relevant reference class is specifiable, it must be so without relying on knowledge.

As I pointed out against Hájek's argument in §4. 1 of Chapter 2, there is no easy route to the denial of the metaphysical tenet of a substantive theory of knowledge from the metaphysical reference class problem, if it is independently motivated. In other words, it is difficult to undermine the interpretation of **Supervenience** given by a certain substantive theory of knowledge, merely on the basis of the metaphysical reference class problem.[2] The implications of the reference class problem for **Supervenience** is discussed at length in §3. 2 of Chapter 7.

§2. Internalism and Externalism

Historically speaking, most epistemological endeavors have been centered on a particular kind of epistemic factor, epistemic justification. The classical analysis of knowledge, dating back to Plato, identifies knowledge with justified true belief. Then, the classical analysis, in our terms, specifies justification as the only relevant epistemic factor. Although justification is still given a pride of place in contemporary epistemology, Gettier-cases have revealed that justification cannot be the only epistemic factor to determine the overall strength of epistemic position. Thus, the prevalent theories *post* Gettier are pluralist, entailing that knowledge is a function of two or more epistemic factors, justification and some anti-Gettiering externalist condition. Some, such as Armstrong (1973), Dretske (1981b), Nozick (1981), and Plantinga (1993a, 1993b), may endorse a monist position, since they go so far as denying that justification is a component of knowledge.

---

[2] This may be close to the point of Schmitt's (1992) defense of type reliabilism against the reference class problem. Schmitt seems to hold that the existence of cases where only type reliabilism can explain our intuitions about the truth or falsity of knowledge attributions is enough to prove that type reliabilism is on the right track. Kappel (2006) argues that Conee & Feldman (1998) wrongly interpret Schmitt as responding directly to the reference class problem.

Some theories discussed below are indeed primarily concerned with justification. A theory of justification can take either of the two forms, internalist or externalist. The internalism/externalism distinction, despite being commonplace in the substantive theory of knowledge, is difficult to draw. There is no single set of properties such that, for internalism and externalism, it is comprehensive of one, and exclusive of the other. Suffice it to say that internalism is, at a first approximation, associated with some, if not all, of the following threads, and externalism is the denial of internalism:

(Mentalism) The justificatory factor(s) relevant for *S*'s justification for *p* is in *S*'s mind in some relevant sense.[3]

(Accessibility) The justificatory factor(s) for *S*'s justification for *p* is accessible to *S*, by reflection alone.[4]

(Propositional Justification) Justification consists in inferential relations between propositions.[5]

---

[3] The term 'mentalism' stems from Feldman & Conee (2001), and they defend their version of internalist evidentialism while partially refusing **Accessibility**. I will discuss their theory in §3. 4. Feldman & Conee also attribute **Mentalism** to Alston (1988), whose 'internalist externalism' only partly endorses **Accessibility**. **Accessibility** entails **Mentalism**, but the converse does not hold.

[4] Proponents of **Accessibility** include BonJour (1985, 1997), Chisholm (1977), Fumerton (1995), and Sellars (1956).

[5] The distinction between propositional and doxastic justification is first drawn by Firth (1978). Proponents of **Propositional Justification** are too many to cite; all traditional coherentists endorse it, and even traditional foundationalists concede that it only fails for foundational mental states.

(Meta-Justification) To be inferentially justified in believing *p* on the basis of another proposition *e*, *S* must be (1) justified in believing *e* and (2) justified in believing that *e* makes probable *p*.[6]

(Epistemic Probability) The degree of justification is measured in terms of epistemic probability of some sort.[7]

Mentalism and **Accessibility** are not essential for internalism, if they are strictly construed as concerning all the relevant epistemic factors. Some internalists may accept the relevance of defeaters one does not possess for justification, and partially deny **Mentalism** (see (Harman, 1980)); on the other hand, if **Mentalism** is construed moderately as only requiring that some justificatory factor is mental, it is consistent with externalism. For even externalists grant that some psychological factor is responsible for justification or being in a good epistemic position; **Accessibility** is too strong a requirement. For example, the coherentist, who holds that justification of a belief is constituted by coherence of the belief system containing it, typically endorses **Accessibility**. But then, **Accessibility** requires *S* to have access to every belief comprising *S*'s belief system. Harman (1986) argues that such access is psychologically unrealistic, and even humanly impossible. Moreover, if **Accessibility** is combined with **Meta-Justification**, the coherentist inevitably faces a dilemma: **Meta-justification** requires that *S* be justified in believing that coherence holds

[6] **Meta-Justification** is taken from Fumerton (1995, p. 36), where he calls it the 'principle of inferential justification.'

[7] Fumerton (1995, 2004) and Swinburne (2001) take **Epistemic Probability** as the central commitment of internalism.

between *p* and every other proposition in *S*'s belief system. Williams (1991, pp. 292-9) points out that, justification for this belief being absent, coherentism collapses into a version of foundationalism, since it is committed to the existence of the epistemically privileged foundation of the belief about coherence.

This problem with **Meta-Justification** is not only for coherentism. Fumerton (1995) argues that **Meta-Justification**, combined with **Propositional Justification**, gives rise to a Carrollian regress in the following way: suppose that the inferential relation *e* bears to *p* is relevant for *S*'s justification for *p*. Now, this inferential relation is a justificatory factor, and hence it needs to be justified. Then, the inferential relation holding between *e and p being inferentially connected* and *e*, on the one hand, and *p*, on the other hand, is another relevant justificatory factor. But this inferential relation, in turn, needs to be justified, and so on, *ad infinitum*.[8] Many internalists choose to give up the full force of **Meta-Justification**, and concede that only the second order justification is required for the first-order justification. But this means that **Propositional Justification** must be partially denied, i.e., there must be high-order justification that is non-inferentially justified. **Accessibility**, **Propositional Justification**, and **Meta-Justification**, thus, cannot be endorsed in their full-force. But partial endorsements of these tenets tend to be associated with internalism.

The divide between internalism and externalism is drawn with **Propositional Justification**, as follows: on the classical analysis, it is necessary for *S*'s knowing *p* that *S*

---

[8] The same problem is raised by Bonjour (1997, pp. 1-6). Note that the Carrollian regress is to be distinguished from the Pyrrhonian regress which arises in the liner model of inferential justification. Fumerton calls them the 'conceptual regress' and the 'epistemic regress,' respectively.

is justified in believing *p*. Being justified in believing *p* does not require believing *p*, let alone justifiably believing *p*. This is the rationale for distinguishing between propositional justification and doxastic justification. It is common among internalists to reckon doxastic justification to consist of two components: (i) propositional justification, and (ii) the basing relation or well-foundedness, i.e., *S* believes *p* on the basis of the epistemic factor(s) relevant for propositional justification for *p*.[9] By contrast, most externalist theories, at the outset, conceptualize justification as involving doxastic justification.[10] For, as we will see in the subsequent sections, the focus of the externalist theories is the process or method by which *S* forms the belief that *p*.

Meta-Justification is akin to the KK principle, though it is a thesis about justification. The general idea behind it is that *S* needs to have justification for how the relevant justificatory factor(s) regarding *p* contributes to *S*'s justification for *p*. The justificatory factor is an epistemic factor, and thereby instantiating it raises the probability

---

[9] I am taking the basing relation to be causal. Some deny this, and claim that doxastic justification consists in meta-recognition of the justificatory force of epistemic factor. On this view, the requirement for doxastic justification amounts to something akin to **Meta-Justification**. For more on the distinction between propositional and doxastic justification, see (Korcz, 1997) and (Turri, 2010).

[10] Comesaña (2006) points out that it is not necessary for internalism to be a theory of propositional justification, and that neither is it necessary for externalism to be a theory of doxastic justification. To use his example, an internalist theory on which *S* justifiably believes *p* iff *S* is completely convinced that p is about doxastic justification; an externalist theory on which *S* is justified in believing *p* iff *p* is objectively probable is about propositional justification. Turri (2010), following Goldman (1986), offers an externalist theory of propositional justification.

of *p*. Then, the contribution it makes to justification for *p* is elevation of the probability of *p*.

But then, what type of probability is it? **Epistemic Probability** is relevant here. By far the majority of internalists hold that it is not objective probability of any sort; it must be epistemic probability of some sort.[11] [12] A rationale for this specification of the relevant kind of probability is the problem of Cartesian skepticism to which any internalist theory is notoriously susceptible. In the face of Cartesian skepticism, it needs to be shown that *S* is justified in believing empirical propositions, without appealing to any alleged piece of empirical knowledge. Objective probability is only knowable *a posteriori* if it is knowable at all. Then, it is required that the relevant kind of probability be epistemic, or at the very least, be not objective. Although this argument is questionable, I postpone the discussion of Cartesian skepticism until §2 of Chapter 6.

§3. Substantive Theories of Knowledge

In what follows, I will discuss three theories – type reliabilism, token reliabilism, and internalist evidentialism – and argue that each of them is vulnerable to one or more variants of the reference class problem. Any version of internalist evidentialism is more or less associated with all the five tenets given above. By contrast, type and token reliabilism

---

[11] One exception is BonJour (1997, pp. 206-16), who suggests that the probability he countenances is a kind of hypothetical frequency. In the meantime, he still maintains that it is *a priori*.

[12] Plantinga (1993a, 1993b) and Alston (2005, ch. 5), in proposing their externalist theories, explicitly deny **Epistemic Probability**.

are both externalist; they deny **Mentalism** and **Accessibility** to a great extent, and **Propositional Justification** has little importance for them. Type reliabilism is a theory of doxastic justification, whereas token reliabilism is designed to be a theory of anti-Gettiering factor to be distinguished from justification. Nevertheless, token reliabilism maintains that the primary target of analysis is the way in which *S* actually believes *p*. What the two theories share is the denial of **Meta-Justification** and **Epistemic Probability**. Indeed, **Epistemic Probability** is most important for our purposes here. For what type of probability theory a theory of knowledge is modeled on is crucial for how the strength of epistemic position is determined.

The reference class problem in the theory of knowledge, until recently, has been taken to be limited to type reliabilism. The founder of type reliabilism, Goldman (1979), formulates the problem and notes the difficulty of dealing with it. In the subsequent literature, it is often recast as a serious challenge to type reliabilism (see (Conee & Feldman, 1998), (Feldman, 1985), and (Pollock, 1984)). Although the name most commonly associated with the problem is 'generality problem,' I will call it the 'reference class problem' throughout. For the generality problem for type reliabilism and its variants for other substantive theories of knowledge are all isomorphic to versions of the reference class problem in probability theory. This is because each substantive theory of knowledge is to be understood on the model of one or another substantive theory of probability.

§3. 1. Type Reliabilism

What I call 'type reliabilism' is championed by many epistemologists, such as Alston (1989, 2005), Goldman (1979, 1986), Greco (1999, 2000, 2010), Plantinga (1993a, 1993b), and Sosa (1991a, 2007).[13] [14] Despite differences in details, they all share the following idea: necessarily, for all *S* and *p*, *S* is justified in believing *p* iff the belief *p* is formed by a reliable process, disposition, faculty, ability, or competence.[15] This definition entails that justification is attributed to a belief rather than a proposition, and hence type reliabilism is primarily a theory of doxastic justification.

Goldman explains the reliability of a process as follows:

> I have characterized justification-conferring processes as ones that have a 'tendency' to produce beliefs that are true rather than false. The term 'tendency' could refer either *actual* long-run frequency, or to a 'propensity', i.e., outcomes that would occur in merely *possible* realizations of the process. Which of these is intended? Unfortunately, I think our ordinary conception of justifiedness is vague

---

[13] I am following Comesaña (2006) in drawing a line between two kinds of reliabilism in terms of the *type/token* distinction. Type and token reliabilism are more commonly referred to as 'global' and 'local' reliabilism, respectively, and these terms stem from McGinn (1984).

[14] Sosa distinguishes between two kinds of knowledge, animal and reflective knowledge, and correspondingly between two kinds of justification, aptness and meta-justification. His type reliabilism is a theory of aptness.

[15] Strictly speaking, reliability is only a necessary condition for justification, and some anti-defeater condition is supplied so as to make it a necessary and sufficient condition. In addition, the degree of justification can be raised or lowered with the reliability of the process that sustains the belief p. I here ignore these complications, and treat the reliability of belief-forming process as if it is the only relevant epistemic factor for justification.

> on this dimension too. For the most part, we simply assume that the 'observed' frequency of truth versus error would be approximately replicated in the actual long run, and also in relevant counterfactual situations, i.e., ones that are highly 'realistic', or conform closely to the circumstances of the actual world. (Goldman, 1979, p. 114; see also (Goldman, 1986, p. 49))

Goldman, thus, permits two models for understanding the reliability of a process. One model is actual frequentist. As we have seen in §3. 2 of Chapter 2, actual frequency of a single event (token) is relativized to a reference class (type). In the exact same way, the reliability of a process token is relativized to a process type: let F be process type {processes leading to true beliefs}. Then, on the actual frequentist model, for all token process g, the reliability of g is the actual frequency of tokens of F within some process type G. That is, it is the ratio of the cardinality of actual process tokens of type $\{F \wedge G\}$ to the cardinality of actual process tokens of type G.

Goldman suggests that the other model is modal, but it is equivocal: modal type reliability may well be assimilated either to long-term propensity or to single-case propensity. However, the most charitable interpretation of the passage cited suggests a simple variant of the actual frequentist model. Let near-actual tokens be those which occur in possible worlds near the actual world. On the 'near-actual' frequentist model Goldman seems to imply, for all process token g, the reliability of g is the near-actual hypothetical frequency of tokens of F within some process type G. That is, it is the ratio of the cardinality of near-actual process tokens of type $\{F \wedge G\}$ to the cardinality of near-actual process tokens of type *G*. Many type reliabilists seem to endorse the near-actual frequency model in defining their notion of type reliability.[16]

Type reliabilism, thus understood, needs to be reformulated in terms of the type/token distinction: necessarily, for all *S* and *p*, *S* is justified in believing *p* iff *S*'s belief *p* is produced by a token exercise of a reliable process, faculty, disposition, or competence type. Type reliabilism shares with frequentism the probabilistic apparatus to determine the degree of *S*'s justification for *p*, and, in turn, the overall strength of *S*'s epistemic position regarding *p*. Let EJ_e be a function for the degree of (externalist) epistemic justification,

---

[16] Virtually all type reliabilists refuse the actual frequentist model, since (i) a process type may have never been actualized, and (ii) actual reliability may be merely accidental. However, it is not clear what alternative model they have. At best, they state the type reliabilist condition in dispositional terms and recognize the need to modalize type reliability. For example, Greco (2000) claims:

> Abilities in general are stable and successful dispositions to achieve certain results under certain conditions. But abilities cannot be defined in actual conditions only. Rather, when we say that someone has an ability we mean that she would be likely to achieve the relevant results in a variety of conditions similar to those that actually obtain. In the language of possible worlds, someone has an ability to achieve some result under relevant conditions only if the person is very likely to achieve that result across close possible worlds. (p. 207)

This is as equivocal as Goldman's explanation. But again, the most charitable interpretation seems to be in terms of near-actual frequency (Greco offers this as a modification of the actual frequency model). Type reliability construed on the near-actual frequency model is similar to token reliability (it is consistent with **Disposition'**, discussed in §3. 3. 1). But it is still frequentist, in that what matters is the numbers, not the closeness, of the relevant possible worlds.

mapping a belief onto a value within the unit interval [0, 1]. This gives the type reliabilist interpretation of **Epistemic Position**: the strength of epistemic position regarding p, $EP(p)$, is a function of $EJ_e(p)$, though EJ_e differs between the actual frequentist model and the near-actual frequentist model. The difference between them can be stated as follows:

(Actual Type Reliabilism) $EJ_e(p) = EJ_e(p \mid g \in G) = x$ iff the actual frequency of true beliefs within beliefs produced by process type $G = x$

(Near-Actual Type Reliabilism) $EJ_e(p) = EJ_e(p \mid g \in G\ ) = x$ iff the near-actual frequency of true beliefs within beliefs produced by process type $G = x$

.

Both definitions are only concerned with degrees of justification. To bring about a theory of justification, they need to be combined with a definition of justification *simpliciter*. Such a definition is typically given by imposing a threshold for justification *simpliciter*:

(Threshold for Externalist Justification) Necessarily, for all S and p, $EJ_e(p)$ for $S \geq 1 - \varepsilon$ iff S's belief p is doxastically justified.

'ε' takes some value within the unit interval. The smaller ε is, the more stringent the epistemic standard for justification is. Reliabilists do not usually specify the exact value of ε, and only suggest that it is less than .5.

§3. 2. The Reference Class Problem for Type Reliabilism

Type reliabilism inherits not only the probabilistic apparatus but also the reference class problem from frequentism. There are infinitely many process types *G*s with varying reliability, and no one stands out as correct for the reliability of *g*. Conee & Feldman (1998) give the following example to illustrate the reference class problem for type reliabilism. Suppose that Smith looks out a window, sees a maple tree, and believes that there is a maple tree nearby. The process causally responsible for Smith's belief is a token of the following process types:

(1) process of a retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby
(2) process of relying on a leaf shape to form a tree-classifying judgment
(3) the visual process
(4) vision in bright sunlight
(5) perceptual process that occurs in middle-aged men on Wednesdays
(6) process which results in justified beliefs
(7) perceptual process of classifying by species a tree behind a solid obstruction

These process types all differ in reliability, and thereby no unique degree of justification can be associated with Smith's belief that there is a maple tree nearby, unless there is a principled way to determine which type is relevant. The reference class problem for type reliabilism takes two somewhat different forms, depending on which theory of probability, actual frequentism or near-actual frequentism, it is modeled on.[17]

In addition, Cruz & Pollock, (2004) and Wallis (1994) point out that there is another dimension along which multiplicity of the problem arises. Their point is anticipated by Popper; remember that he observed of a chance set-up that it may consist of two kinds of factors, the one concerning experimental apparatus and the one concerning environment of experiment. Similarly, a process type may consist of two kinds of factors, the one concerning psychological process and the one concerning external environment. Indeed, these two kinds of factors are not distinguished in (1) through (7); (1) to (3) are purely psychological, whereas the rest involves environmental factors. Corresponding to the two kinds of factors, two versions of the reference class problem for type reliabilism may be differentiated: (i) what psychological process type is relevant in any given case, and (ii) what environment type is relevant in any given case. Let us call the former the 'process type problem,' and the latter the 'environment type problem.'[18]

[17] Sosa (1985) cites the generality problem as one of the pressing problems for his type reliabilism. Similarly, Greco (2009, 2010) accepts the threat of the generality problem raised against his view by Shope (2008). Although Plantinga (1993b) accuses theories, such as epistemic probability theories, Dretske's (1981b) information theory, and Goldman's type reliabilism, of not being able to circumvent variants of the reference class problem, Cruz & Pollock (2004), Feldman (1993), and Sosa (1993) point out that Plantinga's own type reliabilist theory, proper functionalism, falls prey to the generality problem. William Alston, as discussed in §3 of Chapter 4, has long struggled with the generality problem.

[18] Wallis calls them the 'characterization problem' and the 'relevance class problem,' respectively. As he notes, these two problems have been lumped together under the name 'generality problem.' For example, Pollock's (1984) generality problem is primarily the environment type problem.

The process type problem and the environment type problem are not entirely separate. For, as Heller (1995b, p. 504) remarks, "[t]he more we say about the process, the less we need to say about the environment." The process type may be so narrowly specified as to leave no further room for environment type specification, and *vice versa*. But this is not always the case.

Type reliabilism, once modeled on near-actual frequentism, may appear to be capable of dealing with the environment type problem: the relevant type of environment to which reliability is relativized is that which includes only the environments of nearby possible worlds. This solution is at best incomplete. How is the closeness of possible worlds measured? Token reliabilism, as I argue in the next section, inevitably succumbs to this problem, and appeal to the intuitive measure of closeness is far from decisive. Moreover, if what matters is nearby worlds rather than nearest worlds, what range of distance does 'nearby' cover? Unless these questions are settled, the environment type problem looms.

§3. 3. Token Reliabilism

The reference class problem in the theory of knowledge is often reckoned to be a problem only for type reliabilism, just like the reference class problem in probability theory is often reckoned to be a problem only for actual or limiting frequentism. The reason for this myopic partiality is easy to see; on both theories, the ultimate determinant or bearer of the property in question, epistemic position or justification, on the one hand, or probability, on the other hand, is a type rather than a token. Thus, given the multiplicity of types a token can instantiate, the reference class problem manifests itself in these theories in an obvious way. On token reliabilism, on the other hand, the determinant of epistemic position is rather a token: the strength of *S*'s epistemic position regarding *p* is determined by the counterfactual profile of a token exercise of process, faculty, or method. In this regard, token reliabilism is akin to the single-case propensity theory. Indeed, they share the same type of probabilistic apparatus. Thus, notwithstanding the focus on tokens, token reliabilism inevitably runs afoul of a variant of the reference class problem, in a similar way that the single-case propensity theory does.

Token reliabilism comes in two different versions, sensitivity and safety theory. First, I delineate sensitivity theory and its own reference class problem(s), and then, show that safety theory is vulnerable to the reference class problem in much the same way.

§3. 3. 1. Sensitivity Theory

Sensitivity theory is proposed by Nozick (1981), and subsequently borne out by many, such as Becker (2009), Black (2008a), Black & Murphy (2007), Cross (2007), Goldman (1986), and Heller (1999).[19] [20] Their theories differ in details, e.g., Goldman

[19] The term 'sensitivity' is coined by DeRose (1995); Nozick calls the condition for knowledge captured by sensitivity the 'variation condition.' On the ground that the variation condition alone is vulnerable to certain counterexamples, he also requires what he calls the 'adherence condition' for knowledge: if *p* were true, *S* would believe *p*. However, Shatz (1987) argues that the variation condition involves a resource to handle such counterexamples: to individuate *S*'s method narrowly. As I argue below, assuming a certain way of

holds that token reliability, as captured in sensitivity, is an extra condition for knowledge in addition to his type reliabilist justification condition, whereas Nozick does not endorse justification as a necessary condition for knowledge. The idea behind Nozick's sensitivity theory is 'truth tracking': knowing *p* requires exercising a discriminatory ability or method to discriminate between *p* and non-*p*. Such a discriminatory ability or method involves two distinct elements: to believe *p* when *p* is true, and not to believe *p* when *p* is false, over a range of circumstances. Both are intimately related to the dual goals of cognition, truth attainment and error avoidance. Sensitivity theory is particularly focused on the latter.

Nozick, for a first try, fleshes out the sensitivity condition for knowledge in the following way: necessarily, for all *S* and *p*, if *S* knows *p*, then if *p* were false, *S* would not believe *p*. On the Lewis-Kratzer's ordering semantics for counterfactuals, a counterfactual $p \mathrel{\Box\!\!\rightarrow} q$ is true at a world of evaluation *w* iff *q* is true in all the closest possible worlds to *w* in which *p* is true.[21] For the truth condition for a counterfactual to be determinate, two

---

method individuation flies in the face of the reference class problem. Nozick further adds an anti-defeater condition for knowledge, but again, I ignore this complication here.

[20] Dretske (1970, 1971, 1981a) and Goldman (1976) are predecessors of sensitivity theory, each putting forth a comparable counterfactual account of knowledge. Goldman, in his (1983) review of Nozick (1981), notes the similarity between his account and sensitivity theory. I will mention Goldman's version of sensitivity theory in due course.

[21] Lewis (1973) offers a more complex truth condition, in order to accommodate cases in which possible worlds are infinite and the limit assumption, that worlds are not infinitely closer to *w*, does not hold. Stalnaker's (1968) semantics for counterfactuals involves both the limit assumption and the uniqueness assumption that there is the closest possible world to *w*. If any counterfactual account of knowledge is on the right track, the semantics for such an account must deny the uniqueness assumption, if it is possible for *S* to

parameters, what are called a 'modal base' and an 'ordering source,' must be set: a modal base is a domain of possible worlds over which modals quantify over; and an ordering source is a comparative similarity measure of how similar and close worlds in the modal base are to *w*. Nozick, however, does not talk about these parameters. For the time being, I follow him in ignoring complications arising from modal base and ordering source. Then, Lewis-Kratzer's semantics entails that *S*, at the actual world, satisfies the sensitivity condition iff *S* does not believe *p* in all non-*p* worlds closest to the actual world.

Nozick argues that sensitivity in this simple form is subject to a certain counterexample:

> (Grandmother) A grandmother sees her grandson is well when he comes to visit; but if he were sick or dead, others would tell her he was well to spare her upset. Yet this does not mean she doesn't know he is well (or at least ambulatory) when she sees him. (Nozick, 1981, p. 179)

Then, Nozick introduces the notion of *knowing by a method* to deal with this and related cases; grandmother's vision is an efficient method to discriminate grandson's being well from his being not, whereas testimony is not. Sensitivity, thus, needs to incorporate the method *S* actually uses to form the belief *p*.

---

be in the equal epistemic position to two or more propositions. For, as I argue below, the closeness of worlds determines how strong *S*'s epistemic position is. For more on ordering semantics, see (Portner, 2009).

As a matter of fact, however, it is controversial how sensitivity thus methodized is to be formulated. Both Williamson (2000a, ch. 7) and Cross (2007) claim that at least two ways are possible:

(Sensitivity') for all *S* and *p*, if *S* knows *p via* method token *M*, then if *p* were false and *S* were to use *M* to arrive at a belief regarding *p*, *S* would not believe *p via M*.

(Sensitivity*) for all *S* and *p*, if *S* knows *p via* method token *M*, then if *p* were false, *S* would not believe *p via M*.

Cross, in passing, notes that **Sensitivity'** is correct and what is intended by Nozick[22]; whereas Williamson prefers **Sensitivity***, since **Sensitivity'** faces counterexamples of a certain type, which Goldman offers in his (1976) and cites in his (1983) review of Nozick (1981) with a minor modification:

(Dachshund) Oscar sees Dack the dachshund and believes there is a dog before him. If there weren't a dog before him, a hyena would be there instead, which Oscar would misclassify as a dog. (Goldman, 1983, p. 84)

Williamson reckons Oscar's actual method token *M* to be *judging by vision*. This specification of *M* is based on Nozick's internal criterion for method individuation: method tokens M_1 and M_2 are identical iff they are the same 'from inside,' i.e., iff *S* cannot

[22] For the exegetical issue, Cross is right; Nozick formulates sensitivity methodized in terms of **Sensitivity'** in his (1981, p. 179).

recognize differences between them.[23] **Sensitivity'** is false for the Dachshund case with *judging by vision* in place of *M*. But how does the Dachshund case speak in favor of **Sensitivity***? The difference between the two formulations is that **Sensitivity***, not **Sensitivity'**, is satisfied if the actual method *S* uses to believe *p* is not the method *S* uses in the relevant counterfactual worlds. Williamson attempts to externalize methods elsewhere (his 2000a, ch. 4), and also claims that in order to avoid implausible consequences of **Sensitivity***, "one must individuate methods externally rather than internally" (ibid., p. 156). By the external criterion for method individuation, he contends, Oscar's actual method token M_1 is *judging on the basis of perceptual experience as of Dack*, and clearly this is not the method Oscar uses in the counterfactual worlds.[24] Instead, method token M_2 used therein would be *judging on the basis of perceptual experience as of a hyena*. Thus, **Sensitivity*** is true for the Dachshund case with M_1 in place of *M*.

Williamson's argument is problematic. First, as I discussed above, the motivating idea behind sensitivity is the discriminatory power of ability or method. The mere fact that Oscar would use an alternative method if things were different has nothing to do with the discriminatory power of the method he actually uses to form the belief that there is a dog

---

[23] "The method used must be specified as having a certain generality if it is to play the appropriate role in subjunctives. This generality is set by the differences the person would notice; the methods are individuated from the inside." (Nozick, 1981, p. 233)

[24] Once the external criterion is in order, Oscar's actual method token can be specified differently, e.g., as vision, animal discrimination, telling kinds of things by looks, judging on the basis of perceptual experience, judging on the basis of perceptual experience as of Dack, etc. As argued below, to pick up the relevant one out of these is part of the reference class problem for the sensitivity theory.

before him.[25] Second, it is odd to conclude that **Sensitivity*** is better than **Sensitivity'** from the Dachshund case; the problem is in the internal criterion rather than **Sensitivity'** *per se*.[26] What follows if **Sensitivity'** is applied with the external criterion? This, indeed, is the question Goldman raises in his review. His answer is that if the method is individuated externally, **Sensitivity'** is either false, or at least, has no clear application to the Dachshund case. The reason is that M_1 and M_2 being different, the antecedent of **Sensitivity'** is false. So the external criterion seems to be much worse for **Sensitivity'**. Is it? In what follows, I will defend **Sensitivity'** from this line of objection.

Goldman, in his (1976), proposes a token reliabilist theory to address cases like the Dachshund case, which, in his (1983) and (1986, ch. 3), is rendered a version of **Sensitivity**':

[25] Williamson is no fan of sensitivity theory, and offers another type of counterexamples that allegedly works against both **Sensitivity'** and **Sensitivity***. Becker (2009) countenances **Sensitivity***, and tries to save it from the alleged counterexamples, by deploying Goldman's fine-grained way of method individuation mentioned below. This strategy only brings Pyrrhic victory; the reason that the method individuated in too fine-grained a way works against the counterexamples is that it would not be used in most, if not all, other possible worlds. This strategy, then, takes sensitivity further away from the original intention of Nozick.

[26] It is worth mentioning that Williamson's refusal of the internal criterion is in part motivated by his anti-skepticism strategy: there are differences between perceptual evidence of a veridical case and that of a non-veridical case. This strategy commits him to the claim that *S*'s evidence would be different if *S* were a victim of a skeptical hypothesis.

(Goldman Sensitivity) For all *S* and *p*, if *S* knows *p via* token method *M*, then if *p* were false, *S* were to use *M* to arrive at a belief regarding *p*, and any relevant alternative were true, *S* would not believe *p via M*.

Goldman's criterion for method individuation is in terms of 'perceptual equivalence.' It is both internal and external: being internal in the sense that it in part relies on the quality of *S*'s sensory experience, and being external in the sense that it involves the external object as of which *S* has sensory experience (Goldman's criterion is more fully explained in §3. 5 of Chapter 4). In addition to the difference in criterion for method individuation between Nozick and Goldman, **Goldman Sensitivity** differs from **Sensitivity'** as to what possible worlds are relevant for sensitivity – as discussed in the next section, method individuation and selection of the relevant worlds are intimately related. Goldman stipulates of the Dachshund case that the worlds in which Oscar encounters hyenas are closer to the actual world than are the worlds in which he sees other familiar animals, because "[hyenas] are frequenters of this field" (Goldman, 1976, p. 92). That is, he is taking comparative similarity of possible worlds to be proportional to objective probability, in particular, to be frequency of some sort: a possible world w_1 is more similar to w than w_2 is if the objective probability of w_1's existence is higher than the objective probability of w_2's. Once the ordering source for evaluating sensitivity is set by such a probabilistic measure of similarity, as both Goldman and Williamson envisage, **Sensitivity'** fails to yield the verdict that Oscar knows in the Dachshund case. **Goldman Sensitivity** is designed to deal with this difficulty; the relevant possible worlds are those that, among the worlds in which the relevant alternatives are true, are closest to the actual world on the objectively probabilistic ordering source. Thus, a set of possible worlds remote from the actual world may be relevant for sensitivity, when the relevant alternatives are true in such a set of worlds.

As mentioned in §2 of Chapter 1, DeRose is a proponent of the counterfactual robustness theory: "[a]n important component of being in a strong epistemic position with respect to *p* is to have a belief as to whether *p* is true match the fact of the matter as to whether *p* is true, not only in the actual world, but also at the worlds sufficiently close to the actual world" (DeRose, 1995, p. 34). Important here is the idea delineated in this passage, viz., that sensitivity comes in degrees, and hence it is a measure of the strength of epistemic position.[27] One plausible way to conceptualize the degrees of sensitivity, as DeRose suggests it, is to take the degree of *S*'s sensitivity regarding *p* to be proportional to how far out across non-*p* possible worlds *S* does not believe *p*. Moreover, if the strength of epistemic position is a function of sensitivity, the ordering source relevant for sensitivity must be *epistemic*, i.e., comparative similarity of possible worlds must be proportional to the strength of *S*'s epistemic position.[28]

---

[27] DeRose, though being commonly construed to do so, does not embrace sensitivity as a necessary condition for knowledge. Rather, he employs it as a heuristic and explanatory device to account for *S*'s failure of knowledge; I take his point to be that if *S* does not know *p*, *S* fails to satisfy sensitivity with regard to *p*, but the converse and the inverse of this conditional do not hold. For it is enough for *S*'s knowing *p* if *S*'s belief *p* matches the fact regarding *p* in nearby possible worlds, even though it does not in the worlds relevant for sensitivity. Such a case obtains when the closest non-*p* worlds are still remote compared to the nearby *p* worlds, relative to the relevant modal base and ordering source. For more on DeRose's attitude toward sensitivity theory, see (DeRose, 2004a).

[28] This is similar to how Kratzer's ordering semantics for modals conceptualizes the strength of modals.

For example, Oscar has discriminative abilities or methods to tell the difference between dogs and other familiar animals. The existence of such abilities or methods is the very reason that we do not hesitate to think that he is in a sufficiently high epistemic position with regard to the belief that there is a dog before him. This suggests that the possible worlds relevant for the Dachshund case are those where he sees such animals as cats, cows, birds, etc., rather than hyenas, even though it is objectively probable that he sees hyenas. Thus, the epistemic ordering source may be different from the objectively probabilistic ordering source; the relevant epistemic ordering source must be such that, relative to it, the possible worlds where he sees those animals are more similar and thereby closer to the actual world. Even with the internally individuated method token *judging by vision*, **Sensitivity'** is true for the Dachshund case relative to this epistemic ordering source, whereas it is false relative to the probabilistic ordering source.

If this is right, the objection Goldman makes against **Sensitivity'** boils down to the claim that intuitive closeness of possible worlds is not specific enough to give the relevant epistemic ordering source, since many ordering sources are possible. The objection, then, has nothing to do with the question of which formulation of the sensitivity condition is the right one. For all of the formulations proposed face the same problem of what ordering source is epistemically relevant. Provided that the alleged problems with **Sensitivity'** are refuted, I shall henceforth refer to **Sensitivity'** as the sensitivity condition.

The discussions above are intimately related to the version(s) of the reference class problem that sensitivity theory inevitably faces. To begin with, the counterfactual condition for knowledge, manifested in sensitivity theory, can be construed as a dispositional condition, with the assumption that the conditional analysis of disposition is correct, i.e., disposition is analyzed in terms of a counterfactual conditional. Although it is controversial what conditional analysis best captures disposition, a simple analysis will do for the present purposes:

(Disposition) for all *x*, *r*, and *s*, *x* is disposed to give response *r* to stimulus *s* iff *x* would give *r* if *s* were to obtain.

Obviously, the dispositional analysis of knowledge in accordance with the first formulation of sensitivity Nozick gives is obtained, if '*x*,' '*r*,' and '*s*' are replaced with '*S*,' 'not believing *p*,' and '~*p*,' respectively, simply because the right-hand side of **Disposition** is a counterfactual.[29] **Disposition**, as it stands, is inadequate to capture the fact that dispositions come in degrees, and Manley & Wasserman (2007) propose a conditional analysis that incorporates the degrees of disposition:

(Disposition') for all *x*, *r*, and *s*, *x* is more disposed than *y* to give response *r* to stimulus *s* iff there are more *s*-cases in which *x* would give *r* than *s*-cases in which *y* would give *r*.

Manley & Wasserman intend **Disposition'** to be a generic formula and compatible with various ways to fill out details. One way to measure the dispositional strength of *x* by

[29] Gundersen (2010) forcefully argues that sensitivity theory is a kind of the conditional analysis of disposition, and that proposed counterexamples to sensitivity theory are structurally similar to alleged counterexamples to the conditional analysis of disposition.

Disposition', as they mention it, is to combine it with comparative closeness of possible worlds. That is, a set of *s*-worlds are ordered on the basis of some similarity of significance, and *x* and *y* are compared with regard to how far out in such a set of worlds each of them can give response *r*; the more far-reaching has stronger disposition to give response *r* to stimulus *s*.

The single-case propensity theory identifies the probability of *A* with the strength or degree of the propensity or disposition of a token chance set-up to produce *A*. Given **Disposition'** and the closeness-based measure of strengths of dispositions, it follows that the single-case propensity theory and sensitivity theory are structurally alike:

(Sensitivity) $EP(p) = EP(p \mid \text{the belief } p \text{ is produced by method token } M) = x$ iff M has the disposition to produce the belief p to degree x

§3. 3. 2. The Reference Class Problem for Sensitivity Theory

Due to the structural similarity between sensitivity theory and the single-case propensity theory, sensitivity theory faces a variant of the reference class problem, just as the single-case propensity theory does: how the token method *M* is individuated? The value of *EP* varies with the characterization of *M*, and thereby, unless this question is settled, *EP* has no determinate value on the sensitivity interpretation of **Epistemic Position**. Furthermore, another variant of the reference class problem arises for sensitivity theory. As we have seen, even if a method is individuated in some way or another, it may not be fully settled what ordering source is in order. The problem, then, is what ordering source is relevant for sensitivity (alternatively, it may be put as the question of which possible worlds are relevant for sensitivity). As we have seen above, the choice of the ordering source affects whether a belief is sensitive and the extent to which it is. So these two problems must be settled in order to apply sensitivity to any given case. They are akin to the process type problem and the environment type problem for type reliabilism, respectively.[30]

Again, these two problems are interconnected. DeRose (1995, pp. 20-2) argues that it is not necessary for the sensitivity theorist to methodize the sensitivity condition, in order to deal with cases like the Grandmother case. His point is that if the relevant ordering source is properly adjusted, the simple unmethodized formulation of sensitivity can yield the right result of Grandmother case. This is because method individuation and specification of the relevant worlds are interdependent. Both the modal base and the ordering source determine the similarity space relevant for counterfactuals. Thus, if both are fully specified, there is no need to specify the method independently; conversely, specifying the method can contribute to fix the relevant possible worlds, by narrowing down the modal base and/or by setting the relevant similarity measure for the ordering source. Then, it is possible for the ordering source to take care entirely of what possible worlds are close to the actual world, e.g., it may fully determine the relevant closest worlds as those where grandmother sees grandson. Of course, this is not always the case, and the relevant method and the relevant ordering source may need to be specified separately.

---

[30] It is often pointed out that the generality problem for type reliabilism arises for sensitivity theory as a problem concerning method individuation. As far as I know, the first to notice the similarity between the two problems is Pappas (1987).

§3. 3. 3. Safety Theory

Safety theory is a version of token reliabilism. Although it is a relatively new development, it has drawn much attention, and is supported by such figures, as Peacocke (1999), Pritchard (2005a), Sainsbury (1997), Sosa (1999, 2002), and Williamson (2000a, 2009a). Again, their theories differ in detail; for instance, Sosa, just as Goldman does with regard to sensitivity, holds that the full-blooded analysis of (animal) knowledge requires safety as an extra condition for knowledge, in addition to a type reliabilist condition for justification[31]; whereas, others are barely concerned with justification. It is a bit difficult to give a precise formulation of the safety condition, since all the proponents offer slightly different formulations.

Safety is contrasted with sensitivity in that safety is centered on truth acquaintance, and sensitivity on error avoidance.[32] The underlying idea of safety is that knowledge requires *S*'s belief to reflect the way the world is, over a certain range of possible worlds. Thus, Sosa formulates safety in terms of counterfactual conditional[33]: necessarily, for all *S*

---

[31] Sosa, however, has recently abandoned this view in his (2007), where he modifies his type reliabilist account with an explanatory requirement and attributes to it the role of anti-Gettiering device, which his safety theory used to play. In the meantime, safety is still incorporated into his account of reflective knowledge. I discuss his explanatory requirement in §2 of the next Chapter.

[32] Sosa holds that the safety condition is the contrapositive of the sensitivity condition – note that even if so, they are not equivalent, since they are modal statements. DeRose (2004a) objects to the contraposition relation between them.

[33] Sosa, in effect, formulates the safety condition differently, from paper to paper and even within a paper. The formulation of safety here is taken from his (1999, p. 146; 2002, p. 265).

and *p*, if *S* knows *p*, then if *S* were to believe *p*, *p* would be true. He also methodizes safety along the lines of **Sensitivity’**.[34]

One problem with the counterfactual formulation is, as DeRose (2004a, pp. 29-31) and Williamson (2000a, pp. 148-50) note, that when *S* believes *p*, it is a true-true counterfactual. On Lewisian semantics, a counterfactual with true antecedent and true consequent is guaranteed to be true. For this reason, some theorists, like Williamson and Pritchard, prefer to formulate safety directly in terms of possible worlds. Following them, I take it that the safety condition is to be fully formulated as follows: necessarily, for all *S* and *p*, if *S* knows *p via* method token *M* at *w*, then *S* believes *p via M* in all the possible worlds closest to *w* in which *p* is true.[35]

---

[34] However, Sosa, at some points, seems to imply that *S*’s actual belief *p via M* is unsafe if *S* does not believe *p via M* in the close *p*-worlds. Then, the same type of ambiguity, as between **Sensitivity’** and **Sensitivity***, arises for the safety condition.

[35] This is similar to the formulation Williamson gives in his (2000a, p. 149). For his safety theory, see also (pp. 123-8). Pritchard’s formulation is the following:

> (Pritchard Safety) For all agents, φ, if an agent knows a contingent proposition φ, then, in nearly all (if not all) nearby possible worlds in which she forms her belief about φ in the same way as she forms her belief in the actual world, that agent only believes that φ when φ is true. (Pritchard, p. 163)

I don’t know what to make of this formulation. Pritchard attributes a heavy anti-counterexample role to the locutions ‘nearly all (if not all)’ and ‘nearby.’ This blurs what worlds are relevant for **Pritchard Safety**. Moreover, this is odd, given that Pritchard explicitly endorses that the relevant similarity is epistemic rather than objective. If only ‘nearly all’ rather than all worlds are relevant, it seems to entail that the strength of epistemic position is relativized to the two parameters: the comparative closeness of the worlds from the

DeRose (2004a) argues that safety theory can incorporate strengths of epistemic position: "[o]ne's belief that *p* is safer the more remote are the least remote possibilities wherein one believes that *p* without it being the case that *p*" (p. 33).[36] If safety comes in degrees, this is a way to conceptualize such degrees.

§3. 3. 4. The Reference Class Problem for Safety Theory

This much being granted, it is possible to propose a dispositional analysis of safety, and make parallels with the single-case propensity theory, just as we did for sensitivity. But such an analysis would be overly unnatural, entailing something like the claim that the fact that *p* has the disposition to be believed by *S*. Be that as it may, it is clear that safety theory suffers the variants of the reference class problem that correspond to those for sensitivity theory: first, how method *M* is individuated, and second, what the relevant ordering source is.[37] Again, unless these questions are settled, the strength of *S*'s epistemic position has no determinate value on the safety interpretation of **Epistemic Position**.

---

actual world and the number of the worlds at equidistance from the actual world, in which *S*'s belief is true. Perhaps, the principle of parsimony should apply here.

[36] It is not clear whether Sosa accepts DeRose's proposal. In his (2004a) reply to DeRose, Sosa simply offers a Gettier-like counterexample to sensitivity theory, and refuses DeRose's counterfactual robustness theory (as I noted in footnote 27, DeRose does not endorse sensitivity theory in a straight manner). This does not settle the questions of whether safety admits of degrees, and if so, of how the degree is determined.

[37] Becker (2008) argues that the generality problem is generalizable to Pritchard's safety theory.

§3. 4. Internalist Evidentialism

Internalists often raise some variant of the reference class problem against an externalist theory, and attempt to establish theoretical advantages of internalism over it. For example, the most avid critics, Conee & Feldman (1998, p. 140), claim that the prospects for a solution to the reference class problem for type reliabilism are 'worse than bleak,' and therefore type reliabilism looks 'hopeless.' In this section, I argue that internalist theories are no less susceptible to the reference class problem. My target here is primarily internalist evidentialism, since it is the prevailing version of internalism and any internalist theory involves some evidentialist tenet.[38]

As mentioned in §2, the internalist endorses **Propositional Justification**, and takes the primary bearer of justification to be inferential relations between propositions. However, she must accept that the notion of doxastic justification must be integrated into the full-blooded theory of knowledge. Suppose that *S* has good evidence e_1 for *p*, and thereby is propositionally justified in believing *p*. *S* believes *p*, but the very evidence on the basis of which *S* does so is e_2 rather than e_1, where e_2 is bad evidence for *p*. At the intuitive level, *S* is epistemically deficient in some relevant sense, and does not know *p*. To avoid this result, the basing relation or well-foundedness, i.e., the requirement that e_1 be

[38] For this reason, I am not directly concerned with accounts of the structure of justification, such as foundationalism and coherentism. In any case, such accounts presuppose that the relevant epistemic factor is evidence. A controversy in coherentism is whether coherence *per se* is or not an extra epistemic factor; if it is, the more coherent a system of beliefs is, the more justified individual beliefs within that system are. Alternatively, coherence may be reckoned to be a condition for justification *simpliciter*.

causally responsible for *S*'s believing *p*, or that *S* believe *p* on the basis of e_1, must be added as an extra condition for knowledge. Alston (2005, pp. 89-92) holds that the basing relation itself is an epistemic factor in the narrow sense, i.e., satisfying the condition of the basing relation makes *p* more probable. This view seems implausible, since it entails that the degree of doxastic justification for *p* is always higher than the degree of propositional justification for *p*. It is not clear why a causal relation between *p* and e_1 adds to the degree of justification. So I here treat the basing relation as a transmitter of the degree of justification, i.e., it transmits the degree of justification to the belief that *p* from the proposition *p*.[39] This entails, on internalist evidentialism, that the degree of doxastic justification never comes apart from the degree of propositional justification; though they may differ on externalist theories (if they have accounts of both types of justification).

By way of example, Feldman & Conee (1985) formulate their internalist externalism as follows:

> *S*'s doxastic attitude *D* at *t* toward proposition *p* is doxastically justified iff
>
> (i) having *D* toward *p* is justified for *S* at *t*;
>
> (a) *S* has some body of evidence *e* as evidence at *t*;
>
> (b) having *D* toward *p* fits *e*, and

---

[39] Feldman, in his (2005) review of Alston (2005), opposes to Alston's construal of the basing relation. It seems that he alludes to the same direction I'm taking here: the basing relation is a transmitter of the degree of justification. On this construal, the basing relation *per se* is not an epistemic factor in the narrow sense, though it is in the broad sense.

(c) there is no more inclusive body of evidence *e*' had by *S* at *t* such that having *D* toward *p* does not fit *e*';

and

(ii) *S* has *D* toward *p* on the basis of *e*.[40]

(i) is the condition for propositional justification, which consists of an evidentialist thesis (a), an adequacy requirement for evidence (b), and a version of the total evidence requirement (c); whereas, (ii) is the condition for the basing relation.

Any internalist evidentialist theory, as a theory of propositional justification, must have clauses corresponding to (a) through (c). (a) is the central tenet of evidentialism; and the internalist tenet constrains what counts as evidence: **Mentalism** requires it to be in *S*'s mind, and **Accessibility**, at least, requires evidence directly responsible for justification to be accessible to *S*. Feldman & Conee advocate **Mentalism**. (b) is also important for evidentialism. Though Feldman & Conee does not specify the relation 'fitting,' presumably, it is some inferential relation, e.g., *e* supports *p* or *e* is in favor of *p*; unless such inferential relation is in place, there is no sense in which *e* is evidence *for p*.

Because of their commitment to **Epistemic Probability**, internalists grant that the degree of justification is not objective, and reckons the degree of justification to be measured in accordance to some version of the epistemic probability theory. Although the internalists are rarely specific about what type of theory they accept, internalist evidentialism can be combined with any epistemic probability theory discussed in §3. 5 of

[40] This is slightly modified form their own definition in (ibid., p. 93). Oddly, they incorporate (a) through (c) into (ii) rather than (i), but it is clear in their paper that (a) to (c) are conditions for propositional justification.

Chapter 2, e.g., the logical theory, Bayesianism, Kyburgian theory, or the like.[41] In what follows, I use ‘internalist justification’ interchangeably with ‘epistemic probability’, and identify (internalist) justification function EJ_i with one or another epistemic probability function.[42] This means that the degrees of internalist justification are identified with the degrees of rational credence of some sort. Any theory of rational credence must involve the requirement of total evidence, as Feldman & Conee explicitly does in (c). However, they, in the definition above, do not distinguish immediate evidence *e* for *p* from the rest of the total evidence e*, often called the background or auxiliary evidence; e* need not be *S*’s entire evidence set with *e* subtracted. More properly, it is the portion of the evidence set that is inductively relevant for *p*.

Thus understood, internalist evidentialism amounts to the following view:

(Internalist Evidentialism) $EJ_i(p) = EJ_i(p \mid e) = x$ iff $EJ_i(p \mid e \wedge e^*) = \mathrm{x}$

[41] Fumerton (1995, 2004) is explicit that the relevant probability theory for his internalist evidentialism is Keynesian logical probability. Swinburne (2001) gives an excellent account for how an epistemologist chooses the relevant probability theory for justification; it depends on how much of idealization she allows for a theory of justification.

[42] Some internalist evidentialists, such as Conee & Feldman (2008), Lehrer (2000, 2005), and Smith (2010), argue that probability *per se* is not adequate for justificatory relation, and suggest incorporating explanatory relation into justification (see also (Achinstein, 2001)). Many coherentists, in fact, have the same type of requirement; coherence in the relevant sense is not merely probabilistic but explanatory. Bonjour (1985), however, warns not to read too much into such explanatory relation. As I discuss in §2 of the next Chapter, I am in agreement with Bonjour.

This says that the degree of internalist justification for *p* is determined by evidence for *p* and the total evidence. Just like type reliabilism, internalist evidentialism defined here is a theory of gradable justification, and hence needs to be combined with a view on justification *simpliciter*. The majority of internalists appeal to some threshold model of justification; to follow Kyburg (1961), the relevant threshold model for justification *simpliciter* is formulated as follows:

> (Threshold for Internalist Justification) Necessarily, for all *S* and *p*, $EJ_i(p)$ for $S \geq 1 - \varepsilon$ iff *S* is propositionally justified in believing *p*.

§3. 5. The Reference Class Problem for Internalist Evidentialism

From these definitions, it should already be clear that internalist evidentialism cannot evade the reference class problem. The variants of the reference class problem that arise for the epistemic probability theory carry over to internalist evidentialism. Since internalist evidentialism incorporates the requirement of total evidence, the reference class problem may not undermine its metaphysical tenet that there is always a portion of *S*'s evidence set relevant for uniquely determining the degree of propositional justification for *p*. This, however, does not unload the burden of explaining, for any given case, what e^* is, and why it is inductively relevant for *p*, from the internalist's shoulders. In addition, versions of internalist evidentialism are susceptible to other variants of the reference class problem, to which I turn below.

Comesaña (2006) claims that "the problem with the characterization of well-foundedness is intuitively the same as the generality problem, and so any theory of justification that incorporates the notion of well-foundedness into the characterization of justification will have to deal with the generality problem" (p. 31). He goes on to argue, by reference to Feldman & Conee's internalist evidentialism, that type reliabilism can exploit the resources that internalism appeals to in dealing with the problem concerning well-foundedness. The structure of Comesaña's argument seems to be the following:

(1). If the problem concerning well-foundedness for internalist evidentialism is solvable, the reference class problem for type reliabilism is solvable as well.

(2). The problem concerning well-foundedness for internalism is solvable.

(3). Therefore, the reference class problem for type reliabilism is solvable.

It is not clear what Comesaña thinks of as the problem concerning well-foundedness for internalist evidentialism, appearing in (1). But I take it that the problem is related to the relation *p* bears to the total evidence. On Feldman & Conee's model, *S*'s propositional justification for *p* consists in *p*'s fitting the total evidence *e*, where the fitting relation is some sort of inferential relation. But what inferential relation is it? *e* is divided into individual bits of evidence, e_1, …, e_n. Now, the question is which of these pieces of evidence is causally responsible for *S*'s believing *p*. Is *S* believing *p* merely on the basis of e_1, or of e_1 and e_1 *makes p probable*, or more? There are in principle infinite order-rising meta-evidence. Also, even if most constituents of *e* are first-order evidence, multiple pieces of evidence may be causally responsible for *S*'s believing *p*.

Whether *S*'s belief *p* is doxastically justified depends on what evidence *S* bases her belief *p* on. Comesaña seems to suggest that the difficulty concerning determining the base of *S*'s belief *p* is a variant of the reference class problem, because it is tantamount to the problem of specifying the belief-forming process or method *S* exercises in believing *p*, at least if the process or method in question involves evidence. However, the internalist is barely bothered by this problem; she simply specifies a small subset of *e* as the immediate evidence e_1 for *p*, and whatever else in the total evidence as the background or auxiliary evidence. Comesaña's solution to the reference class problem for type reliabilism is, indeed, to specify the relevant process type as {belief-forming process based on such e_1}.

Comesaña, again, is not explicit about why this problem for internalist evidentialism is easy to solve. It seems that the reason lies in (partial) endorsement of **Accessibility**; the relevant immediate evidence is that which *S* has access to by reflection alone, while *S* may not have such access to other part of the total evidence. Once thus construed, Comesaña's 'evidentialist reliabilism' is essentially the same as Alston's (1988) 'internalist externalism.' They are internalist in that justification requires internal access to the primary epistemic factor, the immediate evidence; whereas, they are externalist in that they deny **Epistemic Probability**, and measure the strength of epistemic position or the degree of justification along the lines with type reliabilism. As we will see in §3 of Chapter 4, Alston has long struggled with the reference class problem, though his solution is different from Comesaña's.[43]

[43] Alston offers his solution to the reference class problem in his (1995, 2005).

Both Comesaña and Alston, however, with reluctance, concede that the internalist maneuver does not fully solve the reference class problem; it is only designed to solve the process type problem, and leaves the environment type problem untouched. Furthermore, it is not even clear whether it can fully address the process type problem. Comesaña discusses the problem of how sensory evidence is individuated. Though he does not mention it, the problem he is concerned with is the problem of the speckled hen, first raised by Chisholm (1942). Suppose that two subjects, S_1 and S_2, see a hen with 48 speckles. S_1 has much better identificatory and discriminatory ability about the number of tiny objects than S_2 does. Indeed, S_1 recognizes that the hen has 48 speckles, while, for S_2, it is just a hen with lots of speckles. Intuitively, S_1, not S_2, is justified in believing that the hen has 48 speckles. But there is a sense in which they both have experience with the same content. With the assumption that sensory evidence is all there is to justification for perceptual beliefs, internalist evidentialism entails that S_1 and S_2 are both justified to the same degree.[44]

This problem is just part of the process type problem, focusing on what input type the relevant process type should involve.[45] Comesaña suggests that two solutions are

[44] On both internalist evidentialism and the epistemic probability theory, internalist justification or epistemic probability is dependent on propositional contents. The problem here is an aspect of the difficulty with assigning justificatory role to non-propositional evidence. For attempts to accommodate non-propositionality of sensory evidence into the framework of the epistemic probability theory, see (Jeffrey, 1983, ch. 11) and (Pearl, 1990)).

[45] Roddy (1994) raises this problem against Millar (1991), as a variant of the generality problem. The problem, as he states it, is what experience type is relevant for justification.

available: (a) to individuate the content of experience as of the hen narrowly in such a way that it differs between S_1 and S_2; and (b) to individuate the method S_1 and S_2 use to form the belief in question in such a way that it differs between S_1 and S_2. (a) amounts to using Goldman's internal-external criterion for method individuation, and (b) is just the same strategy in disguise. I will examine the internal-external criterion in more detail in §3. 5 of Chapter 4; to preview it, it has liabilities because the internal-external criterion specifies processes or methods too narrowly. At any rate, Comesaña proposes no principled way to specify the relevant input type, and thereby the relevant process type in any given perceptual case.

At this point, assuming premise (1) of Comesaña's argument, we can easily derive, by *modus tollens,* the denial of (2), viz, that the problem concerning well-foundedness is unsolvable. Indeed, the version of internalist evidentialism that accepts the justificatory role of experience fall prey to the difficulty with experience individuation. This is a form of the reference class problem for internalist evidentialism.

This problem only arises for a version of internalist evidentialism, such as Feldman & Conee's, on which sensory experience is a justificatory factor of perceptual, non-inferential belief. Those who are impressed by Sellars' (1956) attack on the myth of the given, however, do not endorse such a version of internalist evidentialism – indeed, Sellers' argument against the myth of the given, in part, relies on a problem akin to the problem of the speckled hen. What I call 'Sellarsian internalist evidentialism' is the view that only mental states or events with propositional contents can count as evidence.[46] On this view, perceptual belief is formed non-inferentially, but its justification is dependent on what Sellars (1966, p. 211) calls 'trans-level inference':

(Trans-Level Inference)

X's thought that *p* occurred in manner *M*

So (probably) *p*

Sellars, in his (1956), states that in perceptual cases, manner *M* is a reliable manner. That is, in order for *S* to be perceptually justified in believing *p*, *S* needs to be justified in believing that the belief *p* is formed under the condition in which *S*'s perception is reliable. Let us call this requirement the 'requirement of trans-level inference.'

The requirement of trans-level inference can be motivated independently of Sellars's attack on the myth of the given: without it, any version of internalist evidentialism runs afoul of the problem of easy knowledge – I will discuss this problem in detail in §1 of Chapter 6. In the face of the problem of easy knowledge, Feldman & Conee, in the (2004) afterword to their (1985), concede that their internalist evidentialism has to incorporate a requirement like that of trans-level inference (pp. 105-7). Furthermore, the requirement of trans-level inference is simply **Meta-Justification** applied to propositions perceptually believed. Thus, those internalists who are in favor of **Meta-justification** have no reason to refuse it in the first place.

[46] Proponents of Sellarsian internalist evidentialism involve Brandom (1995), Davidson (1983), Rorty (1979), Rosenberg (2002), and Williams (1977, 1991).

If internalist evidentialism is combined with the requirement of trans-level inference, then another variant of the reference class problem inevitably arises. It is natural, especially for Sellarsian internalist evidentialists, to render reliability a measure of internalist justification. For this seems to be the only way to respect the fact that perceptual justification admits of degrees, if the quality of experience plays no role in perceptual justification. Trans-level inference, then, is an expert function in the sense discussed in§3. 6 of Chapter 2:

(Reliability Expert) $EJ_i(p) = EJ_i(p \mid RELI(p) = x) = x$

What this principle does is to specify the reliability function *RELI* for perceptual cases as an expert function such that it matches $EJ_i(p)$ with the reliability of *S*'s perception with regard to *p*.[47] Sellars does not specify what sense of reliability is relevant for trans-level inference, and *RELI* is undefined. But a variety of reliabilist theories, as discussed in §3, can come to help internalist evidentialism to specify *RELI* – *RELI* can be the function EJ_e or *EP* interpreted on one or another externalist theory. That is, internalist evidentialism can piggy-back on some version of reliabilism. However, this piggy-backing, as it usually does, comes with a great cost, the reference class problem; internalist evidentialism inherits the reference class problem(s) from the reliabilist theory with which it is combined. Then,

[47] **Reliability Expert** is similar to **Principal Principle** or **Support**; it transmits the degree of reliability, as known or justifiably believed by *S*, to the degree of internalist justification. For this reason, although I do not belabor the point here, it is possible to construct a Dutch book argument for **Reliability Expert**.

internalist evidentialism has no way-out of the reference class problem, insofar as it requires a reliability expert function for perceptual justification.[48]

The extent to which internalist evidentialism is susceptible to the reference class problem *via* **Reliability Expert** is not as great as externalist theories are. **Mentalism** or **Accessibility** constrains what evidence *S* has, and thereby what candidate reference classes exist for *RELI*; for externalists, such candidates are infinitely many; whereas, for internalists, they are those which are in *S*'s mind or *S* can introspectively identify or recognize. That is, for internalism, the candidates are relatively limited. For this reason, it seems to be relatively easy for internalist evidentialism to specify one or another reference class as relevant. However, Bishop (forthcoming) argues that the reference class problem is no less urgent for internalism than it is for externalism, by appeal to the following case, where 'BFPT' is an abbreviation of 'belief-forming process type':[49]

---

[48] Perceptual justification here is a kind of propositional justification. It, however, is intimately related to doxastic justification. '$RELI(p) = x$' contains two pieces of information: (i) the belief *p* is formed or caused by a certain process or method, and (ii) its reliability is *x*. When the process or method in question involves evidence, the truth of (i) is a condition for doxastic justification. It seems that the truth of (i) and (ii) is required for applications of **Reliability Expert**, and this may be the reason that when (i) and (ii) are false, *S* is not propositionally justified. I argue in §3. 1 of Chapter 7 that this is a way to accommodate our intuition about the Optimist case below.

[49] Bishop's argument is similar to the one I have given on the basis of the requirement of trans-level inference. The difference is that he is only arguing that the generality problem for type reliabilism is transmitted to internalist evidentialism. Given my arguments for the reference class problem(s) for any externalist theory, the reference class problem for internalist evidentialism has broader scope. It can be generated from any notion of reliability *via* **Reliability Expert**.

> (Optimist) After years of neglect, Paula spends much of the day, from noon to 4:55, balancing her check book, performing a few thousand simple arithmetic calculations. She spends the next 5 minutes coming to various absurdly optimistic beliefs about her future. She tries to justify these beliefs by appealing to a BFPT that she knows is reliable. She argues: ''My beliefs about my future will be produced by process tokens that are instances of a highly reliable BFPT." That BFPT consists of the process tokens that share the property of having produced my beliefs between noon and five.

Process type {forming beliefs between noon and five} is, perhaps, one of the few process types Paula believes her process *g* to be a token of. Given the expert assignment relativized to this process type, $EJ_i(p \mid g \in \{\text{forming beliefs between noon and five}\})$, her degree of justification for *p* is very high. But intuitively, she is not justified in believing *p*. The reason we hesitate to grant her justification seems to be that {forming beliefs between noon and five} is not appropriate for the Optimist case. Thus, even if internalist evidentialism manages to reduce the number of candidate process types, the reference class problem is still pressing. Internalist externalism, just like other theories, must offer a principled way to specify the relevant reference class in any given case.

This completes my argument for the ubiquity of the reference class problem in the substantive theory of knowledge. Any theory succumbs to one or more variants of the reference class problem. So, if Conee & Feldman is right that the reference class problem for type reliabilism is 'worse than bleak,' the same can be said of any other theory.

However, I am more sanguine about the possibility of solution than Conee & Feldman is. Much effort to solve the reference class problem has been made both in probability theory and the theory of knowledge. It would be illuminating for the pursuit of the solution to examine whether each of the alleged solution is on the right track, and if not, why it is not. Chapter 4 is devoted to this business.

Chapter 4: Alleged Solutions to the Reference Class Problem

Introduction

I have been emphasizing the structural similarity between (variants of) the reference class problem in probability theory and (variants of) the reference class problem in the theory of knowledge. Given the structural similarity between the two problems, it is natural to expect that solutions, if they exist, would be structurally alike as well. The two problems being structural similar, however, does not mean that they are substantively the same; neither does possible solutions being structurally similar mean that they are substantively the same.

In §1 and 2, I elucidate a substantive difference between objective probability theory and the theory of knowledge, and what ramification it has for the reference class problem. In particular, §1 delineates the constraints imposed on a solution to the reference class problem in general. In §2, a distinction is made with regard to the context of inquiry in which objective probability theory and the theory of knowledge pursue the explanation of their leading question. The difference in context of inquiry pertains to what types of factors are permissible for the explanation in each discipline. Different criteria for permissibility, then, lead to different ways in which objective probability theory and the theory of knowledge seek the solution to the reference class problem. §3 presents and examines alleged solutions to the reference class problem(s) in the theory of knowledge.

§1. Constraints on Solutions

In §4. 2 of Chapter 2, I divided the reference class problem into three sub-problems: the problem of conflicting reference class, the problem of trivializing reference class, and the problem of non-projectible reference class. Some of these problems have been prominent both in probability theory and the theory of knowledge. For example, Feldman (1985), in formulating the reference class problem for type reliabilism, dubs the problem of trivializing reference class the 'Single-Case Problem,' and adds a further problem:

> A very broad account of relevant types of belief-forming processes leads to what we may call "The No-Distinction Problem." This arises when beliefs of obviously different epistemic status are produced by tokens that are of the same (broad) relevant type. For example, if the relevant type for every case of inferring were the type "inferring," then [type reliabilism] would have the unacceptable consequence that the conclusions of all inferences are equally well justified (or unjustified) because they are believed as a result of processes of the same relevant type. (p. 161)

Thus, Feldman characterizes the reference class (generality) problem as that of providing "an account of relevant types that is broad enough to avoid The Single Case Problem but not so broad as to encounter The No-Distinction Problem" (ibid.).[1]

[1] Goldman (1979, pp. 115-6) already formulated the reference class problem for type reliabilism in the same way.

In probability theory, Salmon (1971b) characterizes the reference class problem in much the same way as Feldman does in the theory of knowledge:

> The reference class must, therefore, be broad enough to provide the required number of instances for examination to constitute evidence for an inductive inference [from observed frequency to relative frequency]. At the same time, we want to avoid choosing a reference class so broad that it includes cases irrelevant to the ones with which we are concerned. (pp. 41-2)

I include the No-Distinction Problem in the problem of conflicting reference class, since, as Salmon states it, it amounts to the problem of excluding irrelevant reference classes.

Both Feldman and Salmon require any solution to the reference class problem to avoid the two extremes, selecting too broad a class and too narrow a class. This follows from the characterization of the problem *per se*. What is more, Conee & Feldman (1998) argue that any solution to the reference class problem for type reliabilism must satisfy three constraints: (i) it must offer a principled, rather than ad hoc, case-by-case, way to select the relevant reference class in any given case; (ii) the relevant reference class it assigns must result in justification attribution or denial that is co-extensional with our intuitive justification attribution or denial; and (iii) it must stay in the naturalistic reductionist spirit of reliabilism. (ii) and (iii) are directed merely at type reliabilism, or more narrowly, naturalistic type reliabilism, as a theory of justification, and hence lack general applicability. For more generality, I modify (i) through (iii) as follows:

(Non-Arbitrariness) A solution must offer a principled, rather than ad hoc, case-by-case, way to select the relevant reference class in any given case, if there is the one therein.

(Accordance) The relevant reference class a solution assigns must result in justification/knowledge attribution or denial that is co-extensional with our intuitive justification/knowledge attribution or denial.

(Reductionism) A solution must not rely on concepts or properties that are conceptually and explanatorily posterior to knowledge or justification.

Non-Arbitrariness requires any solution to the reference class problem to offer a way to determine the factors pertaining to the relevant reference class for any given case, if there is the one.[2] Note that it is consistent with the possibility that there is no uniquely relevant reference class; this may be the case, e.g., when justification or knowledge ascription is vague. **Accordance** is a specific form of the conceptual constraint for substantive theories of knowledge. **Reductionism** follows from the reductionist tenet of the substantive theory of knowledge.[3] For this reason, the knowledge-first approach is not

---

[2] **Non-Arbitrariness** may be read more strongly, e.g., a solution must offer one or a set of rules or procedures to determine the relevant reference class in any given case. Conee & Feldman imply this reading of their (i) at some points. How this reading differs from the reading I am giving depends on what is meant by a rule or procedure.

[3] Williamson, a representative proponent of the knowledge-first approach, is well aware of this consequence. Against Conee & Feldman (1998), he demurs as follows:

susceptible to **Reductionism**, if it purports to solve the reference class problem in some way or another – as I remarked in §2 of Chapter 2, the epistemological reference class problem is pressing even for the knowledge-first approach.

These constraints may be equally applicable to solutions to the reference class problem in probability theory, insofar as they are circumscribed by the conceptual constraint and reductionism. So the solution to the reference class problem in the theory of knowledge may be able to appeal to the structural similarity between the reference class problems in the two disciplines. Such appeal, however, is doomed to be limited; the solution to the reference class problem substantively differs across the disciplines, in at least one respect.

§2. Epistemic Context of Inquiry

> Let us concede for the sake of argument that the generality problem is indeed insoluble. It does not follow that appeals to reliability in epistemology should be abandoned. For the insolubility of the generality problem means that the concept of reliability cannot be defined in independent terms; it does not mean that the concept is incoherent. Most words express indefinable concepts; 'reliable' is not special in that respect. Irrespective of any relation to the concept *knows*, we clearly do have a workable concept *is reliable*; for example, historians sensibly ask which of their sources are reliable. The concept is certainly vague, but most words express vague concepts; 'reliable' is not special in that respect either. The concept *is reliable* need not be precise to be related to the concept *knows*; it need only be vague in ways that correspond to the vagueness in *knows*. No reason has emerged to doubt the intuitive claim that reliability is necessary for knowledge. (Williamson, 2000a, p. 100; cf. also his (2009b))

Attempts to solve the reference class problem have been endeavored in probability theory, as well as in the theory of knowledge. The first promising line along which a solution is pursued is drawn by Reichenbach (1949). He proposes to identify the relevant reference class with "the narrowest class for which reliable statistics can be compiled" (p. 374). This idea consists of two components: the relevant reference class in any given case is (a) the narrowest among candidate reference classes for which (b) reliable statistical data exists.

(a) allows for two different readings, depending on how a class is defined. A class can be defined, either by enumerating all the class members or by specifying the predicate of which all the class members are true. There are cases where only one way of definition is possible, e.g., when language lacks the relevant predicate, or when it is intensional. I here put aside such cases. Then, 'the narrowest reference class' refers to the class that contains the minimal number of elements or the extension of the strongest predicate.

What Reichenbach means by (b) is not clear.[4] But again, for present purposes, I assume that the standard interpretation of (b) is right, viz., the single or token event in question is a random member of the reference class it is assigned to.

---

[4] If 'reliable' in (b) is construed in some reliabilist way, then, obviously, this is no solution to the reference class problem, merely replacing one reference class problem with another. The passage in which Reichenbach offers the clearest exposition of (b) is that "the reference class that, on the basis of our present knowledge, will lead to the greatest number of successful predictions" (ibid., p. 377). This passage contains several controversial points: first, it relativizes the choice of the relevant reference class to our present knowledge. As we will see below, once thus relativized, definite probabilities are not objective (this is motivated by Reichenbach's positivist idea that statements of definite probabilities are not directly verifiable

As a matter of fact, virtually all the existing solutions to the reference class problem in probability theory are developed along the lines of Reichenbach's idea. As one of the early attempts, Carl Hempel, in a series of writings (1962, 1965, 1968, 1977), works out Reichenbach's idea as part of the I-S (inductive-statistical) model of scientific explanation. Wesley Salmon, in his (1971a, 1971b, 1984), contrasts the I-S model with his alternative, the S-R (statistical-relevance) model, which contains his own solution to the reference class problem on the basis of Reichenbach's idea.

First, according to Salmon, the two models are formulated as follows:

> I-S model: An explanation of a particular fact is an *inductive argument* that confers upon the fact-to-be-explained a *high inductive probability*.
>
> S-R model: An explanation of a particular fact is an assemblage of facts *statistically relevant* to the fact-to-be-explained *regardless of the degree of probability* that results. (Salmon, 1971a, p. 11)

Second, Fetzer (1972, 1977) points out that a difference between the two models consists in what why-question each purports to answer: the I-S model asks why it is reasonable to believe that the explanandum has occurred or will occur; whereas the S-R model asks why the explanandum has or will occur. In other words, the two models differ as to in what *context of inquiry* they seek the explanatory relation between an explanans

---

by experience); second, whose knowledge matters? Reichenbach does not answer this question; third, it may be a pragmatic matter what count as 'successful predictions.' If so, utilities are important for the choice of the relevant reference class.

and an explanandum. To use Fetzer's terminology, the I-R model has its place in an *epistemic*, '*reason-seeking*' context, and the S-R does in a *causal*, '*explanation-seeking*' context.[5] Fetzer makes this point as an objection to the S-R model originally proposed in Salmon (1971b), since not every statistically relevant factor is causal; Salmon (1984) accepts this objection and modifies the S-R model accordingly. These two differences are reflected in the ways in which each model addresses the reference class problem.

Hempel's solution to the reference class problem is to choose the *maximally specific reference class* (called the 'requirement of maximal specificity' or 'RMS'), and RMS is reducible to the requirement of the total evidence. According to Hempel, an explanans *Ga* explains an explanandum *Fa* only if $Ep(Fa \mid Ga \wedge e) > r$, for some sufficiently high r, where e is the total evidence, and *Ep* is an epistemic probability function akin to the confirmation function of Carnap's logical theory.[6] A predicate '*G*' is relevant, i.e., maximally specific relative to e iff (i) '*Gx*' entails neither '*Fx*' nor '*~Fx*,' for all x, (ii) e contains a probabilistic lawlike statement connecting F and G, and (iii) '*G*' denotes the narrowest class that satisfies (i) and (ii).[7] For Hempel, a lawlike statement is tantamount to

[5] Hempel himself draws a similar distinction in context (1965, pp. 370-4; 1968, p. 151). But what he calls 'explanation' is simply to give explanatory reason in the form of argument.

[6] Hempel changes over time his view on the relation of the requirement of the total evidence to the I-R model of explanation. In his (1968), he revises the old view and emphasizes that not every evidence is relevant for explanation; the total evidence is restricted to law-like statements and statements about single events. For other points of revision, see the next footnote.

[7] This is Hempel's formulation of RMS, taken from his (1962) and (1965). He modifies RMS in his (1968), by incorporating the notion of statistical relevance. I ignore this complication here, merely because those who

a statement of the limiting frequency of F within G. This much being given, the explanation of the I-R model relies on an expert probability function, since it identifies limiting frequencies with epistemic probabilities.

By contrast, Salmon's solution to the reference class problem is to use the *broadest objectively homogeneous class* (henceforth, I call this the 'requirement of broadest homogeneity' or 'RBH'). The homogeneity of a class is defined in the following way: a partition $\{G \wedge C_i\}$ of a class G is a set of mutually exclusive and exhaustive subclasses of G, and a cell is a subclass in the partition $\{G \wedge C_i\}$. Now, G is homogeneous if the objective indefinite probability of F relative to each cell $G \wedge C_i$ of G is not different from each other cell, i.e., if $P(F \mid G \wedge C_i) = P(F \mid G \wedge C_j)$, for all $i \neq j$, where Salmon interprets the objective probability function P along the lines with the long-run propensity theory. Class G is objective iff every factor included in G is causally relevant for F. Partition $\{G \wedge C_i\}$ is the broadest iff further division of it gives rise to a shift in objective probability, i.e., $P(F \mid G \wedge C_i) \neq P(F \mid G \wedge C_i \wedge C_j)$, for any j. When adding C_j to or subtracting it from $\{G \wedge C_i\}$ does not make a difference in value to $P(F \mid G \wedge C_i)$, i.e., when $P(F \mid G \wedge C_i \wedge C_j) = P(F \mid G \wedge C_i) \neq P(F \mid G \wedge C_j)$, C_i 'screens off' C_j from F in reference class G, making C_j statistically and causally irrelevant. Note again that not every statistically relevant factor is causal, and for this reason, the objectivity of a class is crucial insofar as the S-R model is a model for causal explanation. Unlike the I-R model, the S-R model does not require high probability for causal-statistical explanation. That is, *Ga* explains *Fa* only if G is causally and

appeal to the maximal specificity for a solution to the reference class problem in the theory of knowledge pay no attention to the modified RMS.

statistically relevant for F with regard to a, by boosting the probability of F; put differently, only if $P(F \mid G) > P(F)$, even though $P(F \mid G)$ is low.

As discussed in §3. 3, an attempt at solution to the reference class problem in the theory of knowledge appeals to Salmon's RBH. Although it is barely noticed, however, such brute appeal cannot work. For the context in which the S-R model operates is not the same as the context in which the theory of knowledge belongs. A context of inquiry is determined by interests and purposes of inquiry; in particular, it is determined by what question is to be addressed thereby. The S-R model and the theory of knowledge substantively differ as to what question they purport to address.

Lehrer (2000, p. 198) makes a distinction between the question of *why S believes p* and the question of *how S knows p*. The former question, which he reckons to be about the basing relation, is causal, and hence is not the primary subject matter of the theory of knowledge. Lehrer's point, then, is that the theory of knowledge is not primarily concerned with causal questions about knowledge and its components. In addition, Lehrer (2005, p. 417), from the vantage point of internalist evidentialism, elucidates that it is because appeal to *S*'s evidence fails to explain *why S's belief that p is true* that *S* does not know *p* in Gettier cases. This question may be more properly expressed as the question of why *S believes p truly*, since it is equivocal as to whether the truth of the content of *S*'s belief or *S*'s act of believing truth matters, and Lehrer has the latter in mind; it is of little epistemic importance why proposition *p* is true. Then, answering the latter question is a precondition for answering the question of how *S* knows *p*; for example, Lehrer's internalist evidentialism amounts to the view that *S* knows *p* only if *S*'s evidence for *p* explains *S*'s believing *p* truly.

This point is generalizable to other theories of knowledge; *S* knows *p* only if the relevant epistemic factor(s) explains why *S* believes *p* truly; or put differently, only if *S* believes *p* truly *because S* instantiates the relevant epistemic factor(s). Indeed, recently, Greco (2002, 2003, 2004a, 2007a, 2008, 2009, 2010) and Sosa (2003, 2007) incorporate such an explanatory requirement into their versions of type reliabilism. Greco, however, mistakenly construes the explanation in question to be causal (2008, pp. 419-20; 2009, p. 18; 2010, pp. 106-7)). This cannot be right; as Lehrer notes, it may be a causal matter whether *S believes p* truly because *S* instantiates a set of epistemic factors, if the event of *S*'s believing *p* truly is identified with the event of *S*'s believing *p*; but much more is required for the explanation of why *S* believes *p truly*. There is no reason to think that every epistemic factor is causal.[8]

---

[8] Greco argues that when luck is salient, it deprives type reliability of causal explanatory power, i.e., it fails to causally explain why *S* believes *p* truly. This may be the affirmation of the Aristotelian thesis that explanation excludes luck. Causal explanation, presumably, excludes 'causal luck,' e.g., deviant causation. But there is no reason that causal explanation excludes epistemic luck, as I discuss below. Greco argues for his 'subject-sensitive' contextualism by appeal to the context-sensitivity of causal explanation. Though he is not explicit, his version of contextualism is non-standard, in that what varies with subject's or attributor's context is the strength of epistemic position, in addition to the epistemic standard. Greco does not mention it, but perhaps, the best way to derive the context sensitivity of causal explanation is van Fraassen's (1980, §2, ch. 4). He argues that causal explanation is relative to a why-question, which is of contrastive form. On his account, the choice of contrast class is context-sensitive, and it determines the relevant reference class for causal explanation in a given context. If epistemic explanation is similar to causal explanation in the relevant respects, a similar argument for context-sensitivity may be constructible for the former.

Let us call this non-causal explanatory relation between epistemic factors and *S*'s believing truly the 'epistemically explanatory relation.' What type of relation is it? Multiple answers are possible, of which two are at-hand. Remember the two senses of epistemic factor. In the broad sense, an epistemic factor is that which contributes to the turning of *S*'s belief into knowledge. **Epistemic Position** is a way to flesh out this sense of epistemic factor. With **Epistemic Position**, an answer to the question is the following; one or a set of *S*'s epistemic factors is *epistemically explanatory* if it makes the strength of *S*'s epistemic position regarding *p* sufficiently high. This is similar to the I-S model's high probability requirement for statistical explanation. Another answer follows from the narrow sense of epistemic factor: a factor is epistemic iff it make *S*'s epistemic position regarding *p* higher than otherwise; that is, a factor is *epistemically relevant* iff instantiating it makes a difference in strength of epistemic position.[9] This is similar to the S-R model's statistical-causal relevance. Gettier cases are precisely cases where certain epistemically relevant factors fail to be epistemically explanatory.[10] I will elucidate this point in more

[9] Note that my notion of epistemic relevance is different from Salmon's. He defines an epistemically homogeneous class as one which, for all one knows, is objectively homogeneous. In my definition, any reference class, used in the epistemic context, is epistemic, because it is an epistemic factor. It may or may not be defined by reference to one's knowledge or evidence.

[10] For this reason, I do not think that the epistemically explanatory relation is an epistemic factor. Its absence in Gettier cases just means that some epistemic factor, i.e., an anti-Gettiering factor, is missing, or *S* fails to instantiate it, in such cases. A similar point is made in Fantl & McGrath (2009b, p. 174). Sosa (2007), unlike in his preceding writings, does not give safety a pride of place as an anti-Gettier device, since, it seems, he thinks that the explanatory requirement suffices to deal with Gettier cases. However, he does not specify any condition under which the explanatory relation holds between *S*'s type reliability and *S*'s believing truly. It

detail in §3. 7; to preview, in one of the Gettier cases, the Barn Façade case, the reason that epistemically relevant factors fail to explain why *S* believes truly that there is a barn before her is that *S* is surrounded by convincing barn facades. The proximity of such barn facades is far from having causal effects on the production of *S*'s belief and its truth.

If this is right, it is plausible that the locus of the theory of knowledge is not objective-causal context, rather it is epistemic. For this reason, Salmon's solution, as it is, is not applicable to the reference class problem in the theory of knowledge; at the very least, the notion of statistical-causal relevance needs to be replaced with the notion of epistemic relevance defined above. On the other hand, Hempel's solution is thoroughly epistemic in the sense that it identifies the relevant reference class with the one that one is rational to or justified in choosing. Theorists of epistemic probability have proposed their solutions to the reference class problem along the lines of RMS.[11] However, it is not without problems.

---

remains a possibility that it holds if or/and only if the safety condition is satisfied by *S*. Furthermore, our intuitive sense of epistemic explanation is not decisive to figure out the condition for the epistemically explanatory relation; for example, Sosa (2007) holds that *S* satisfies the explanatory requirement in the Barn Façade case, whereas Greco (2003, 2009, 2010) holds that *S* does not. Both judgments are based on their intuitions, and apparently they have different intuitions about the same case.

[11] As noted in footnote 34 of Chapter 2, Kyburg's solution to the reference class problem for the epistemic probability theory is of this type, though it also relies on his idiosyncratic concept of probability, viz., that probabilities take an interval-value rather than a number-value. Another prominent figure in this field is Pollock (1990), whose solution is similar to Kyburg's in some but not all respects. His solution hinges on what he calls 'nomic probability.' I am not concerned with their solutions, mainly because, besides the problems discussed in Chapter 2 and 3, (i) they are dependent on *sui generis* notions of probability, (ii) I

I will discuss, in more detail, the solutions to the reference class problem for type reliabilism based on RMS and RBH, in §3. 2 and 3. 3, respectively.

§3. Alleged Solutions to the Reference Class Problem

The solution to the reference class problem has been sought both in probability theory and the theory of knowledge. In the theory of knowledge, since the problem has been thus far mainly associated with type reliabilism, most alleged solutions are designed to militate against the variants of the reference class problem that are specific to type reliabilism. Such solutions are, *mutatis mutandis*, applicable to other theories. In this section, I examine whether the existing solutions succeed at capturing the epistemically relevant reference class. The solutions to be examined are: (a) psychological realism, (b) maximal specificity, (c) broadest homogeneity, (d) normalcy, (e) perceptual equivalence, (f) task sensitivity, and (g) contextualism. (a) is based on Reichenbach-Hempel's proposal, and (b) on Salmon's.

§3. 1. Psychological Realism

The reference class problem for type reliabilism concerns which process type among many is assigned to a process token in a given case. Then, the problem itself

---

doubt that they can evade the problems put forth below, and (iii) they are inadequate to deal with the epistemological problems to be discussed in Chapters 5 and 6.

presupposes some way in which a process token is typed. Goldman (1979) proposes that processes are functionally typed:

> We need to say more about the notion of a belief-forming ‘process’. Let us mean by a ‘process’ a *functional operation* or procedure, i.e., something that generates a *mapping* from certain states – ‘input’ – into other states – ‘output’. The outputs in the present case are states of believing this or that proposition at a given moment. (p. 115)[12]

A belief-forming process token is characterized as a series of causally connected predoxastic state tokens. To follow Adler & Levin (2002), assume that the *i*th predoxastic state token in the causal chain involved in the process token belongs to a type, and it is represented as a function F_i. Then, $<F_1, \ldots, F_n>$ is a function corresponding to a belief-forming process type, and it is realized by a process token resulting in a belief token *b* iff, for some function f, *b* is $f<F_1, \ldots, F_n>$; if a process type involves a belief type as output, it is represented as $<F_1, \ldots, F_n, f<F_1, \ldots, F_n>>$. Alston (1995, 2005, ch. 6) holds that there is a unique function that any belief-forming process token psychologically realizes, and that it is the task of psychology to find such a unique function. He offers this idea as a solution to the reference class problem, under the heading of ‘psychological realism’:

[12] Similarly, Sosa (1991b) defines an intellectual faculty or virtue type as a function of propositions, fields, and environments.

> (Psychological Realism) For any belief-forming process token *g*, the relevant process type for *g* is the function *f* that *g* psychologically realizes.

Alston combines **Psychological Realism** with Hempelian RMS, and we will discuss it in the next section. But even before proceeding, it must be clear that **Psychological Realism** itself faces some difficulties: first, its success as a solution hinges on psychological research, but it is doubtful that psychology alone is adequate to fulfill the purpose of function identification (I will make a case for this doubt in §3. 6); second, as Alston is aware, it, at best, is a solution to the process type problem, not touching on the environment type problem[13] (he appeals to some sort of normalcy condition for a solution to this problem. This type of solution is discussed in §3. 3).[14]

---

[13] This may be a consequence of Goldman's version of type reliabilism that Alston is committed to. Goldman, in proposing his type reliabilism, makes explicit that the process of interest is 'cognitive':

> Clearly, the causal ancestry of beliefs often includes events outside the organism. Are such events to be included among the 'inputs' of belief-forming processes? Or should we restrict the extent of belief-forming processes to '*cognitive*' events, i.e., events within the organism's nervous system? I shall choose the latter course, though with some hesitation. (Goldman, 1979, p. 116)

Goldman, however, does not strictly abide by this restriction in this and subsequent writings (he concedes his violating the restriction in his (2008).

[14] Vahid (2009) raises another objection against **Psychological Realism**: Alston is merely replacing the reference class problem with a problem concerning indeterminacy of perceptual content. The latter problem is identical with the problem of the speckled hen. However, it is not quite clear why **Psychological Realism**

§3. 2. Maximal Specificity

Alston's solution to the reference class problem for type reliabilism is two-tiered, in that **Psychological Realism** is designed to cooperate with RMS as Hempel proposed it: "the functions in question are maximally specific, in that any difference in input that is registered by the function indicates a different function" (Alston, 1995, p. 26). Similarly, Goldman identifies the relevant process type with 'the *narrowest* type that is *causally operative* in producing the belief token in question' (Goldman, 1986, p. 50). The Hempelian strategy, then, may be formulated as follows:

(Maximal Specificity) For any belief-forming process token *g*, the relevant process type for *g* is the maximally specific type.

It is not clear how **Maximal Specificity** works against the reference class problem for type reliabilism. As we have seen above, the Hempelian maximally specific type is epistemic, and relativized to the total evidence. Hempel does not specify whose total evidence is relevant. Presumably, it is the total evidence of a scientific community. But, even if it is assumed that the total evidence relevant for the typing of cognitive process is that which is shared by the whole community of psychologists, it is not guaranteed that each process token is associated with a unique type – psychologists may disagree on what

---

cannot extend to this problem, since, as discussed in §3. 3. 2 of Chapter 3, it is part of the problem concerning the typing of a belief-forming process token.

the unique type is. And, if the notion of total evidence is personalized, as in epistemic probability theories, some problems still remain.

First, the Hempelian construal of **Maximal Specificity** entails that type reliabilism must abandon its objective orientation. As elucidated above, the why-question with which Hempel's RMS is concerned is why it is rational to assign a value to definite probability by selecting a certain reference class. This strategy entails that unless subjective factors, such as one's knowledge or evidence, are invoked, the relevant reference class cannot be determined. If this is the case, the relevant reference class may vary from person to person. Thus, this strategy is devastating for the objective tenet of type reliabilism.

Second, Gillies (2000, p. 121) points out that **Maximal Specificity** does not work when more than one maximally specific type exists[15]; one may have in one's total evidence statistical data for each of $Freq(F \mid G_1)$ and $Freq(F \mid G_2)$, but not for $Freq(F \mid G_1 \wedge G_2)$, where the sizes of G_1 and G_2 are equal. In such a case, the relevant reference class is not determined.

It may be argued that these problems stem from the epistemic rendering of **Maximal Specificity**, but it can be rendered objectively. Granted, then, the maximally specific type is simply either the type containing the minimal number of tokens, or the extension of the strongest predicate in a language, perhaps, psychological language. This way of developing **Maximal Specificity**, however, is a *non sequitur*. It straightforwardly falls prey to the problem of trivializing reference class; there is no distinction between a type and a token, since the maximally specific type in this sense is that which only the

[15] Gillies attributes this problem to David Corfield and Jon Williamson.

token in question can instantiate.[16] Hempel's RMS contains clause (i), viz., that '*Gx*' entails neither '*Fx*' nor '~*Fx*,' precisely to avoid this problem. Even if its *ad hoc* flavor is erased, (i) does not exclude many narrow reference classes relative to which the frequency of true beliefs is close to 0 or 1.

§3. 3. Broadest Homogeneity

Beebe (2004) proposes a two-tiered strategy against the reference class problem, one component of which is Salmon's RBH (another is **Task Sensitivity**, the topic of §3. 6)[17]:

(Broadest Homogeneity) For any belief-forming process token *g*, the relevant process type for *g* is the broadest homogeneous type.

[16] Conee & Feldman (1998) argue against Alston and Goldman along this line. Although Adler & Levin (2002) attempt to refute their argument, it is not clear at all how their appeal to variables or functions rather than values is helpful here. This problem for **Maximal Specificity** is classical anyway (see, e.g., (Kyburg, 1970a)). Conee & Feldman address Adler & Levin's refutation in Feldman & Conee (2002). The debate between Conee & Feldman and Adler & Levin is reviewed at length in §3. 6 below and §2. 2 of Chapter 7.

[17] Adler (2005) utilizes **Broadest Homogeneity** as the best available strategy for type reliabilism, though his intention is to show that even with the best strategy, type reliabilism fails to handle his counterexample. In other words, his counterexample does not hold if there is good reason to deny **Broadest Homogeneity**. His counterexample is a lottery case, which is the topic of §2 of the next chapter.

Beebe uses **Broadest Homogeneity** to narrow down the range of candidate process types in Conee & Feldman's (1998) list mentioned in §3. 2 of Chapter 3:

(1) process of a retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby

(2) process of relying on a leaf shape to form a tree-classifying judgment

(3) the visual process

(4) vision in bright sunlight

(5) perceptual process that occurs in middle-aged men on Wednesdays

(6) process which results in justified beliefs

(7) perceptual process of classifying by species a tree behind a solid obstruction

Conee and Feldman include (5), an obviously irrelevant process type, in order to evince the predicament of type reliabilism.

For a preliminary stage, Beebe defines the properties causally involved in processes (1) through (7) as follows:

D: the property of occurring up close to a perceived object

O: the property of occurring when the surface of a perceived object with the greatest area is oriented perpendicularly to the line between subject and object

L: the property of occurring in a high degree of ambient lighting, the property of occurring after less than one second of exposure to a perceived object

W: the property of occurring on Wednesdays

M: the property of occurring in middle-aged men

B: the property of producing a true belief.

Then, he suggests that that the following are true of these properties:

$$P(B \mid A \wedge D) \neq P(B \mid A \wedge D \wedge O) \neq P(B \mid A \wedge D \wedge O \wedge L)$$

$$P(B \mid A \wedge D) \neq P(B \mid A \wedge D \wedge E)$$

$$P(B \mid A \wedge D \wedge O) \neq P(B \mid A \wedge D \wedge O \wedge E)$$

$$P(B \mid A \wedge D \wedge O \wedge L) \neq P(B \mid A \wedge D \wedge O \wedge L \wedge E)$$

$$P(B \mid A \wedge D \wedge O \wedge L) \neq P(B \mid A \wedge {\sim}D \wedge {\sim}O \wedge {\sim}L)$$

$$P(B \mid A \wedge D \wedge O \wedge L) = P(B \mid A \wedge D \wedge O \wedge L \wedge W \wedge M)$$

This means that "factors D, O, L and E are statistically relevant but not that W and M are" (Beebe, 2004, p. 189), and Beebe concludes that type reliabilism, with **Broadest Homogeneity**, has a resource to rule out process type (5).

This may seem a promising line of response at first glance. However, it has some flaws. First, it is not clear why **Broadest Homogeneity**, construed as Salmonian RBH, is required to rule out (5). W is further divided to W_1: occurring and W_2: being on Wednesdays, and W_2 is not a causal factor. Then, W is not objective in Salmon' sense, and hence to be excluded at the outset in the S-R model. As I argued above, not every causal factor is epistemic, but **Broadest Homogeneity** may be still important, once it is combined with the notion of epistemic homogeneity rather than that of objective homogeneity. Second, more importantly, **Broadest Homogeneity**, just as **Maximal Specificity**, has the notorious difficulty with the problem of trivializing reference class. Salmon (1984) attempts to circumvent the problem by requiring that every factor in the relevant reference class *G* be temporary prior to those in the target class *F*. This requirement may make sense in the context of causal-statistical explanation, but may not in the context of epistemic explanation. Even worse, Salmon (1989) concedes that **Broadest Homogeneity** cannot evade the problem of trivializing reference class so easily, since many causally relevant but trivializing factors are up for grabs.[18] Similarly, epistemically relevant but trivializing factors are abundant. For example, the spatio-temporal location of *S*'s believing *p* or its neighborhoods would be such factors.

§3. 4. Normalcy

Alston (1986, 2005), Goldman (1986), and Sosa (1991c) suggest that the reliability of *S*'s process type is relativized to a set of normal worlds or environments, regardless of what world or environment *S* is in. This suggestion can be reckoned to be a solution to the environment type problem for type reliabilism.

> (Normalcy) For any belief-forming environment token *g*, the relevant environment type for *g* is the normal type.

[18] Salmon (1984) also makes use of von Mises' selection function, in order to avoid the problem of trivializing reference class. For problems with this strategy, see White (1988).

How is the normal world or environment type defined? Goldman regards the set of normal worlds as those "consistent with our general beliefs about the actual world" (ibid., p. 107); Sosa (ibid., p. 142) specifies the normal environment relative to fellow members of the epistemic group to which *S* belongs.[19] It seems what Sosa is getting at is not so different than Alston's idea that the relevant environment is "a wide range of situations *of the sort we typically encounter*" (2005, p. 124). That is, the relevant environment type is that which members of *S*'s community normally or typically encounter. All of these figures are aware that what counts as normal is vague (Goldman), context-dependent (Sosa), and may vary with cultural, technological, or other changes (Alston).[20] Then, one problem with **Normalcy** is that it introduces irredeemable subjectivity into the typing of environments, and may well run afoul of naturalism central for Goldman's and Alston's type reliabilism.[21]

---

[19] Sosa also offers a different line of response to the reference class problem: the relevant process/environment type is what *S* is aware of as exercising (Sosa, 1991a, p. 274). This is based on the same sprit as Comesaña's solution, as I interpret it in §3. 5 of Chapter 3; both exploits the internalist element, **Accessibility**, to narrow down the range of candidate reference classes. Greco (2004c) demurs that it is psychologically unrealistic that *S* always has access to what process/environment type is being exercised. Moreover, as with Comesaña's, Sosa's strategy fails to deal with the Optimist case.

[20] If typicality here is construed probabilistically, e.g., as a high encounter rate, the vagueness of normalcy is a manifestation of the reference class problem.

[21] Goldman (1986) maintains that even if this is the case, it is still an empirical question whether a process is reliable in the normal world. To defend naturalism, Goldman, in his (1992a, 2008), makes a further move, by distinguishing between the theory of knowledge and the theory of knowledge attribution. He seems to think that subjective elements only enter the latter. I am not sure what this distinction amounts to. Perhaps, what he means by the theory of knowledge attribution is a psychological account of what factor induces the

As with **Maximal Specificity**, then, **Normalcy** results in the partial denial of the objective orientation of type reliabilism.

§3. 5. Perceptual Equivalence

Goldman's notion of perceptual equivalence is already mentioned in §3. 3. 1 of Chapter 3. More precisely, Goldman defines perceptual equivalence as follows: "[a] *perceptual equivalent* of an actual state of affair is a possible state of affairs that would produce the same, or a sufficiently similar, perceptual experience" (Goldman, 1976, p. 92). He, then, further adds that "[i]f the percept produced by the alternative state of affairs would not differ from the actual percept in any respect that is causally relevant to *S*'s belief, this alternative situation is a perceptual equivalent for *S* of the actual situation" (ibid., p. 94).[22] Perceptual equivalence is relativized to a person and a time, since what counts as a perceptual equivalent may vary inter-subjectively and intra-subjectively over time.

---

attributor to utter a knowledge attribution or denial, and has little to do with the truth value of knowledge attributions. But then, the reference class problem for type reliabilism as a theory of knowledge is completely left out.

[22] This passage is followed by Goldman's exposition of the original Dachshund case in which wolves, not hyenas, are abundant nearby:

> Consider now the dachshund wolf case. The hypothetical percept produced by a wolf would differ from Oscar's actual percept of the dachshund in respects that are causally relevant to Oscar's judgment that a dog is present. Let me elaborate. There are various kinds of objects, rather different in shape, size, color, and texture, that would be classi fied by Oscar as a dog. He has a number of

Goldman's idea is originally part of his theory of token reliability, but can be utilized for the sake of type reliabilism:

> (Perceptual Equivalence) For any perceptual belief-forming process or method token *g*, the relevant process or method type for *g* is the type containing as input all the perceptual equivalents of the perceptual experience involved in *g*.[23]

For example, in the simplest case, a process token, which leads to a token belief *b* from perceptual experience token *o* as of an object, belongs to type <*O*, *B*>, where *O* is the type containing all the perceptual equivalents of *o*, and *B* is a belief type instantiated by *b*.[24]

> visual 'schemata', we might say, each with a distinctive set of features, such that any percept that 'matches' or 'fits' one of these schemata would elicit a 'dog' classi fication … Now although a dachshund and a wolf would each produce a dog-belief in Oscar, the percepts produced by these respective stimuli would differ in respects that are causally relevant to Oscar's forming a dog-belief. Since Oscar's dachshund schema includes such features as having an elongated, sausage-like shape, a smallish size, and droopy ears, these features of the percept are all causally relevant, when a dachshund is present, to Oscar's believing that a dog is present.

[23] McEvoy (2005) proposes this as a solution to the reference class problem for type reliabilism. Becker (2008) appeals to it in solving the reference class problem for token reliabilism.

[24] Goldman more finely characterize *o* as consisting of three parameters: an object, its non-relational properties, and what he calls a 'distance-orientation-environment relation' (DOE). A DOE is 'a conjunction of relations or properties concerning distance, orientation, and environmental conditions' (ibid., p. 93). He, however, gives nothing more than a rough sketch of how to specify a DOE in any given case.

This way of specifying the relevant process type or individuating the method token is what I called the 'internal-external criterion' for method individuation.[25]

It is not obvious whether **Perceptual Equivalence** circumvents the problem of the speckled hen; everything hinges on what experience counts as 'the same, or a sufficiently similar, perceptual experience.' But even if sufficient similarity is elucidated in some way or another, it still faces a problem: **Perceptual Equivalence** is too narrow a criterion for typing or individuating perceptual experience. Gendler & Hawthorne (2005) offer several examples to illustrate this point, one of which is the following:

> (Red Paper) Agent Orange is insensitive to certain subtleties of shading whereas Colonel Mustard is not; there will be cases where Colonel Mustard will know that he is seeing a red piece of paper, whereas Agent Orange will not – even though Mustard and Orange never disagree about whether a sheet of paper is red and thus even though neither is more easily deceived, neither more reliable in redness verdicts, than the other. Suppose that the two are sitting side-by-side. In front of them is a piece of paper of the shade red-36, surrounded by pieces of white paper that have been illuminated to look as if they are of the shades red-32, red-34, and

[25] The internal-external criterion is more congenial to externalism than to internalism, in two respects: first, *o* contains a distal object of perception and its properties, and thereby the internal-external criterion partially denies **Mentalism**; second, it also partially denies **Accessibility**, not only because the denial of **Mentalism** entails the denial of **Accessibility**, but also, on Goldman's account, *S* does not need to know what process or method *S* is using.

> red-38. Casting their gazes on the red-36 sheet, Colonel Mustard and Agent Orange both form the judgment: there is a red piece of paper before me. (p. 343)

Shades red-32, red-34, red-36, and red-38 are equivalent percepts for Agent Orange, whereas shade red-36 is not perceptually equivalent to the other shades for Colonel Mustard. Then, Agent Orange's perception of the paper of the shade red-36 is a token of process type <{shades red-32, red-34, red-36, red-38}, there is a red piece of paper>, whereas Colonel Mustard's is a token of process type <shade red-36, there is a red piece of paper>. Relative to this particular situation, the former is much less reliable than the latter. But the strengths of their epistemic positions seem the same, since, on the intuitive level, no epistemically relevant difference exist between them, and neither Agent Orange nor Colonel Mustard knows that there is a red piece of paper. That is, **Perceptual Evidence** does not satisfy **Accordance**.

§3. 6. Task Sensitivity

As mentioned above, Beebe's solution to the reference class problem is a two-pronged approach. In particular, he tries to handle Feldman's Single-Case Problem with **Broadest Homogeneity**, and the No-Distinction Problem with what he calls the 'tri-level condition':

(Tri-Level Condition) The reliability of a cognitive process type *T* determines the justification of any belief token produced by a cognitive process token *t* that falls under *T* only if all of the members of *T*:

a) solve the same type of information-processing problem *i* solved by *t*;

b) use the same information-processing procedure or algorithm *t* used in solving *i*; and

c) share the same cognitive architecture as *t*. (Beebe, 2004, p. 180)

(a) to (b) correspond to David Marr's three levels of theory or explanation of cognitive process: (a') computational, (b') algorithmic, and (c') implementation level. Since Marr's theory is important for my own solution to the reference class problem offered in Chapter 7, I here expound it in my own terms. According to Marr, cognitive process is identified with information-processing or computational process, and any computational explanation of a cognitive process must proceed top-down, (a') through (c') *via* (b'). Each level is characterized by what question the explanation of the cognitive process purports to answer at that level:

a') Computational: What is the goal of the computation, why is it appropriate, and what is the logic of the strategy by which it can be carried out.

b') Algorithmic: How can this computational theory be implemented? In particular, what is the representation for the input and output, and what is the algorithm for the transformation?

c') Implementation: How can the representation and algorithm be realized physically? (Marr, 1982, p. 25; each name is slightly changed)

Advantages of the top-down approach may be easy to see by comparing it with the bottom-up approach. A cognitive process type can be multiply realized in different physical systems, and thereby starting with the implementation level does not contribute to identifying the cognitive process type. The identification of process type at the algorithmic level is functional, i.e., it specifies the process type as the function realized by the cognitive system. The problem is that the algorithmic level alone cannot identify what function is realized by a physical system (ibid. pp. 25-6)). Sprevak (forthcoming) nicely illustrates this point by the following example[26]:

> Consider two computational systems that perform the same numerical calculation. Suppose that one system takes ink-marks shaped like Roman numerals (I, II, III, IV, …) as input and yields ink-marks shaped like Roman numerals as output. Suppose that the other system takes ink-marks shaped like Arabic numerals (1, 2, 3, 4, …) as input and yields ink-marks shaped like Arabic numerals as output. Suppose that the two systems compute the same function, for example, the addition function. What could this equivalence consist in? There need be no physical or functional similarity between their respective inputs and outputs. (pp, 23-4, online version)

[26] For more examples, see (Shagrir, 1999).

Sprevak argues that there is no way to determine whether the two systems realize the same function or not, without fixing what task each system purports to perform. In other words, the functional identity of process is only determined at the computational level, taking account of what problem-solving task it purports to perform. According to Marr, the task of computational system is only specified in terms of semantic or intentional representation (note that what he calls 'representation' in the quote above is merely syntactic, and I do not use 'representation' in his idiosyncratic sense). When two systems manifest the same psychological and behavioral discriminatory patterns, the process typing or the functional identification of each system must take account of what is represented by the inputs and the outputs of the process or function (see (Shagrir, 1999, 2001) and (Shapiro, 1995, 1997)). This is why any computational theory must start from the top, computational level.[27]

Thus, Beebe's tri-level condition amounts to the following proposal against the reference class problem:

---

[27] What has been labeled 'the received view' is that computation must involve representational or intentional content; more precisely, the computational identity of process type cannot be determined without representational contents involved in process. Opponents of the received view, such as Egan (1991, 1992, 1994, 1995) and Piccinini (2008), argue that only syntactically specified function and algorithm enter the computational identity of process type. The debate concerning the received view, in part, concerns exegetic issues of Marr's theory, and it is discussed by the following participants of the debate: besides Egan and Piccinini, Burge (1986), Butler (1998), Davies (1992), Kitcher (1988), Morton (1993), Segal (1989, 1991), Shagrir (1999, 2001), Shapiro (1995, 1997), and Sprevak (forthcoming). I assume the received view throughout this thesis, since it is implicit in epistemological theories.

> (Task Sensitivity) For any belief-forming process token *g*, the relevant process type for *g* is determined in part by what problem-solving task *S* purports to perform by *g*.

Task Sensitivity is no doubt relevant for the solution to the reference class problem. Consider again the list of processes (1) through (7). A debate pertaining to the reference class problem concerns the question of whether they are of the same type; on the one hand, Feldman & Conee (2002), who propose the list in Conee & Feldman (1998), hold that the description of each process encodes a different function, and thereby each corresponds to a different type; on the other hand, Adler & Levin (2002) argues that the list merely describes the same function and process type in different levels of detail. Beebe points out that this problem cannot be settled without considering what task each process is performing. If the processes are taken to carry out the same task, they are categorized as being of the same type. Beebe is definitely right; this is a moral to be drawn from the Two System case.

Similar solutions to the reference class problem are proposed by Baergen (1995) and Wallis (1994).[28] Conee & Feldman make a case against Baergen's solution, to the effect that the computational type, revealed by psychological research, may not satisfy **Accordance**. The psychologist may choose to type a cognitive process narrowly as a result of fine-grained specification of the relevant task. Then, type reliabilism may fail to yield intuitive predications of knowledge or justification for cases like the Red Paper case.

---

[28] Baergen explicitly mentions Marr's theory of vision. Wallis's argument is based on the adaptive nature of cognitive system (cf. (Kornblith, 2002)). On his view, the relevant task is such that *S* encounters it repeatedly and develops a strategy or method against it in a certain range of environments. Then, what strategy *S* is using in what environment is relevant for both the process type problem and the environment type problem.

Moreover, appeal to **Task Sensitivity** may not determine a unique type for a given process token. This is indeed another moral Sprevak draws from his Two System case; there is no reason to prefer the typing of the two systems as {number-processing} to the typing of them as {Roman numerals-processing} and {Arabic numerals-processing}, respectively. He offers a solution to this problem, and it is discussed at length in §2. 2 of Chapter 7 in connection with my contextualist solution to the reference class problem.

§3. 7. Contextualism

Mark Heller offers a type of solution to the reference class problem by relying on contextualism. In fact, he constructs two different versions of contextualist solution, one for type reliabilism in his (1995b) and another for a variant of sensitivity theory in his (1999).[29] The arguments for both versions are based on linguistic considerations.

First, he remarks about the ordinary concept of reliability:

> "Reliable" is a perfectly ordinary word that is in perfectly ordinary situations is applied to tokens which are instances of several types, where those types have different degrees of reliability. Yet we somehow manage to use this word without difficulty in ordinary discourses. Just as our use of the term in ordinary discourse is context relative, reflecting the different concerns of different speakers on different

[29] Baumann (2005) also proposes a contextualist solution to the reference class problem for type reliabilism, while he argues that token reliabilism has no way-out of the reference class problem. I do not see any reason that the contextualist apparatus only works for type reliabilism, but not for token reliabilism.

> occasions of use, "reliable" is also context dependent in epistemological discourse. (Heller, 1995b, p. 502)

Since this argument is linguistic, the context in question is the context of use, where an attributor utters a reliability statement. Heller generalizes this point to the contexts of knowledge attribution: insofar as type reliability is an epistemic factor, i.e., '*S* knows *p*' entails '*S*'s process token leading to believing *p* is sufficiently reliable,' the context where an attributor utters a knowledge ascription or denial contains the contextual factors that determine what process type and/or environment type is relevant in that context.[30]

Second, Heller (1999) correctly points out that the modal base and the ordering source for a counterfactual are both context-sensitive. As a matter of fact, this is built into Lewis-Kratzer's ordering semantics for modals. Hence, it is quite plausible that the counterfactual condition of sensitivity is determined by contextual factors. The difference between the two solutions is of little importance here, and so I put them together in the following way:

> (Contextualism) For any belief-forming process token, method, or environment *g*, the relevant type for *g* is contextually determined relative to a context of use.

The contexts of use in which knowledge attributions or denials are uttered all belong to the epistemic context defined in §2, but are more finely individuated. The

[30] Greco (2009, 2010) notes that Heller's solution is not restricted to contextualism, and SSI can militate against the reference class problem. It seems that Greco misses the fact that Heller's arguments are linguistic.

epistemic context is the context of inquiry the purpose of which is to evaluate one's epistemic performance or practice; whereas, a particular context of knowledge attribution may be governed by more local interests and purposes, e.g., an attributor in a context of use may be interested in evaluating only a certain aspect of *S*'s epistemic score.

In §2, I distinguished the epistemic context from the causal context; the latter is guided by interests in objective probabilities and probabilistic explanations. Heller does not pay much attention to the difference between the two contexts of inquiry. In his (1989) and (1999), Heller argues that Nozick's sensitivity theory is vulnerable to Goldman-Ginet's Barn-Façade case[31]:

> (Barn-Façade) Henry spots a barn in front of him while he is driving. Then he believes that there is a barn before him. However, unbeknownst to him, Henry is in the Barn Façade County that contains numerous barn-façades. He would mistakenly take these to be barns if he were looking at them. But the barn he spotted is the only real barn therein.

The majority of proponents and even critics of sensitivity theory reckon sensitivity to yield the right verdict that Henry does not know that it is a barn, on account of his failure to satisfy the sensitivity condition. Heller, however, disagrees, because, he contends, the counterfactual relevant for the Barn-Façade case, 'if there were not a barn there, *S* would not believe that there was,' is true: "[i]f the barn had not been there, the facsimiles would

[31] This case is first presented in Goldman (1976), and he attributes it to Carl Ginet. I will discuss the Barn-Façade case in more detail in §3 of Chapter 5.

still have been just where they are. The most similar worlds in which the barn is not there are ones in which there is no facsimile there either" (Heller, 1989, 26). In the way he envisages the Barn-Façade case, the non-*p* worlds closest to the actual world are those worlds in which no structure exists in front of Henry; and, of course, Henry does not believe in those worlds that there is a barn before him. If this is correct, it follows that Henry satisfies the sensitivity condition.

On the basis of this consideration, Heller proposes to modify the sensitivity condition, in the following way: necessarily, if *S* knows *p* at *w*, then *S* does not believe *p* in the non-*p* worlds sufficiently close to *w*. The set of sufficiently close non-*p* worlds is more inclusive than the set of closest non-p worlds, and hence the sensitivity condition thus modified works against the Barn-Façade case.[32] This is precisely the move Goldman (1986) makes against counterexamples to **Sensitivity'**, as we have seen in §3. 3. 1 of Chapter 3. Goldman's move, indeed, is motivated by his considerations of the Barn-Façade case and other similar cases. As I argued, the moral to be drawn from these cases is that the epistemic ordering source is to be different from other, more objective ordering sources, not

[32] The Barn Façade case may be described as involving Henry's belief that *it* is a barn, not that there is a barn. Since 'it' is a demonstrative and a rigid designator, the referent of 'it' is fixed across possible worlds. But, in order for sensitivity to apply to the Barn Façade case, Henry's (more precisely, his counterparts') beliefs of the form 'it is a barn' must be about different objects in different worlds. Otherwise, the sensitivity condition is trivially satisfied, merely because Henry does not believe the same proposition in other worlds as he does in the actual world. Given this, the sensitivity condition (and the safety condition as well) needs to be reformulated in terms of *p and similar propositions* rather than *p* alone: roughly, if *p* were false, *S* would not believe *p* or other similar propositions. Since this point is not of much importance, I freely switch between 'it is a barn' and 'there is a barn,' when I discuss the Barn Façade case and its variants.

that the sensitivity condition, as Nozick formulates it, is flawed.[33] It is strange that Heller makes this move, since he endorses the idea articulated in DeRose's counterfactual robustness theory, viz., the closeness of the possible worlds across which *S*'s belief accords with the fact regarding *p* is the measure of *S*'s strength of epistemic position regarding *p*. It seems that he is still captivated in the objective, non-epistemic context of inquiry.

With this caveat, **Contextualism** seems to be on the right track. One advantage it has over other proposed solutions is that it can integrate the virtues of them. On contextualism, the attributor factors that contribute to the determination of the relevant reference class include what is salient to the attributor and what interests she has, with regard to the choice of the relevant reference class. When more detailed description of *S*'s

---

[33] Another way to make this clear is to compare the sensitivity theory with the counterfactual analysis of token causation. On the counterfactual analysis, roughly, token event *A* causes token event *B* iff *B* would not happen, if *A* were absent. Now, suppose that we want to know whether the presence of the barn before Henry has caused his believing that there is a barn. Obviously, this causal relation holds. The counterfactual analysis of token causation gets the right result here, since the counterfactual 'Henry would not believe that there is a barn if there were not a barn' is true in this context, with the ordering source set by our interest in the token causation. But most epistemologists regard the very same counterfactual to be false, and claim that Henry's belief is insensitive. This strongly suggests that the context they are in is governed by different contextual factors than in the causal context. Heller is conflating between the two contexts. Moreover, his ignoring the contextual difference leads to a far-reaching result: Heller (1995a) proposes the same type of modification against Lewis' (1979) counterfactual analysis of token causation as he does against the sensitivity condition. But the modified analysis entails that the counterfactual analysis of token causation fails in the case at hand. He cannot have a cake and eat it too, unless he accepts the difference in context. Cohen (1988) and Swain (1978) also appeal to objective probability to deal with the Barn Façade case. This objection is equally applied to their theories of knowledge.

situation is available, narrower reference classes become salient, and hence the contextualist solution accommodates the main idea of **Maximal Specificity** and **Broadest Homogeneity**. Attributor's interests are likely to be governed by the sense of normalcy, unless she has interests quite different than usual; so **Normalcy** is congenial to the contextualist solution. Moreover, the problem facing **Maximal Specificity**, **Broadest Homogeneity**, and **Perceptual Equivalence**, that they may end up with too narrow a reference class, can be overcome if they are combined with contextualism; the determinants of the relevant reference class articulated in these solutions are constrained by attributor's interests; too narrow classes are usually out of her interests, and even if they are, contextualism can accommodate such classes – this line of thought is developed in §2. 3 of Chapter 7.

I will offer my thorough defense of a version of contextualist solution to the reference class problem in the next three chapters. The gist of my solution consists of the combination of **Task Sensitivity** and **Contextualism**; the difficulties with the former are surmounted once they are jointly put to work. The solution appeals to psychological attributor factors, and hence it involves subjective elements. Such subjective elements, in part, undermine the objective orientation of any substantive theory of knowledge to which the contextualist solution is applied. In §3. 2 of Chapter 7, I will inquire as to whether or to what extent this is crucial for the theory of knowledge in general.

Chapter 5: Shifts in Reference Class (i): the Lottery Case and the Gettier Case

Introduction

The reference class problem tells us that there are many reference classes to which the strength of *S*'s epistemic position is relativized. This means that even if one or a set of *kinds* of epistemic factors *F*s, such as type reliability, token reliability, or internalist evidence, are selected by a substantive theory of knowledge as determinants of the strength of *S*'s epistemic position, it is indeterminate how strong an epistemic position *S* is in, since the relevant reference class still needs to be specified. Also, we have seen, in Chapter 4, that many solutions to the reference class problem indicate that the choice of the relevant reference class is at least in part subjective. If this is true, the objective orientation of the theory of knowledge is undermined.

Furthermore, if the choice of the relevant reference class is (even in part) a subjective matter, it may not be the case that there is a unique legitimate reference class for a given triple of *S*, *p*, and *F*. For what reference class counts as relevant may vary inter-personally or even intra-personally. The non-uniqueness of the relevant reference classes can be evinced, at least *prime facie*, if (i) multiple reference classes determine different strengths of epistemic position for a certain triple of *S*, *p*, and *F*, and (ii) the strengths of epistemic position thus determined are each intuitively correct. (ii) should be treated with care: we should not have two different intuitions about *S*'s strength of epistemic position regarding *p* at the same time (if we do, we contradict ourselves); (a) it suffices if we have different intuitions from case to case while a triple of *S*, *p*, and *F* is held fixed. In addition, given that the notion of epistemic position is a theoretical posit, it is not to be expected that our intuitions track precisely the strength of epistemic position. Rather, (b) the relevant cases are those across which our intuitions about knowledge and/or justification attribution regarding the same triple vary.

(a) and (b) constitute a minimal condition for (ii), but it may be too minimal; satisfying (a) and (b) does not exclude the possibility that our varying intuitions are induced by some other shift than shift in reference class. Standard contextualism and SSI appeal to shifts in epistemic standard to explain why our intuitions about similar cases change while a triple of *S*, *p*, and *F* is held fixed. Thus, what is needed to reach (ii) *via* (a) and (b) is an explanation of the relevant cases according to which shifts in reference class underlie our differing intuitions. As a preliminary step for my argument, fully given in Chapter 7, for a contextualist solution to the reference class problem, I will show that such an explanation is possible for a wide range of epistemologically interesting cases, even though, for some of them, SSI or standard contextualism is reckoned to offer the best explanation.

This and the next chapter are devoted to the diagnosis of four kinds of cases: lottery cases, Gettier cases, cases of easy knowledge, and cases of Cartesian skepticism. I argue that shifts in reference class are responsible for our varying intuitions about some cases of these kinds. This chapter focuses on the former two. First, in §1, I introduce a variety of closure and transmission principles, since they play important roles in the cases. Then, §2 and §3 offer diagnoses of lottery cases and Gettier cases, respectively, with accounts of how our intuitions about them are affected by the choice of the relevant reference class. Standard contextualism and SSI have competed for the best explanation of our intuitions about lottery cases. I show that the shift in standard which is, albeit differently, posited by the standard contextualist and the SSIist explanation cannot fully explain our intuitions about certain lottery cases. Also, my diagnosis reveals that one prominent type of Gettier case, the Barn Façade case, is a variant of the lottery case involving a shift in reference class.

§1. Closure Principles and Transmission Principles

Epistemic closure principles play important roles in the two epistemological problems to be discussed in what follows: the lottery problem and the problem of Cartesian skepticism. Closure principles, in general, say that some important epistemic property, e.g., knowledge, warrant, justification, or being in an epistemic position with certain strength, is closed under known entailment. Having this feature in common, however, specific closure principles can be distinguished from each other by variations in three factors: (i) what epistemic property it is about; (ii) whether the known entailment involved contains a single premise or multiple premises; and (iii) whether it only concerns closure of a certain epistemic property under known entailment or it, in addition, concerns transmission of the epistemic property across known entailment.[1]

Let us begin with the two principles differing in (i), one of which is about knowledge, and the other of which is about warrant; they are on the same footing in terms of (ii) and (iii): both merely involve closure, rather than transmission, under known entailment with a single premise:

[1] Another variation is in what property knowledge is closed under. The strongest principle in this regard is that knowledge is closed under entailment.

> (SCK: Single-Premise Closure of Knowledge) Necessarily, if *S* knows *p*, and that *p* entails *q*, then *S* knows *q*.
>
> (SCW: Single-Premise Closure of Warrant) Necessarily, if *S* has warrant for *p*, and knows that *p* entails *q*, then *S* has warrant for *q*.

'Warrant' is jargon in epistemology, coined by Plantinga (1993a, 1993b). He defines it as a single property that turns true belief into a case of knowledge, i.e., as the epistemic factor in Stanley's (2005) broad sense. By definition, then, *S* has warrant for *p* iff *S* is in a sufficiently strong epistemic position with regard to *p*. So, SCW can easily read as the closure principle for sufficiency of epistemic position. But it is not necessary to define 'warrant' in Plantinga's fashion: first, it can be identified with some epistemic factor in Stanley's narrow sense, e.g., internalist or externalist justification, and accordingly, a variety of closure principles can be generated with regard to different epistemic factors; second, each epistemic factor being gradable, the concept of warrant in the minimal sense can be defined as a property of raising the strength of epistemic position: on this weakest definition, SCW amounts to the principle that necessarily if *S* instantiates some epistemic factor with which the strength of *S*'s epistemic position regarding *p* is raised to some degree, and *S* knows that *p* entails *q*, then *S*'s epistemic position regarding *q* is equally raised; third, 'warrant' can be used in the broadest sense to cover all of these definitions. In what follows, I specify the relevant sense of the term only when it is necessary.

By rendering 'single premise' in SCK and SCW 'multiple premises,' the principles differing in (ii) are generated:

(MCK: Multi-Premise Closure of Knowledge) Necessarily, if *S* knows each of propositions p_1, …, p_n, and knows that the conjunction of p_1, …, p_n entails *q*, then *S* knows *q*.

(MCW: Multi-Premise Closure of Warrant) Necessarily, if *S* has warrant for each of propositions p_1, …, p_n, and knows that the conjunction of p_1, …, p_n entails *q*, then *S* has warrant for *q*.

The four closure principles so far mentioned take a conditional form, and only have to do with a metaphysical or conceptual relation between the antecedent and the consequent of the conditional. For this reason, they are not principles that specify a means by which *S* comes to know *q*, nor are they principles that state a pre-condition that must obtain in order for *S* to know *q*. In order to obtain principles pertaining to such a means or a pre-condition, the basing relation needs to be incorporated to the principles above, and once it is done, the principles differing in (iii) are produced – henceforth, I refer to the former as the 'closure principles', and the latter as the 'transmission principles':

(STK: Single-Premise Transmission of Knowledge) Necessarily, if *S* knows *p*, competently deduces *q*, and thereby comes to believe *q*, while retaining knowledge of *p* throughout, then *S* knows *q*.[2]

[2] The formulations of STK and MTK are taken from Hawthorne (2004a, p. 34 and p. 33, respectively), though he calls the former 'Single-Premise Closure of Knowledge (SPC)'and the latter 'Multi-Premise Closure of Knowledge (MPC).' Hawthorne's formulations are influenced by Williamson (2000a, p. 117).

(STW: Single-Premise Transmission of Warrant) Necessarily, if *S* has warrant for *p*, competently deduces *q*, and thereby comes to believe *q*, while retaining warrant for *p* throughout, then *S* acquires warrant for *q*.[3]

(MTK: Multiple-Premise Transmission of Knowledge) Necessarily, if *S* knows p_1, …, p_n, competently deduces *q*, and thereby comes to believe *q*, while retaining knowledge of p_1, …, p_n throughout, then *S* knows *q*.

(MTW: Multiple-Premise Transmission of Warrant) Necessarily, if *S* has warrant for p_1, …, p_n, competently deduces *q*, and thereby comes to believe *q*, while retaining warrant for p_1, …, p_n throughout, then *S* acquires warrant for *q*.

These principles, roughly, predicate that an epistemic factor *S* instantiates, knowledge or warrant, regarding one or a set of entailing premises, is transmitted across known entailment, if *S* believes the conclusion of entailment on the basis of the premise(s) *via* deduction. Thus, the transmission principles are focused on the way in which *S* acquires knowledge or warrant for the conclusion of entailment. As a result, the transmission principles of warrant, STW and MTW, require that the notion of warrant

[3] As far as I know, Lehrer (1965) is the first to propose the justification version of STW as a replacement of the justification version of SCW; he argues that the justification version of SCW is inadequate, on the two grounds: first, even if *S* satisfies its antecedent, *S* fails to satisfy its consequent, when *S* believes *q* on the basis of a reason other than *p*. The notion of justification Lehrer appeals to is clearly doxastic; second, *S*'s belief *q* may not be justified, even when *S* believes *q* on the basis of *p*; such a case obtains if *S* randomly bases the belief *q* on *p*. The requirement that *S* has deduced *q* from *p* is in need. Note that the second requirement does not entail the first, since *S* may deduce *q* from *p*, but *S* may not believe *q* on the basis of *p*; this is the reason why 'thereby' in the formulations of the transmission principles is important.

therein be doxastic rather than propositional, whereas the closure principles of warrant, SCK and MCK, do not specify which notion of warrant is relevant.[4] Notice that there may be cases where a closure principle of warrant holds but the corresponding transmission principle of warrant fails to hold; such cases obtain if warrant for *q* is a pre-condition for warrant for *p*; or, temporally speaking, if it is because *S already* has propositional warrant for *q* that the belief *q* that *S* forms by deduction is doxastically warranted. Such cases constitute counterexamples to the transmission principles of warrant, but not to the closure principles of warrant.[5]

§2. Lottery Cases

§2. 1. The Lottery Problem

What I call the 'lottery problem' and the 'lottery cases' in epistemology are most clearly stated by Vogel (1990) and, following him, Hawthorne (2004a), though the problem

---

[4] I do not define the notions of propositional and doxastic warrant here. I hope that they are nonetheless clear enough, from the distinction between propositional and doxastic justification.

[5] For this reason, Crispin Wright, in a series of papers (2000, 2002, 2003, 2004, 2007), emphasizes the difference between the closure and the transmission principle (it is already implicit in his early papers (1985, 1991). On his view, the difference is important for the problem of Cartesian skepticism. I discuss his view in §2. 3 and 2. 4 of Chapter 6.

was recognized before them.[6] [7] Both Vogel and Hawthorne mainly focus on fallibilist internalist evidentialism, but the lottery problem is generalized to most fallibilist theories of knowledge.[8] In Chapter 1, fallibilism is formulated in terms of **Fallible Knowledge**:

---

[6] For example, Harman (1968) proposes a problem similar to the lottery problem as Vogel presents it. Harman offers different solutions in this paper and subsequent books (1974, 1986). For critical discussions of his solutions, see (Hawthorne, 2004a, pp. 149-56).

[7] The lottery problem in epistemology, though different in several respects, has similarity with, and partly relies on the lottery paradox in probability theory, first raised by Kyburg (1961). Kyburg points out that a paradoxical consequence follows from the two plausible epistemic principles, **Threshold for Internalist Justification** and the justification version of MCW (Kyburg uses the notion of rational acceptability instead of that of internalist justification, but they are equivalent in my definition, because both are defined in terms of epistemic probability in the same way):

Suppose that ε is $1/n$, where $1- \varepsilon$ is the threshold for justification, and that there is a fair lottery containing n tickets, $t_1, \ldots, t_n$. Also suppose that S knows and thereby is justified in believing that one of the n tickets will win. With these suppositions, we can construct the following inconsistent triad:

(1). The proposition that one ticket out of the n tickets will win is justified for S. [Assumption]

(2). For any $i \in [1, n]$, the proposition that ticket i will not win is justified for S. [from **Threshold for Internalist Justification**]

(3). Therefore, the proposition that no ticket will win is justified for S. [from (2) and MCW]

(2) follows from **Threshold for Internalist Justification**, because Kyburg's **Support** predicates that the epistemic probability of *i will not win* equals the known objective probability of *i will not win*, and the latter is equal to the threshold, $1 - \varepsilon$, for justification. Kyburg's own solution to the lottery paradox is to deny the justification version of MCW; another typical response is to deny the sufficiency part of **Threshold for Internalist Justification**, viz., that $Ep(p) \geq 1 - \varepsilon$ is sufficient for justification for p.

(Fallible Knowledge) S can know p even though the strength of S's epistemic position with regard to $p \geq 1 - \varepsilon$, where $\varepsilon \neq 0$,

where '$1 - \varepsilon$' refers to the threshold for the strength of epistemic position enough for knowledge. The smaller ε is, the more demanding the epistemic standard for knowledge becomes, and *vice versa*. However small ε is, if it is not 0, the lottery problem arises.

The lottery problem consists in the fact that many alleged mundane cases of knowledge are analogous to the following case (henceforth, I call it the 'original lottery case'). Suppose that ε is $1/n$, and that there is a fair lottery containing n tickets, $t_1, \ldots, t_n$, one of which is owned by S. Also suppose that S knows that one of the n tickets will win. Now, we can derive the conclusion that S knows that a particular ticket will win:

(1). S knows that one ticket out of the n tickets will win. [Assumption]

(2). S knows that ticket i will not win, for any $i \in [1, n - 1]$. [from (1) and **Fallible Knowledge**]

(3). Therefore, S knows, of some particular ticket, that it will win. [from (2) and MCK]

[8] Hawthorne remarks that sensitivity theory is an exception, but points out certain problems of the sensitivity solution to the lottery problem (2004a, pp. 9-14). His target is DeRose (1996a) who employs sensitivity to deal with the lottery problem.

Most theories of knowledge, jointly with (1) and **Fallible Knowledge**, lead to (2); they, however, differ as to how (2) obtains, since they give different interpretations of the notion of epistemic position. On internalist evidentialism, any expert objective probability function predicates that *S*'s degree of internalist justification for *ticket i will not win* is 1 - 1/n = 1 - ε, for any $i \in [1, n]$, insofar as (1) is equivalent to the premise that *S* has in her evidence set *the indefinite objective probability of a ticket's winning is* ε.[9] Then, *S* is internalistically justified in believing, and other things being equal, knows that ticket *i* will not win[10]; on type reliabilism, presumably the relevant process type is something like

---

[9] Vogel and Hawthorne appeal to some expert probability function. Given his reference to Kyburg, it seems that Vogel has Kyburg's **Support** in mind; Hawthorne is explicit that it is Lewis's **Principal Principle** (2004a, pp. 92-3); the derivation of (2) from **Principal Principle** requires that *S* has in her evidence set the definite probability of *ticket i will not win*, for any $i \in [1, n - 1]$. The lottery problem can be constructed without relying on any expert objective probability function, if the following principle is in place:

> (Principle of Indifference) If the sample space is partitioned into $p_1, \ldots, p_n$, and *S* has no reason to favor any of them over others, then, for any i, $Ep(p_i) = 1/n$.

This principle or its sophistication is espoused by many theorists of probability, regardless of what type of probability they are concerned with. Note that it is explicitly incorporated into feature (iii) of lottery proposition below. In lottery cases, the sample space consists of propositions $p_1, \ldots, p_n$, each of which is to the effect that each of ticket$_1$, …, ticket$_n$ will win. Then, $Ep(p_i)$ is $1/n = \varepsilon$, because *S* has no reason to favor any *p* over every other *p*. However, **Principle of Indifference** is not generally applicable, and its application is often restricted to the game of chance like the original lottery case we are discussing here.

[10] As mentioned in footnote 7, many deny that *S* is justified in this case. I am not concerned with this type of response, since this response does not work against other lottery cases.

{believing that ticket i will lose: for any $i \in [1, n]$}, and it has $1 - \varepsilon$ reliability (For simplicity, I omit the environment type). Then, *S*'s belief that ticket i will not win is externalistically justified, and, other things being equal, amounts to knowledge.[11] Sensitivity and safety theory may evade (2), but they still faces the lottery problem – I discuss these theories in §2. 3.

(3) is hard to accept at an intuitive level, at least in most situations or contexts. In order to block (3), at least four responses are logically possible: (a) deny that internalist justification or externalist justification (type reliability) is the only epistemic factor, and to espouse a pluralist theory of knowledge[12]; (b) deny **Fallible Knowledge**[13]; (c), for

[11] Alan Goldman, in his (1988, pp. 52-3), points out that type reliabilism has this consequence, though, as is often the case in the literature, he does not specify the reference class he is appealing to (cf. also (Adler, 2005)). The reference class used here is not the only one to lead to (2); among others, it can be a type of direct inference process to infer that ticket i will lose on the basis of the probabilistic information that the indefinite objective probability of a ticket's winning is ε.

[12] (a) means treating the original lottery case as similar to Gettier cases. The Gettier problem is the topic of §3.

[13] (b) might be attributed to Williamson (2000a), since he holds that knowledge requires evidential probability 1. But his notion of evidential probability 1 is idiosyncratic, and, as he emphasizes, should not be identified with Cartesian infallibility; if *S* has evidence that p, the evidential probability of p is 1, simply because the evidential probability of $(p \mid p)$ is 1. What makes his view non-trivial is the E=K thesis, viz., that only knowledge counts as evidence. For Williamson, then, the lottery problem amounts to the problem of how *S* knows or fails to know the propositions in (2) and similar propositions in other cases. Williamson (2005a, 2009a) suggests an answer to this question by making use of safety theory.

internalist evidentialism, deny the relevant version of expert objective probability function[14]; and (d) deny MCK.[15]

The responses (a) through (c) entail the denial of (2), but they come with great cost[16]: Vogel's lottery problem shows that it has as radical, unpalatable ramifications as those of traditional skepticism. Vogel argues that there is a vast set of propositions about which one intuitively judges that one does not know, even though they are known to be consequences of everyday propositions about which one intuitively judges that one knows. Such a set of propositions are called 'lottery propositions,' since they share certain features with the propositions in (2) of the form *ticket i will not win*. For example, consider the following case Vogel puts forth:

> (Car Theft) Suppose you own a car which you parked a few hours ago on a side street in a major metropolitan area. You remember clearly where you left it. Do you know where your car is? We are inclined to say that you do. Now it is true that every day hundreds of cars are stolen in the major cities of the United States. Do

[14] Something like (d) is put forth by Cohen (1988, 2004), Dudman (1992), and Luper-Foy (2007). They argue, roughly, that *S* cannot know any proposition on probabilistic or statistical ground. In particular, Cohen combines this type of response with his contextualism. I will discuss his response to the lottery problem in detail in the next section.

[15] Note that even if Kyburg is right about the denial of the justification version of MCW, the denial of MCK does not follow. They are about different properties, justification and knowledge.

[16] Although (d), albeit unpopular, has some air of plausibility, I put it aside here. For the criticisms of (d), see Hawthorne (2004a), pp. 46-50.

you know that your car has not been stolen? Many people have the intuition that you would not know that. (Vogel, 1990, p. 15)

Vogel claims that lottery propositions, such as *my car is not stolen*, constitute *ostensible* counterexamples to SCK. In the Car Theft case, you know that your car is parked on the street, and you know that *your car is parked on the street* entails *your car is not stolen*.[17] But intuitively, you do not know that your car is not stolen.

Vogel points out that there are several respects in which the proposition that your car has not been stolen is similar to the proposition that a particular ticket will not win: (i) *S*'s epistemic position with regard to each proposition is considerably high[18]; (ii) nevertheless, it is not abnormal in some sense if it happens to be false. It is not abnormal if you win a big lottery by sheer luck, and similarly, it is not abnormal if your car has been stolen by bad luck after you park it, given the existence of car thieves[19]; (iii) The item in

[17] Strictly speaking, the former does not entail the latter; it is possible for a stupid thief to have stolen a car and left it in the same place as it is originally parked. So, the entailed proposition must be stronger, e.g., *your car is not stolen away*. Here, however, I stick to the way in which Vogel proposes his case.

[18] Vogel's original characterization of (i) is that the *objective* probability of its being true is considerably high. This characterization does not make it immediately evident that lottery cases are *epistemically* problematic. I hope that my characterization better captures his intention.

[19] The sense of abnormality involved in (ii) cannot be that which is defined in terms of the low probability of the falsity of the lottery proposition; if so, (ii) is simply redundant. It seems that (ii) is relevant for differentiating lottery cases from similar but different cases Vogel invents elsewhere, one of which is the following:

question is a member of a reference class, and it is indistinguishable to *S* from other members with respect to its chance of being chosen. The reference class in the Car Theft case is presumably {cars in the city}. This reference class consists of cars c_1, …, c_n, and among them c_i is your car ,where *i* is some number ∈ [1, n]. Corresponding to c_1, …, c_n, there is a set of propositions p_1, …, p_n, each of which is to the effect that each of c_1, …, c_n is not stolen, and p_i is a member of this set. You have no reason to favor *your car is not stolen* over *every other car in the city is not stolen*, i.e., you have no reason to favor p_i (and any particular *p*) over other *p*s. Again, a particular ticket t_i is a member of the reference

---

> (Hole-In-One) Sixty golfers are entered in the Wealth and Privilege Invitational Tournament. The course has a short but difficult hole, known as the 'Heartbreaker". Before the round begins, you think to yourself that, surely, not all sixty players will get a hole-in-one on the 'Heartbreaker". (Vogel, 1999, p. 165)

Vogel claims that, unlike in lottery cases, it is intuitive in this case that you know that not all sixty players will get a hole-in-one on the Heartbreaker. The proposition that all sixty players will get a hole-in-one on the Heartbreaker corresponds to a particular reference sequence r_i of the results of the sixty golfers' first shots on the Heartbreaker, and it is also belong to other reference sequences. r_i seems too abnormal to be relevant for your epistemic probability of the proposition or its denial. The relevant sense of abnormality or normalcy has yet to be specified, and it is here where **Normalcy**, proposed as a solution to the reference class problem, may be helpful. But note that even if the range of the relevant reference classes or sequences is restricted to normal ones, it does not mean that what counts as normal is uniquely determined; the relevant sense of normalcy may vary. For this reason, I do not believe that there is a principled difference between lottery cases and cases like the Hole-in-One case. We can make silent the sense of abnormality in which it is abnormal that one wins a big lottery; "What are you talking about? You're not the guy who can win such a big lottery. Winning the lottery is like getting a hole-in-one many times in a row!"

class {tickets of the lottery} and one has no reason to favor *its winning* over *every other ticket's winning*. More precisely, there is a set of propositions $p_1, \ldots, p_n$, each of which is to the effect that each of $t_1, \ldots, t_n$ will not lose, and you have no reason to favor p_i (and any particular p) over other ps; (iv) there is some probabilistic evidence that that one of the propositions $p_1, \ldots, p_n$ is false. In the case at hand, there is probabilistic or statistical evidence that some car is stolen, just as there is statistical evidence that some ticket will win.[20] (iv) does not specify who possesses such probabilistic or statistical evidence. As we will see in the next section, if it is possessed by attributors rather than subjects, lottery cases may count in favor of contextualism.

The lottery propositions are those propositions which satisfy features (i) to (iv). Vogel shows that a lottery proposition is constructible for most, if not all, alleged cases of inductive knowledge.[21] Given the ubiquity of lottery propositions and SCK, it follows that one does not know many things one thinks one does. Thus, the denial of knowledge of lottery propositions summons unpalatable skepticism. The lottery problem is the problem of how to avoid this skepticism.

---

[20] (iv) is often stated in the literature in such a way that it requires that *S* know that one of the propositions $p_1, \ldots, p_n$ is false. Vogel does not require this. Moreover, if it is required, the range of the lottery problem is unnecessarily restricted. For a lottery case can be constructed in such a way that there is no guaranteed winner of the lottery, as in Lotto 6/49. For more on this point, see (Hawthorne, 2004a, pp. 15-6); Neufeld & Goodwin (1998) point out the same for the lottery paradox.

[21] Hawthorne (2004a) illustrates by way of example that a lottery proposition can be generated for any perceptually believed proposition. But I do not find his example compelling.

Indeed, neither Vogel nor Hawthorne regards lottery propositions to constitute *genuine* counterexamples to SCK.[22] The reason is the following: all alleged counterexamples involve two sub-cases, in one of which one intuits that *S* knows an ordinary proposition, and in one of which one intuits that *S* does not know a lottery proposition corresponding to that ordinary proposition. SCK is preserved, since, as far as our intuitions go, *S* knows both the lottery proposition and the everyday proposition in the first sub-case, and *S* knows neither in the second sub-case. However, this entails that some epistemic shift occurs between the two sub-cases. Then, the remaining problem is to figure out what shift it is.

§2. 2. Shift in Reference Class in the Lottery Case for Internalism

Vogel (1990), in passing, suggests that there are two kinds of shifts to be postulated to explain our seemingly conflicting intuitions about the lottery cases: shift in probability assignment and shift in epistemic standard. He, in his (2004b), adds that his favored account is in terms of the former shift, and explains that the shift is induced by different reference classes being assigned to the lottery proposition. He seems to assume Kyburg's **Support**, according to which objective indefinite probability is *epistemically* relevant for determining the strength of epistemic position regarding the lottery proposition (I follow Vogel in using **Support** for explanation in what follows, but nothing important hinges on this choice of the relevant expert probability function; though, with other expert probability

[22] Hawthorne has in mind STK rather than SCK. For a time, I ignore the difference between them.

functions, the explanation would be more complex). More generally, the former shift is expressed as shift in strength of epistemic position.

Then, the two kinds of shifts represent two different ways of thinking about epistemic positions in lottery cases: between the two sub-cases of a lottery case, (a) *S*'s epistemic position regarding the lottery proposition shifts, and (b) the epistemic standard for how high epistemic position is required for knowing the lottery proposition shifts. Appeals to (a) and (b) account for our varying intuitions about lottery cases in fundamentally different ways. I call the two accounts postulating these shifts the 'shift-in-reference-class account' (SRC) and the 'shift-in-epistemic-standard account' (SES), respectively. Cohen (1988, 2004) offers a contextualist SES of lottery cases, and Hawthorne (2004a) does a SSIist SES. In this section, I will argue that there are cases that only SRS can account for. The existence of such cases will give a reason in favor of a different version of contextualism than standard contextualism espoused by Cohen and others – the argument along this line is fully offered in §2. 5 of Chapter 7.

Vogel (2004b, p. 507) uses his SRC to save our knowledge of the everyday proposition in the Car Theft case: *S* may extrapolate the everyday proposition that *S*'s car is parked on the street from a past record of invariably finding *S*'s car where it was parked. This is tantamount to *S*'s using, as the relevant reference class in estimating the actual frequency of a car's being on the street, {occasions o_1, …, o_n: *S*'s cars is parked on the street}, which has among its members individual occasions on which *S*'s car is parked on the street in the past. Relative to this reference class, the indefinite probability of *S's car is on the street on an occasion* is virtually 1. Since *S* knows this indefinite probability and this particular occasion o_{n+1} belongs to the same reference class as o_1, …, o_n, it follows by

Support that *Ep*(the car is on the street on o_{n+1}) for S is virtually $1 > 1 - \varepsilon$. Then, S knows on o_{n+1} that her car is where it was parked, and by SCK, that her car is not stolen. Indeed, once relativized to the same reference class, the *Ep*(the car is not stolen on o_{n+1}) has the same value as *Ep*(the car is on the street on o_{n+1}), and this is the reason that SCK is maintained.

This is how SRC explains why we often, albeit not always, do not hesitate to attribute knowledge of the everyday propositions to S. Our actual frequency assignment, or more generally, indefinite objective probability assignment, is relative to a choice of reference class; such a choice, given **Support** or other expert objective probability functions, is also responsible for the degrees of internalist justification. Though Vogel is careful to claim that what shift underlies the variance in intuition is an empirical question, it seems hard to deny that our reasons to believe the everyday propositions are often those which he describes; we tend to say things like "I always park here, and it never ever happened that my car was stolen. I'm pretty sure that it won't."

By contrast, SES explains lottery cases as follows: **Epistemic Position** says that S knows p only if the strength of S's epistemic position regarding $p \geq 1 - \varepsilon$. When we are ready to attribute knowledge of an everyday proposition, a lax standard may be in play so that $1 - \varepsilon$ is relatively small. Then, S can easily satisfy **Epistemic Position**, and thereby knows the everyday proposition. When the every proposition is paired with the lottery proposition it entails, the standard is raised to the point where $1 - \varepsilon$ surpasses the strength of

S's epistemic position regarding the lottery proposition.[23] Then, *S* fails to satisfy **Epistemic Position**, and thereby does not know the lottery proposition. SES, as well, preserves SCK. For, relative to the lax standard, *S*'s epistemic positions regarding the everyday proposition and its corresponding lottery proposition are high enough for knowledge; whereas, relative to the stringent standard, they are not.

SRS and SES equally save our varying intuitions and SCK. And yet, the two accounts are significantly different, for they exploit different features of lottery proposition: on the one hand, SRC appeals to feature (iii) to explain why recognizing a proposition as a lottery proposition tends to induce shift in reference class; on the other hand, SES appeals to the feature (iv) to explain why recognizing a proposition as a lottery proposition tends to raise the epistemic standard; on account of feature (iv), a lottery proposition is associated with a possibility that it is false. Cohen (1998, 2004) combines this point with contextualism.[24] The gist of his contextualist SES is nicely epitomized in the following passage:

> … one cannot avoid thinking about the 1/*n* probability that one will win if one bases one's belief on the fact that the odds of losing are 1 - 1/*n*. So the chance of error

---

[23] Just above the strength of epistemic position? If this sounds *ad hoc*, **Epistemic Position** can be re-formulated in terms of '>' rather than ' ≥,' though the settings of lottery cases would need to be modified accordingly.

[24] He is a proponent of internalist evidentialism, and combines it with standard contextualism. Since he takes internalist evidentialism as a theory of justification, his contextualism is that how high a degree of justification is required for knowledge varies with context. For more on his contextualism, see Cohen (1999).

> will be salient in this case. But matters are different in case where I base my belief on the newspaper report. In such a case, we typically do not infer that what the newspaper says is true, based on the *m*/*n* probability that if the newspaper says it, it is true. Rather, insofar as we think of it at all, we think of a scenario: The reporter determines the result, perhaps by witnessing the drawing, sends in the report and then it is printed in the newspaper. In this scenario, the chance of error is not salient. (Cohen, 2004, p. 431)

Cohen is comparing the original lottery case in which *S* believes that *S*'s ticket will not win as a result of direct inference, with a similar case in which *S* believes the same proposition after *S* reads the result of the lottery in a newspaper. The objective probability of a newspaper's telling truth, or more simply, its (general) reliability, is lower than the probability of a ticket's losing, if the lottery is very big, as is stipulated in the original lottery case. Nevertheless, one tends to judge of the original lottery case that *S* does not know, whereas one tends to judge of the newspaper case that *S* knows. And, more importantly, both judgments seem true. Thus, these two cases are analogs of the sub-cases of the Car Theft case, though the proposition in question is the same between the cases.

Cohen accounts for the asymmetry of our intuitions and their truth in terms of error possibilities accompanying the original lottery case: thinking of probabilistic evidence brings one or more possibilities of error into salience on the attributor's side, in most, if not all, circumstances.[25] It is worth pointing out that the two cases, as he describes in the quote, are not merely different as to whether error possibilities are salient to the attributor; they also differ with regard to the base on which *S* believes that her ticket is a loser: in the original lottery case, *S* bases her belief on the probabilistic evidence that the chance of a ticket's winning is $1/n$ or ε; whereas, in the newspaper case, *S* bases her belief on the evidence that the newspaper says that a ticket other than hers has won.

These two cases alone do not motivate contextualism: first, in Cohen's original lottery case, error possibilities are salient to both the subject and the attributor, and in his newspaper case, they are salient to neither. In other words, attributor factors play no essential role in Cohen's explanation. Hawthorne (2004a) offer a variant of SES that is similar to Cohen's except in one respect[26]: he combines it with SSI, and thereby salience of error possibilities to the subject is responsible for the shift in epistemic standard.[27] For the reason just given, SSI can save our intuitions about the two cases as much as Cohen's contextualism can, but is less committal in that it needs to posit no explanatorily idle factor, at least in these cases.

---

[25] He concedes that salience is defeasible. So it is not always the case that a salient error possibility induces shift in epistemic standard.

[26] Hawthorne is a proponent of some version of internalist evidentialism, and regards internalist evidence as a determinant of the strength of epistemic position. But it is not clear whether he takes his internalist evidentialism to be a theory of justification.

[27] It should be emphasized that Hawthorne accepts salience as a shift-inducing factor with reluctance (see (ibid., pp. 158-72). He, along with other SSIists, puts more emphasis on pragmatic importance or stakes as such a factor.

Second, more importantly, the two cases are not parallel. As Cohen presents them, the following pieces of evidence are involved in each case besides *S*'s total evidence:

The Original Lottery Case		The Newspaper Case	
a	The indefinite probability of a ticket of the lottery's winning is ε.	a'	(The indefinite probability of a newspaper article's telling truth is γ.)
b	*S*'s ticket is a ticket of the lottery.	b'	The newspaper article tells that the number of *S*'s ticket has not been drawn.

S's base for believing *S's ticket will not win* differs between the original lottery case and the newspaper case: (a) in the former and (b) in the latter. Thus, the difference between the two cases is primarily in subject factor rather than in attributor factor, i.e., they differ in type of evidence *S* bases her belief on. Why can't we simply say that *S* cannot know a lottery proposition on the basis of probabilistic evidence like (a), while *S* can on the basis of non-probabilistic evidence like (b')? Moreover, this asymmetry concerning *S*'s base is crucial for Cohen's contextualist account. For he holds that it is because the description of the original lottery case includes the specification of S's base as (a) that the error possibility associated with (a) becomes salient to the attributor as she reads the description.[28]

[28] DeRose (1996a, pp. 576-8) raises two problems against Cohen's account: first, *S* may believe that her ticket is a loser, simply on the basis of evidence like (b). To use DeRose's example, *S*'s belief that her ticket is a loser may be based on "it's a Super Lotto ticket (for heaven's sake!)." Thus modified, the original lottery case is precisely symmetrical to the newspaper case. Then, Cohen's account must predicate that *S* knows that her ticket is a loser. DeRose claims that this result is hard to accept; second, there are cases where *S* bases her

Too many differences being involved, the pair of the original lottery case and the newspaper case is not adequate for settling the question of which account is the best explanation of our intuitions; it is indifferent between the contextualist SES and the SSIist SES, and between SES and SRC: in the original lottery case, the reference class used in *S*'s direct inference is held fixed as (a) in the description of the case because it is stipulated as *S*'s base. What about the newspaper case? As a matter of fact, it is underspecified: unlike (a), (a') is not the base of *S*'s belief. Cohen assumes that (a') is still relevant for the strength of *S*'s epistemic position regarding *S's ticket will not win*. But (a') may not be the only probabilistic evidence *S* has about *S*'s ticket. For example, suppose that *S* has never won any big lottery like the one at stake. Then, *S*'s ticket of this lottery is a member of the reference class {the big lottery tickets *S* has ever bought}, and presumably, *S* has in her evidence set the corresponding probabilistic evidence (a'') that a lottery ticket in this reference class has probability as nearly as 1 of losing. Then, relative to which evidence, (a') or (a''), is the strength of *S*'s epistemic position to be determined? Furthermore, if (a'') is *S*'s base, is it true that *S* knows that *S*'s ticket is a loser? [29]

---

belief on (b'), but about which one does have the intuition that *S* knows. DeRose reckons this a trouble for Cohen, since his account cannot permit such cases. For Hawthorne's response to the first problem, see the next footnote. The second problem is not as serious as DeRose thinks. For, as mentioned in footnote 25, Cohen accepts that salience may be trumped by some other factor, and when that is the case, the epistemic standard does not rise as a result of (b')'s salience.

[29] Hawthorne points out that people are, on some occasions, willing to attribute knowledge in the original lottery case: "Try raising the possibility of lottery success to people who are planning out their lives. Very often, they will respond with 'You know that's not going to happen' or 'I know full well I'm not going to get that lucky'" (ibid., p. 18). Similarly, in deliberating whether to buy a lottery ticket, ordinary people often

This is precisely the reference class problem for internalist evidentialism. One may have the intuition that the answer to the last question is affirmative. SES, whether Cohen's variant or Hawthorne's, cannot explain this case, let alone the case Vogel makes of the Car Theft case in proposing SRC: insofar as the reference class *S* uses is set as (a'') or some similar one, the strength of *S*'s epistemic position is nearly maximal, leaving no room for the epistemic standard going beyond it. In fact, I am inclined to think that the reference

---

claim "You know you are wasting your money." Hawthorne suggests that in such cases, "[t]he operative division of subcases is into 'I win; one of the other million or so ticket-holders wins', the former of which is in turn disregarded, even to the extent that one will be willing to flat-out assert of the lottery victory 'There is no chance of *that* happening'" (pp. 18-9).

Hawthorne, thus, grants that the division of the sample space into *n* partitions each of which is of the form 'ticket *i* will not win' is not necessary, and the division may be simply into the two possibilities, *S will not win* and *some other ticket-holder will win*. This, in my terms, is equivalent to assigning the lottery proposition to a different reference class. However, it is not clear whether Hawthorne conceptualizes the cases as involving *S*'s using such a reference class or as the attributor is using it in evaluating *S*'s epistemic position regarding the lottery proposition. Even if the former is Hawthorne's intention, it is still unclear whether the reference class he appeals to has the result he wants it to do. Given that only two propositions are relevant relative to the reference class specified, isn't it the case that the lottery proposition has .5 probability? It, however, seems that Hawthorne is implying that *S will not win* has probability 1.

I take it, then, that Hawthorne is getting close to the point Vogel made. Relative to the reference class {the lottery tickets the loser has ever bought}, the objective probability of, and, by **Principal Principle** or the like, the epistemic probability of, the lottery proposition are virtually 1. Suppose that *S* bases her belief in the lottery proposition on the proposition corresponding to this reference class, viz., that the indefinite probability of *a ticket out of my lottery tickets will not win* is nearly 1. The strength of *S*'s epistemic position with regard to the lottery proposition, then, is nearly maximal.

class in (a'') is not the correct reference class, unlike the reference class Vogel appeals to in his SRC of the Car Theft case. What is missing is a general recipe for determining which reference class is relevant or correct. I postpone my full-blown discussion until §3. 1 of Chapter 7. But for now, suffice it to say that there are cases SES cannot explain, because it lacks resources to take account of reference classes.

§2. 3. Shift in Reference Class in the Lottery Case for Externalism

I have so far focused on SRC and SES for internalist evidentialism, mainly because Vogel, Cohen, and Hawthorne are all proponents of internalist evidentialism of one or another sort. It is possible to accommodate each of SRC and SES into externalist theories, resulting in a variety of SRCs and SESs. In this section, I delineate what form SES and SRC take for type reliabilism and token reliabilism.

First, type reliabilism. Consider again the Car Theft case. As we have seen, both Cohen and Vogel use feature (iii) of lottery proposition to explain that the degree of internalist justification for each of the propositions $p_1, \ldots, p_n$, to the effect that each of $c_1, \ldots, c_n$ is not stolen, is internalistically justified to the degree $1 - \varepsilon = 1 - 1/n$. Cohen's SES implies that this is also the degree of internalist justification for each of the propositions to the effect that each of $c_1, \ldots, c_n$ is on the street. (iii) is generalizable by requiring that externalist justification for (and more generally, the strength of epistemic position regarding) each of these propositions be equal. In order for externalist justification to satisfy (iii) thus generalized, the process type and the environment type to which type reliability is relativized must be {remembering the car's location} and {places in the city} respectively, or a pair of similar ones.[30] Indeed, these types are quite natural for us to use in estimating type reliability in the case at hand. Relative to this pair of reference classes, *S*'s degree of externalist justification is 1 - ε. Then, type reliabilism is easily combined with SES.

In addition, SRC can be developed for type reliabilism on the model of the internalist evidentialist version of SRC. {places in the city} is not the only environment type against which *S*'s reliability is estimated. Other types, e.g., {the environments in which *S* parked the car on the street}, are equally in order. Relativized to this environment type, the reliability of {remembering the car's location} would be nearly maximal, insofar as *S* has never experienced any accident of car removal. Then, however high the epistemic threshold is, *S* knows both that the car is on the street and that the car is not stolen. The second sub-case shifts the relevant reference class in such a way that externalist justification satisfies (iii), and lowers the degree of externalist justification. The type reliabilist version of SRC is consistent with SCK.[31]

Second, token reliabilism. Since it is not clear how sensitivity theory militates in favor of our intuitions, I focus on safety theory (my points equally hold for the counterfactual robustness theory).

---

[30] The environment type here includes hypothetical environment tokens in which the car is stolen. So, type reliability must not be based on actual frequency. Near-actual frequency or long-term propensity is required to capture the kind of type reliability necessary for satisfying (iii).

[31] Greco (2003, 2004a, 2008, 2010) proposes a type reliabilist version of SRC as a solution to the lottery problem. It has some similarity with the version of non-standard contextualism I am developing. I, however, doubt that his account is successful (see footnotes 8 of Chapter 4).

Safety theory can be combined with both SES and SRC. The safety version of SES goes as follows: *S*'s beliefs in the lottery proposition and in the corresponding everyday proposition are both safe, relative to the ordinary, lax standard. The lax epistemic standard only requires that *S*'s beliefs match the fact regarding these propositions within a narrow range of nearby worlds. *S*'s beliefs satisfy this requirement, since they are true in the nearby possible worlds; whereas, the stringent epistemic standard enlarges the relevant range of possible worlds and requires that *S*'s beliefs match the fact regarding these propositions within the enlarged range of nearby worlds. *S*'s beliefs do not satisfy this requirement.

Thus, SES appeals to the shift in epistemic standard governing how wide a range of possible worlds is across which *S*'s beliefs must match the fact, in order for *S* to know. This account presupposes that the world ordering relevant for lottery cases is such that lottery propositions are only false in remote worlds. What enforces this world ordering? The natural candidate is objective probability: a possible world w_1 is farther away from w than w_2 is iff the probability of w_1 being realized at w is lower than the probability of w_2 being realized at w. There are several problems with the objectively probabilistic ordering source, as mentioned in §3. 7 of Chapter 4. At least some cases require the epistemic ordering source to differ from the objectively probabilistic ordering source – one of such cases, the Barn Façade case, is discussed in detail in §3. More importantly, objective probability is relative to a reference class. Vogelian SRC has the point here: the safety version of SES presupposes that a certain reference class is chosen.

The safety version of SES entails that *S* knows in the original lottery case that her ticket will lose. Pritchard (2005a, 2008a) and Williamson (2009a), in developing a safety-based account of the original lottery case, deny that the ordering source relevant for the case is objectively probabilistic.[32] Even though *S's ticket will lose* has a very low probability of being false, the possible worlds in which it is false are near the actual world. It follows that *S*'s belief in the original lottery proposition is not safe at all, resulting in a very low strength of epistemic position. If the same account is generalizable to other lottery cases to save our shifty intuitions, the result would be the safety version of SRC: the relevant ordering source differs, and hence how strong *S*'s epistemic position regarding the lottery and the corresponding ordinary proposition varies, between the two sub-cases of a lottery case.[33] [34]

---

[32] Pritchard (2008a) develops the solution proposed in Pritchard (2005a) to meet Greco's (2007b) criticisms. But the essential point is the same:

> … the very attraction of a fair lottery lies in the fact that the possible world in which one wins is very like the actual world, even though it is in fact unlikely that such a possible world should be the actual world. This point highlights the sense in which the similarity ordering of possible worlds is not tantamount to an ordering in terms of probability. (Pritchard, 2005a, p. 128)

[33] Both Pritchard and Williamson are invariantists, and committed to the position that *S* knows lottery propositions in other lottery cases. Pritchard (2008a) offers an account to explain the difference between the original and the other lottery cases. His account, however, requires a very fine-grained ordering source for each case, and is dangerously close to Giere and Miller's single-case propensity theory, on which the probability of an event token supervenes on the entire universe.

[34] Sensitivity theory is generally taken to lead to the denial of SCK (cf. (Dretske, 2005a, 2005b) and (Hawthorne,2005)). It is also controversial whether safety theory is consistent with SCK. Williamson (2009a) discusses this issue in connection with the lottery problem.

§3. Gettier Cases

§3. 1. The Gettier Problem

The Gettier problem is named after Edmund Gettier (1963), who offered two counterexamples to the classical analysis of knowledge, attributed by some to Plato. An enormous number of similar counterexamples have been since offered in the literature, and they are all subsumed under the label of 'Gettier cases' or 'Gettier counterexamples.' The classical analysis identifies knowledge with justified true belief, and, in the terminology used here, is tantamount to the monist thesis that justification is the only epistemic factor. Although Gettier only attacks the internalist notion of justification, type reliabilism as an account of justification is equally vulnerable to Gettier counterexamples.[35]

Not all the Gettier cases constitute counterexamples to the classical analysis of knowledge in the same way. Of course, it trivially follows that they are all cases in which *S* has justified true belief *p*, but fails to know *p*. And also, it is commonplace to hold that Gettier cases include accidentally or luckily true belief.[36] But it is not the case that every kind of luck induces Gettier cases, and it remains a problem to specify the relevant kind of luck, *epistemic luck* or *knowledge-preventing luck*.[37] Though what epistemic luck is

---

[35] See (Harper, 1996) and (Zagzebski, 1994, 1999).

[36] This analysis of Gettier cases is first offered by Unger (1968).

[37] For kinds of luck, see Pritchard (2005a). His safety theory is proposed on the basis of his analysis of epistemic luck: a belief is safe iff it is not true by epistemic luck. It is controversial whether safety theory, whether his particular version or in general, can deal with every Gettier case.

remains a problem, the standard analysis of Gettier cases, at the very least, tells that Gettier cases are constructible because it is possible for *S*'s justified belief to be true merely by epistemic luck. This analysis requires that **Fallible Knowledge** be true for justification: *S* can know *p* even though the degree of justification for *p* is less then maximal. For the notion of luck is deeply related to a possibility of error; if infallibilism about justification is true, then it is impossible for justified belief to be false, and hence it is impossible for justified belief to be true by luck.[38]

It is controversial whether non-justificatory theories, typically variants of token reliabilism, are able to exclude all the cases involving epistemic luck as instances of non-knowledge. The Gettier problem, then, is the challenge to find a necessary and sufficient condition for knowledge that is immune to such counterexamples. For this reason, it is the problem central for the theory of knowledge.

Furthermore, even if it is granted that a particular type of luck is involved in Gettier cases, there may be multiple ways in which *S* is susceptible to such a type of luck. In general, at least two kinds of Gettier cases are distinguished: ones relevantly involving false evidence and ones involving no such evidence. In this section, I am only concerned with a particular instance of the latter kind, Ginet-Goldman's Barn Façade case. If my diagnosis of the Barn Façade case is right, the only way to explain our intuition about it is to postulate

[38] Merricks (1995, 1997) and Zagzebski (1994, 1999) argue that the existence of Gettier cases leads to the denial of infallibilism for any epistemic factor. Their argument is sound only if neither of the following is the case: (i) the epistemic standard gets higher in Gettier cases, and (ii) the strength of *S*'s epistemic position is lowered in Gettier cases. In what follows, I will show that (ii) rather than (i) is true, at least, for the Barn-Façade case.

a shift in reference class. That is, SRC is required to explain our intuition about the Barn Façade case. In particular, I make three points: (a) the Barn Façade case is parallel to lottery cases in every epistemically relevant respect; (b) it is because the relevant reference class shifts in the Barn Façade case that *S*'s epistemic position is lowered to the point where *S* does not know; and (c) unlike for lottery cases, SES is not possible for the Barn Façade case.

§3. 2. The Barn Façade Case

The Barn Façade case is described in subtly different ways. When I introduced this case in §3. 7 of Chapter 4, I used the canonical way to describe the case, which is the following:

> (Barn-Façade) Henry spots a barn in front of him while he is driving. Then he believes that there is a barn before him. However, unbeknownst to him, Henry is in the Barn Façade County that contains numerous barn-façades. He would mistakenly take these to be barns if he were looking at them. But the barn he spotted is the only real barn therein.

This is merely the bare bones of the case, and is much simpler than Goldman's original description of the case. Most prominently, Goldman seems to take the mere proximity of many convincing barn facades to be enough to induce the intuition on the reader's side that Henry does not know that it is a barn. Thus, in the original description of the Barn Case, Henry has not encountered any of the fake barns. However, DeRose (2009, ch. 1, fn. 24, p. 23; ch. 2, fn. 2, pp. 49-50) reports that when the case is thus described, many epistemologists including himself do not have the strong intuition that Henry does not know. DeRose, then, suggests a modification of the Barn Case. Since his modification is relevant for my argument below, I cite it at length:

> (Modified Barn Façade) Imagine that Henry *has* encountered many of the fake barns (let's say 19 of them), and has been fooled by all of them that he has encountered, believing them to be real barns. Now, he's encountering the only real barn in the region (so he is 1-for-20 in his barn judgments), but it seems the same to him as all the other objects he has taken to be barns: He is not in any way more confident of his identification of the object as a barn in the case as compared with all the other instances. ('Yet *another* barn,' he thinks.) Make sure you're imagining the case so that the fakes present an extremely convincing real appearance – just like the real thing – and Henry should have the same visual evidence as he would have had if all the barns around were real, and should be just as justified in thinking that what he is presently looking at is a barn as he would have been if all the barn-like things he had encountered had all been real. (DeRose, 2009, ch. 1, fn. 24, p. 23)

With this modification, DeRose is confident that Henry does not know that it is a barn. Once this intuition is taken for granted, the Barn Case must be reckoned a Gettier case: Henry has justification for true belief but fails to have knowledge.

I agree with DeRose (and presumably, the other epistemologists he queried) that the mere presence of fake barns in the vicinity is not sufficient to pump the intuition that Henry does not know. Just imagine that Henry correctly identified many barns as such, while he was driving along a normal county, and then just crossed the border between that county and the Barn Façade County; and that the very first structure he sees in the Barn Façade County is the only real barn there. It seems counterintuitive if Henry does not know that it is a barn. Thus, some modification along the line DeRose suggests must be made.

Why does the original case have less intuition-driving force, compared with DeRose's modified case? The key to the answer to this question is the following: first, it seems that DeRose's emphasis on Henry's having encountered fake barns is misplaced. If this is all that matters for the intuition that Henry does not know, it becomes mysterious why at least some people have that intuition even for the original case; second, DeRose, in effect, spends more on emphasizing something different: Henry is in the same epistemic position to the proposition he believes of each of the barn-like structures in the Barn Façade County. In DeRose's setting of the case, there are twenty barn-like structures $b_1, \ldots, b_{20}$, of each of which Henry believes one of propositions $p_1, \ldots, p_{20}$, to the effect that it is a barn. Among these only b_{20} is a real barn, and thereby only p_{20} is true. Nevertheless, as DeRose accentuates, Henry is equally confident of, and has equal qualitative evidence for $p_1, \ldots, p_{20}$, since all the structures, $b_1, \ldots, b_{19}$, are extremely convincing fakes. To pump the intuition that Henry does not know, it need not to be the case that the real barn in the Barn Façade County is the very last barn-like structure he encounters therein. Simply, it is enough that b_i is a real barn, for some i, and Henry is in the same epistemic position with regard to p_i as with regard to other *p*s. Once the case is understood this way, its similarity to lottery cases is obvious: it satisfies the relevant features of lottery proposition.

First, just like lottery cases, there is a set of propositions $p_1,\ldots, p_{20}$, to the effect that $b_1,\ldots, b_{20}$ are barns; second, feature (ii) is satisfied: it is not abnormal at all if p_i turns out to be false, since b_i could have been a fake barn just like others; third, as exposed above, feature (iii) is easily satisfied. *S*'s epistemic positions with regard to $p_1, \ldots, p_{20}$ are equal; fourth, feature (iv) is satisfied, for it is known on the reader's side that at least one (indeed, all but one) of $p_1,\ldots, p_{20}$ is false. In this regard, the Barn Façade case is like the case of the lottery which has only one loser. As a result of this, feature (i) is not satisfied. But features (i) to (iii) are enough to imply significant ramifications of the Barn Façade case that have never been offered in the full-blood sense.

If *S* knows that the lottery contains only one loser, even if *S* truly believes that a particular ticket t_i will not win, *S* does not know it. The reason is, presumably, that *S*'s epistemic position with regard to *ticket t_i will not win* is too low for knowledge or even for justification. For, relative to the reference class {tickets of the lottery}, *S*'s epistemic position regarding *t_i will not win* is very low, *via* some expert objective probability function. This is how the internalist version of SRC explains the lottery case involving only one loser. Though my diagnosis of the Barn Façade case as such a lottery case indicates that Henry's epistemic position regarding *b_i is a barn* is low, this type of explanation, as it stands, is not applicable to the Barn Façade case. There are two important differences between the lottery case with one loser and the Barn Façade case.[39]

First, unlike in the lottery case, Henry's belief is an instance of perceptual, non-inferential belief, and not based on any probabilistic evidence. Second, notwithstanding his not knowing that the structure before him is a barn, Henry is well justified in believing it. What accounts for the justification for his belief? The answer to this question would differ between internalist evidentialism and type reliabilism. Goldman, as a type reliabilist, appeals to the fact that the environment focused on in the Barn Façade case is normal for perception – Henry is looking at a barn just as usual, and the process type <perception>, of which Henry's process responsible for his belief that it is a barn is a token, is normally reliable.[40] The internalist can appeal to this reliability by **Reliability Expert** if the relevant reliability expert function is construed in the type reliabilist manner. That is, Henry has in his evidence set *{perception} is normally reliable*, and thereby the degree of the internalist justification for his belief conforms to the reliability of {perception} in {normal environments}. Henry's belief being justified is generally taken to show that the Barn

---

[39] As noted in footnote 7, some, in response to the lottery paradox, deny that *S* is justified in believing lottery propositions, by adding some condition for justification. But the lottery case with one loser is distinguished from the other lottery cases thus far discussed, in that *S*'s epistemic position regarding the lottery proposition is quite low in the former. Thus, insofar as high epistemic positions are necessary for justification, no extra-condition for justification is of relevance here.

[40] This is the view Goldman endorses in his (1986). As mentioned in §3. 4 of Chapter 4, Goldman specifies the relevant environmental type in terms of normalcy. Some type reliabilists, such as Greco (2003) and Kornblith (2002), claim that Henry's perception is not reliable. This means that they are using a different environment type than Goldman does. Presumably, it is the one I discuss below.

Façade case, just like other Gettier cases, requires some extra-condition for knowledge besides justification to be met. Token reliabilist conditions are designed to offer such an extra-condition.

As I noted in §3. 2 of Chapter 5, both sensitivity and safety theories are reckoned to work against the Barn Façade case: Henry's belief that the structure before him is a barn is insensitive, since he would believe so were it not a barn; his belief is unsafe, since, in some of the nearest worlds in which he believes so, his belief is false. I argued that these explanations presuppose that the relevant ordering source is such that the possible worlds in which Henry is seeing a barn façade instead of a real barn are near the actual world. It is plausible to assume that these worlds are remote unless Henry is not in the Barn Façade County. Therefore, the key to elucidating our intuitions varying between the Modified Barn Façade case and the normal perceptual cases is a shift in reference class in the form of ordering source. And presumably, it is what underlies some people's intuition that Henry knows in the original Barn-Façade case: they are using the ordering source tied to the normal perceptual environment. Suppose that Henry is seeing a real barn in a normal perceptual situation. His belief that it is a barn is sensitive: the nearest possible worlds where it is not a barn would be the worlds in which Henry is not seeing any structure or seeing some other type of structure he normally encounters, and Henry does not believe that there is a barn or it is a barn; similarly, his belief is safe in such a situation: the nearest possible worlds where Henry believes that there is a barn or it is a barn would be the worlds in which his beliefs to this effect are true.

Indeed, type reliabilism can offer a similar explanation of the Barn Façade case. Lepock (2009) and Luper-Foy (1987a) argue that two kinds of type reliability are necessary for knowledge, generic reliability tied with normal environments and specific reliability tied with the particular environments *S* is in. According to Lepock and Luper-Foy, generic reliability, which is usually the focus of type reliabilism, is responsible for the justificatory status of *S*'s belief; and specific reliability is the condition that allows for type reliabilism to deal with Gettier cases. That is, in Gettier cases, *S*'s process is not reliable relative to a specific type of environment *S* is currently in. Of course, a variant of the reference class problem arises for this type of reliability. Nevertheless, if the diagnosis of the Barn Façade case I have given above is right, it is easy to specify what environmental type underlies our intuition of this particular case that Henry does not know.

In the Barn Façade case, presumably, the relevant process type is {perception} or {barn-identification}. What is the relevant environmental type? It must be such that it satisfies feature (iii) of lottery case, and thereby *S*'s epistemic positions with regard to $p_1,\ldots, p_n$ are the same, or at least not significantly different from each other. The type to meet this condition is {the presence of b_i: for any i, $i \in [1, 20]$}, or more intuitively {the presence of the barn-like structure in the Barn Façade County}. Relative to this pair of process and environment types, Henry's epistemic position regarding p_i, for any i (the reliability of the token process Henry exercises in forming p_i) is .05. This value is too low to be sufficiently high for any realistic epistemic standard. It follows by **Epistemic Position** that Henry does not know p_i.[41] [42]

[41] Millikan (1984), in passing, mentions that reference class is the key to understanding the Barn Façade case, although she does not specify what she means by 'reference class.' Brandom (1994, pp. 206-13) argues that shift in reference class is responsible for our intuition about the Barn Façade case, and he identifies reference class with environment type. He begins by endorsing what he calls the 'reliabilist insight,' viz., that type

§3. 3. Shift in Reference Class in the Barn Façade Case

If my diagnosis of the Barn Façade case is right, externalist theory of any sort entails SRC on which underlying our shifting intuitions between the Barn Façade case and normal perceptual cases is a shift in reference class. On the other hand, the internalist tends to adopt SES in order to explain our shifting intuitions; Brueckner (1994a), Fogelin (1994), and Williams (2001b) argue that the epistemic standard is raised in Gettier cases.[43] In what

---

reliability is, at least in some cases, sufficient for justification, or entitlement in his terminology. But, in the final analysis, Brandom construes reliability to be attributed to *S* by the attributor when she endorses the inference from *S*'s reliability regarding *p* to the truth of *p*. Reliability in this sense, then, is not type reliability. I take it that Brandom is rather broadening the range of total evidence so as to include not only *S*'s evidence but also the attributor's evidence. If the attributor has in her evidence set a proposition regarding the reliability of *S*'s perceptual report on *p*, and legitimately infers the truth of *p* from this evidence and *S*'s report on *p*, then *S* is justified in believing *p*. More concisely, the idea is that trans-level inference can be performed inter-subjectively. I will discuss the variants of the Barn Façade cases Brandom appeals to, in footnote 45 of Chapter 7.

[42] This result is compatible with pluralist theories of knowledge according to which the strength of epistemic position is a function of specific reliability and some other factor, such as generic reliability or internalist justification. In such theories, *S*'s epistemic position regarding p_i may be assigned some higher value than .05 as a function of multiple factors; but the contribution the specific reliabilist condition makes to the strength of *S*'s epistemic position remains all but nothing.

[43] Williams endorses Brandom's analysis of the Barn Façade case mentioned in footnote 41. But in his study of Gettier cases in general, he heavily relies on Fogelin's SES.

follows, I will give an argument to show that even the internalist must endorse SRC as the best explanation of our varying intuitions concerning the Barn Façade case.

The arguments for the elevation of the epistemic standard, offered by the internalists, are very similar. They are all based on comparative considerations. Suppose that the same Henry is seeing a barn in a normal perceptual environment, the barn just like the one he is seeing in the Barn Façade County, and he perceptually forms the true belief that it is a barn. Let EP_n be the strength of Henry's epistemic position regarding *it is a barn* in this normal case, and EP_b be the strength of his epistemic position regarding the similar proposition he forms in the Barn Façade case. There seems – though I will argue that this seeming is misleading – no epistemically relevant difference between the normal case and the Barn Façade case: he is just alike with regard to empirical evidence for his belief, total evidence, and perceptual resources accessible to him, between the two cases. Then, $EP_n = EP_b$. Henry has knowledge in the normal case, and so $1 - \varepsilon_n \leq EP_n$, where '$1 - \varepsilon_n$' denotes the threshold for knowledge imposed by the epistemic standard operative in the normal case. But Henry fails to have knowledge in the Barn Façade case, and so $EP_b < 1 - \varepsilon_b$, where '$1 - \varepsilon_b$' denotes the threshold for knowledge imposed by the epistemic standard operative in the Barn Façade case. The only explanation for his epistemic asymmetry is that the epistemic standard for knowledge, what counts as a sufficiently high epistemic position, shifts between the two cases, i.e., $1 - \varepsilon_n \leq EP_n = EP_b < 1 - \varepsilon_b$.

This argument relies on internalist evidentialism as a monist theory of knowledge. Since the issue is whether internalist evidentialism alone has resources to explain our shifting intuitions, externalism cannot be assumed in order to criticize the argument. But the question of whether SES or SRC best explains our judgment that Henry does not know in the Barn Façade case can be settled without assuming internalism or externalism.[44] I will propose a non-partisan or theory-neutral argument that Henry's epistemic position is lowered in the Barn Façade case.

My argument, too, goes *via* a comparative consideration: first, consider what is necessary for Henry to know in the Barn Façade case. Not much is required. He can easily have knowledge, and thereby get EP_b as high as $1 - \varepsilon_b$, by going behind the barn or being told by a local that it is a real barn, i.e., by getting information that it is not a barn façade.[45] This is possible even when his epistemic position with respect to such information is still less than maximal (if it is not, the Barn façade case simply boils down to a case of infallibilist skepticism). So, ε_b must not be so small; next, consider how Henry can improve his epistemic position in the normal case. There are a number of ways, e.g., by having better eyesight, or collecting more information on the surface structure of a barn, and so forth. Since ε_b is not small, this suggests that it is possible that Henry's epistemic position in the normal case gets as high as $1 - \varepsilon_b$ while his epistemic position is less than maximal ($1 - \varepsilon_b \leq EP_n < 1$). But it does not seem sufficient for Henry to know in the Barn Façade case that Henry has much better eyesight or that he is an expert on the structure of barns. Still, Henry does not know in the Barn Case ($EP_b < 1 - \varepsilon_b$). It follows that $1 - \varepsilon_b \leq$

[44] DeRose (1996b) offers two related reasons against Bruckner's objection, but they both beg the question against Bruckner, since DeRose assumes externalism at the outset; DeRose, in his (2009), adds another reason, but it is merely based on his intuition.

[45] It is because this step cannot be taken that this argument is not applied to show that the strength of epistemic position is lowered in lottery cases; in lottery cases, the strength of epistemic position is set at a very high value.

EP_n, and that $EP_b < 1 - \varepsilon_b$. Therefore, $EP_n \neq EP_b$. This conclusion is consistent with the possibility that $1 - \varepsilon_n < 1 - \varepsilon_b$. That is, it is possible that both the epistemic standard is raised and the epistemic standard is lowered in the Barn Façade case. However, once $EP_n = EP_b$ is denied, there is no reason to postulate the two kinds of shifts here.

This argument strongly supports the claim that SRC is required to explain our intuition about the Barn Façade case, since, among resources available to theories of knowledge, reference class is the only factor to shift the strength of epistemic position between the normal perceptual cases and the Barn Façade case. Even if it is granted, as the internalists contends, that the degree of internalist justification does not vary between the normal perceptual cases and the Barn Façade case, a pluralist position must be in order; at least one determinant of the strength of epistemic position must be an externalist factor to which SRC is combined.

Chapter 6: Shifts in Reference Class (ii): the Case of Easy Knowledge and the Case of Cartesian Skepticism

Introduction

This chapter is a continuation of the last chapter, and offers diagnoses of two further kinds of epistemologically interesting cases: cases of easy knowledge and cases of Cartesian skepticism. Again, I argue that shifts in reference class are responsible for our varying intuitions about these cases. In §1, I argue against the problem of easy knowledge that we can acquire second-order warrant for the reliability of a source, when independent sources mutually function as the checker of each other's reliability. It is part of the choice of the relevant reference class(es) what sources are available to *S* and whether they count as independent of each other. §2 analyzes several arguments for Cartesian skepticism. Standard contextualism explains skeptical cases in terms of shifts in epistemic standard, but it fails to deal with the full-blooded version of Cartesian skepticism. This version of skepticism shares the essential structure with the problem of easy knowledge, and hence the same type of diagnosis is applied to both. My diagnosis, then, gives further reason for non-standard contextualism, because contextual variance in reference class needs to be postulated to save our everyday intuition that we know and skeptical intuition that we do not.

§1. Cases of Easy Knowledge

§1. 1. The Problem of Easy Knowledge

What is called the ‘problem of easy knowledge’ comes in different forms, and philosophers disagree as to what it, even in a particular form, really amounts to. As regards to the form of the problem, I am here only concerned with the one concerning track-record arguments or bootstrapping, first raised in clear fashion by Alston (1986).[1] As regards to the substance of the problem, I take it to consist in the difficulty of finding some other way to acquire knowledge of the reliability of alleged basic sources than bootstrapping.

Alston raises the problem of easy knowledge as a problem for type reliabilism, and such internalists as Fumerton (1995) and Vogel (2000) argue that despite type reliabilists’(including Alston’s own) efforts to handle it, no good solution on behalf of type reliabilism is in the offing. Cohen (2002), however, points out that the problem of easy knowledge arises not only for type reliabilism but also for any other theory of knowledge that allows for ‘basic knowledge.’ In his definition, a theory allows for basic knowledge or has a ‘basic knowledge structure’ (BKS) iff it denies the following requirement:

[1] Besides the bootstrapping version of the problem of easy knowledge, the ‘problem of easy knowledge’ also refers to the closure version of the problem of easy knowledge and the problem with dogmatism as an anti-skeptical strategy. Cohen (2002) offers an alleged case to illustrate the former, but I do not see why it is a problem; there is no clear sense in which his case involves failure of knowledge by deduction. I briefly mention dogmatism in footnote 40.

> (KR): A potential knowledge source *K* can yield knowledge for *S*, only if *S* knows *K* is reliable. (Cohen, 2002, p. 309)[2]

KR is a principle that states a pre-condition for first-order knowledge, and weaker principles pertaining to warrant or justification are possible. Once being weakened so as to be about justification, KR is similar to **Meta-Justification**, that first-order justification requires second-order justification. Although the problem of easy knowledge can be formulated differently in terms of different principles, differences between them are not of much interest here. [3]

KR is intended to be consistent with a broad range of theories, and it applies to each theory differently. In particular, depending on what theory it is applied to, (a) what source *K* is, and (b) what it is to be reliable, can vary. For internalist evidentialism, Cohen suggests, *K* may be propositional evidence for *p*, and to know that *K* is reliable is to know that *K* makes *p* probable in the sense of epistemic probability. He implies that internalist evidentialism is congruent to KR if it has something like the requirement of trans-level inference (so he seems to ground KR on **Meta-Justification**). For type reliabilism, *K* is a process type, and for token reliabilism, it is a method token.

KR is inconsistent with the existence of 'basic' sources of knowledge. Such sources, as perception, memory, and introspection, are typically reckoned *basic*, on the

---

[2] Cohen does not specify what 'KR' stands for. I take it that it abbreviates the principle or the requirement of 'knowledge of reliability.'

[3] The problem of easy knowledge, in any form, does not arise from the gap between knowledge and its epistemic components.

grounds, not only that knowledge from any of them is produced non-inferentially, but also that it can be held without having knowledge (warrant, or justification) of its reliability. For example, how can we know that perception is reliable without knowing its deliverances? The negative answer to this question leads to the denial of KR.

No BKS theory can evade the problem of easy knowledge, whether it is externalist or internalist.[4] BKS theories have to accept cases where *S* achieves knowledge too easily by means of a 'track-record argument' (Alston) or 'bootstrapping' (Vogel).[5] It is difficult to characterize precisely what bootstrapping is, given the widespread usage of the term in the recent literature. Here, I simply follow Alston, and take it to be inductive inference of the following form, where 'basic belief' refers to belief that is knowable in ways inconsistent with KR:

(1) At t_1, *S* formed, by *K*, the basic belief that p_1, and p_1.

(2) At t_2, *S* formed, by *K*, the basic belief that p_2, and p_2. (and, so on.)

(C) Therefore, *K* is a reliable source of belief.

Why is bootstrapping so problematic? Vogel illustrates this point by way of example:

---

[4] As noted in §2 of Chapter 3, an internalist theory is possible without endorsing **Meta-Justification** or the requirement of trans-level inference.

[5] Vogel takes the term 'bootstrapping' from Williams (1991). The basic setting of the Gas Gauge case below is also taken from the case Williams presents (p. 347).

(Gas Gauge) Roxanne, in driving her car with the reliable gas gauge, believes what it says, without knowing that it is reliable. So, when the gauge reads '*F*,' she believes both that, at time t_1, the gauge reads '*F*,' and that, at time t_1, the tank is full. Moreover, she combines these beliefs, and believes:

(1) At time t_1, the gauge reads '*F*' and *F*.

Now, Roxanne does this over and over again: at various times, she reads what the gauge says, and accordingly forms beliefs of the form 'at time t_i, the gauge reads '*X*' and *X*.' By putting such beliefs together, Roxanne concludes:

(C) The gauge is reliable.[6]

Three different belief-forming processes are involved in Roxanne's bootstrapping: (i) seeing where the gauge needle is, (ii) detecting the state of the tank on the basis of (i), and (iii) statistical induction leading to (C) *via* repeated applications of (i) and (ii). All (i) through (iii) are reliable processes even though Roxanne has no warrant or justification about the reliability thereof. The focus of the case is (ii), since the reliability of the gauge in (C) is tantamount to the reliability of the gauge's telling the truth about the state of the

[6] In Vogel's description, the inferential process Roxanne goes through is more complex, involving more intermediate steps. The case here is a bit simplified so that Roxanne's inference is of the exact form given by Alston.

tank. In what follows, I refer to (ii) simply as {reading the gauge}. (iii) is causally and epistemically dependent on (ii), in that (ii) is involved in the process of (iii), and the reliability of (iii) obtains conditional on the reliability of (ii). (i) is not of much relevance, since it does not affect the Gas gauge case whether Roxanne knows the reliability of (i).

At the very least, intuitively, the bootstrapping process Roxanne uses in the Gas Gauge case does not seem to be a legitimate means to know, have warrant for, or even be justified in believing (C). With this case, Vogel argues that type reliabilism faces the problem of easy knowledge, since it permits knowledge of *K*'s reliability by bootstrapping. Cohen further adds that internalist BKS theory has the same result. And, I may add, safety theory may lead to the same consequence, since (C) is safe, if the gas gauge is safe in the sense that it tells truth in the nearest possible worlds where Roxanne reads it.[7]

Bergmann (2004, p. 168) derives what he calls the 'no self-support principle' from our intuition about the Gas Gauge case, and uses it to explain how the denial of KR leads to the problem of easy knowledge:

> (NSS): One cannot obtain a justified or warranted belief that a belief source *K* is trustworthy by relying on *K* or *K*-dependent sources alone.[8]

---

[7] Cohen claims that sensitivity theory can circumvent the closure version of the problem of easy knowledge. The same may be said of the bootstrapping version of the problem. For, given the setting of the Gas Gauge case, if the gas gauge of Roxanne's car were not reliable, she would still believe that it is.

[8] The formulation of NSS here is taken from Cohen's (2002, p. 319) reformulation of Bergmann's, but it is a bit modified. Bergman's is in terms of epistemic properties; and both Bergmann's and Cohen's involve the locution 'even in part relying on source *K*.' I take it that this locution is introduced because Roxanne's

What NSS does is to require an independent checker of *K*'s reliability for acquiring justified or warranted belief about *K*'s reliability. *K* or *K*-dependent sources alone cannot function as such an independent checker.

The Gas Gauge case generates the problem of easy knowledge in the form of a dilemma: the first horn of the dilemma is the denial of KR, and it is motivated by the existence of basic sources; NSS constitutes the second horn of the dilemma: if KR is denied for basic sources, it must be possible that one can know their reliability without merely relying on them or sources depending on them, i.e., NSS must be false for basic sources; at least, insofar as it is true that we know their reliability. Bergmann (2004) and Van Cleve (2003), in construing the problem of easy knowledge as a version of skepticism, argue that no other means than bootstrapping seems available for the purpose of gaining knowledge of the reliability of basic sources.[9] Then, a source being basic entails that we have no means

---

bootstrapping is dependent on (ii). NSS being given here may better capture Bergmann's and Cohen's intention. For the aim of introducing NSS is not to block *K* from being involved, but to require some independent source, in acquiring warranted belief about *K*'s reliability. Bergmann attributes NSS to internalists, and in particular, to Fumerton (1995). Black (2008b) and Vogel (2000, 2008) also endorse one or another principle akin to NSS.

[9] Bergmann, Van Cleve, and Sosa (2009, ch. 10) argue that a version of Cartesian skepticism follows from KR and the existence of basic knowledge: NSS prohibits bootstrapping as a means to know the reliability of perception. Then, from KR, it follows that we do not have perceptual knowledge. If bootstrapping is the only way to know the reliability of perception, the denial of KR leads to the denial of NSS. In §2. 4, I explicate in detail what presuppositions are involved in this argument for Cartesian skepticism.

to acquire knowledge of its reliability other than bootstrapping. Given that we know the reliability of basic sources in everyday situations, there must be a way-out of the dilemma.

A gloss as to how the Gas Gauge case is related to the dilemma may be required. The Gas Gauge case *per se* does not exclude the possibility that in ordinary cases, we can know the reliability of the gas gauge from a different source than reading it, say, by reading a magazine article on or having someone else tell about the functional quality of the brand of the gas gauge. But even when such means are unavailable, it seems that we can still know that the gauge is reliable while driving a car.[10] More importantly, the Gas Gauge case is designed to be parallel to cases in which we must rely on a basic source, typically perception, in order to know its reliability; (ii) is assimilated to the process of detecting the way the external world is on the basis of perceptual or sensory experience.

Most alleged solutions to the problem of easy knowledge in this form are either to deny NSS or to argue that there is some *a priori* way to acquire knowledge of the reliability of the gauge.[11] But I argue, in what follows, that a solution is possible without denying

---

[10] The belief that the new gauge is reliable may be warranted on the ground that it is a new gauge, and new gauges are generally reliable. I assume that this evidence is not available to us, e.g., we do not know that the gauge is new.

[11] Alston (1986), Bergmann (2004), Markie (2005), Van Cleve (2003), Williams (1991) use the problem of easy knowledge as a support for one or another version of externalist BKS theory. They argue that the cases of easy knowledge are those in which we rely on some basic source in acquiring first-order knowledge, and the only way to allow for such knowledge is to deny NSS and KR, i.e., to endorse bootstrapping as a legitimate way to acquire knowledge of reliability. On the other hand, Cohen (2002) and Sosa (2009, ch. 10) argue that the reliability of a basic source is *a priori* knowable, or at the very least, justifiable for *S*, on account of the coherence of *S*'s belief system. Van Cleve objects to this solution that it makes no practical

NSS or the *a posteriority* of knowledge of the gauge's reliability. My focus here is on the Gas Gauge case, and I extend my solution to perceptual cases in §2. 4.

§1. 2. Shift in Reference Class in the Case of Easy Knowledge

Although the solution to the problem of easy knowledge is not usually sought this way, it may be worthwhile contrasting Roxanne's situation with its ordinary counterpart. For the dilemma arising from the Gas Gauge case is acute only if: (a) Roxanne does not know that the gas gauge is reliable, and (b) we, in everyday situations analogous to Roxanne's, know that the gas gauge is reliable.

It might be tempting to account for the difference between (a) and (b) by different epistemic standards being associated with them. Cohen (2005), however, denies the relevance of varying epistemic standards for the problem of easy knowledge, on the ground, roughly, that there is a principled reason that knowledge is not to be gained too easily. Applied to the Gas Gauge case, his point is that bootstrapping is no legitimate way to acquire knowledge, regardless of what epistemic standard is invoked.[12] This consideration indicates that there is an epistemic difference between Roxanne and us. Such a difference is only explained in terms of epistemic resources available to Roxanne and us. Roxanne's

---

difference to the Gas Gauge case: it may be incorporated into Roxanne's story that her belief system is perfectly coherent, but this would not affect our intuition that she does not know the reliability of the gauge.

[12] Cohen is concerned with the closure version of the problem of knowledge he raised in his (2002). But some of his reasons for the irrelevance of epistemic standard are applicable to the bootstrapping version as well.

epistemic resources are even more limited than ours, since the setting of the Gas Gauge case restricts the range of processes she can exploit in forming the belief that the gauge is reliable. The Gas Gauge case is indeed a peculiar case in which only one process is allowed to be put to work.[13]

Suppose that you just had a new gas gauge installed into your car, and you have no evidence for or against its reliability, even though it is in fact reliable. After the installation, you have driven your car for days. Do you acquire warrant or come to know that the new gas gauge is reliable? If your situation is just like Roxanne's, you must not.

Notice that in the everyday situation of driving a car, multiple processes are in play when you believe (1) and other similar propositions, of the form 'at time t_i, the gauge reads '*X*' and *X*.' While you are driving a car, you avail yourself of a variety of types of information from different sources, and they all contribute to your epistemic position with regard to (C). For example, you *see* the gas needle go up and down, *see* how fast or slow the scenery is passing, *hear* how hard or dull the engine is working, *feel* how shaky or stable the car is moving, and so on. The cognitive processes by which you acquire information of these types can be classified as {reading the gauge}, {seeing the passing scenery}, {hearing the engine sound}, {feeling the car's movement}, respectively. As a

[13] So, this solution can be combined with non-standard contextualism I develop in the next chapter. We are in a certain context when we correctly judge that Roxanne does not know, whereas we are in a different context when we correctly judge that one knows in situations similar to Roxanne. Cohen accepts that standard contextualism has no application to the problem of easy knowledge. Van Cleve (2003) seems to be aware that something like a contextualist solution to the problem of easy knowledge is at least logically possible; he mentions that an implicit premise involved is that "the term 'knowledge' has been used univocally…" (p. 50).

matter of fact, one can and does use tokens of these process types to check the reliability of the gas gauge.

Insofar as the gauge is reliable, these types of information bring about further information on how the gauge co-varies with the car, viz., that the faster or harder the car runs, the farther the gas gauge goes down to '*E*.' It is safe to assume that, prior to the occasion of driving the car with the new gas gauge, you know general facts about cars, such that cars run faster or heavier when scenery passes faster, or engines sound louder, or car bodes move more wildly, and that the faster cars run, the more gas they take.[14] Given the antecedent knowledge of these general facts, you have a means by which you can check on the reliability of the gauge, independently of process {reading the gauge}. Such a means consists in comparing how the gauge is working with how the car is running. The information of how the gauge is working is, presumably, gained *via* process {reading the gauge}, and the information of how the car is running is *via* some or all of processes {seeing the passing scenery}, {hearing the engine sound}, and {feeling the car's

[14] Of course, there are more general truths about the functions of car, gas gauge, and gas tank. For example, you may know that the more gas is in tanks, the further gas gauges go up. But this piece of information is irrelevant for the case at hand. Roxanne is stipulated not to have charged gas. For doing so gives independent evidence for believing that when the tank is full, the gauge reads '*F*,' which, in turn, supports that the gauge is reliable. The case here is designed so as to be as close to the Gas Gauge case as possible.

movement}. Put differently, you can check whether the gauge is reliable, by appeal to how the gas gauge coordinates with the state of the car.[15]

The information from these different sources is implicitly involved when you form beliefs of the form 'at time t_i, the gauge reads '*X*' and *X*.' More importantly, it allows you to know (C) without relying on the single source {reading the gauge} alone. This is not a mere intuitive matter; it is supported by the following probabilistic consideration.

Let us define the notion of coordination more precisely: sources $K_1, \ldots, K_n$ coordinate iff their deliverances $e_1, \ldots, e_n$ indicate the occurrence of the same event *p*. For example, if tokens of {hearing the engine sound} deliver that the engine is working heavily, and tokens of {reading the gauge} deliver that the gauge is going down fast, then the two process types coordinate, because they indicate the same event that the car is running fast. Our warrant or strength of epistemic position regarding the coordination of sources admits of degrees, and is described probabilistically; e.g., on frequentism, it is the degree of the reliability determined by how often process tokens coordinate, and, on the epistemic theory, it is the degree with which *S* rationally expects the coordination of sources. While leaving open what '*EP*' denotes, here, I assume that the strength of epistemic position conforms to the probability calculus, and identify the degree of coordination of $K_1, \ldots, K_n$ with the value of $EP(e_1 \wedge \ldots \wedge e_n)$, where each of $e_1, \ldots, e_n$ is to the effect that a deliverance of each of $K_1, \ldots, K_n$ indicates *p*.

Let *h* be 'K_1 and K_2 are reliable.'[16] The following is true by Bayes' theorem:

---

[15] A similar point is made by Vogel (2008, pp. 531-2), though he states it in terms of coherence of belief system. Vogel, however, argues that this is not enough for a solution to the problem of easy knowledge; for him, its full-blown solution requires elucidating which step of bootstrapping is illegitimate and why.

(1) $EP(h \mid e_1 \wedge e_2) = (EP(h) / EP(e_1 \wedge e_2)) \times EP(e_1 \wedge e_2 \mid h)$

h makes it more probable that both e_1 and e_2 are true than $\sim h$ does, since K_1 and K_2 are more likely to deliver the same result p if they are both reliable than if at least one of them is unreliable; put differently, unreliable processes are more likely to deliver diverse results than reliable ones. Then,

(2) $EP(e_1 \wedge e_2 \mid h) > EP(e_1 \wedge e_2 \mid \sim h)$

In cases of the problem of easy knowledge and their ordinary counterparts, before having deliverances of the source(s) in question, S has no evidence for or against its reliability of, and neither can S tell whether it is reliable or not. Hence, (3) $EP(h) = EP(\sim h)$. From premises (1) through (3), it follows:

$EP(h \mid e_1 \wedge e_2) > EP(\sim h \mid e_1 \wedge e_2)$

---

[16] Since reliability comes in degrees, h is more precisely 'K_1 and K_2 are reliable to degree r.' The value of r affects the probabilities involved in the proof below, but I omit the complication arising from it.

Therefore, when K_1 and K_2 coordinate, their deliverances increase the strength of epistemic position with regard to the reliability of each of K_1 and K_2 more than with regard to the unreliability thereof.[17] [18]

This is not enough to show that S knows or even is justified *simpliciter* in believing that K_1 and K_2 are reliable. $EP(h \mid e_1 \wedge e_2)$ may be still less than the threshold for knowledge or justification. The value of $EP(h \mid e_1 \wedge e_2)$ is negatively proportional to $EP(e_1 \wedge e_2)$, and positively proportional to $EP(e_1 \wedge e_2 \mid h)$, i.e., it depends on $EP(e_1 \wedge e_2 \mid h)$ - $EP(e_1 \wedge e_2)$. For this reason, it is crucial for EP to be as high as the threshold that $EP(e_1 \wedge e_2)$ is sufficiently low. $EP(e_1 \wedge e_2)$ seems to depend on the degree of the independence of K_1 and K_2, in that it gets lower as they are expected to function more independently of one another, prior to gaining e_1 and e_2.

When are K_1 and K_2 expected to be independent? It is difficult to state precise conditions for two sources being independent. This is indeed part of the reference class

[17] This proof is an instance of the schema, proved by Myrvold (2003), to the effect that if theory h unifies phenomena $e_{1,}$ …, e_n that are independent on prior grounds, the degree to which h is supported by $e_{1,}$ …, e_n depends on how unifying h is with regard to $e_{1,}$ …, e_n. His point is that $e_{1,}$ …, e_n provide direct, empirical support for h, as opposed to indirect, non-empirical support. This feature of the schema differentiates the solution to the problem of easy knowledge given here from the coherentist solution. For other differences, see footnote 52.

[18] Having e_1 and e_2 may require warrant or justification for them. It seems possible to acquire first-order warrant or justification for a belief and second-order warrant or justification that its source is reliable at the same time. To work out this point requires more inquiry. But if this is the case, the proof here does not violate KR. Note that accepting KR does not solve the problem of easy knowledge unless the question as to how the reliability of alleged basic sources is known is settled.

problem. First, there are reasons to regard, say, {hearing the engine sound} and {reading the gauge} as independent to a low degree. Auditory sense and visual sense causally work independently in some relevant sense, because the functioning of your vision does not affect the functioning of your hearing; however, there are reasons to regard them as dependent to a high degree. They both belong to the general type {perception} or {sensory experience}. When this broad type is employed, their success and failure are taken to go hand-in-hand with one another. We will see in §2. 4 that the extraordinary broad typing underlies the problem of Cartesian skepticism. Second, the degree of the independence of K_1 and K_2 depends on the environment type to which they are relativized. {reading the gauge} may be reckoned to be independent of {seeing the passing scenery}, if the former is relativized to {right in front of eyes} and the latter to {at a distance}. Visual sense may or may not be reckoned to work differently between near distance and middle distance.

When the reference classes are properly selected, $EP(h \mid e_1 \wedge \ldots \wedge e_n)$ can rise as high as the threshold for knowledge or justification. In addition, $EP(e_1 \wedge \ldots \wedge e_n)$ decreases, and hence $EP(h \mid e_1 \wedge \ldots \wedge e_n)$ increases, as n adds up. It seems that we type or individuate sources narrowly in ordinary situations, and many *e*s are available.

This is a solution to the dilemma generated by the Gas Gauge case: it accepts NSS and denies that bootstrapping gives knowledge that alleged basic source K is reliable; but it also denies that the only resource available for us to acquire such knowledge in ordinary situations is K itself. Knowledge about the reliability of K is not easy; we need to exercise multiple sources, and empirically check on the reliability of them by reference to how they coordinate with each other (how this is done may differ from theory to theory). The reliability knowledge thus gained, then, is still *a posteriori*. The intuition that Roxanne does not know the reliability of the gauge seems to be driven by the setting of the Gas Gauge case that excessively restricts the range of the processes available to Roxanne. To describe Roxanne's process as bootstrapping has this effect, since, by definition, bootstrapping is the inferential process ultimately relying on {reading the gauge} alone.

Nothing thus far said prevents *S* from performing inference of the bootstrapping form. But when $EP(h \mid e_1 \wedge \ldots \wedge e_n)$ is high enough for knowledge or justification, it is not the process of inference that yields knowledge or justification, rather only making explicit what *S* already has.[19]

The solution to the dilemma is a variant of SRC in that it explains our varying intuitions regarding (a) and (b) by postulating a shift in reference class. When the choice of reference class is such that our belief in *K*'s reliability is a product of many independent processes, we can know the reliability of them; when the choice of reference class is such that it is primarily based on a single process, we cannot.

§2. Cases of Cartesian Skepticism

§2. 1. The Problem of Cartesian Skepticism

[19] In other words, inference of the bootstrapping form does not transmit warrant from first-order propositions to second-order proposition about *K*'s reliability. Wright (2007) and Neta (2007a) use a similar point against dogmatism. For internalism, bootstrapping may be still important, since it renders propositional justification or warrant doxastic.

The 'problem of Cartesian skepticism' refers to a family of problems rather than a single problem, though they may not be distinguished from each other with enough care. What they have in common are the following features: first, each variant of Cartesian skepticism consists of an argument that purports to establish the impossibility of *knowledge about the external world* or *empirical knowledge*; second, it is *global*, in that it puts into doubt every alleged piece of empirical knowledge with the same argument; third, it involves a *hypothetical scenario* in which *S* is subject to global epistemic deception. Each variant, then, questions how knowledge about the external world is possible against the skeptical doubt.[20]

Cartesian scepticism is generated *via* different forms of arguments. The standard form of argument, which DeRose (1995) calls the 'argument from ignorance,' is formulated as follows, where 'HAND' refers to the proposition that *S* has a hand, and 'BIV' a skeptical hypothesis incompatible with HAND, e.g., that *S* is a handless brain-in-a-vat:

(1). *S* does not know ~BIV.

(2). If *S* does not know ~BIV, and knows that HAND entails ~BIV, then *S* does not know HAND.

(3). Therefore, *S* does not know HAND.

[20] Conant (2004) points out that a variant of Cartesian skepticism, what he calls the 'Kantian problem,' questions how it is possible for mental acts to have objective imports, or to be directed at the external world. I am not concerned with this problem.

By contraposing (2), we can easily run a *modus tollens* on this argument, which is the so-called Moorean proof:

(1'). *S* knows HAND.

(2'). If *S* knows HAND, and knows that HAND entails ~BIV, then *S* knows ~BIV.

(3'). Therefore, *S* knows ~BIV.

The BIV hypothesis can be adjusted so as to be incompatible with any other everyday proposition about the external world. From the argument from ignorance repeated for each such proposition, the global skeptical conclusion that *S* does not know anything about the external world follows. But, at the same time, the Moorean proof is possible for any everyday proposition.

The difficulty of Cartesian skepticism, manifest in these two arguments, consists in the fact that we have strong intuitive pull toward (1) and (1'), and they contradict each other. Moreover, the argument from ignorance (and the Moorean proof) is valid; once our intuitions of (1) and (1') are taken for granted, the only way to refute it is to deny (2) and its contrapositive (2'). They, however, are instances of SCK, viz., that knowledge is closed under known entailment with a single premise.[21]

§2. 2. Shift in Reference Class in the Case of Cartesian Skepticism for Externalism

[21] As noted in footnote 34 of Chapter 5, Dretske and Nozick endorse that sensitivity entails the denial of SCK, and for the very reason, the sensitivity solution to skepticism is highly contentious.

The contextualist response to the problem of Cartesian skepticism has been reckoned to have at least two advantages over other responses: first, it saves the phenomena, our shifting intuitions about (1), (1'), (3), and (3'); second, it is consistent with SCK, at least in the sense that it holds for any knowledge ascription in a context. The idea behind the first point is that the soundness of the argument of ignorance and its accompanying intuitions of the truth of (1) and (3), on the one hand, and the soundness of the Moorean proof and its accompanying intuition of the truth of (1') and (3'), on the other hand, are a function of different contexts with epistemic standards differing in stringency, and the contextual shift in standard is caused by the salience of the skeptical hypothesis to the attributor.[22]

Being proponents of different theories of knowledge, standard contextualists, such as Cohen (1988, 1999, 2000), DeRose (1995), and Lewis (1996), combine the contextualist apparatus with different theories of knowledge; as a result, their diagnoses of why each

---

[22] Stanley (2005, pp. 125-30) argues that SSI can mimic the contextualist solution to the skeptical problem, by replacing the notion of context of use with that of *S*'s non-epistemic situation. On Hawthorne's SSI, *S*'s non-epistemic situation includes a certain psychological feature of *S*, what error possibilities are salient to *S*. But it fails to accommodate our intuitions about their-person knowledge denials about skeptical cases. Stanley and Fantl & McGrath focus exclusively on *S*'s pragmatic situation, how much is at stake on *p*. As Stanley is aware, it is difficult to apply this version of SSI to the first-person skeptical cases; how is it that *S*'s considering skeptical possibilities raise stakes on everyday propositions like HAND? (for this point, see (Williams, 2004, p. 188)) Notwithstanding this difficulty, he contends that when stakes are low, his SSI predicates that *S* knows *p*. Fantl & McGrath, in view of this problem, countenance that contextualism may have advantages over SSI as a response to skepticism and to the lottery problem (Fantl & McGrath, 2009b, p. 33).

premise seems correct in a context and their accounts of how SCK is maintained differ from each other. Here is DeRose's:

> Consider my belief that I have hands. I believe this at the actual world and it's true. What's more, in the other nearby worlds in which I have hands I believe that I do. There are also, at least in my own case, some alarmingly close worlds in which I don't have hands. These include worlds in which I lost my hands years ago while working on my uncle's garbage truck. In the closest of these not-*P* worlds, I'm now fully aware of the fact that I'm handless, and my belief as to whether I have hands matches the fact of the matter. My belief as to whether I have hands doesn't match the fact in various worlds in which I'm a BIV, of course, but these are *very* distant. (DeRose, 1995, pp. 34-5)

On DeRose's counterfactual robustness theory, the epistemic standard for knowing *p* prescribes how far out in the possibility space *S*'s belief in HAND must match the fact as to whether HAND is true. In quotidian contexts where a relatively lax standard is in play, '*S* knows HAND' is true at *w* iff *S*'s belief as to whether HAND matches the fact concerning HAND in possible worlds relatively close to *w*. If *S* satisfies this condition, DeRose argues, *S* also satisfies a similar condition for knowing ~BIV, and thereby '*S* knows ~BIV' is true in these contexts. That is, SCK is preserved on the counterfactual robustness theory; whereas, in skeptical contexts where skeptical alternatives like the BIV hypothesis are salient to the attributor, a stringent standard comes into play. Then, '*S* knows HAND' is true at *w* iff *S*'s belief as to whether HAND matches the fact regarding

HAND in possible worlds very remote from *w*. Insofar as such worlds include a BIV world, *S* fails to satisfy this condition, since, of course, the BIV worlds are the ones in which *S* falsely believes that *S* is not a BIV. Again, the result is consistent with SCK, since, in the skeptical contexts, *S* also fails to satisfy the condition for knowing BIV.

Is this a satisfactory solution to the problem of Cartesian skepticism? Klein (2000), Kornblith (2000), and Williams (2001a) contest that it underestimates the genuine threat of Cartesian skepticism; according to them, Cartesian skepticism is construed, to use Kornblith's terms, as either a high-standard skepticism or a full-blooded skepticism: high-standard skepticism, even if its argument is successful, only establishes that one does not have knowledge by a stringent, perhaps, infallibilist standard; full-blooded skepticism, on the other hand, attempts to prove that one does not have knowledge *simpliciter*, i.e., one does not have knowledge even relative to the ordinary, lax standard. These critics argue that the real threat of Cartesian skepticism consists in its potential for being a full-blooded skepticism. Thus, standard contextualism, because it is SES, does not even address the genuine problem of skepticism, only responding to high-standard skepticism. Stanley (2005, pp. 125-30) likewise accepts that SSI, being SES, is in no better position than contextualism to respond to full-blooded skepticism.

Furthermore, the critics claim that full-blooded skepticism concerns internalist justification: thus, the conclusion of Cartesian skepticism, in its full-blooded form, must be not so much that *S* is not internalistically justified in believing HAND to some unusually high degree, as that *S* is not internalistically justified at all. However, it is one thing to claim that the contextualists, whether internalist or externalist, are mainly concerned with high-standard skepticism – which is true. It is another to claim that the full-blooded version of Cartesian skepticism is necessarily an attack on internalist justification – which is false, or so shall I argue.

As a matter of fact, not only DeRose, the contextualist, but also those invariantists who address Cartesian skepticism by means of token reliabilism, reckon Cartesian skepticism to be high-standard skepticism. For they are all committed to the assumption that the skeptical worlds are distant from the actual world[23]; note that DeRose, in the quote above, explicitly states that the skeptical worlds are 'very distant.' In a similar vein, Sosa, in offering a safety-based response to skepticism, claims:

> In the actual world, and for quite a distance away from the actual world, up to quite remote possible worlds, our belief that we are not radically deceived matches the fact as to whether we are or are not radically deceived. (Sosa, 1999, p. 147)

---

[23] Also, proponents of the near-actual frequency model of type reliabilism are committed to this assumption. For example, Greco (2000) claims:

> … it is clear that the world in which I am a brain in a vat is a very far world from the actual world, given that the actual world is anything like we think it is. (p. 217)

Greco explains the remoteness of the BIV worlds in terms of the difference in *S*'s constitution: they are remote because *S*'s constitution in the BIV worlds is very different than in the actual world. The 'constitution' in Greco's sense is the characteristics of *S* that are relevant for *S*'s abilities of belief-formation. It is not clear at all how Greco distinguishes *S* from BIV in constitution. Moreover, I suspect that he begs the question against the skeptic. For the point of Cartesian skepticism is that *S* and BIV are not relevantly different in ability of belief-formation, since they are able to have the same sensory experience.

And here is what Pritchard, another proponent of safety theory, notes:

> Safety thus allows the possibility of everyday knowledge just so long as sceptical possible worlds are, as a matter of fact, modally remote. (Pritchard, 2005a, p. 72)

Just as with the counterfactual robustness theory, safety and sensitivity theory, if they are designed to militate against Cartesian skepticism, presuppose that the skeptical worlds are far away: for safety, it is because the nearest worlds where *S* believes HAND are those where HAND is true that *S*'s belief in HAND is safe; similarly, for sensitivity, it is because the nearest possible worlds where HAND is false are those where *S* does not believe HAND that *S*'s belief in HAND is sensitive. This entails that, in order for these theories to work, the skeptical worlds must be farther out than the worlds with regard to which *S*'s belief is safe or sensitive; otherwise, safety or sensitivity does not hold for HAND. Perhaps, the skeptical worlds are located at the margin of possibility space. On these theories, the strength of *S*'s epistemic position is proportional to the distance to the relevant worlds with regard to which *S*'s belief is safe or sensitive. Thus, the safety and the sensitivity response are only available against high-standard skepticism. They amount to the position that knowledge is possible by a lax, fallibilist standard, even though knowledge is impossible by a stringent, infallibilist standard. Thus, the point of the invariantist response from sensitivity and safety is that if **Fallible Knowledge** is true, i.e., if knowledge is consistent with the margin of error, then knowledge is possible.

How are those theorists entitled to assume that the skeptical worlds are modally remote? Curiously, they do not say much of this, and merely suggest, at best, that the skeptical hypotheses are too far-fetched, or that the skeptical worlds are very different from actuality. This is hardly an answer to the question; as I have emphasized at several points, every world is similar, and hence can be close, to each other in some respect. It might be true that certain skeptical worlds, e.g., the BIV worlds, are considerably different from the actual world on some *objective* ordering source – for example, over-technology in the BIV worlds makes a difference in an objective respect –, and even on some objectively probabilistic ordering source – the objective probability of the BIV hypothesis being realized is quite small, and thereby the BIV worlds are far away.[24] I do not deny that objective respects of similarity may be *epistemically* relevant to some extent.[25] But it is worth mentioning again that these objective ordering sources may not be the ordering source of primary interest in the problem of Cartesian skepticism, if they are not epistemic to a legitimate extent. Indeed, David Lewis is aware of the importance of this point:

---

[24] Sosa (2007) holds that the BIV skepticism is to be differentiated from dreaming skepticism. On his view, the BIV worlds are remote, whereas the worlds in which *S* is dreaming HAND are close; dreaming of HAND often happens, and possibilities of dreaming HAND are much easier to be realized than possibilities of being a BIV. I do not see any relevant difference between the two types of skepticism; the BIV and the dreaming worlds seem to be epistemically equidistant from the actual world. Sosa's distinction is, again, a product of his prior commitment to an objectively probabilistic ordering source.

[25] It cannot be the case that the way the actual world objectively is has no influence over the ordering source.

> Actuality is a possibility uneliminated by the subject's evidence. Any other possibility *W* that is likewise uneliminated by the subject's evidence thereby resembles actuality in one salient respect: namely, in respect of the subject's evidence. (Lewis, 1996, p. 556)[26]

Lewis, thus, reluctantly concedes that his Rule of Resemblance, which is a rule to determine the relevant ordering source, has an *ad-hoc* exception: the similarity in *S*'s evidence does not count. *S*, in the skeptical worlds, has the same, or at the very least, considerably similar sensory experience as of HAND's being the case to the experience *S* does in the actual world.[27] If the ordering source only takes account of the similarity based

[26] Lewis is usually regarded to be a proponent of standard contextualism. However, his position allows for multiple interpretations. Along the lines suggested here, it is possible to construe his position as a version of non-standard contextualism, on which the epistemic position rather than the epistemic standard varies with context.

[27] Williamson (2000a, 2000b, 2001) dismisses this: on his view, *S* has different evidence between the veridical cases and the non-veridical cases: two bits of evidence being phenomenally indistinguishable to *S* entails neither that they are the same evidence, nor that their evidential value are the same. Williamson uses this point to argue that the BIV worlds are not epistemically accessible to the actual worlds: in our terms, the BIV worlds are remote on the relevant epistemic ordering source. McDowell (1982, 1986, 1995, 2008), who has a similar idea to Williamson's on evidence, may agree with him. I do not fully address their view here, but only mention two problems for it: first, this view calls for a different treatment of dreaming skepticism, since dreaming skepticism can be generated by the hypothesis that HAND is true but *S* is dreaming as if HAND is true; second, the view may yield the unintuitive verdict that Henry knows in the Barn Façade case. McDowell (1982, fn. 37, p. 390), in passing, mentions the second problem, and denies that Henry is

on phenomenology of sensory experience, the BIV worlds must be very close to the actual world.[28] [29]

Note that the BIV hypothesis stipulates that *S*'s sensory experience is very similar to BIV's, and it is not possible to deny this respect of similarity. If the relevant ordering source is set in terms of this respect of similarity, the token reliabilist theories result in the

---

experiencing the presence of a barn. I do not see any reason whatsoever to deny this. For problematic ramifications of Williamson's view for the Gettier problem, see (Comesaña & Kantin, 2010).

[28] A similar case is made by Pritchard (2005a), albeit for a different purpose. He makes a distinction between two oft-confused types of 'epistemic luck':

> (Veritic Epistemic Luck) It is a matter of luck that the agent's belief is true. (p. 146)
>
> (Reflective Epistemic Luck) Given only what the agent is able to know by reflection alone, it is a matter of luck that her belief is true. (p. 175)

According to Pritchard, these two types of luck differ as to what ordering source is relevant for them: the ordering source pertinent to veritic luck is "in terms of the way in which she *in fact* formed her belief,' and the ordering source pertinent to reflective luck is "in terms of the way in which the agent *believes* (or would believe) she formed her belief in the actual world" (p. 175). Pritchard argues that *S*'s belief that *S* is not a BIV is reflectively lucky but not veritically lucky, since the BIV worlds are nearby only on the ordering source for reflective luck. This response to the problem of Cartesian skepticism is peculiarly asymmetric to his response to the lottery problem discussed in §2. 3 of Chapter 5. He denies that the relevant ordering source for the original lottery case is objectively probabilistic.

[29] Williams (1991, 2001b) emphasizes this point: the skeptical worlds are as close to the actual world as non-skeptical worlds are. His ground for this is that sensory evidence makes the skeptical possibilities as likely as non-skeptical possibilities. I do not think that this is quite right, as I discuss in the next section. But sill, the similarity in phenomenal property of sensory experience will do for putting both worlds at equidistance.

verdict that *S* knows neither HAND nor ~BIV. On the other hand, when the invariantist token reliabilist theories are combined with the Moorean response, it is presupposed that the relevant ordering source is set in terms of a different respect of similarity. This presupposition may be no problem insofar as they are designed to work against high-standard skepticism. But, against full-blooded skepticism, it begs the question unless it is established that the phenomenal similarity of sensory evidence does not count.

Of course, the difficulty in deciding which ordering source is relevant is the reference class problem for token reliabilism all over again. Insofar as the argument from ignorance and the Moorean proof are about full-blooded skepticism, they appeal to different reference classes. If our skeptical intuitions are full-blooded, in the sense that they imply that *S* has no strong epistemic position regarding HAND or ~BIV, a way to save both our ordinary and skeptical intuitions, then, is to relativize the truth or falsity of our intuitive knowledge ascriptions or denials to a reference class; our ordinary knowledge attributions are true when an ordinary reference class is used, while they are false when the skeptical reference class set by the phenomenal similarity of sensory evidence is used. This response to the full-blooded version of Cartesian skepticism is only possible when SRC is combined with contextualism, and it preserves SCK, at least within a context. Non-standard contextualism I put forth in Chapter 7 will be developed along this line.

§2. 3. The Underdetermination Argument for Cartesian Skepticism

For internalism, Cartesian skepticism appears as a problem of internalist justification. And, as is the case with externalism, the two versions of skepticism are to be distinguished. Again, they give different accounts of why premise (1) is true: high-standard skepticism predicates that it is because the degree of *S*'s justification for ~BIV is not close enough to the maximal value; whereas, the full-blooded skepticism predicates that it is because *S* is not justified at all in believing ~BIV.

Cohen (1988, 1999, 2000) proposes a contextualist SES response to high-standard skepticism about internalist justification: different epistemic standards, or more precisely, different standards for how high the degree of *S*'s justification must be for *S* to be justified *simpliciter* are invoked in different contexts. On the one hand, relative to the extravagantly stringent standard of skeptical contexts, *S* is not justified in believing ~BIV, even given every evidence *S* has against the BIV hypothesis; in probabilistic terms, the stringent standard decreases ε to some infinitesimal value, so that $Ep(\sim\text{BIV} \mid e) < 1 - \varepsilon$, where '*Ep*' is a function for internalist justification, and *e* is *S*'s total evidence. In addition, assuming that the justification version of SCW, it follows that, relative to the same standard, *S* is not justified in believing HAND, because $Ep(\text{HAND} \mid e) < 1 - \varepsilon$. Thus, (1) and (3) are both true in skeptical contexts, and the truth of them is elucidated in terms of lack of the high degree of justification. On the other hand, relative to the lax standard of everyday contexts, *S* is justified in believing HAND; when ε is increased to some relatively large value, $Ep(\text{HAND} \mid e) \geq 1 - \varepsilon$. Again, from the justification version of SCW, it follows that *S* is also justified in believing ~BIV, because $Ep(\sim\text{BIV} \mid e) \geq 1 - \varepsilon$. Thus, (1') and (3') are both true in everyday contexts.

Pritchard (2005a, 2005b) and Williams (1991, 2001a, 2001b, 2004) argue that full-blooded Cartesian skepticism for internalist evidentialism is generated by an underdetermination argument.[30] First, they point out that premise (1) of the argument of ignorance is supported by the fact that sensory experience alone underdetermines whether *S* is a BIV or not. Pritchard, following Brueckner (1994b), formulates the principle that underlies (1) as follows:

> (Underdetermination Principle, UP) For all *S*, *p*, *q*, if *S*'s evidence for believing *p* does not favor *p* over some hypothesis *q* which *S* knows to be incompatible with *p*, then *S*'s evidence does not justify *S* in believing *p*. (Pritchard, 2005a, p. 108)[31]

Pritchard goes on to claim that skeptical conclusion (3) is derivable from UP without relying on SCK or SCW, since UP, with other trivial conditions, entails that *S* has no justification at all for HAND. The UP-based skeptical argument goes as follows, where

---

[30] Both Pritchard and Williams refer to Brueckner (1994b), who offers a formalization of the underdetermination argument for Cartesian skepticism. Brueckner, however, does not hold that the underdetermination argument is for full-blooded skepticism, because he notes that UP given below may require infallibilism. It seems that Pritchard and Williamson ignore this aspect of Brueckner's UP. For more on this, see footnote 34.

[31] Brueckner's original formulation includes the phrase 'incompatible with' in place of '*S* knows to be incompatible with.' Pritchard's modification is quite appropriate, because the relevant notion of justification here is internalist.

'EXP' refers to visual experience as of HAND being the case or visual experience that it *looks* that HAND is the case[32]:

(U1). *S*'s evidence EXP for believing HAND does not favor HAND over BIV.

(U2). If *S*'s evidence EXP does not favor HAND over BIV, then *S*'s evidence EXP does not justify *S* in believing HAND.

(U3). Therefore *S*'s evidence EXP does not justify *S* in believing HAND (and therefore *S* is not justified in believing HAND, and therefore *S* does not know HAND).[33]

Unfortunately, neither Brueckner nor Pritchard explicates the notion of 'favoring *p* over *q*,' even though it is of crucial importance for UP and the UP-based argument for skepticism. Moreover, it is not clear at all that the UP-based argument amounts to an argument for full-blooded skepticism.[34] Pritchard notes of the notion of favoring that "the

---

[32] Vogel (2004a) also formalizes a knowledge version of the underdetermination argument, and it involves an underdetermination principle about knowledge. His underdetermination principle is as vague as Brueckner-Pritchard's, and my objection to the latter below may well be, *mutatis mutandis*, applied to the former.

[33] This argument is taken from Pritchard (2005a, p. 108), only with a modification on propositional letter.

[34] Indeed, Brueckner denies this. UP requires *S*'s evidence to favor *p* over any hypothesis incompatible *q*, in order for *S*'s evidence to justify *p*. Brueckner proposes to weaken UP so as to require only that *S*'s evidence favor *p* over any *relevant* hypothesis incompatible with *q*. Brueckner refers to the relevant alternatives theory, and it is usually reckoned a way to save our ordinary knowledge from high-standard skepticism. However, the relationship between the relevant alternatives theory and infallibilism is more complex than it is usually thought to be (see footnote 43 of Chapter 7).

sceptic maintains that our evidence is of its nature *inconclusive* as regards sceptical hypotheses, since there is nothing to tell between the circumstances that I take myself to be in and being, say, a BIV" (ibid., p. 108; emphasis is added). If (U1) is equivalent to the claim that *S* only has inconclusive evidence against BIV, the version of Cartesian skepticism resulting from the UP-based argument is high-standard skepticism.

Pritchard also implies that EXP does not discriminate between HAND and BIV. Williams (2001b, p. 75) makes the same point, and then construes this indiscriminability in terms of probability: EXP underdetermines that HAND is true, because, conditional on EXP, BIV is just as likely to be true as HAND. Although it is plausible that evidence *e* does not favor *p* over *q* if $Ep(p \mid e) = Ep(q \mid e)$, (i) $Ep(\text{HAND} \mid \text{EXP}) = Ep(\text{BIV} \mid \text{EXP})$ cannot be true. HAND entails ~BIV, and *S* is supposed to know this entailment. By the probability calculus, it follows (ii) that $Ep(\text{HAND} \mid \text{EXP}) \geq Ep(\text{BIV} \mid \text{EXP})$. Also, BIV entails ~HAND, and *S* is supposed to know this entailment. So, (iii) $Ep(\text{BIV} \mid \text{EXP}) \leq Ep({\sim}\text{HAND} \mid \text{EXP})$. Given (ii) and (iii), (i) entails that $Ep(\text{HAND} \mid \text{EXP}) \leq Ep({\sim}\text{HAND} \mid \text{EXP})$. This is absurd – why does experience as of a hand make the absence of a hand equally or even more probable than the presence of a hand? Then, a *reductio* is easily constructed against (i).[35]

[35] At the very least, it must be true that EXP boosts the epistemic probability of HAND and reduces that of ~HAND. This is not enough for the denial of $Ep(\text{HAND} \mid \text{EXP}) \leq Ep({\sim}\text{HAND} \mid \text{EXP})$, if the prior probability of ~HAND is considerably low. It seems difficult to decide the prior probability of HAND or ~HAND, in view of Cartesian skepticism. It would be the degree of rational credence on *a priori* grounds alone, and it is not clear how *a priori* grounds make a significant difference in rational credence with which *S* believes empirical propositions.

Although Pritchard and Williams fail to fully articulate UP, it has some air of plausibility. Perhaps, the best interpretation of the notion of favoring *p* over *q* would be the following:

(Favoring) Evidence *e* favors *p* over an (known to be) incompatible hypothesis *q* iff $Ep(p \mid e) > Ep(p)$ but $Ep(q \mid e) \not> Ep(q)$.[36]

As a matter of fact, however, a principle weaker than **Favoring** is orthodoxy in the epistemic probability theory. Carnap (1962) makes a distinction between confirmation as 'firmness' and confirmation as 'increase in firmness': they being rendered the notion concerning justification rather than confirmation, the former is the essentially the same as **Threshold for Internalist Justification**, and the latter as the first conjunct of the right hand side of **Favoring**. That is, orthodoxy requires the following as a necessary condition for *e* to give justification for *p*:

---

[36] Kvanvig, in his (2008) review of Pritchard (2005a), points out that two different readings of the notion of favoring *p* over *q* is possible: (a) *e* favors *p* over *q* iff *e* "is a reason for believing *p*, and isn't a reason, one way or the other, about *q*"; and (b) *e* favors *p* over *q* only if *e* is a reason for *p*, and is "is also a reason against *q*, i.e., a reason in favor of ~*q*" (p. 281). Kvanvig argues that although Pritchard's intended reading is (b), (a) is more plausible. For, Kvanvig submits, from (a) and UP, it follows that *S* is justified in believing everyday propositions. That is, (a) is not susceptible to skepticism. My elucidation of the notion of favoring *p* over *q* is modeled on (b). I doubt that any rendering of this notion is in accordance with (a), insofar as it is construed probabilistically. For what it says is that even though HAND entails ~BIV and EXP raises the probability of HAND, it has no bearing on BIV or ~BIV. As I discuss shortly, given plausible assumptions, EXP raises the probability of BIV.

(Favoring') Evidence *e* favors *p* only iff $Ep(p \mid e) > Ep(p)$.[37]

The difference between **Favoring** and **Favoring'** is suggestive of what presuppositions are involved in premise (1) of the argument from ignorance, in several ways: (i) once the notion of favoring is defined in terms of epistemic probability, it is the case that EXP satisfies **Favoring'**, but not **Favoring** with regard to HAND. That is, *Ep*(HAND | EXP) > *Ep*(HAND), but *Ep*(~BIV | EXP) > *Ep*(~BIV); (ii) this point explains why EXP fails to transmit justification to ~BIV; (iii), given that EXP satisfies **Favoring**, the explanation of *S*'s unjustifiedness for HAND in the skeptical case needs to be explained by some other necessary condition for justification. I will consider these points one by one in what follows.

(i) is supported by many, such as Cohen (2005, pp. 424-5), Hawthorne (2004a, pp. 73-77), Neta (forthcoming), Silins (2008, pp. 123-28), White (2006), and Williamson (2005b). They all make more or less the same arguments, the gist of which is as follows: the prior probability of EXP is not maximal (*Ep*(EXP) < 1), and if *S* has normal sight and is rational, *S* should give more credence to EXP given HAND than otherwise. Then, the probability of EXP is raised conditional on HAND:

Ep(EXP | HAND) > *Ep*(EXP)

[37] This condition is sometimes treated as a condition for something to count as evidence for *p*. This distinction makes no relevant difference, insofar as justification for *p* requires evidence for *p*.

It follows by Bayes' theorem that

Ep(HAND | EXP) > Ep(HAND)[38]

Hence, EXP favors HAND, i.e., EXP satisfies **Favoring'** with regard to HAND. This does not entail that EXP satisfies **Favoring'** with regard to ~BIV, even though HAND entails ~BIV. Note that BIV entails EXP, since, as the BIV hypothesis goes, *S* is a BIV having sensory experience as of a hand, just as *S* is in the actual world. So, by the probability calculus,

Ep(EXP | BIV) = 1

Then, with the assumptions that Ep(BIV) > 0, and that 1 > Ep(EXP) > 0, it follows by Bayes' theorem that

Ep(BIV | EXP) > Ep(BIV)[39]

[38] This derivation is taken from White (2008). Although he does not state it, the proof for the derivation runs as follows:

(1). Ep(EXP | HAND) > Ep(EXP) [assumption]

(2). Ep(EXP | HAND) = (Ep(EXP) / Ep(HAND)) × Ep(HAND | EXP) [Bayes' theorem]

(3). (Ep(EXP | HAND) / Ep(EXP)) × Ep(HAND) = Ep(HAND | EXP) [from (2)]

(4). (Ep(EXP | HAND) / Ep(EXP)) > 1 [from (1)]

(5). Therefore, Ep(HAND | EXP) > Ep(HAND) [from (3) and (4)]

Again, by the probability calculus,

$$Ep(\sim\text{BIV} \mid \text{EXP}) < Ep(\sim\text{BIV})$$

---

[39] These two assumptions are plausible enough; BIV and EXP have non-zero prior probability, and the prior probability of EXP cannot be maximal, given that there is some, albeit small, probability that *S* can rationally expect not to have hands. From these assumptions, the proof goes as follows:

(1). $Ep(\text{EXP} \mid \text{BIV}) = 1$ [from the BIV hypothesis]

(2). $Ep(\text{BIV} \mid \text{EXP}) = (Ep(\text{BIV}) / Ep(\text{EXP})) \times Ep(\text{EXP} \mid \text{BIV})$ [Bayes' theorem]

(3). $Ep(\text{BIV} \mid \text{EXP}) = (Ep(\text{BIV}) / Ep(\text{EXP}))$ [from (1) and (2)]

(4). $Ep(\text{EXP}) < 1$ [assumption]

(5). $Ep(\text{BIV} \mid \text{EXP}) > Ep(\text{BIV})$ [from (3) and (4)]

Neta (forthcoming) holds that this argument *per se* is inadequate, and adds an anti-defeater condition. This move is not necessary, or even worse, misses the point: first, he seems to forget the fact that many avoid wordiness by using '$Ep(p \mid e)$' as an abbreviation of '$Ep(p \mid e \wedge$ total evidence$)$.' If *S*'s total evidence includes a defeater for EXP or ~BIV, obviously, the probability assignment in the proof is different; second, more importantly, Neta wrongly holds that the total evidence is relevant for the probabilistic consideration at stake. In fact, since full-blooded Cartesian skepticism is designed to show that sensory experience *alone* does not give a reason in favor of ~BIV, the total evidence should not be taken account of. For this reason, it cannot be that those who offer or support the probabilistic consideration are using $Ep(p \mid e)$ as the abbreviation mentioned above. However, I agree with Neta that general principles about favoring, such as **Favoring** and **Favoring'**, need to incorporate some anti-defeater clause.

Hence, not merely does EXP not favor ~BIV, but also favors BIV.[40]

The consequence is that EXP fails to satisfy the condition UP imposes for justification for *p*, if it is construed in terms of **Favoring**. As I argued above, it does not immediately follow from UP that UP gives rise to the full-blooded version of Cartesian skepticism rather than the high-standard version. Nevertheless, UP tells something important: EXP fails to transmit internalist justification in the minimal sense to ~BIV. For, no matter what epistemic contribution EXP makes to HAND, EXP contributes nothing good to ~BIV, provided that EXP lowers the epistemic probability of ~BIV. More generally, the probabilistic consideration and **Favoring** suggest the following: for any *e*, *p*, *q*, if *e* does not favor *p* over *q*, *e* does not transmit warrant in the minimal justificatory sense to ~*q* from *p*, even if *p* entails ~*q*. In other words, cases of UP failure are those of warrant transmission failure; if UP does not hold, STW (the single-premise warrant transmission principle) does not hold for the increase of internalist justification. As a result, EXP alone does not transmit internalist justification to ~BIV from HAND: if *S* believes HAND on the basis of EXP, and, in turn, believes ~BIV on the basis of HAND, then, for *S*, EXP does not give justification for ~BIV, regardless of whether EXP justifies HAND.

[40] Those who offer this proof tend to criticize dogmatism, advocated by Pryor (2000, 2004). Dogmatism is, roughly, the position that sensory evidence alone is sufficient to justify beliefs about the external world, and they, in turn, are sufficient to justify beliefs in the denial of skeptical hypotheses. Pryor, however, severally restricts the range of application of his dogmatism. It is not clear whether he accepts its application to the problem of Cartesian skepticism, since his main concern is to defend the Moorean proof of the existence of external world, not the Moorean proof of the denial of the BIV hypothesis or the like.

§2. 4. Shift in Reference Class in the Case of Cartesian Skepticism for Internalism

EXP's failure to transmit internalist justification to ~BIV from HAND does not exclude the possibility that *S* is justified in believing HAND, if UP is construed in terms of **Favoring'**. I argue, in what follows, that the denial of this possibility, supported by our skeptical intuitions, is to be explained by a shift in reference class.

Crispin Wright, in a series of writings (2000, 2002, 2003, 2004, 2007), emphasizes that the failure of the internalist justification version of STW does not entail the failure of the internalist justification version of SCW.[41] In particular, he argues that EXP's failure to transmit internalist justification to ~BIV from HAND is consistent with the claim that

[41] The view I attribute to Wright here may contain oversimplifications; his view is filled with idiosyncrasies, some of which may be of interest. First and foremost, he calls warrant or justification for ~BIV 'entitlement,' and distinguishes it from the standard notion of internalist justification. Entitlement is the type of warrant *S* needs not to earn, i.e., it does not require evidence. On Wright's view affected by Wittgenstein (1969), entitlement concerning *p* is present in a certain information context where *S* takes for granted or trusts *p*. Thus, entitlement is borne by a distinct mental attitude of taking for granted or trusting rather than believing; Wright distinguishes entitlement to trust *p* from justification for believing *p*. It is not clear how this is congenial to his preservation of the justification version of SCW. Furthermore, trusting seems to be quite similar to believing. Tucker (2009, p. 449) points out that Wright reckons trusting *p* to be incompatible with disbelieving or withholding judgment about *p*. Fantl & McGrath (2009b, p. 110) define justification for *p* in such a way that if *S* should neither withhold nor disbelieve *p*, *S* is justified in believing *p* (see also (pp. 117-24) for their discussions on Wittgenstein). In addition, it is doubtful that entitlement in Wright's sense is epistemic. Jenkins (2007) argues that it is not.

internalist justification for ~BIV is a pre-condition for internalist justification for HAND.[42] If this is the case, the internalist justification version of SCW is preserved. For, then, whenever *S* is justified in believing HAND, *S* is also justified in believing ~BIV. This line of thought may be supported on two grounds.

First, White (2006, p. 533) defends Wright's idea in the following way: it is a theorem of probability theory that necessarily, for any e, p, q, if p entails q, then $Ep(p \mid e) \leq Ep(q \mid e)$. Then, $Ep(\text{HAND} \mid \text{EXP}) \leq Ep(\sim\text{BIV} \mid \text{EXP})$. As we have seen, $Ep(\sim\text{BIV} \mid \text{EXP}) < Ep(\sim\text{BIV})$. It follows that $Ep(\text{HAND} \mid \text{EXP}) < Ep(\sim\text{BIV})$. So, given **Threshold for Internalist Justification**, *S*'s belief in HAND is justified on the basis of EXP, only if *S* is justified in believing ~BIV.[43]

Second, we have seen that EXP favors HAND in the sense of **Favoring'**: $Ep(\text{HAND} \mid \text{EXP}) > Ep(\text{EXP})$. Isn't this enough for EXP to justify HAND for *S*? The answer is indeterminate, since UP is only a necessary condition for justification. EXP fails to give justification for HAND, if one of the following conditions is not satisfied: first, **Threshold for Internalist Justification**: $Ep(\text{HAND} \mid \text{EXP})$ may be too low to reach the threshold for justification, or put differently, it may be the case that $Ep(\text{HAND} \mid \text{EXP}) < 1 - \varepsilon$; second, **Meta-Justification**: *S* may not be justified in believing that EXP gives

[42] Wright (2002) explains transmission failure in terms of failure of a counterfactual condition (see also (Brown, 2004) for an attempt to refine it). It, however, unclear that internalist justification has any bearing on counterfactuals. Smith (2009) offers a thorough account of transmission failure in the framework of safety and a counterfactual account of type reliabilism. Smith's account may be relevant for elucidating why and when externalist warrant or justification fails to transmit.

[43] White takes this argument to be supportive of Wright (2002).

justification for HAND, or that $Ep(\text{HAND} \mid \text{EXP}) \geq 1 - \varepsilon$; third, NSS: *S* may have no independent means to have justification that the source of the belief in HAND is reliable.

When applied to non-inferential empirical belief, **Meta-Justification** takes the form of the requirement of trans-level inference, and a plausible construal of the requirement of trans-level inference is **Reliability Expert**. Applied to the case at hand, **Reliability Expert** predicates that the degree of internalist justification for EXP conforms to the reliability of EXP assigned by one or another reliability expert function. Let *RELI*(EXP) be the reliability assignment to the effect that EXP is a reliable indicator of HAND. For **Reliability Expert** to be applied to *RELI*(EXP), it needs to be in *S*'s evidence set. Assuming that *e* is evidence for *S* only if *S* is justified in believing *e*, **Reliability Expert** entails that it is a pre-condition for HAND to be justified for *S* that *S* is justified in believing that the source of the belief in HAND is reliable, or more particularly, *RELI*(EXP).[44]

It is crucial here whether *S* is justified in believing *RELI*(EXP), since the lack of justification for *RELI*(EXP) may explain why *S* fails to satisfy **Threshold for Internalist Justification** for HAND, i.e., $Ep(\text{HAND} \mid \text{EXP})$ is lower than the threshold $1 - \varepsilon$ for justification, even if it is set by an ordinary, lax standard. Suppose that you have no idea at all how reliable an indicator your particular visual experience, EXP, is for HAND. Since

---

[44] Wright (2004, pp 168-9) seems to hold that it is the case that $Ep(\text{HAND} \mid \text{EXP}) \leq Ep(RELI(\text{EXP}))$. If this is the case, it is easy to show that justification for *RELI*(EXP) is a pre-condition for EXP to justify HAND. But Wright does not offer any reason for this, and it is not clear what amounts to such a reason, especially because reliability comes in degrees. All I can agree is that Wright is right if *RELI*(EXP) says that EXP is a perfectly reliable indicator of HAND, but this is far from important for full-blooded skepticism.

you have no idea about the reliability of EXP, rationality prescribes that you should not believe HAND with much credence. The content of *RELI*(EXP), in full detail, is that EXP is a reliable indicator of HAND to the degree *r*. Even if EXP alone cannot boost *Ep*(HAND) much, *r* may be equal or higher than 1 - ε. If so, **Reliability Expert** is the only way for EXP to justify HAND. For, then, it follows from **Reliability Expert** that $Ep(\text{HAND} \mid \text{EXP} \wedge RELI(\text{EXP}) = r) = r \geq 1 - \varepsilon$.

Is *S* justified in believing *RELI*(EXP)? To answer this question, NSS is relevant. *RELI*(EXP) is an application of the reliability expert function *RELI* to a particular sensory evidence, EXP. However *RELI* is interpreted, definite reliability assignments are relativized to a source, process, or method type *K*.[45] So, *RELI*(EXP) in *S*'s evidence set are indeed divided into two pieces of evidence: that process *K* is reliable (at *r*), and that EXP is produced by a token of *K*. NSS is satisfied only if *S* has a means to justify that *K* is reliable independently of *K* itself. Insofar as NSS is a condition for **Meta-Justification**, NSS alone

[45] Also, if necessary, it is relativized to an environment type. I ignore this complication here, but see footnote 48. I also ignore the token-reliabilist interpretation of *RELI*. On the token reliabilist interpretation, *RELI*(EXP) means that beliefs based on EXP is sensitive or safe, and the relevant method token is already individuated as EXP. But the problem is how *S* is justified in believing that beliefs based on EXP are sensitive or safe. As discussed in §3. 6 of Chapter 2, it is difficult to explain the epistemic role of single-case propensity, and the single-case propensity theorist tends to accept that known or justified single-case propensity assignments are based on frequency data. Applied to the theory of knowledge, then, sensitivity or safety assignment are known or justified on the basis of type reliability data. Hence, the conclusion I draw below would equally apply even if *RELI* is interpreted in the token reliabilist fashion.

is capable of generating the problem of Cartesian skepticism for the version of internalist evidentialism that endorses **Meta-Justification**, in the following way[46]:

(N1). If *S* is justified in believing HAND, then *S* is justified in believing *RELI*(EXP). [from **Meta-Justification**]

(N2). *S* is justified in believing *RELI*(EXP), then *S* is justified in believing that *K* is a reliable source. [from relativization of definite reliability to indefinite reliability]

(N3). If *S* is justified in believing that *K* is a reliable source, then there is a means for *S* to justify this without relying on *K* or *K*-dependent sources alone. [from NSS]

(N4). There is no means to justify that *K* is a reliable source other than *K*.

(N5). Therefore, *S* is not justified in believing HAND.[47]

---

[46] This is so, even if EXP alone justifies HAND. Moreover, if externalism is combined with **Meta-Justification**, the same argument is generated for it.

[47] This argument captures most, albeit not all, important features of the argument for Cartesian skepticism constructed by Williams (1991, ch. 2; 2001b, pp. 75-7) (cf. also (Wright, 2004, pp. 169-70)). However, the notion of underdetermination is not essential for this argument. Williams holds that EXP underdetermines HAND because HAND and BIV are equally probable conditional on EXP. As I argued, this is false. Both Williams and Wright mention the similarity between the argument for Cartesian skepticism and the argument for Humean skepticism about induction; the reliability of sensory experience in the former is analogous to the uniformity of nature in the latter, with regard to its role in argument.

The crucial point is whether premise (N4) is true. In §1. 2 of Chapter 6, I have shown against the problem of easy knowledge that in ordinary perceptual cases, (N4) is false. Does the same consideration apply to the problem of Cartesian skepticism? Let K_1, …, K_n be process types and e_1, …, e_n be outputs of their tokens, such that they indicate HAND. For (N4) to be false, it is crucial that K_1, …, K_n are independent of each other. The skeptic denies this, since all K_1, …, K_n are empirical processes, and thereby they are subject to the skeptical doubt all at once.

To see this, consider process type K_v {visual experience} and Kt {tactual experience} such that the outputs of tokens of K_v and Kt, visual experience as of a hand, EXP, and tactual experience as of a hand, EXP_t, respectively, happen to indicate the same event HAND. In order for S to be justified in believing that each of K_v and K_t is reliable, Ep(each of K_v and K_t is reliable | EXP ∧ EXP_t) needs to be at least as high as 1 - ε. This is possible only if the prior degree of justification, Ep(EXP ∧ EXP_t), is sufficiently low. S's belief in EXP ∧ EXP_t may be supported by evidence concerning the independence between K_v's function and K_t's function; such evidence is typically statistical or probabilistic evidence on how frequent K_v and K_t coordinate relative to some environment type.

However, this frequency is high relative to the environment type depicted by the BIV hypothesis. For it is the environment in which all empirical processes available to S is manipulated in the coherent, wholesale way; or put differently, in which all empirical processes are not independent of each other.[48] In the skeptical environment, then, that K_v

[48] This relativization is possible only if S's actual environment is a token of this environment type. As I argue below, skeptical hypotheses enforce an overly broad reference class, and so the relevant environment type is

and K_t happen to coordinate does not make it probable that K_v and K_t are reliable. This means that, in view of skeptical hypotheses like the BIV hypothesis, there is little reason to treat K_v and K_t separately. In the type of environment made salient by skeptical hypotheses, all empirical processes are on a par with regard to their function, in the sense that their success and failure are uniformly controlled. This feature of skeptical hypotheses forces us to classify K_v, K_t and any other empirical process as being of the one and the same type. Presumably, such an excessively general type is {sensory experience}. Once process type *K* is thus set, premise (N4) can be nothing other than truth; any empirical means to justify that *K* is reliable falls under *K*, and hence it is bootstrapping. Thus, the full-blooded skeptical conclusion (N5) looms: *S* is not justified at all in believing HAND (and other empirical propositions). (N5) holds regardless under what process type *S* classifies EXP or relative to what environment type *S* estimates the frequency of the coordination of K_1, …, K_n. In the face of skeptical hypotheses, other types than {sensory experience} become irrelevant.[49]

In his response to full-blooded skepticism resulting from this argument, Wright denies (N4) and concedes that it is *a priori* justified for *S* that sensory experience is reliable.[50] In a similar vein, Cohen (1999) holds that ~BIV is *a priori* justified. I will

---

{occasions of sensory experience}, which include as its token any occasion on which *S* has sensory experience.

[49] At least, this is the only way to accommodate the data that in the face of full-blooded skepticism, our third-person denials of knowledge or justification seem true, even relative to the ordinary, lax epistemic standard.

[50] More precisely, Wright argues that *sensory experience is reliable* is justified, or in his terminology, entitled for *S* relative to an information context. He is relying on the Wittgensteinian idea that the reliability of

pursue a different line of response here, which is already predicted in my response to the problem of easy knowledge.

If I am right, insofar as there is the full-blooded version of Cartesian skepticism, it amounts to a generalized problem of easy knowledge. For this reason, the type of response to the problem of easy knowledge I have proposed is equally possible for the problem of Cartesian skepticism: it presupposes a certain reference class, but there is no reason that it is the only reference class instantiated by an individual experience when one forms empirical belief. The point I am making is conditional: insofar as we have the intuition about the skeptical cases that *S* does not know HAND, and our intuition is to be explained by *S*'s being not justified *at all* in believing HAND, the only way of explanation is to pick up an excessively wide reference class. On the other hand, we can explain our everyday intuition that *S* knows HAND, by picking up narrower reference classes, such as K_v and K_t. This shift in reference class offers a nice explanation of why we are inclined to agree on premise (1) of the argument from ignorance, and why we are inclined to agree on premise (1') of the Moorean proof, even if the argument is for full-blooded Cartesian skepticism.

---

sensory experience is a hinge, i.e., it is taken for granted or trusted in the context of empirical inquiry, without any evidence for it. It is not clear whether Wright endorses that the context in which the reliability of sensory experience is trusted is broad so as to include any individual occasion of empirical inquiry, or that it only includes some but not all occasions. If the former is the case, Wright's view is tantamount to the reliability of sensory experience being *a priori* justified; if the latter is the case, his view is essentially the same as 'contextualism' as Williams (1991) proposes it. Williams' contextualism entails that the distinction between the *a priori* and the *a posteriori* is only drawn within a context of inquiry, not across contexts.

In my response to the problem of easy knowledge, I argued that coordination of independent processes can be used as a checker of the reliability of each process involved, and that we type process tokens narrower than {sensory experience} in everyday contexts. Processes or sources specified narrowly are usually reckoned to belong to different functional orders. This is why the frequency of (EXP $\wedge$ EXP_t) is usually reckoned to be low. Having visual experience EXP alone cannot contribute to establishing the reliability of K_v. But we can also touch the object, and have EXP_t, by exercising a token of K_t. The coordination of process types K_v and K_t of different functional orders raises the probability of the reliability of each process type. And more importantly, we have registered such coordination on many occasions. With information on coordination accumulated, we are justified in believing K_v is reliable, and *RELI*(EXP), viz., that a particular output, EXP, of a token of K_v is reliable. Even within one sensory modality, narrow specifications of the relevant reference classes are possible: you have EXP at a time t_1, and have it again at a later time t_2. If you remember that nothing has happened to your hand between these times, it gives some justification for the reliability of process types {K_v at t_1} and {K_v at t_2}. Unlike Wright's solution, this solution is consistent with our natural notion that we can *a posteriori* have knowledge of or justification for the reliability of empirical processes. [51] [52]

---

[51] The skeptic may try to refute this response on two grounds. The first ground: the point I made amounts to some justification being given for *RELI*(EXP) when two or more processes provide coordinated information about one and the same object. This presupposes that there is such an object, and begs the question against skepticism. Reply: the problem the skeptic is pushing is the Kantian problem in Conant's sense mentioned in footnote 20. I am not addressing that question. The second ground: it is true that all empirical processes belong to a single reference class {experience}. Then, unless this choice of the relevant reference class is

An argument for standard contextualism goes as follows: it offers the best explanation of varying intuitions about certain kinds of cases – lottery cases and skeptical cases. In the last chapter, I have argued that standard contextualism has limited resources to deal with some cases of the kinds they purport to do. For what underlies the cases is shifting reference class, and standard contextualism, being SES, cannot accommodate it. Furthermore, varying reference class is the key to other kinds of cases to which standard contextualism does not apply – (some of) Gettier cases and cases of easy knowledge. Only SRC, if properly constructed, is capable of offer a unifying account of our shifty intuitions about these types of cases. In the next chapter, I will propose and defend a version of non-standard contextualism as such an account, with several arguments.

---

legitimately excluded, appeal to multiple reference classes begs the question. Reply: this is a version of high-standard skepticism in disguise. The skeptic does not show that it is false that multiple empirical processes function independently, and thereby belong to different reference classes. The choices of the relevant reference classes on her side and on my side are equally possible. My point is that if there is any argument for full-blooded Cartesian skepticism, the only way to generate it is by employing the reference class that covers every empirical means of inquiry into the world. To require excluding the reference class chosen by the skeptic is simply to raise the epistemic standard for those who respond to the skeptical problem. The equal burden of proof is on the skeptic's shoulder. How can she exclude the multiple reference classes we use in our ordinary practice? SRC, if combined with contextualism, can save both skeptical and anti-skeptical intuitions.

[52] The inference to the best explanation response to Cartesian skepticism is that the systematicity or coherence of experience makes the non-skeptical hypothesis *a priori* more probable than the skeptical hypotheses. This response is similar to mine in a way, but significantly different: what I call coordination of different processes is the systematicity or coherence of information gained from independent sources; the inference to the best explanation response shares with skepticism the relevance of the excessively wide reference class.

Chapter 7: Task-Sensitive Contextualism and the Reference Class Problem

Introduction

In this chapter, I propose a version of non-standard contextualism as a solution to the reference class problem in the theory of knowledge. While it is similar to the version of contextualism put forth by Heller, its details and arguments are different. The two arguments Heller offers for his contextualism are simple: in his (1995b), he suggests that it is a linguistic fact that ‘reliable’ is always associated with a reference class, and what class it is is contextually determined; and in his (1999), he points out that on Lewisian semantics for counterfactuals, the modal base and the ordering source for a counterfactual are contextually determined. These arguments are intended to offer solutions to the versions of reference class problems for type reliabilism and for sensitivity theory, respectively. Heller does not say much about how his contextualism is adequate as a solution to the reference class problem generally. My arguments for non-standard contextualism are more general in two senses: (i) they are meant to demonstrate the adequacy of non-standard contextualism as a general solution, and (ii) they do not hinge on any idiosyncratic feature of a particular substantive theory of knowledge.

The range of applications of non-standard contextualism is restricted to those variants of the reference class problem that face substantive theories of knowledge. It is not part of my intention to solve the reference class problems for objective probability theories. And also, the full-blown solution to the problem of non-projectible reference class is beyond the scope of the discussions here. Furthermore, even though the reference class problem for epistemic probability theories is epistemological, I have no general recipe for specifying the relevant reference class for any given direct inference or application of any given expert function.[1] Be that as it may, non-standard contextualism has strong force as a solution to the reference class problem for both externalist and internalist theories of knowledge, or so shall I argue.

The specific version of non-standard contextualism I advocate is called 'task-sensitive contextualism,' on which the choice of the relevant reference class is constrained by, among other things, what task the attributor assigns to the subject. In §1, I will outline the semantic framework for non-standard contextualism. The semantic framework for standard contextualism will not do, since it does not posit a reference class as a contextual parameter. Then, I will offer five different arguments for task-sensitive contextualism. Each sub-section of §2 offers one argument, with an explanation of how task-sensitive contextualism contributes to the solution to the reference class problem. Arguments 1 and 2 are psychological: argument 1 purports to show the relevance of contextual factors for the choice of reference class, by appeal to the psychological study of category-based inference; argument 2 appeals to the way in which computational theorists of cognition and

[1] The rationality of the choice of a reference class is a beast really hard to capture – I discuss this topic in §2. 1, in connection with psychological studies of category-based inference. Part of this difficulty seems to consist in the fact that direct inferences are used for many different purposes in different contexts. For example, a specific form of the reference class problem for the epistemic probability theory arises in legal contexts: if the probative value of probabilistic evidence is relativized to a reference class, what reference class is to be used in legal adjudication? For the reference class problem in legal decision making, see (Allen & Pardo, 2007) and (Colyvan, Regan & Ferson, 2001). *International Journal of Evidence and Proof*, 11(4), 2007, is devoted to discussions on Allen & Pardo's paper, including their replies.

philosophers of computation circumvent a problem structurally similar to the reference class problem. Argument 3 demonstrates that the trivializing reference class is not a problem once task-sensitive contextualism is adopted. Argument 4 is based on a function of knowledge in our epistemic practice: knowledge attributions are made to certify an answer to the question explicitly or implicitly asked in context. The choice of the relevant reference class is sensitive to the contextual variance in question. Argument 5 adds the final brush on the considerations in Chapters 5 and 6, with which I conclude that task-sensitive contextualism best explains our intuitions about many cases.

§3 is devoted to figuring out what ramification task-sensitive contextualism has for both the substantive theory (§3. 1) and the formal theory of knowledge (§3. 2). A significant ramification for externalist theories is that strength of epistemic position is a function of both subject and attributor factors. A similar ramification obtains for internalist evidentialism. As a result, no substantive theory can satisfy **Supervenience** – the axiom central for the enterprise of the theory of knowledge –, if it is read in a certain way.

§1. The Outline of Non-Standard Contextualist Semantics of 'Know'

Before proposing arguments for the version of non-standard contextualism required for solving the reference class problem in the theory of knowledge, I will discuss what semantic framework it embodies, since doing so best captures differences between standard and non-standard contextualism. Through this discussion, I delineate the outline of the non-standard contextualist semantics for 'know.'

Proponents of standard contextualism, such as Cohen, DeRose, and Lewis, model the semantics of 'know' on that of gradable adjectives.[2] As Stanley (2005, pp. 17-8) points out, however, they do not specify how the main tenet of standard contextualism, that the semantic content of 'know' varies across contexts of use in which different epistemic standards are in play, is semantically implemented; standard contextualism is silent on what logical form 'know' has or how the denotation of 'know,' the knowledge relation, is structured.[3] After all, contextualism in epistemology is not a formal semantic theory.[4] As Stanley correctly submits, standard contextualism aims no more than to make the case that knowledge ascriptions are context-sensitive in a distinctively epistemological way.[5]

---

[2] Stanley (2004, 2005) argues against contextualism that no good analogy holds between 'know' and gradable adjectives. Though his arguments are strong, I do not respond to them here. For objections to Stanley's arguments, see (Blome-Tillmann, 2008), (DeRose, 2009, ch. 5), and (Halliday, 2007).

[3] Hawthorne, (2004a, ch. 2, fn. 84, p. 81) reports that Cohen suggests in conversation that 'know' invariably expresses a three-place relation between a subject, a proposition, and a standard, and the value of the standard index varies with context. This way of semantic implementation is contrasted with that on which 'know' expresses different relations in different contexts. The former is quite similar to the standard analysis of gradable adjectives, as we will see below; the latter treats 'know' as if it is an indexical. The framework I propose for non-standard contextualism is more congenial to the former than the latter.

[4] This is not necessarily vicious; for this reason, the contextualist idea is open to many ways of semantic implementation. For the details of such ways, see Ludlow (2005).

[5] Stanley also claims that contextualism in epistemology is better taken to be a view about the semantics of 'know *p*' rather than 'know.' The reason is that there are respects in which 'know,' just like every other verb, is trivially context-sensitive: for example, it is indexed to a time, and what time it is indexed to is context-sensitive. I put aside this complication here, and treat both standard and non-standard contextualism as being about 'know.'

Even though the proponents of standard contextualism endorse similarities in semantic behavior between 'know' and gradable adjectives, they do not derive contextualism directly from such similarities.[6] As a matter of fact, gradable adjectives differ from 'know' in many, though not all, respects. The point of appeal to similarities is rather to emphasize that the context sensitivity of 'know' is not too outlandish.

At the very least, the contextualist accepts the basic idea of the standard analysis of gradable adjectives, which, according to Kennedy (2007), consists of the following two assumptions:

(a) Gradable adjectives map their arguments onto abstract representations of measurement, or DEGREES.

(b) A set of degrees totally ordered with respect to some DIMENSION (height, cost, etc.) constitutes a SCALE. (p. 4)

On the basis of (a) and (b), a gradable adjective '*F*' is analyzed as a function *f* whose argument is an individual that has a property lexically tied with *F* and whose value is a degree – a point or an interval – on the scale for that property.

The context-sensitivity of gradable adjectives appears when they are put in the 'positive form,' the form that involves no overt degree morphology. To follow Kennedy and many others, I assume that the positive form involves an implicit slot for or index to the pertinent standard of comparison, and with Stanley (2000), that such a slot is filled by a

[6] This type of argument is what DeRose calls an 'indirect positive argument for contextualism.' He denies that he is offering such an indirect argument; see (DeRose, 2009, pp. 66-9; ch. 5, fn. 8, p. 169).

variable. The context-sensitivity of a gradable adjective in the positive form amounts to the value of the standard of comparison variable being determined by context (and semantic properties of the other co-occurring expressions).

In the case of 'know,' the axioms of the formal theory of knowledge yield the two assumptions corresponding to (a) and (b):

(a'). 'Know' maps its arguments onto an abstract representation of measurement.

(b'). A set of degrees totally ordered with respect to epistemic position constitutes an epistemic scale.

(a') and (b') are semantic expressions of **Supervenience** and **Degree**, respectively. Likewise for gradable adjectives, 'know' is analyzed as a function whose argument is a proposition with regard to which *S* instantiates some epistemic factor(s) and whose value is a point or an interval on the epistemic scale. The task of any substantive theory of knowledge, as we have seen in Chapter 3, is to determine what such a function, *EP*, really is. One prominent difference between (b) and (b'), however, is that the scales for some gradable adjectives, such as 'tall' or 'heavy,' have no upper- or lower-bound; a thing can be infinitely taller or heavier and shorter or lighter[7]; whereas, due to **Limitation**, the strength

[7] Following others' works, Kennedy (pp. 32-6) points out that there are four types of structures of scale for a gradable adjective: (i) it has no upper- or lower-bound, (ii) it has the upper-bound but no lower-bound, (iii) it has the lower-bound but no upper-bound, and (iv) it has both the upper- and the lower-bound.

of epistemic position is standardized, i.e., the epistemic scale has the maximal and the minimal value.[8]

Non-standard contextualism – a specific version of which is offered as a solution to the reference class problem in epistemology in the subsequent sections – is consistent with but significantly different from standard contextualism. The main tenet of non-standard contextualism is the following:

> (Non-Standard Contextualism) For all subjects *S* and propositions *p*, the truth of sentences of the form '*S* has strength of epistemic position *EP* regarding *p*' is relative to a reference class *G*, and the value of *G* relevant for *S* and *p* is, at least in part, determined by context of use.

On the one hand, according to standard contextualism, 'know' in the form of '*S* knows *p*' has an implicit slot for or index to an epistemic standard. What context does in fixing the semantic content of a particular tokening of 'know' is to determine the value of the epistemic standard variable. The epistemic standard variable, once standardized, is expressed by '$1 – \varepsilon$.' Standard contextualism, then, entails that the truth value of knowledge ascriptions or denials is context-sensitive, since it follows from **Epistemic Position** that '*S* knows *p*' is true only if $EP(p) \geq 1 - \varepsilon$.

---

[8] I am assuming with the standard contextualists that the implicit variable for the standard of comparison ranges over degrees or thresholds. This assumption is controversial, but even if it is false, it is not fatal to contextualism. For discussions, see (DeRose, 2008), (Fara, 2000), (Hawthorne, 2004a, pp. 53-4), (Kamp, 1975), (Kennedy, 2007), (Klein, 1980), (Ludlow, 1989), and (Stanley, 2000, 2002, 2004, 2005).

On the other hand, according to non-standard contextualism, ‘know’ in the form of ‘*S* knows *p*’ has an implicit slot for or index to a reference class. For, however *EP* is interpreted by one or another substantive theory of knowledge, the value of *EP*(*p*) is indeterminate unless it is relativized to a reference class. Insofar as **Epistemic Position** is a necessary condition for knowledge, the truth of ‘*S* knows *p*’ always presupposes that a certain reference class is chosen for *S* and *p*. Non-standard contextualism assigns a different role than standard contextualism does to context in fixing the semantic content of a particular tokening of ‘know.’ The role of context is to determine the relevant reference class. Standard and non-standard contextualism are consistent with each other. ‘Know’ may have two implicit slots to be contextually filled. Furthermore, the two contextual parameters, epistemic standard and reference class, are independent, in that fixing one does not fix the other.

The basic semantic framework proposed by the standard analysis of gradable adjectives is not applicable to non-standard contextualism. On the standard analysis of gradable adjective, it is a context-sensitive matter whether *S* is tall, or more properly, whether ‘*S* is tall’ is true. But it is not a context-sensitive matter how tall *S* is, or more properly, whether ‘*S* is of height *d*’ is true. The position of John on the scale for tallness is constant across contexts. On the other hand, on non-standard contextualism, not only whether ‘*S* knows *p*’ is true but also whether ‘*S* has epistemic position *EP* regarding *p*’ is true is a context-sensitive matter. The position of *S* with regard to *p* on the epistemic scale varies with context. Thus, non-standard contextualism, as opposed to standard contextualism, cannot be built on the model of the semantics of a wide variety of gradable adjectives.

However, there are gradable adjectives that have a reference class slot or index. Remember that Heller's (1995b) contextualist solution to the reference class problem relies on the fact that the ordinary gradable adjective 'reliable,' as applied to a token process, is relativized to a reference class. Of course, the reference class for 'reliable' may be (in part) explicit in such constructions as 'John's perception on this occasion is reliable.' or 'that process is reliable in good lighting conditions.' etc. Even though 'reliable' may be relativized to both a process type and an environmental type, I assume for simplicity that a reliability sentence has a single slot for a reference class. Thus, the logical form of reliability sentences with explicit expressions for the pertaining type(s) must have a slot for a reference class. Given this, it is plausible that the positive form of 'reliable' is analyzed as having an implicit slot for a reference class. In addition, 'probable' in the sense of definite probability must have a similar logical form. What function 'probable' or 'reliable' encodes and what type of things saturates the reference class slot depend on what theory of probability or reliability is correct.

The semantics of 'know' can be modeled on that of these specific gradable adjectives. Indeed, it is natural to expect that the semantic behaviors of 'know' and its cognates are comparable to those of 'probable' and its cognates, since they are analyzed similarly in the formal and substantive theory. Then, it is plausible to conceptualize that 'know' has an implicit slot for a reference class, where what type of thing saturates the reference class slot depend on what substantive theory of knowledge is correct.

This model is not yet contextualist. The contextualist element comes in when the model is combined with the claim that context determines the value of the reference class variable. An important difference between standard and non-standard contextualism is this:

not only is the value of the epistemic standard variable '$1 - \varepsilon$' independent of the value of the reference class variable 'G,' but also it has a different semantic function. The value of $1 - \varepsilon$ does not have effects on the degree onto which EP maps p; on the other hand, the value of $EP(p)$ depends on the value of G. That is, unlike functions encoded by most gradable adjectives, EP takes three arguments: a subject, a proposition, and a reference class.

Standard contextualism does not entail that any choice of epistemic standard is permissible in a context; when 'know' is used with a modifier for the epistemic standard, e.g., 'really' or 'with certainty,' it compositionally constrains the value of the epistemic standard variable. Even when 'know' is used without such a modifier, certain features of the context do the same job. Similarly, on non-standard contextualism, the value of the reference class variable is contextually constrained. This is the sense in which context determines the 'relevant' reference class.

§2. Five Arguments for Task-Sensitive Contextualism

In this section, I will adduce five different arguments for a version of non-standard contextualism as a solution to the reference class problem in epistemology. If any one of them is compelling, it gives indirect support for the semantic framework of non-standard contextualism. Although the arguments are independent, and some of them are weak, they jointly give strong support for a line of response to the reference class problem: the choice of the relevant reference class is sensitive to what cognitive task the attributor of putative knowledge assigns to the subject. The version of non-standard contextualism that integrates task assignment is called 'task-sensitive contextualism.'

Standard contextualism typically postulates two types of attributor factors as responsible for the contextual shift in epistemic standard: (a) psychological features of the attributor: in particular, what error possibility incompatible with *p* is salient to her; and (b) pragmatic features of the attributor: in particular, how much is at stake on *p* for her. I do not deny, and even take it to be plausible, that these psychological and pragmatic factors contribute to and constrain task assignment on the attributor's side. For the sake of clarity, however, I focus on a distinctive type of attributor factor, attributor's interests in the task she thinks *S* is engaged or ought to engage in.

§2. 1. Argument 1: Pragmatics of Inductive Inference

The first argument for the context-sensitivity of the relevant reference class comes from psychology of the evaluation of inductive inference.[9] Here, I propose to use psychological studies of a type of induction, *category-based inference*, as a model on which our practice of choosing the relevant reference class is to be understood.[10] It seems that category-based inference is analogous to direct inference in many relevant respects. In particular, the proposal I am offering consists of two hypotheses: (i) the relevance of a

[9] In the psychological literature, *evaluating* inductive inference and *performing* it are rarely distinguished. The assumption behind this seems to be that when subjects of an experiment are presented an inference and asked to judge its strength, they go through, and thereby perform it on their own.

[10] What I call 'category-based inference' is more commonly referred to as 'category-based induction.'

reference class is assessed in terms of the strength of direct inference involving that reference class; and (ii) our evaluation of the strength of direct inference is psychologically analogous to that of category-based inference.

(i) is plausible enough: the relation definite probability bears to a reference class is easily represented in the form of direct inference. For example, consider again Vann's case of John Smith, and compare the two different direct inferences one may make about the probability of John Smith living up to the age of 61:

John Smith visited Madeira

P(x lives to the age of 61 | x visited Madeira) = r_1

∴ P(John Smith lives to the age of 61) = r_1

John Smith is a consumptive Englishman

P(x lives to the age of 61 | x is a consumptive Englishman) = r_2

∴ P(John Smith lives to the age of 61) = r_2[11]

Intuitively, the second direct inference makes the conclusion more strong or rational than the first one does; the reason is no doubt that the minor premise of the first inference, even if true, is less relevant, than that of the second inference is, for the definite probability

[11] The reference classes in these inferences, {x: x is a consumptive Englishman} and {x: x visited Madeira}, are taken from Hájek(2007b).

P(John Smith lives to the age of 61); in other words, the reference class {*x*: *x* visited Madeira} is not of much relevance for the definite probability in the conclusion.[12] [13]

(ii) is more controversial. Unfortunately, as far as I know, there has been no psychological study of direct inference, and, no empirical evidence is directly available for the purpose of working out our practice of evaluating direct inference. Be that as it may, there are two reasons that psychological studies of category-based inference provide clues to the psychology of direct inference, and given (i), to the relevance of reference classes.

First, category-based inference (with a single premise) is argument of the following form:

Category *A* has the property *X*

∴ Category *B* has the property *X*

For category-based inference to be strong, it is crucial for *A* to be an appropriate category, type, or class for the conclusion. In this regard, category-based inference is similar to direct

[12] The epistemological reference class problem presupposes that one or another reference class assignment is rational, and hence is relevant. There may be no relevant reference class in some cases.

[13] Since both relevance and inductive strength come in degrees, (i) may be read as the hypothesis that the strength of direct inference co-varies with the relevance of the reference class involved in it. This is not my intended reading; for the present purposes, it is enough that any inductive inference is either strong or weak, and correspondingly, any reference class is either relevant or irrelevant. Indeed, most psychological models for category-based inference are indifferent to the degree of inductive strength.

inference; strong direct inference requires that the chosen reference class be appropriate for deriving the conclusion.

Second, the strength of category-based inference seems to be judged by similar criteria to those pertaining to direct inference. There is ample evidence about psychological properties of category-based inference that category-based inference is judged to be strong on account of similarity, homogeneity, and typicality (for an overview, see (Heit, 2000)). Such properties seem to be shared by direct inference. We have seen in Chapter 4 that alleged solutions to the reference class problem in the theory of knowledge assess the relevance of reference class by the criteria of similarity, homogeneity, and typicality. These criteria are also found in our practice of evaluating direct inference. For example, if $P(x$ lives to the age of 61 | x is a consumptive Englishman who visited Madeira$) = r_3$ is available to us, the direct inference on the basis of this indefinite probability may be judged to be more strong. For the class {consumptive Englishmen who visited Madeira} is narrower, i.e., has more in common with John Smith, compared with {consumptive Englishmen} or {visitors of Madeira}.

If the overall similarity, homogeneity, or typicality of categories is all that matter, the strength of category-based inference would not vary with *X* once *A* and *B* are held fixed. However, there are cases in which the strength of category-based inference is not in accordance of the criteria of similarity, homogeneity, and typicality.

Here are two such divergent cases. First, Heit & Rubinstein (1994) experimentally confirm that different respects of similarity come into play in assessing inferences with different properties (cf. also (Ross & Murphy, 1999) and (Smith, Shafir & Osherson, 1993)). For example, when *X* is the anatomical property *having a liver with two chambers*, subjects tend to judge that the argument with (*A*: chicken; *B*: hawk) is stronger than argument with (*A*: tiger; *B*: hawk); whereas, where *X* is the behavioral property *preferring to feed at night*, subjects tend to judge that the argument with the former pair of categories is weaker than the one with the latter pair.[14]

Second, one tends to attribute high inductive strength to the argument in which the same type of causal relation holds between *A* and *B*, on the one hand, and *X*, on the other hand. Effects of causal relation on the strength of category-based induction are observed when subjects have knowledge about it (Sloman, 1994, 1997), or the settings of the experiment make it particularly salient (Lassaline, 1996). To use Sloman's (1994) example, compare the argument

(a) Many ex-cons are hired as bodyguards
∴ Many war veterans are hired as bodyguards

with the one

(b) Many ex-cons are unemployed
∴ Many war veterans are unemployed.

[14] To minimize influences from other factors than *X*, these 'blank' properties – properties known but unfamiliar to subjects – are used in their experiments.

One's judgments of this pair of reasoning are asymmetrical: the former is strong, and the latter is weak. Heit (2000) draws the general moral from the examples like this that "inductive inference with meaningful properties critically depends on determining which known characteristics of the categories are causally related to or predictive of the property to be projected" (p. 583). Furthermore, Rehder (2006) finds that similarity in the respect of causal relation outweighs overall similarity (see also (Medin, Coley, Storms & Hayes, 2003)).[15]

Early psychological models for category-based inference have been built around the similarities and hierarchical relations encoded by categories *A* and *B*. The divergent cases just mentioned, however, imply that the relevant respect of similarity differs from inference to inference. In order to account for the variability in relevant respect of similarity, recent models are developed on the basis of the idea that information encoded in *A* and *B* is inadequate for explaining our varying assessments of category-based inference; rather, a good model must incorporate extra-linguistic or pragmatic features of category-based inference.

---

[15] These two types of examples, then, constitute counterexamples to similarity-based models, such as the similarity-coverage model (Osherson, Smith, Wilkie, López & Shafir, 1990) and feature-based induction model (Sloman, 1993, 1998). To be fair, however, Osherson et al. are aware of some of these difficulties for their model.

Among several recent models, Heit's (1997) contrast model is suggestive of constructing a model for direct inference in general, and more particularly, for the choice of the relevant reference class in knowledge attribution[16]:

The contrast model, as Heit applies it to category-based inference, is originally proposed by Tversky (1977), as a model of similarity judgment. On Tversky's contrast model, the relevant respect of similarity varies with mainly two contextual factors: explicit or implicit instructions and the salience of features. In an experimental setting, typically, the instruction for subjects states in what regard the objects in question are to be compared with each other; and subjects may come to get a grip on a more specific instruction in the process of experimentation, as the explicit instruction is not fully specified. The salience of one or a set of features of objects also determines the relevant respect of similarity.

For Tversky, the salience of a feature depends on two types of properties, called 'intensive' and 'diagnostic,' respectively. The former refers to properties with high degrees, e.g., the brightness of a light, the loudness of a tone, the saturation of a color, the size of a letter, and so on. The latter is "the classificatory significance of features, that is, the importance or prevalence of the classifications that are based on these features" (ibid., p. 342). Unlike intensive properties, diagnostic properties are sensitive to what set of objects is under consideration, and for this reason, yields the 'extension effect': taking the union of two sets of objects makes salient the feature shared by all the members of the union even

[16] The relevance model proposed by Medin et al. (2003) has the virtues of both the contrast model and the hypothesis-based model mentioned below. Though it can offer the same type of explanation of the cases at hand, I focus on the contrast model here, merely to avoid undue prolixity. The relevance model is explicitly pragmatic; it is based on Sperber & Wilson's (1996) relevance theory.

when that feature is already shared by the members of one set. Thus, contrasting a pair of objects to be compared with a third object tends to bring a certain feature to salience, and thereby, an otherwise irrelevant respect of comparison becomes relevant for that pair.[17]

Heit argues that the contrast model has straightforward application in accounting for the first divergent case: when the premise contains the anatomical property *having a liver with two chambers*, anatomical features are made salient. In the respect of anatomy or internal structure, hawks are more similar to chickens than tigers to chickens. Thus, the contrast model yields the correct verdict that the argument from chickens to hawks is stronger than argument from tigers to hawks. By contrast, when the premise contains the unfamiliar behavioral property *preferring to feed at night*, more salience is added to behavioral features. Then, features about the habits and ways of preying determine the relevant measure of comparison. With regard to such features, hawks and tigers would seem particularly similar, and hawks and chickens would seem less similar. So, the contrast model, again, correctly implies that the argument from hawks to tigers is stronger than that from hawks to chickens.[18]

Heit does not deal with the second divergent case, and it is not obvious how the contrast model applies to it. It, however, seems possible for the contrast model to take advantage of the hypothesis-based model, developed by McDonald, Samuels & Rispoli

---

[17] For more detailed mechanism of how the measure of similarity varies with context, see (Goldstone, Medin & Halberstadt, 1997) and (Medin, Goldstone & Gentner, 1993).

[18] Psychological models of induction are in general designed to accommodate violations of the norms of inference, and explain why, in certain cases, one commits oneself to fallacies like the conjunction fallacy and the inclusion fallacy. Heit notes that the contrast model fails to accommodate the inclusion fallacy.

(1996) and Sloman (1994, 1997): on this model, roughly, the strength of category-based induction depends, in part, on whether the same type of causal explanation suggested by the premise applies to the conclusion. Thus, (a) is reckoned strong, because, in judging its strength, one readily comes up with the same explanation for the premise and the conclusion, viz., that war veterans and ex-cons are hired as bodyguards because they are tough and experienced fighters; whereas, (b) is reckoned weak, because alternative explanations are available for the premise and the conclusion. The point of the hypothesis-based model is that one is guided by an implicit causal-explanatory instruction when given inferences involving causal relations. As Tversky emphasizes, the contrast model posits the implicit instruction as a determinant of the relevant measure of similarity, and thereby, it can incorporate the virtue of the hypothesis-based model.

The contrast model is doubly pragmatic, in that it incorporates as determinants of the relevant respect of similarity the two contextual factors: (i) non-linguistic information embedded in context and (ii) the purpose for making or evaluating inference in context. Hence, the contrast model explicitly endorses the context-sensitivity of the judged strength of category-based inference. This gives some support for the claim that the strength of direct inference, and the relevance of reference class, as we judge them, is context-sensitive. In addition, relying on the contrast model, it may be possible to further hypothesize that the relevance of a reference class is determined by instruction and salience.

Even if it is granted that the relevance of a reference class for a given direct inference is context-sensitive, it does not yet show that non-standard contextualism about ‘know’ is right. But in §2. 3 and 2. 4, I will show that manipulation of (i) and (ii), in effect, results in differences in knowledge attribution. So, if these two factors are context-dependent, this gives us reason to think that 'know' is context-sensitive as well.[19]

§2. 2. Argument 2: The Top-Down Approach to Computation

In §3. 6 of Chapter 4, I discussed **Task-Sensitivity** as a solution to the reference class problem, and raised two problems against it: (i) computational-level specification of the task of a token process, given in psychology, may not uniquely determine the relevant process type, and (ii) it may not be in accordance with our intuitive specification. In this section, I outline how computational theorists and philosophers of computation overcome the problem similar to (i). The way they do is to combine computational-level specification of the relevant process type with an interpreter's perspective. Task-sensitive contextualism can follow the same path in solving (i), and it also suggests a way to deal with (ii).

A problem quite similar to the reference class problem has been a focus of the study of computation. To begin with, the study of computation, or more narrowly, the study of a certain kind of computation, say, visual computation, aims to identify the process type

[19] Heit & Rubinstein (1994) regard their findings of divergent cases as the first step toward the solution to the new riddle of induction, being guided by the questions like why the anatomical property *having a liver with two chambers* is more projectible from the property *being chickens* to the property *being hawks* than from *being tigers* to *being hawks*. Even though this is still in the state of development, psychological research surely sheds significant lights about the solution to the new riddle of induction. Given the similarity between the new riddle of induction and the problem of non-projectible reference class, I believe that the same should be said about the latter. Holland, Holyoak, Nisbett & Thagard (1989) offers a similar pragmatic approach to the new riddle of induction (see also (Holyoak & Thagard, 1996)).

relevant for the purpose of the study; a computational theorist of vision, Wildes (1990), notes,

> The key idea behind the computational approach is to view the study of vision as an inquiry into the type of information processing that is necessary to make useful inferences about properties of a three-dimensional world from corresponding visual images. This point of view allows one initially to abstract away from the algorithmic and implementational details of how any particular vision task is performed by any particular visual system. Instead, initial effort is placed into establishing formal relationships that serve to define exactly what problems are being solved and into specifying why those particular problems are being solved. (p. 332)

Wildes, along with Marr (1982), pursues the study of vision *via* the top-down approach, i.e., from the computational level, through the algorithmic level, to the implementation level. The computational level explains what task a system purports to perform, and why it has the task it does. Specification of the task of a system requires investigation into the external environment surrounding the system and the interaction between them. Shagrir, (2001), Shapiro (1995, 1997), and Sprevak (forthcoming) interpret Marr to show that task specification is to be in the semantic or representational vocabulary, i.e., the computational level specifies what semantic or representational content information processed by the system has. The algorithmic level is constrained by the computational level, since the algorithm or function specified at the algorithmic level is that whose realization contributes to the achievement of the task specified at the computational level.

The bottom-up approach is inadequate for constructing the computational theory of a system, mainly on two grounds: first, different physical systems are capable of implementing the same algorithm or realizing the same function. This means that starting at the implementation level cannot settle the question central for the implementation level; second, a system is capable of implementing different algorithms or realizing different functions, while taking in the same inputs and yielding the same outputs. The first point is simply *multiple realizability*, and the second point is a consequence of *universal realizability*.[20]

---

[20] The term 'universal realizability' comes from Searle (1992, pp. 207-212), but commonly applied to the Putnam's thesis that every open system realizes every abstract finite automaton. Putnam, in an appendix of his (1998, pp. 121-5), offers a proof of this thesis. His proof starts with finite-state automaton without input and output (henceforth, inputless FSA), and then he extends the result of the proof to FSA with input and output. Searle, on the basis of an argument similar to Putnam's, claims that multiple realizability entails universal realizability.

Chalmers, in his (1994), and more thoroughly in his (1996), argues that Putnam's proof presupposes certain weak notions of causal transition and causal dependence, and they are inadequate to capture causal complexity required for the realization relation. Putnam's notions are weak, since they only require the actual sequence of physical state-transition (by an inputless FSA) or the dependence of the actual outputs on the actual inputs (by a FSA with input and output) to correspond to one sequence or a small set of input-output sequences of formal-state transition. Put differently, with Putnam's notions, it is enough for a physical system to implement g that the sequence or sequences produced by the physical system accords with the sequences of actual values or a limited range of values of g. Indeed, I often see these weak notions being used in the discussions on the reference class problem for type reliabilism. Chalmers proposes stronger, Before proceeding, it is necessary to explicate what it is for a physical system to implement an algorithm or realize a function. It is commonplace in philosophy of mind to define the realization relation in terms of isomorphism between types of physical states and types of algorithmic states; for example, here is David Chalmers' definition:

> A physical system implements a given computation when there exists a grouping of physical states of the system into state-types and a one-to-one mapping from formal states of the computation to physical state-types, such that formal states related by an abstract state-transition relation are mapped onto physical states-types related by a corresponding causal state-transition relation. (Chalmers, 1994, p. 392)

Let $I_1, \ldots, I_l$ be input types, $O_1, \ldots, O_m$ be output types, and $S_1, \ldots, S_n$ be algorithmic state types, of algorithm M. Similarly, Let $i_1, \ldots, i_l$ be output types, $o_1, \ldots, o_m$ be input types, and $s_1, \ldots, s_n$ be physical state types, of physical system S. A function g characterizes every state transition $[S, I] \rightarrow [S', O]$. Chalmers' definition is that S implements M, or equivalently, S realizes g, iff (i), for every state transition $[S, I] \rightarrow [S', O]$, there is a physical state transition $[s, i] \rightarrow [s', o]$, and (ii) isomorphism holds between input, output, and state types of M, on the one hand, and those of S, on the other hand. More

---

counterfactual notions as substitutes, and what they do is, roughly, to broaden the domain of g so as to include non-actual states. As Chalmers admits, this does not much change the results of Putnam's proofs; a weak form of universal realizability still holds: every physically complex physical system realizes every finite automaton. For rejoinders to Chalmers' objection to Putnam, see (Bishop, 2002) and (Scheutz, 2001).

precisely, (ii) means that there is a bijective (*onto* and *one-to-one*) function f such that $f(s) = S$, $f(i) = I$, $f(s') = S'$, and $f(o) = O$.

Universal realizability, as Chalmers accepts it, is the thesis that a physical system realizes many different computational functions from inputs to outputs at the same time. The problem it poses for computational theories, then, is to determine *the* realized function. This problem endangers **Psychological Realism**, discussed in §3. 1 of Chapter 4, as a solution to the reference class problem for type reliabilism. **Psychological Realism** is the position that there is a unique computational function for any process token, and the relevant function is that which constitutes a natural kind, as psychological inquiry reveals it. The faith in psychology that **Psychological Realism** presupposes is not shared among computational theorists and philosophers of computation.

The multiplicity of computational functions for a system to be in a position to realize, indeed, arises on several different levels. First, identity of computational function depends on how algorithmic states are typed. Suppose that function f is defined in such a way that $f(s_1 \vee s_3 \vee s_5 \vee s_7) = S_1$ and $f(s_2 \vee s_4 \vee s_6) = S_2$. f violates the condition for isomorphism, if each of $s_1, \ldots, s_7$ corresponds to a physical type. For, if so, f is not bijective, simply because it is not a one-to-one function. An easy fix is available: to assign s_1, s_3, s_5, s_7 to a single type, and s_2, s_4, s_6 to another.[21] Moreover, it is commonplace in

[21] Scheutz (2001) emphasizes the types specified by physics in circumventing the problem from disjunctive states. This does not mitigate the force of the problem to a great degree. Physical states of human brain may differ from person to person at the micro-physical level. In order to theorize the computation of human brains in a united way, otherwise disjunctive types may be typed as one. Either way, typing by physics is of little

computational theories to hold that computational functions, specified at the algorithmic level, are described in virtue of syntactic properties. But it is not necessary to assign each syntactic property to a single type, because such assignment may be too fine-grained.

Second, an algorithm often comprises sub-algorithms, but some of them may be irrelevant. Chalmers (1996) points out that most cognitive or computational systems have a combinatorial structure, i.e., a structure divided into components which causally interact locally and globally in complex ways. Thus, Chalmers argues that an internal state of a system with combinatorial structure is to be characterized as a vector $[S_1, ..., S_n]$ rather than a single state *S*, where the *i*th component of the vector can take on a finite number of different values; put differently, the computational function *g* for a complex system takes sub-computational functions as its domain values. However, it may be unnecessary to include every possible sub-computational function in the characterization of *g*. If so, multiple *f*s are possible, depending on which sub-computational function is relevant.[22]

Third, and most importantly, algorithmic states of a system are typed in virtue of what they represent. Recall the Two System case, mentioned in §3. 6 of Chapter 4. It involves the two systems such that the functions realized by them may be characterized differently, as {Roman numerals-processing} and {Arabic numerals-processing}, or identically, as {number-processing}. Both characterizations specify the computational function in semantic or representational terms: the former characterization specifies the representational objects of each system as numerals of a certain sort, whereas the latter

---

importance for the purpose of the theory of knowledge; our judgments about knowledge and ignorance do not track such fine-grained typing.

[22] Shagrir (2001) offers a good illustration of this point by way of example.

characterization specifies them as numbers. This difference in representation affects the typing of inputs; although Roman and Arabic numerals are syntactically distinct, the latter characterization treats them as falling under the same type. Hence, identity of computational function depends on identity at the representational level.

In addressing the problem from universal realizability, those philosophers who are influenced by Marr, such as Shagrir, Shapiro, and Sprevak, emphasize the importance of the computational level for the algorithmic level: task specification imposes a constraint on what function a system realizes.[23] To borrow a useful notion from Shagrir (2001), the function a system realizes is *the function associated with a task*, i.e., the function whose realization contributes to the achievement of the task of the system. According to Marr, the task of vision is to derive three-dimensional representations of shapes and positions of things from two-dimensional images.[24] The computational theory of vision, then, purports to describe the function associated with this task at the algorithmic level.

Although Marr is not explicit about this, his words suggest that the task he assigns to vision is determined by the facts about human adaptation to environment. Sprevak, however, adds that the computational theorists may assign narrower tasks than Marr does to

---

[23] Pylyshyn (1980, 1984) addresses the problem of functionally identity from a different perspective. He argues that time complexity and cognitive penetrability are key to specifying the relevant function.

[24] See, especially, (Marr, 1982, pp. 36-8). In the terminology of Shapiro (1997, p. 134), relative to this 'chief task' of vision, each part of vision has a 'service task,' i.e., one whose completion contributes to the achievement of the chief task.

a certain system, for some specific purpose.[25] He argues that task assignment is a matter of one's interpretation, being constrained by the interpreter's purpose and the linguistic convention, in addition to the interaction between the system and its environment. His Two System case is designed to show this point: whether the two systems realize one function {processing numbers} or different functions {Roman numerals-processing} and {Arabic numerals-processing} depends on how one interprets the systems. And, room for various interpretations would remain even if every fact about the two systems and their environments is determined. In the face of the problem concerning universal realizability, other philosophers, such as Chalmers (1994, 1996) and Searle (1992), also accept that what automation a physical system implements is 'observer-relative' or 'interest-relative.'[26] Task-assignment at the computational level gives substance to the idea of interest-relativity of implementation or realization.

---

[25] He argues that the process type specified independently of external environment is relevant for the purpose of inquiring how the system would perform the same task if evolved in a different environment. The reference class induced by Cartesian skepticism is similar to this internally specified process type, in that it does not involve any environmental feature; but they are also different: Sprevak identifies the internal type with a type characterized in the syntactic vocabulary; whereas the reference class the skeptic appeals to involves phenomenal properties, which are representational at least on the standard view. Sprevak's point is somewhat similar to the claim made by those who respond to Cartesian skepticism with the weapon of externalism about mental content, by holding that victims of the skeptical scenario cannot have normal representational beliefs, since it lacks the history of interaction with environment.

[26] Searle often uses these terms in contending that what algorithm a system implements cannot be determined by reference to its intrinsic features alone. See, e.g., (Searle, 1992, pp. 209-10)

As mentioned in §3. 1 of Chapter 4, the problem concerning functional identity is the focus of the debate between Adler & Levin (2002) and Feldman & Conee (2002). Consider, again, some process types in the list of Conee & Feldman (1998): (1) {process of a retinal image of such-and-such specific characteristics leading to a belief that there is a maple tree nearby}, (2) {process of relying on a leaf shape to form a tree-classifying judgment}, and (3) {the visual process}. Adler & Levin hold that (1) through (3) describe the same function at increasing levels of generality; while Feldman & Conee accept that process tokens are typed in virtue of what functions they realize, they reply to Adler & Levin that these processes need not involve the same function. If (1) to (3) are of the same type, the reliability of them has the same value, and no reference class problem arises for type reliabilism.

Insofar as Marr is right about the task of vision, the function defined by (3) requires such typing as to categorize a broad range of shapes and positions as being of different types. The function, then, is highly complex, and hence, in order to realize the function, a physical system must have causal structure complex enough to mirror the relations among input, output, and states types that the function encodes. However, the task assignment relevant for the purpose of computational theory may not be necessary for different purposes. If type reliabilism is right, knowledge attributions or denials involve type reliability assessments. For the purpose of assessing type reliability in attributing or denying knowledge, only coarse-grained typing may be required.

Furthermore, attributors may assign different tasks for different purposes; a way to make a function out of (1) is by dividing input types only into two, a specific feature of maple tree – seven leaf points, say – and other leaf points, and dividing output types correspondingly. In order for a system to realize such a function g_1, only the transition [I: seven leaf points ∨ non-seven leaf points] → [O: belief that it is a maple tree ∨ belief that it is not a maple tree ∨ withholding] is relevant. There is no reason to rule out this function as irrelevant, especially if one aims to judge whether a subject reliably discriminates maple tree from other tree species. A function g_2 for (2) is more complex: it involves the transition (I: shapes of leafs) → (O: belief that it is of species k ∨ withholding), where the domain of I is different shapes of leafs, and the domain of O includes beliefs about a variety of tree species k. g_2 may count as the same as g_1, if the attributor uses it for the purpose of judging about the reliability of a subject's discriminatory ability regarding maple true and other tree species; for a different purpose, say, of judging whether a subject reliably discriminates among various tree species, they count as different.[27]

Conee & Feldman (1998) object to Heller's contextualist solution to the reference class problem for type reliabilism, on the ground that contextual factors do not always determine the unique process type for a given process token. Suppose that John sees Smith look at a maple tree nearby, and asserts "Smith knows that there is a maple tree nearby." Since John may not have in mind any salient process type, Conee & Feldman claim that the

[27] A similar point can be made about the Red Paper case presented in §3. 5 of Chapter 4. What count as perceptual equivalents cannot be settled independently of whether the relevant function encodes them as being of the same type. Again, the types of perceptual experience relevant for psychology may not coincide with the types relevant for knowledge attributions, and knowledge attributions may be made for different purposes in different contexts. In general, however, our ordinary typing for the purpose of figuring out what one perceives is likely to constrain on typing for knowledge attributions (see (Vahid, 2009) and (Brandom, 1994, p. 213)).

process token can be associated with any of (1) through (3). This objection is not even *prim facie* strong. First, the case like this does not show that it is in principle impossible to obtain the relevant process type from contexts. It seems that it is much harder to say that John's knowledge attribution is not driven by any interest in what process or ability Smith exercises in identifying a maple tree as such, whether John is aware of it or not at the moment. Second, this case summons the reference class problem only if (1) to (3) differ in type reliability. But this is not necessarily the case.[28]

The theoretical privilege of the computational level over the algorithmic level entails that the problem of what type a token process exemplifies cannot be settled unless the relevant task of the system in question is specified. Provided that task specification is interest-relative, this provides further ground for task-sensitive contextualism. Insofar as

---

[28] This objection is designed to emphasize that reference classes have nothing to do with the truth condition of '*S* knows *p*.' Conee & Feldman appeal to this case in defending internalist evidentialism, since they hold that the degree of internal justification does not depend on the choice of reference class. As we know, this is not the case. Another objection Conee & Feldman makes against Heller is that attributors may have wrong ideas about the process token in question. Suppose that Jones sees Smith identify a bird as being of a certain species when Smith has had only a brief glimpse of it under poor lighting conditions. By thinking that Smith's process token is of {bird classification on the basis of a brief glimpse under poor lightning}, John may deny Smith knowledge; while Smith's bird identification is based on hearing the tweet of the bird. I am not sure how this is meant to be an objection. As Conee & Feldman claim, the process exemplified by Smith is not a token of unreliable process type {bird classification on the basis of a brief glimpse under poor lightning}. Then, John's knowledge denial is false, since type reliabilism entails that '*S* knows *p*' is true only if the process type the token of which has produced *S*'s belief *p* is reliable. Having the same input and output is not enough to categorize two or more process tokens as being of one and the same type. More complex causal conditions need to be met, as noted in footnote 20.

type reliabilism is right, knowledge attributions or denials require specification of the relevant process type, and what type is relevant is determined by attributor's interests or purposes in attributing or denying knowledge.

§2. 3. Argument 3: No Problem of Trivializing Reference Class

The third argument provides another reason to adopt task-sensitive contextualism: if task-sensitive contextualism is adopted, the problem of trivializing reference class for the theory of knowledge is solved, or more properly, resolved. The problem of trivializing reference class is the problem of how the reference class involving the individual in question as its sole member is excluded; the reference class is trivializing, since, relative to it, the type-token distinction is lost. Indeed, as mentioned in §3. 2 and §3. 3 of Chapter 4, one main reason for the failures of **Maximal Specificity** and **Broadest Homogeneity** is that they lead to the trivializing reference class. Once relativized to a trivializing reference class, a proposition or a token belief *p* may be assigned the maximal value of epistemic position merely because it is true.

This may seem too bewildering at first glance, given epistemologists' long efforts to provide an epistemic necessary condition for knowledge independent of truth and belief. Goldman (1999, 2001, 2002), however, argues that there are two uses or senses of 'knowledge' that conform to "some standard, ordinary sense of the term in colloquial English (as judged by what epistemologists who attend to ordinary usage have identified as such)" (2002, p. 183): in one standard sense, knowledge is "justified, warranted, or acquired in some suitable fashion (e.g., by reliable methods)", and, in the other sense, "to know that *p* is simply to possess the information that *p*, where "information" entails truth but "possession" merely entails belief, not necessarily justified or warranted belief" (ibid., p. 185).

Goldman, in putting forth his main argument for the use of knowledge as possessing information, relies on the example given by John Hawthorne:

> (Six Children) I give six children six books and ask them each to pick one of the books at random. All but one contains misinformation about the capital of Austria. I ask the children to look up what the capital of Austria is and commit the answer to memory. One child learns 'Belgrade', another 'Lisbon', another 'Vienna', and so on. I ask an onlooker who has witnessed the whole sequence of events (or someone to whom the sequence of events is described) 'Which one of the schoolchildren knows what the capital of Austria is?' or 'How many of the children know what the capital of Austria is?' It is my experience that those presented with this kind of case will answer, not by saying 'None of them', but by selecting the child whose book read 'Vienna' – even though that child was only given the correct answer by luck. (Note in this connection that if I make a five dollar bet on a certain child knowing that Vienna is the capital of Austria, you will pay up as soon as you are convinced that the child believes the capital to be Vienna. You will not inquire further about how the child came by that information – whether by dumb luck or from an

> informant that normally lies – even if you have reason to suspect such an unreliable source.)[29]

Hawthorne intends this case to show that '*S* know *p*' has a usage equivalent to '*S* possesses the information *p*.' But, unlike Goldman, Hawthorne does not claim that 'know' is ambiguous; rather, he argues that this case counts *prime facie* in favor of standard contextualism, on the basis of the following diagnosis:

> When an onlooker is asked to make a knowledge judgment in these contexts, she is not very worried about the possibility that the person to whom knowledge is ascribed has made a mistake because she is very sure that he hasn't; and it is clear, moreover, that the questioner is also not very interested in how easily the knowledge candidate might have erred. Thus, in effect, the standards associated with the verb 'know' in such a context are extremely low. (Hawthorne, 2004a, pp. 69-70)

Thus, for Hawthorne, the Six Children case involves a limiting context in which '*S* know *p*' requires no strong epistemic position with regard to *p*, in order for it to be true.

---

[29] Hawthorne first offers this type of example in his (2000, pp. 202-3), and then develops it in his (2002, pp. 253-55) and (2004a, pp. 68-9), with the help of Tamar Gendler. The example Goldman refers to is the 2002 version as presented at the Rutgers Epistemology Conference in 2000. I am here citing the latest, 2004 version.

Goldman and Hawthorne have different diagnoses of the same case: (a) it supports the ambiguity of 'know' and (b) it supports the context-sensitivity of 'know.'[30] Although it is not clear to what degree (a) differs from (b), if the ambiguity in (a) is lexical ambiguity; (b) is favorable, simply for the reason of theoretical parsimony: it need not posit another type of sense or use, other than the sense being of interest for epistemologists. However, the context sensitivity in (b) must be accounted for by non-standard contextualism rather than standard contextualism. Standard contextualism squares with the Six Children case, only if (i) *S*'s epistemic position is significantly low, and (ii) an extremely lax standard is in play in the context. I will argue that neither is the case, and that non-standard contextualism offers a better explanation of the Six Children case and the like.

What grounds does Hawthorne have for (i)? He seems to intend the Six Children case to be an example of a case in which *S* satisfies no condition for knowledge imposed by any theory of knowledge: *S*'s degree of internalist justification for the belief that Vienna is the capital of Austria is much lower than usual, and so is the degree of reliability of *S*'s process type or method token by which *S* has that belief. The general reason for *S*'s low epistemic position, as Hawthorne suggests it, is that a type of luck is involved in the case. However, the type of luck involved is not *epistemic* in the proper sense. Pritchard (2005a) distinguishes epistemic luck from other benign types of luck, one of which, called

[30] Hawthorne (2002) mentions, though does not examine, two more WAM options: (c) 'know' always means processing information, but the context-sensitive assertibility condition of 'know' does not allow to assert '*S* knows' in many contexts other than the context of the Six Children case, and (d) 'knows' always means having warranted belief, but the context-sensitive assertibility condition allows to assert '*S* knows' in the context of the Six Children case even though it is false.

'evidential epistemic luck,' is such that "[i]t is lucky that the agent acquires the evidence that she has in favor of her belief" (p. 136). As Pritchard points out, evidential epistemic luck has nothing to do with the strength of epistemic position. Thus, it is far from maintainable that, in this case, *S* knows despite her epistemic position being unusually low.

Furthermore, it is doubtful that any particular theory of knowledge entails that *S*'s epistemic position is low in the Six Children case. First, *S* has no negative evidence against either *Vienna is the capital of Austria* or *the book she is given is reliable*; if she does, other things being equal, *S* is rather justified in disbelieving the information of the book. It is too big a stretch to read the Six Children case as endorsing the possibility of knowledge as negatively justified true belief. Second, not only does *S* have no negative evidence, but also she has some positive evidence that Vienna is the capital of Austria – it is what the book says! Then, a charitable interpretation may be that *S* lacks second-order evidence, i.e., evidence about the reliability of the book. If this is the case, *S* violates **Meta-Justification**. But does she? It is natural to read the case in such a way that *S* does know or at least justifiably believes that books in general are reliable sources of information.[31] And, even

[31] In believing this, *S* is assigning the book at hand to a reference class {books}. This choice of a reference class should not be a problem; as I said above, *S* has no negative evidence against the reliability of books in general, and this fact makes it natural for *S* to use this reference class. If *S* finds some feature of the book fishy, she would not trust the book, because, then, she would have evidence that books in the class {books with such a feature} are unreliable. Still, some may claim that *S* has no second-order justification because *S* does not know or justifiably believe that the particular book she is given is reliable. This is a *non-sequitur*. It is hard to make good sense of the reliability of a particular book in other ways than just mentioned – how to make sense of the reliability of a particular thing is precisely the reference class problem. In a sense of

worse, insofar as **Meta-Justification** is in place, its failure entails that *S* is not justified at all; third, Hawthorne assimilates the Six Children case to Goldman-Ginet's Barn Façade case, and sums up the lesson of them that "in many contexts, Gettierized true belief *is* knowledge" (Hawthorne, 2000, p. 203). Being parallel to Gettier cases, then, *S* is internalistically justified in the Six Children case.[32]

Of course, having internalist justification is not enough to put *S* in a strong epistemic position: taking a pluralist view of epistemic factor, *S* may still be in a weak epistemic position if *S* does not do well along the dimension of some type reliabilist or token reliabilist condition. Hawthorne describes the Six Children case in terms of an 'unreliable source,' by which he seems to mean the following: first, the process by which *S* comes to believe that Vienna is the capital of Austria is type unreliable, since the book *S* is given is one of the six books selected at random, among which only one contains correct information; second, *S*'s method token is not sensitive, since, given the book is randomly selected, there are nearby possible worlds in which *S* falsely believes that Vienna is the capital of Austria.

---

reliability, *the particular book is reliable* is equivalent to *what the book says is true*. Then, this response entails that *S* does not know that Vienna is the capital of Austria.

[32] Hawthorne always pairs the variants of the Six Children case with the Stopped Clock case, in which *S* believes truly about the current time *t* by having a glance at a stopped clock because it stopped at the time of the day at which *S* looks at it. But I doubt that it is true in the Stopped Clock case that *S* knows that it is *t*, even though Hawthorne reports that many do have such intuition. It seems more appropriate to describe the content of *S*'s belief as *the clock reads t*. Hawthorne (2002) offers several variants of the Barn Façade case, some of which is to be analyzed in the same way as the Six Children case, and some of which is more similar to the New Barn Façade case I describe in the next section.

Again, these ideas are based on confusions: the randomness at stake merely concerns the way in which *S* comes to exercise the process or method of reading the book *S* does; it has little to do with the reliability of the process or method. The type of randomness Hawthorne alludes to is indeed quite similar to another type of benign luck pointed out by Pritchard, which he calls 'doxastic epistemic luck': "[i]t is lucky that the agent believes the proposition" (Pritchard, 2005a, p. 138).[33] It, however, is possible to build doxastic epistemic luck into the typing of the relevant process or specification of the relevant method. That is, to specify the process type or the method token as {reading one of the six books by random selection} yields the result that it is unreliable in the sense of both type and token reliability. This reference class would be in place, if the Six Children case were similar to the Barn Façade case; as analyzed in §3. 2 of Chapter 5, the reference class implicit in the Barn Façade case is alike in the relevant respects.

However, the reference class {reading one of the six books by random selection} does not seem to be the one we assign in the context where we judge that *S* knows; nor does it seem to be relevant in that context. To see this, let us consider why Hawthorne holds (ii) that the standard in the context is extremely lax. He claims that "insofar as we think ourselves into contexts where the point of the knowledge question is information possession, our "intuitions" are going to be very liberal in connection with our disposition to ascribe knowledge" (Hawthorne, 2000, p. 255). That is, it is the attributor's interest only in the truth about the capital of Austria that sets the epistemic standard at some extremely low point. Is this really the case? If the attributor's interest is only in whether *S*'s belief is

[33] Pritchard is explicit about implications of epistemic doxastic luck for sensitivity. And such implications are equally applicable to safety.

true or not, why does she need to choose {reading one of the six books by random selection}? This choice puts the attributor in a curious mindset: she cares about what *S* would believe if *S* were given a different book among the six books, while being interested only in the truth of *S*'s belief. Insofar as the attributor is driven by this interest, it is more natural for her to assign a reference class that only involves the very process or method that has led *S* to the belief about the capital of Austria. Such a reference class cannot be other than {reading the book *S* is given}, and it is trivializing.[34] [35] But if my diagnosis is right, the moral to be drawn from the Six Children case is that there are contexts where trivializing reference classes are relevant.[36]

---

[34] I am omitting the process-environment distinction, since this process type naturally fixes the relevant environment type. An alternative way to state my point is to fix the environment type as {this occasion}. Whatever process type is selected, this environment type is trivializing.

[35] Even with this reference class, it is a bit unclear whether *S* satisfies sensitivity. The answer depends on how we conceptualize the details of the case. Does the book keep telling that Vienna is the capital of Austria in the nearest possible worlds where it is not? If it does, *S*'s belief is insensitive. But we may think that the book itself is reliable in the sense of sensitivity, i.e., it does not tell that Vienna is the capital of Austria in the nearest possible worlds. Then, *S*'s belief, in turn, is sensitive.

[36] The Six Children case alone may not be decisive for establishing the existence of such contexts: the attributor may be taken to be using {reading books that say that Vienna is the capital of Austria}. This reference class is not trivializing, but sill results in the verdict that *S*'s epistemic position is nearly maximal. To fully settle what context the attributor is in fact in, psychological inquiry into what interest drives the attributor's epistemic evaluation regarding the Six Children case is required. But, for my point here to hold, it is enough that when the attributor chooses the trivializing reference class, the resulting knowledge attribution is not counterintuitive.

If non-standard contextualism is true, the problem of trivializing reference class is no problem. For a trivializing reference class being relevant in a context does not mean that such a class is relevant across the board. Furthermore, we tend to attribute knowledge in cases like the Six Children case. Task-sensitive contextualism offers a compelling explanation of such cases, because they are cases in which attributors' interests lead to selecting trivializing reference classes as relevant. Hence, the existence of such cases militates in favor of task-sensitive contextualism.

Goldman and Hawthorne share the idea (ii) that the epistemic standard is significantly lowered in the Six Children case, when attributors apply 'know' in the sense of information possession. This is not right, since trivializing reference classes maximize the strength of *S*'s epistemic position rather than lower the epistemic standard. Then, task-sensitive contextualism entails the denial of (i) *S*'s being in a low epistemic position , but this is what we should expect. There is a sense in which the epistemic standard implicit in the Six Children case is lower than usual: the same token process is unreliable if a wider reference class is used. But this is a different sense of epistemic standard with which strengths of epistemic positions can be compared over reference classes. The concept of epistemic standard currently received in epistemology has no such application; the strengths of epistemic positions are only comparable within a reference class, since different reference classes yield different strengths of epistemic position even for the same

S, time, and place. It seems that Goldman and Hawthorne conflate between the two senses of epistemic standard.[37]

§2. 4. Argument 4: Support from Question-Sensitivity

The third argument emphasizes a role of attributor's interests in fixing the relevant reference class. The fourth argument defends the same idea more generally, on the basis of a practical function of knowledge attribution. The value turn in epistemology has motivated inquiry into the normative question of *why knowledge is valuable*. It is commonplace in recent epistemology to approach the normative question by investigating the practical role or function of knowledge attribution, i.e., it is addressed in connection with the role knowledge attribution plays, or the point it has, in our epistemic practice. One role of knowledge attribution, among others, is pointed out by Jonathan Schaffer:

> (Question-Sensitivity) Knowledge ascriptions certify the subject's ability to answer the question. (Schaffer, 2005, p. 236)[38]

Although Shaffer maintains that **Question-Sensitivity** supports *contrastivism* **–**

---

[37] Perhaps, this is because Hawthorne's notion of epistemic standard is quite liberal, as noted in footnote 22 of Chapter 1. Indeed, Hawthorne (2002, pp. 262-7) alludes to shift in reference class; referring to Lewis's (1996) account, he claims that the world ordering is what varies with context. This may be what he means by (ii).

[38] Schaffer notes that **Question-Sensitivity** is found in Castañeda (1980): he claims that "knowledge involves essentially the non-doxastic component of a power to answer a question" (p. 194).

his version of non-standard contextualism –, I will argue that it speaks more in favor of task-sensitive contextualism.[39]

Schaffer glosses **Question-Sensitivity** as follows: (i) 'certify' denotes the illocutionary act, performed by the attributor, of conferring warrant on *S*'s answer to the question.[40] An answer to the question may be formed and stored in *S* as a belief, and needs not to be uttered; (ii) 'able to answer' denotes an epistemic capacity, so that *S* may guess correctly without having an ability in the sense relevant here; (iii) the question needs not to be explicitly posed, but it is always recoverable from context. Schaffer suggests that the

---

[39] There are versions of contrastivism. Here, I only mean Schaffer's brand. Contrastivism is not merely a version of contextualism; it consists of two central ideas: (i) 'know' expresses a ternary relation among a subject, a proposition, and a contrast class, rather than a binary relation between a subject and a proposition; and (ii) what contrast class 'know' has as its relatum is context-dependent. (i) entails that contrastivism should be reckoned a position about the semantics of 'knowing that,' not of 'know,' since it concerns what semantic content 'know' has within its embedded that-clause. The contextualist element of contrastivism is (ii), and in this regard, contrastivism has some similarities with the version of contextualism I am proposing: both add another contextual parameter, contrast class or reference class, to the semantics of 'know' or 'know that.' Beyond this point, it is difficult to compare them, especially because Schaffer does not explain how the contrastivist apparatus is related to epistemic standards. For more on his contrastivism, see (Schaffer, 2004, 2005, 2007a, 2007b, 2008, 2009).

[40] Schaffer uses the notion of 'entitlement,' taken from Brandom (1994), instead of that of warrant. But none of Schaffer's points regarding **Question-Sensitivity** hinges on Brandom's idiosyncratic account of entitlement. Brandom, in fact, holds that one form of entitlement conferral is reliability assessment, and the relevant reference class for reliability assessment is determined by the interests and goals of those who perform the speech act of knowledge attribution (pp. 212-3). Brandom's case for the interest-relativity of reference class is discussed a bit more in footnote 45.

question is determined by what Stalnaker (1978) calls 'pragmatic presuppositions' of context – a proposition is pragmatically presupposed iff it is true in "the set of possible worlds recognized by the speakers to be the 'live options' relevant to the conversation" (pp. 84-5)[41]; (iv) Schaffer further adds that questions always have the contrastive form '*wh*: *p* or *q*,' where '*q*' denotes what he calls a 'contrast class,' i.e., a class of alternatives to *p*. From here, he argues that since knowledge relation logs a question, it also has the contrastive form 'knowing *p* rather than *q*.' Contrastivism is the position that 'know' expresses different ternary relations in different contexts, depending on what contrast class saturates its contrast class slot.

Question-Sensitivity, *pace* Schaffer, can be accommodated more naturally into task-sensitive contextualism – more natural in the sense that it is less committal in the following respects: (a) task-sensitive contextualism can be combined with a wider variety

[41] Stalnaker's notion of pragmatic presupposition is more thoroughly developed in Stalnaker (1974). Stalnaker comments on Schaffer (2004) in his (2004).

Blome-Tillmann (2009) proposes a version of contextualism which directly relies on the notion of pragmatic presupposition: on his contextualism, the epistemic standard for knowledge is, in part, determined by what is pragmatically presupposed in the context in which it is made. This brand of contextualism still remains in the framework of standard contextualism, and is to be distinguished from Schaffer's or my version of contextualism; it does not add another contextual parameter to the apparatus of standard contextualism, and standard contextualism is tolerant about what contextual factor determines the epistemic standard.

of substantive theories of knowledge than contrastivism can be, and (b) it need not incur the extra commitment to (iv).[42]

(a) is easy to see. One odd thing about Schaffer's argument is that he construes (ii) so as to identify having an epistemic ability to answer a question with having compelling ground for the answer to it. From his reference to Wittgenstein (1961), it is clear that Schaffer presumes internalist evidentialism of some sort at the outset. A variety of substantive theories of knowledge, however, can implement **Question Sensitivity** in one way or another. Reliabilist theories in general are to be reckoned accounts of epistemic abilities or competence: such terms, as 'process,' 'ability,' 'capacity,' 'competence,' and 'method,' are used interchangeably in the literature, whether they refer to tokens or types. In addition, even internalist evidentialism needs to appeal to some of these notions, if it incorporates **Reliability Expert**. Thus, (ii) can be construed in terms of these notions, but it is not obvious how **Question-Sensitivity** leads to contrastivism on such construals.[43]

---

[42] Schaffer's argument proceeds by showing how contrastivism 'fits' **Question-Sensitivity** and other data. The argument here purports to show that a closer fit exists between task-sensitive contextualism and **Question-Sensitivity**.

[43] This must be the case, especially because Schaffer takes **Question-Sensitivity**, with (iv), to have motivated the relevant alternatives theory, first proposed by Austin (1946): *S* knows *p* only if *S* can exclude every relevant alternative to *p*. Despite the name commonly attributed to it, the relevant alternatives theory is no substantive theory, since it does not define the important notion of 'exclude an alternative'; on the standard interpretation, excluding an alternative is identified with knowing that it is false, and hence the relevant alternatives theory does not offer a reductive analysis of knowledge. Many combine the relevant alternatives theory with their favored theories of knowledge, e.g., Austin endorses internalist evidentialism, and Dretske (1970, 1971, 1981a) and Goldman (1976, 1986, ch. 3) develops variants of sensitivity theory with reference to

This is a problem for contrastivism since it, just like contextualism, is meant to be a semantic view on ‘know that,’ being neutral on the substantive question of what the relevant epistemic factor(s) is.[44]

(b) is also related to the relation between **Question-Sensitivity** and the substantive theory of knowledge. Each theory of knowledge focuses on some of the notions just mentioned, and it is necessarily relativized to a reference class. Indeed, the relevant reference class may vary with what question is explicitly or implicitly asked in context. In addition, the contextually defined question pertinent to the choice of the relevant reference

---

the relevant alternatives theory. It is sometimes taken for granted that the relevant alternatives theory results in standard contextualism if what alternatives count as relevant is a context-sensitive matter. This is wrong; it is based on the naïve thought that the epistemic standard goes up and down, depending on how large the set of the relevant alternatives is; the larger the set is, the standard goes up, and *vice versa*. For token reliabilism, the relevant alternatives theory contextualized does not have this consequence; on token reliabilism, the set of the relevant alternatives amounts to the set of the relevant possible worlds, and the strictness of epistemic standard is determined by how far, not how many, the relevant worlds must be (cf. (Heller, 1989) and (Shatz, 1981)). For type reliabilism, the relevant alternatives can be taken to the relevant situations relative to which process types are evaluated. The choice of the relevant situations is independent of how high the epistemic standard for knowledge or justification is. These points are, *mutatis mutandis*, applicable to internalist evidentialism, insofar as it is combined with **Reliability Expert**. That is, as I argue below, **Question-Sensitivity** may result in shift in reference class, which changes the epistemic position rather than the epistemic standard.

[44] In addition, this move is even textually incoherent: Schaffer claims that an epistemic ability is “a capacity insofar as one need not actually speak or otherwise exercise the ability in order to possess it” (Schaffer, 2005, p. 236). He clearly has in mind some counterfactual notion of epistemic ability or capacity, which is precisely the notion widely shared among reliabilists.

class needs not to be of the contrastive form. To make this point explicit, a pair of variants of the Barn Façade case is useful:

> (Normal Barn Façade): As Henry passes by the Barn Façade County, he runs into John, a local, and starts chatting with him. John often gloats that visitors to the County are tricked by barn facades. So he asks Henry, pointing by chance to the real barn "what is it?" Henry replies "it's a barn." Finding a new victim of the hoax, John thinks, "Ha ha! He doesn't know it's a barn."

> (New Barn Façade): As Henry passes by the Barn Façade County, he runs into John, a local, and starts chatting with him. John often gloats that visitors to the County are tricked by barn facades. At the moment, however, it occurs to him that Henry might not have proper discrimination power about barns. If so, the trick would not work for Henry and John's ill-natured desire would not be satisfied. So he asks Henry, pointing by chance to the real barn "what is it?" Henry replies "it's a barn." With a little relief, John thinks, "OK, he knows it's a barn."[45]

---

[45] Brandom (1994) argues that attributor's goals or interests regarding knowledge attribution determine the relevant reference class, by appeal to variants of the Barn Façade case. Greco (2009, p. 24; 2010, p. 80) offers a similar argument, though it is designed to be in favor of his 'subject-sensitive' contextualism). Brandom (pp. 209-12) asks to suppose that the Barn Façade County is part of a state the majority of the counties of which contain no barn façade, and that state is part of a country the majority of the states of which contain many barn facades, and so on. Each of the county, the state, and the country constitutes a reference class. Brandom merely points out that reliability takes different values, depending on which reference class

The two cases are precisely the same in all the relevant respects except the epistemic interest John has in mind when he asks the question “what is it?”; they are about the same Henry, the same John, and the same barn in the Barn Façade County (the same time, and the same world, if necessary). The normal Barn Façade case is simply the Barn Façade case involving an attributor, John, and the details of its setting are set in accordance with DeRose’s modified Barn Façade case. Thus, insofar as it is true in the modified Barn Façade case that Henry does not know that the structure he is seeing is a barn, John’s utterance “he doesn’t know it’s a barn” must be true in the normal Barn Façade case. On the other hand, there is strong intuitive pull toward the thought that John’s utterance “he knows it’s a barn” is true in the new Barn Façade case.[46]

---

the attributor chooses. But this case does not show that the reference class *relevant* for the truth of knowledge attributions varies with the attributor’s interest and goal. I cannot come up with any case in which the state or the country becomes the relevant reference class in attributing or denying knowledge to Henry who is seeing a barn in the Barn Façade County. The reference class comprising the Barn Façade County is a sub-class of the ones comprising the state or the country. It seems that this is the case in which narrower reference classes trump larger ones. Note that the relevant reference classes in the Normal and the New Barn Façade case are not such that one is a subset of the other, or one is narrower than the other.

[46] Cohen (1998) argues that standard contextualism cannot be applied to Gettier cases. The reason is simply that he does not find any context for Gettier cases in which ‘*S* knows *p*’ is true. As mentioned in §3. 2 of Chapter 5, however, DeRose reports that many do not have the intuition that Henry does not know. Though I share part of Cohen’s worry, variants of the Barn case are very different from other Gettier cases, and need not to be treated as parallel to them.

This is not a crude matter of intuition: the argument in §3. 3 of Chapter 5 shows that what the modified Barn Façade case does is to shift the relevant reference class (environmental type) from a normal one, say, {the presence of structures in countryside}, to {the presence of the barn-like structures in the Barn Façade County}. If so, it seems possible that the shift is induced in the opposite direction; the reference class can shift even when every fact about *S* is held fixed. The new Barn Façade case embodies at least one of such reverse procedures; a normal perceptual reference class is assigned to Henry's token environment. Thus, insofar as Henry knows in normal cases, John's knowledge attribution must be true in the new Barn Façade case.[47]

In both cases, John asks the question of the form 'what is it?' on the surface grammar. Are his questions contrastive in logical form? Perhaps, the question in the normal Barn Façade case may be analyzed as being of the form 'what is the structure Henry is seeing: a barn or a barn façade?'; and the question in the new Barn Façade case may be as being of the form 'what is the structure Henry is seeing: a barn or other structures in countryside?' If this analysis is right, the contrastivity of question can affect the difference in relevant reference class between the two cases. Even if contrastivity does not hold for questions, there is also an easier way to account for the difference. John's questions of the form 'what is it?' are driven by his explicit interests in Henry's discriminative abilities. The point of asking 'what is it?' in the normal Barn Façade case is to inquire as to whether

[47] Or alternatively, the new Barn Façade case may well be accounted for by the invariance of the relevant reference class, i.e., it is held fixed at the normal one. As DeRose reports, some do not have the intuition that Henry does not know in the original Barn Façade, presumably because the shift in reference class does not occur among them. The new Barn Façade may capture what is going on behind their intuition.

Henry is able to tell barns from barn facades; whereas, the point of asking the same form of question in the new Barn Façade case is to inquire as to whether Henry is able to tell barns from other structures in countryside. So, John's different interests in Henry's discriminative abilities are enough to set the relevant reference class differently between the two cases.

The contrastivity of question is not only unnecessary but also insufficient to account for the varying truth value of John's knowledge attributions (and denials). If Schaffer is right about (ii), **Question-Sensitivity** entails that knowledge attributions involve assigning an epistemic ability to answer the question. Such an ability is one to form or keep a belief as the answer to the question. Then, insofar as abilities are construed in the type or token reliabilist manner, the reference class must be specified for the ability in question. The variants of the Barn Façade case show (iii') that the interests of the attributor – in the broad sense that may involve the question as elucidated in (iii) – contribute to determining the reference class *relevant for the ability to answer the question*, not merely *the relevant question*. So **Question-Sensitivity** can be modified as follows.

> (Question-Sensitivity') Knowledge ascriptions certify the ability (type or token) to answer the question, relative to the environment (type or token).

The ability-environmental pair pertaining to a question may be determined by the question *per se*, or more generally, by the epistemic task the attributor assigns to *S* such that success or failure of the task assigned is important for the satisfaction of the attributor's interests in *S*'s epistemic score. For example, John, in the New Barn Façade case, assigns to Henry the task of discriminating barns from barn facades, and, in the normal Barn Façade case, assigns the task of discriminating barns from structures in countryside. John's interests in the two cases consist in seeing whether these tasks are achievable by Henry.

Question-Sensitivity', along with (i), (ii), and (iii'), constitutes an argument for task-sensitive contextualism. Just like questions, tasks may be implicit in context; the New Barn Façade case is parallel to the Modified Barn Façade case, but our denial of knowledge in the latter is not driven by any explicit task assignment. Be that as it may, the task is recoverable from context. Such factors, as salience, similarity, and normalcy, implicitly define the relevant task. For example, the description of the Modified Barn brings to salience the similarity between Henry's situation and those situations where Henry is seeing barn facades. This would have the same effect on the choice of the relevant reference class as John's explicit task assignment does. Of course, it requires further studies to elucidate what mechanism underlies implicit task assignment. This, however, is, at least in part, a psychological question, and as we have seen in the arguments in the preceding sections, psychological studies do support task-sensitive contextualism.

§2. 5. Arguments 5: The Best Explanation

Now, it is time to finalize the argument for task-sensitive contextualism that started in Chapters 5 and 6: it offers a united solution to the epistemologically difficult cases, such as some of the lottery cases, (variants of) the Barn Façade case, cases of easy knowledge, and cases of full-blooded Cartesian skepticism. How task-sensitive contextualism is applied to these cases is straightforward: Chapters 5 and 6 established how they are to be diagnosed in terms of SRC (shift-in-reference-class account) rather than SES (shift-in-standard account), and task-sensitive contextualism incorporates SRC.

That task-sensitive contextualism incorporates SRC constitutes an argument for itself. The fifth argument for task-sensitive contextualism is modeled on the master argument for standard contextualism, which, in the simplest form, goes as follows:

(1). One has the intuitions about a range *A* of cases where *S* knows *p*.

(2). One has the intuitions about a range *B* of cases where *S* does not know *p*, where each *B* case is paired with an *A* case.

(3). For each pair of *A* and *B* cases, *S*'s epistemic position with regard to *p* is held constant between them.

(4). Standard contextualism provides the best explanation of the variability of one's intuitions about *A* and *B*.

(5). Therefore, standard contextualism is true.

Similarly, the argument for task-sensitive contextualism is made in a similar fashion:

(1'). One has the intuitions about a range *C* of cases where *S* knows *p*.

(2'). One has the intuitions about a range *D* of cases where *S* does not know *p*, where each *D* case is paired with a *C* case.

(3'). Each pair of *C* and *D* cases involves the same *S*, the same token environment, and the same time.

(4'). Task-sensitive contextualism provides the best explanation of the variability of one's intuitions about *C* and *D*.

(5'). Therefore, task-sensitive contextualism is true.

'*D*' refers to the epistemologically challenging cases named above, and '*C*' their ordinary counterparts. Range *B* includes some of the lottery cases and the cases of high-standard local and global skepticism.[48] Note that '*B*' and '*D*' (and correspondingly '*A*' and '*C*') are not co-extensional, though their extensions partly overlap, because, as I argued in §2. 2 and 2. 3 of Chapter 5, the standard contextualist attempts to explain the lottery cases in general. But some lottery cases are those to which SRC rather than SES applies, and *C* contains such cases. Standard contextualism competes with its rivals, such as SSI, contrastivism, relativism, and WAM (warranted assertability maneuver), for the *best* explanation of our intuitions about *A* and *B*, and it may well be questioned whether standard contextualism is the best explanation.[49]

The explanation by task-sensitive contextualism has general theoretical virtues, such as consistency, simplicity, unification, and so on. In addition, the 'best' here is meant to be assessed, in particular, in terms of the theoretical virtue of 'intuitive adequacy,' i.e., how well an explanation saves the phenomena, i.e., our intuitive verdicts on cases. The

[48] More precisely, those lottery cases to which SES applies. By local skepticism, I mean the cases like the Bank case (DeRose, 1992) and the Airport case (Cohen, 1999).

[49] Fantl & McGrath (2009b, ch. 2) discuss in detail comparative adequacy of the explanations given by standard contextualism and SSI. Their discussions contain a more thorough reconstruction of the master argument for standard contextualism.

argument for task-sensitive contextualism is *prime facie* strong, since it does explain our intuitions well (and, there is no competing account at the moment).

§2. 6. Summary of Arguments

Arguments 1 through 5 jointly constitute the master argument for task-sensitive contextualism as the solution to the reference class problem in the theory of knowledge. Does this solution satisfy the constraints discussed in §1 of Chapter 4, **Non-Arbitrariness**, **Accordance**, and **Reductionism**? It, obviously, satisfies **Accordance**, since it amounts to intuitive adequacy. As regard to **Non-Arbitrariness**, task-sensitive contextualism specifies certain factors as determinants of the relevant reference class. Such factors are not arbitrarily picked up; they are entrenched in our interests concerning knowledge attributions and denials. It is a subtle question whether task-sensitive contextualism passes the test for **Reductionism**. After all, **Reductionism** is a constraint for substantive theories of knowledge with the partisan tenet. Being independent of substantive questions concerning knowledge, it is not clear at all why contextualism is susceptible to this constraint. Task-sensitive contextualism does not settle the question of how the concept of knowledge is analyzable non-circularly, and then it is open to both partisan and non-partisan approaches. This is an advantage of the contextualist solution, not a defect.

Compared to other alleged solutions, the contextualist solution is surely in a better position, not only because it is better off in terms of **Accordance** or intuitive adequacy, but also because it can take advantage of the virtues of other solutions. As mentioned through this thesis, there are serious problems in other places that are structurally similar to the reference class problem. In particular, arguments 1 and 2 are based on the contextualist solutions to some of these problems. Thus, it is natural to expect a structurally similar strategy to work against the reference class problem in epistemology, even though they are not substantively the same in the sense delineated in §1 of Chapter 4. Arguments 3, 4, and 5 give epistemological grounds for task-sensitive contextualism. In addition, argument 4 shows that task-sensitive contextualism incorporates a function or role of knowledge attribution in our epistemic practice. Even though none of the arguments is knock-down, the fact that they all indicate the same direction makes a strong case for the claim that task-sensitive contextualism is the best candidate for a solution to the reference class problem in the theory of knowledge.

This completes my argument for task-sensitive contextualism. In what follows, I will outline its ramifications for both the substantive and the formal theory of knowledge.

§3. Ramifications of Task-Sensitive Contextualism

§3. 1. Ramifications for the Substantive Theory of Knowledge

Task-sensitive contextualism, just like standard contextualism, can be combined with any substantive theory of knowledge. Its application to externalist theories is straightforward, as it is easy to see by now: on type reliabilism, the reliability of *S*'s token process is relativized to a pair of process type and environment type, which are determined contextually; similarly, on token reliabilism, the degree to which *S*'s belief is sensitive or safe is relativized to a pair of ordering source and modal base, which are contextually set.

Thus, any externalist theory entails that the strength of *S*'s epistemic position is, in part, a function of contextual factors.

By contrast, it may not be obvious that internalism has the same consequence. When **Reliability Expert** is in place, the degree of internalist justification is relativized to an expert assignment by the reliability expert function *RELI*. A *RELI* is a function that takes a proposition about type or token reliability as its argument and delivers a certain degree of internalist justification as its value. The domain of any *RELI* must be reliability data in *S*'s evidence set, i.e., the set of known or justified propositions that a token process or method has certain type or token reliability relative to a certain reference class. Let us call reference classes involved in such propositions 'propositional reference classes.' The range of propositional reference classes are more restricted than candidate reference classes for externalist theories.

As far as internalist evidentialism is concerned, it is useful to draw a distinction between a propositional reference class and a doxastic reference class. The degree of propositional justification for *p* amounts to epistemic probability of some sort, which is a measure of personal rationality, i.e., how rational *p* is for *S* to believe. The relevant reference class for propositional justification, then, is rationally constrained. It is here where the epistemic probability theorist ties the choice of the relevant reference class with the requirement of total evidence. However, as Comesaña (2006) points out, any full-blown theory of internalist justification must involve an account of both propositional and doxastic justification. A propositionally justified proposition fails to yield doxastic justification for *p* when *S* does not believe *p* on the basis of it. Insofar as the basing relation is causal, then, the doxastic reference class is that which is involved in the base of *S*'s belief *p*.

Given this distinction, there are many different kinds of cases where *S*'s belief fails to be an instance of knowledge *via* **Reliability Expert**. Such cases are summarized as follows:

(i) The doxastic reference class is not a propositional reference class.

(ii) The propositional-doxastic reference class is irrelevant.

(ii-a) irrelevant: the proposition involving the propositional reference class is false.

(ii-a-a) the proposition is false about the reference class.

(ii-a-b) the proposition is false about the reliability.

(ii-b) irrelevant: the proposition involving the propositional reference class is externalistically unconstrained.

Cases of (i) obtain when *S*'s belief is based on propositionally unjustified reliability data, whether it is true or false. Cases of (ii-a) are ones where *S*'s evidence set includes justified but false reliability data and *S* bases her belief *p* on it. Such cases are further divided into two sub-cases: in (ii-a-a), reliability data is false because it tells wrong type or token reliability; and in (ii-a-b), reliability data is false because it tells wrong type-token relationship, i.e., it assigns a process token to a process type to which it *does not* belong, or it contains a *false* counterfactual for safety or sensitivity assessment.

It seems that the Optimist case, mentioned in §3. 5 of Chapter 3, belongs to (ii-a-a). Paula types the process token by which she forms an optimistic belief about the future as {belief-forming from noon to 5:00}. She has a daily habit to do arithmetic calculation from noon to 4: 55, and then to think optimistically about her future. So the type {belief-forming from noon to 5:00} is indeed disjunctive, being equivalent to the type {arithmetic calculation from noon to 4: 55 ∨ optimistic belief-forming about the future after 4:55 till 5:00}. As far as belief-forming process tokens are concerned, they must be typed in terms of causality rather than temporality. Even if Paula is justified in believing that this disjunctive type is reliable, we do not judge that her beliefs about the future are propositionally justified. The reason seems to be that {arithmetic calculation from noon to 4: 55} and {optimistic belief-forming about the future after 4:55 till 5:00} are causally distinct process types.[50] I take it that this amounts to the type-token relationship not holding in the way Paula believes it to be. This case suggests that *RELI* only applies if reliability data involved in its application contains true information on the type-token relationship. Then, insofar as applications of *RELI* are relevant for *S*'s satisfiability of **Meta-Justification**, cases of (ii-a-a) (and perhaps, cases of (ii-a-b) as well) are those where *S*'s belief *p* is not propositionally justified. For, if *S* fails to satisfy **Meta-Justification** with regard to *p*, *S* is not justified in believing *p*.

[50] As I argued in §2 of Chapter 4, not every epistemically relevant reference class is causal. The point here is that process types, distinguished from environmental types, need to be at least causal. The type Paula appeals to does not satisfy Chalmers' causal complexity requirement for complex systems involving sub-computational functions – the same type of requirement for FSAs are mentioned in footnote 20.

Similar considerations are applied to the original lottery case discussed in Chapter 5: *S*'s degree of internalist justification for *her ticket will lose* might be nearly maximal if *S* bases her belief on the belief that it is a member of {the big lottery tickets *S* has ever bought}. A ticket's winning is a causal matter, and determined by how the winning tickets are drawn by whatever selection system. The problem of this reference class, then, is that it contains tokens which are subject to causally independent selection systems. As far as the objective probability of *S*'s winning a particular lottery is concerned, it seems that the choice of the relevant reference class must be causally constrained.[51]

In addition, causality is the key for *S*'s justification in cases of the problem of easy knowledge including cases of full-blooded Cartesian skepticism. As I argued in Chapter 6, in such cases, it is crucial for *S*'s justification for reliability data, *a process type is reliable to a certain degree*, that *S* exercises multiple causal-functionally independent processes. Relative to the environment type made salient by those cases, however, this condition is not met. Insofar as the third-personal knowledge denials in the cases are taken seriously, the only way to explain them is appeal to *S*'s failure to satisfy **Meta-Justification**. If *S* fails to satisfy **Meta-Justification**, it is a case of (i). Again, this suggests that cases of (i) are those

[51] So this class is irrelevant. But I am reluctant to claim that this case is categorized as (ii-a-a), since my point presupposes that the membership relation relevant for the probability of a ticket's winning is causal, not merely statistic. Some theory of probability would deny this. I do not commit myself to settle the question of what theory of probability is appropriate for what purpose. This case, then, might fall under (ii-b). Either way, I have no general account of the choice of reference class for direct inference. For this reason, I cannot determine whether *S* is justified in believing that *S*'s ticket will lose, in the original lottery case.

where *S*'s belief *p* is not propositionally justified.[52] If my diagnosis of cases of the problem of easy knowledge is right, the selection of the relevant reference class on the attributor's side affects whether *S* is justified *via* **Meta-Justification**.

The lottery cases (besides the original lottery case) and the Gettier cases to which SRC is applied are instances of (ii-b), though that is difficult to see because they do not specify the doxastic reference class. For example, in the Barn Façade case, Henry is internalistically justified in believing that he is seeing a barn, while what meta-justification he has for this belief is not specified. What class is the doxastic reference class is a causal matter, and hence, even if unspecified, it may be objectively determined (insofar as *S* believes *p* on the basis of some reliability data). The point of the case rather consists in the fact that there is no way for *S* to rationally access to the relevant reference class; Henry does not have information that he is in the Barn Façade County, and he cannot rationally select the relevant reference class, specified in my diagnosis of the case in §3. 3 of Chapter 5. Hence, whatever doxastic reference class it is, it does not constitute knowledge, even though it does internalist justification, since the relevant referent class affects whether *S* satisfies externalist conditions for knowledge. Lottery cases are amenable to the same type of explanation.[53]

---

[52] This result is peculiar, since it seems to hold even if *S* is justified in believing that she exercises multiple causal-functionally independent sources. Perhaps, the reason is that causal factors trump other factors, and put a strong constrain on the degree of internalist justification. This seems to be because reliability data is about a causal process or method.

[53] Remember that Hájek (2003, 2007b) introduces expert functions as *objective* constrains on epistemic probabilities, since, he argues, the connection with the objective world is severed if they only need to be

Thus, in cases of (i) and (ii-a-a), contexts affect the strength of *S*'s epistemic position, at least in the sense that they are crucial for whether *S* is justified or not in believing *p*. Moreover, if my analysis of the cases of (ii-b) is on the right track, the relevant reference class regarding externalist conditions for knowledge is also that which is subject to contextual variance. Thus, even in cases where the degree of *S*'s internalist justification for *p* is not a function of a contextually determined reference class, it still determines whether *S* knows *p* or not.

§3. 2. Ramifications for the Formal Theory of Knowledge

As a result of its consequence for substantive theories of knowledge, task-sensitive contextualism has a significant ramification for the formal theory of knowledge as well. Any substantive theory of knowledge aims to identify one or more truth-conductive factors *F*s as the determinants of the strength of *S*'s epistemic position with regard to *p* in any given case. This point is articulated by one of the axioms of the formal theory, **Supervenience**:

> (Supervenience) The strength of *S*'s epistemic position with regard to *p* supervenes on epistemic factors *F*s *S* instantiates regarding *p*.

In view of task-sensitive contextualism, this thesis allows for two different readings:

---

probabilistically coherent. This objective sense of rationality is missing in cases of (ii-b), and it seems to be best captured by some externalist theory.

> (Supervenience A) If *S* and *S*' both instantiate epistemic factors *F*s with regard to *p*, then *S* has epistemic position *EP* with respect to *p* iff *S*' does.
>
> (Supervenience B) If *S* and *S*' both instantiate epistemic factors *F*s with regard to *p*, then any utterance of '*S* has epistemic position *EP* with regard to *p*' is true iff any utterance of '*S*' has epistemic position *EP* with respect to *p*' is true.

Task-sensitive contextualism entails that **Supervenience B** is false, whereas it is consistent with **Supervenience A**. Notice that reference classes are epistemic factors, since they are responsible for strengths of epistemic position. Moreover, any *S* instantiates many different reference classes at the same time.

The distinction between *S*'s epistemic condition and epistemic position, drawn by Heller (1995b, pp. 513-4, 1999, p. 127), best serves the purpose of capturing the difference between **Supervenience A** and **B**. Heller's notion of *S*'s epistemic condition with respect to *p* can be defined as the value of a function *EC*: for all proposition *p*, environment token *e*, and time *t*, *EC*(*p*, *e*, *t*) for *S* is the sum of *EP*(*p*, given r_1), …, and *EP*(*p*, given r_n), where each of 'r_1', …, 'r_n' denotes a reference class among the *n* reference classes *S* instantiates regarding *p* in *e* at *t*. There are many reference classes for any set of *S*, *p*, *e*, *t*. Of course, what factors count as reference classes are dependent on which substantive theory of knowledge is advocated; to take an example of type reliabilism, *S*'s epistemic position regarding *p* is identified with the value of type reliability relative to a pair of process and environment types. There are many pairs of process and environment types exemplified by a process token at a time, each of which results in a value of type reliability. The epistemic condition is the sum of all the values of such type reliability. The concept of epistemic condition reflects the fact that the totality of *S*'s epistemic factors regarding *p* – what process, ability, disposition to believe, etc. *S* has with respect to *p* – does not change across contexts, times and environments, unless there is a factual difference about *S*.[54]

With the concept of epistemic condition, (3) in Argument 5 can be re-formulated: (3*) *S*'s epistemic condition with regard to *p* is held constant between *C* and *D*. This is a specific form of the following thesis:

> (Supervenience') The strength of *S*'s epistemic condition with regard to *p* supervenes on the totality of *S*'s epistemic factors regarding *p*.

As Heller submits, it is a fully objective matter how strong *S*'s epistemic condition is, since it supervenes on facts about *S*.[55]

---

[54] Wunderlich's (2003) solution to the reference class problem is to use the epistemic condition rather than the epistemic position, as the object of epistemic appraisal. Wunderlich points out that if *r* is infinite, there is no way to assign a single value to *EC*. Then, he offers several ways to restrict the range of *r*. As a consequence, Wunderlich submits, the use of *EC* only permits comparative epistemic appraisals. The epistemic condition, as construed along lines with type reliabilism, is essentially the same as 'transglobal reliability,' put forth by Henderson & Horgan (2006, 2007) and Henderson, Horgan, & Potrč (2007). They argue that transglobal reliability is required to capture our intuitions about a family of cases, some of which are variants of the Barn Façade case. Task-sensitive contextualism can deal with these cases more naturally.

[55] Greco (2008, 2010, ch. 7), from a different perspective, addresses two worries about contextualism and SSI: they make the truth condition of knowledge attributions subjective, and they deprive epistemology of its proper subject matter.

Supervenience' does not entail **Supervenience B**. For type and token reliabilism, if *S* and *S*' (or more properly, their process or method token) instantiate the same epistemic factors relative to the same set of reference classes, they are both in the same epistemic condition. This does not determine what *EP*(*p*) is truly attributable to *S* or *S*'. Task-sensitive contextualism is the position that any sentence of the form '*S* has epistemic position *EP*' has an implicit slot to be saturated by a reference class, and what class saturates it varies from context to context. Then, the sentence of '*S* has epistemic position *EP*' may be true if uttered in one context, but not if uttered in another, when the contextual factors of the two contexts assign different reference classes (unless, of course, they accidentally result in the same value of type or token reliability).

This result is generalizable: no full-blown substantive theory can satisfy **Supervenience B**. The moral to be drawn from Gettier cases is that any full-blooded substantive theory must be, in part, externalist, i.e., knowledge has at least one externalist epistemic factor as its component. As a consequence, even if some internalist theory captures one epistemic factor, it must incorporate some externalist factor, in order to fully capture the epistemic necessary conditions for knowledge or conform to **Epistemic Position**. Then, task-sensitive contextualism entails that the contribution any putative externalist factor makes to the overall strength of the epistemic position an attributor can truly attribute to *S* is a function of that externalist factor and attributor factors.[56] Insofar as

[56] The strength of internalist justification is determined by one of the propositional reference classes. Which one is relevant may depend on the subject's interests, just as the relevant reference class in the sense of (ii-b) depends on the attributor's interests. One problem with this view about internalist justification is that it violates the requirement of total evidence, since, on this view, the degree of internalist justification depends

pluralism about epistemic factor is true, the truth of '*S* has epistemic position *EP* with regard to *p*' does not supervene on the totality of factors about *S*, let alone her epistemic factors. Therefore, the denial of **Supervenience B** follows. In addition, as pointed out in the last section, even monist internalist evidentialism leads to the same result, insofar as it accepts **Reliability Expert** and **Meta-Justification**.

On the other hand, task-sensitive contextualism is consistent with **Supervenience A**. **Supervenience A** is true iff *S* and *S*' always satisfy the relation denoted by 'have epistemic position *EP*,' if *S* and *S*' both instantiate epistemic factors *F*s with regard to *p*. What relation 'have epistemic position *EP*' denotes depends on context, since, if task-sensitive contextualism is true, the relation always involves implicit reference to a reference class to be contextually determined. However, once the denotation of 'have epistemic position *EP*' is held fixed, whether *S* satisfies it is not a context-sensitive matter. This indeed follows from **Supervenience'**; that *S* and *S*' are in the same epistemic condition entails that they do not differ in epistemic position relative to each single reference class they instantiate (unless, of course, different reference classes they instantiate accidentally result in the same

---

on a subset of *S*'s total evidence to be picked out by *S*'s interests. A way to evade this problem is to retain the requirement of total evidence by stipulating that the subject's interests affect what beliefs or mental states are in *S*'s total evidence. Neta (2007b, 2008) and Williams (1991) independently propose this type of view as a version of subject-centered 'contextualism' – though it is not about context of use. Neta (2008), in making an argument for his subject-centered contextualism, grants that it is based on the contention that the total evidence requirement is true. One problem with this version of contextualism is that it fails to capture third-person skeptical intuitions.

epistemic condition). Task-sensitive contextualism is consistent with **Supervenience'**, and hence, is with **Supervenience A**.

The strength of epistemic condition is always an objective matter. But the same does not hold for epistemic positions: task-sensitive contextualism entails that the value of *EP*(*p*) is in part a subjective matter, in that non-epistemic attributor factors inherent in context are determinants of the strength of *S*'s epistemic position; what context does is to pick up a particular *r*, on which the value of *EP* is partly dependent. This explains why task-sensitive contextualism is not a solution to the metaphysical reference class problem in a strong sense that its solution must be consistent with **Supervenience B**. Indeed, if my arguments for task-sensitive contextualism are successful, no solution is possible for the metaphysical reference class problem in this sense. But there is still a sense in which it is not irredeemably subjective how strong an epistemic position *S* is in: **Supervenience B** is maintained.[57] Against the epistemological reference class problem, task-sensitive contextualism gives an explanation of how context imposes constraints on what reference class is appropriate for the purpose of a certain knowledge attribution or denial.

Besides the denial of **Supervenience B**, task-sensitive contextualism deprives neither the substantive nor the formal theory of knowledge of anything important: the epistemologist still needs to address significant epistemological questions, such as what the relevant kind of epistemic factor is, why knowledge is valuable, and what role or function

---

[57] The supervenience theses here are all to be construed in terms of local supervenience. Task-sensitive contextualism entails that neither local nor global supervenience holds for strengths of epistemic position insofar as **Supervenience** reads as **Supervenience B**.

knowledge has in our epistemic practice. These problems still remain and deserve serious inquiry.

Bibliography

Achinstein, Peter. (2001). *The Book of Evidence*. New York, NY: Oxford University Press.

Adler, Jonathan E. (2005). "Reliabilist Justification (or Knowledge) as a Good Truth-Ratio." *Pacific Philosophical Quarterly* 86(4): 445-458.

Adler, Jonathan E & Levin, Michael. (2002). "Is the Generality Problem Too General?" *Philosophy and Phenomenological Research* 65(1): 87-97.

Allen, Ronald J. & Pardo, Michael S. (2007). "The Problematic Value of Mathematical Models of Evidence." *The Journal of Legal Studies* 36(1): 107-140.

Alston, William P. (1980). "Level-Confusions in Epistemology." *Midwest Studies in Philosophy* 5: 135-150. Reprinted in Alston (1989): 153-171.

-----. (1985). "Concepts of Epistemic Justification." *The Monist* 62(2): 57-89. Reprinted in Alston (1989): 81-114.

-----. (1986). "Epistemic Circularity." *Philosophy and Phenomenological Research* 47(1): 1-30. Reprinted in Alston (1989): 319-349.

-----. (1988). "An Internalist Externalism." *Synthese* 74(3): 265-283. Reprinted in Alston (1989): 227-245.

-----. (1989). *Epistemic Justification: Essays in the Theory of Knowledge*. Ithaca, NY: Cornell University Press.

-----. (1995). "How to Think about Reliability." *Philosophical Topics* 23(1): 1-29.

-----. (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca, NY: Cornell University Press.

Armstrong, David M. (1973). *Belief, Truth and Knowledge*. Cambridge: Cambridge University Press.

Austin, John L. (1946). “Other Minds.” *Proceedings of the Aristotelian Society* Suppl. 20: 149-187.

Ayer, Alfred J. (1963). “Two Notes on Probability.” In A. J. Ayer, *The Concept of a Person and Other Essays*. London: MacMillan: 188-208.

Baergen, Ralph. (1995). *Contemporary Epistemology*. Fort Worth, TX: Hartcourt Brace.

Baumann, Peter. (2005). “Varieties of Contextualism: Standards and Descriptions.” *Grazer Philosophische Studien* 69: 229-246.

Becker, Kelly. (2008). “Epistemic Luck and the Generality Problem.” *Philosophical Studies* 139(3): 353-366.

-----. (2009). “Margins for Error and Sensitivity: What Nozick Might Have Said.” *Acta Analytica* 24(1): 17-31.

Beebe, James R. (2004). “The Generality Problem, Statistical Relevance and the Tri-Level Hypothesis.” *Noûs* 38(1): 177-195.

Beebee, Helen & Papineau, David. (1997). “Probability as a Guide to Life.” *The Journal of Philosophy* 94(5): 217-243.

Bergmann, Michael. (2004). “Epistemic Circularity: Malignant and Benign.” *Philosophy and Phenomenological Research* 69(3): 709-727.

Bishop, Mark. (2002). “Counterfactuals Cannot Count: A Rejoinder to David Chalmers.” *Consciousness and Cognition* 11(4): 642-652.

Bishop, Michael A. (forthcoming) “Why the Generality Problem is Everybody’s Problem.” *Philosophical Studies*.

Black, Tim. (2008a). “Defending a Sensitive Neo-Moorean Invariantism.” In Pritchard & Hendricks (2008): 8-27.

-----. (2008b). “Solving the Problem of Easy Knowledge.” *The Philosophical Quarterly* 58(233): 597-617.

Black, Tim & Murphy, Peter. (2007). “In Defense of Sensitivity.” *Synthese* 154(1): 53-71.

Blome-Tillmann, Michael. (2007). “The Folly of Trying to Define Knowledge.” *Analysis* 67(295): 214-219.

-----. (2008). “The Indexicality of ‘Knowledge’.” *Philosophical Studies* 138(1): 29-53.

-----. (2009). “Knowledge and Presuppositions.” *Mind* 118(470): 241-294.

BonJour, Laurence. (1985). *The Structure of Empirical Knowledge*. Cambridge, MA: Harvard University Press.

-----. (1997). *In Defense of Pure Reason: A Rationalist Account of A Priori Justification*. Cambridge: Cambridge University Press.

Brandom, Robert B. (1994). *Making it Explicit: Reasoning, Representing, and Discursive Commitment*. Cambridge, MA: Harvard University Press.

Brogaard, Berit. (2006). “Can Virtue Reliabilism Explain the Value of Knowledge?” *Canadian Journal of Philosophy* 36(3): 335-354.

Brown, Jessica. (2004). “Wright on Transmission Failure.” *Analysis* 64(281): 57-67.

Brueckner, Anthony. (1994a). “The Shifting Content of Knowledge Attributions.” *Philosophy and Phenomenological Research* 54(1): 123-126.

-----. (1994b). “The Structure of the Skeptical Argument.” *Philosophy and Phenomenological Research* 54(4): 827-835.

Burge, Tyler. (1986). “Individualism and Psychology.” *The Philosophical Review* 95(1): 3-45.

Butler, Keith. (1998). “Content, Computation, and Individuation.” *Synthese* 114(2): 277-292.

Carnap, Rudolf. (1950). *Logical Foundations of Probability*. Chicago, IL: University of Chicago Press.

-----. (1962). *Logical Foundations of Probability*, 2nd Ed. Chicago, IL: University of Chicago Press.

Carrier, Leonard S. (1993). “How to Define a Nonskeptical Fallibilism.” *Philosophia* 22(3): 361-372.

Castañeda, Hector-Neri. (1980). “The Theory of Questions, Epistemic Powers, and the Indexical Theory of Knowledge.” *Midwest Studies in Philosophy* 5: 193-238.

Chalmers, David J. (1994). “On Implementing a Computation.” *Minds and Machines* 4(4): 391-402.

-----. (1996). “Does a Rock Implement Every Finite-State Automaton?” *Synthese* 108(3): 309-333.

Chisholm, Roderick M. (1942). “The Problem of the Speckled Hen.” *Mind* 51(204): 368-373.

-----. (1977). *Theory of Knowledge*, 2nd Ed. Englewood Cliffs, NJ: Prentice Hall.

Coffman, E. J. (2008). “Warrant Without Truth?” *Synthese* 162(2): 173-194.

Cohen, Stewart. (1988). “How to be a Fallibilist.” *Philosophical Perspectives* 2: 91-123.

-----. (1998). “Contextualist Solutions to Epistemological Problems: Scepticism, Gettier, and the Lottery.” *Australasian Journal of Philosophy* 76(2): 289.

-----. (1999). “Contextualism, Skepticism, and the Structure of Reasons.” *Philosophical Perspectives* 13: 57.

-----. (2000). “Contextualism and Skepticism.” *Philosophical Issues* 10: 94-107.

-----. (2002). “Basic Knowledge and the Problem of Easy Knowledge.” *Philosophy and Phenomenological Research* 65(2): 309-329.

-----. (2004). “Reply to Baumann.” *Erkenntnis* 61(2-3): 429-433.

-----. (2005). “Why Basic Knowledge is Easy Knowledge.” *Philosophy and Phenomenological Research* 70(2): 417-430.

Colyvan, Mark., Regan, Helen M & Ferson, Scott. (2001). “Is It a Crime to Belong to a Reference Class.” *Journal of Political Philosophy* 9(2): 168-181.

Comesaña, Juan. (2006). “A Well-Founded Solution to the Generality Problem.” *Philosophical Studies* 129(1): 27-47.

Comesaña, Juan, & Holly Kantin. (2010). “Is Evidence Knowledge?” *Philosophy and Phenomenological Research* 80(2): 447-454.

Conant, James. (2004). “Varieties of Skepticism.” In D. McManus (ed.), *Wittgenstein and Skepticism*. London: Routledge: 97-136.

Conee, Earl & Feldman, Richard. (1998). “The Generality Problem for Reliabilism.” *Philosophical Studies* 89(1): 1-29. Reprinted in Conee & Feldman (2004): 135-159.

-----. (2004). *Evidentialism: Essays in Epistemology*. Oxford: Oxford University Press.

-----. (2008). “Evidence.” In Q. Smith (ed.), *Epistemology: New Essays*. Oxford: Oxford University Press, 83-104.

Cross, Troy. (2007). “Comments on Vogel.” *Philosophical Studies* 134(1): 89-98.

Cruz, Joe & Pollock, John L. (2004). “The Chimerical Appeal of Epistemic Externalism.” In R. Schantz (ed.), *The Externalist Challenge*. Berlin, Germany: de Gruyter, 125-141.

David, Marian. (2001). “Truth as the Epistemic Goal.” In M. Steup (ed.), *Knowledge, Truth, and Duty: Essays on Epistemic Justification, Responsibility, and Virtue*. Oxford: Oxford University Press: 151-169.

-----. (2005). “Truth as the Primary Epistemic Goal: A Working Hypothesis.” In Sosa & Steup (2005): 296-312.

Davidson, Donald. (1983). “A Coherence Theory of Truth and Knowledge.” In D. Henrich (ed.), *Kant oder Hegel?* Stuttgart, Germany: Klett-Cotta: 423-438.

Davies, Martin. (1992). “Perceptual Content and Local Supervenience.” *Proceedings of the Aristotelian Society* 92: 21-45.

DePaul, Michael R & Grimm, Stephen R. (2007). “Review Essay on Jonathan Kvanvig’s *The Value of Knowledge and the Pursuit of Understanding*.” *Philosophy and Phenomenological Research* 74(2): 498-514.

DePaul, Michael R & Zagzebski, Linda T. (eds.). (2003), *Intellectual Virtue: Perspectives from Ethics and Epistemology*. Oxford: Oxford University Press.

DeRose, Keith. (1992). “Contextualism and Knowledge Attributions.” *Philosophy and Phenomenological Research* 52(4): 913-929.

-----. (1995). “Solving the Skeptical Problem.” *The Philosophical Review* 104(1): 1-52.

-----. (1996a). “Knowledge, Assertion and Lotteries.” *Australasian Journal of Philosophy* 74(4): 568.

-----. (1996b). “Relevant Alternatives and the Content of Knowledge Attributions.” *Philosophy and Phenomenological Research* 56(1): 193-197.

-----. (1999) “Contextualism: An Explanation and Defense.” In Greco & Sosa (1999): 187-205.

-----. (2000). “How Can We Know that We’re not Brains in Vats?” *The Southern Journal of Philosophy* Suppl. 38: 121-149.

-----. (2004a). “Safety, Sensitivity, and Sceptical Hypotheses.” In Greco (2004b): 22-41.

-----. (2004b). “The Problem with Subject-Sensitive Invariantism.” *Philosophy and Phenomenological Research* 68(2): 346-350.

-----. (2008). “Gradable Adjectives: A Defence of Pluralism.” *Australasian Journal of Philosophy* 86(1): 141-160.

-----. (2009). *The Case for Contextualism: Knowledge, Skepticism, and Context*, Vol. 1. Oxford: Oxford University Press.

Dretske, Fred I. (1970). “Epistemic Operators.” *The Journal of Philosophy* 67(24): 1007-1023.

-----. (1971). “Conclusive Reasons.” *Australasian Journal of Philosophy* 49(1): 1-22.

-----. (1981a). “The Pragmatic Dimension of Knowledge.” *Philosophical Studies* 40(3): 363-378.

-----. (1981b). *Knowledge and the Flow of Information*. Cambridge, MA: MIT Press.

-----. (2005a). “Reply to Hawthorne.” In Sosa & Steup (2005): 43-46.

-----. (2005b). “The Case against Closure.” In Sosa & Steup (2005): 13-26.

Dudman, Victor H. (1992). “Probability and Assertion.” *Analysis* 52(4): 204-211.

Eagle, Antony. (2004). “Twenty-One Arguments against Propensity Analyses of Probability.” *Erkenntnis* 60(3): 371-416.

Egan, Frances. (1991). “Must Psychology Be Individualistic?” *The Philosophical Review* 100(2): 179-203.

-----. (1992). “Individualism, Computation, and Perceptual Content.” *Mind* 101(403): 443-459.

-----. (1994). “Individualism and Vision Theory.” *Analysis* 54(4): 258-264.

-----. (1995). “Computation and Content.” *The Philosophical Review* 104(2): 181-203.

Fantl, Jeremy & McGrath, Matthew. (2002). “Evidence, Pragmatics, and Justification.” *The Philosophical Review* 111(1): 67-94.

-----. (2007). “On Pragmatic Encroachment in Epistemology.” *Philosophy and Phenomenological Research* 75(3): 558-589.

-----. (2009a). “Critical Study of John Hawthorne’s *Knowledge and Lotteries* and Jason Stanley’s *Knowledge and Practical Interests*.” *Noûs* 43(1): 178-192.

-----. (2009b). *Knowledge in an Uncertain World*. Oxford: Oxford University Press.

Fara, Delia Graff. (2000). “Shifting Sands: An Interest-Relative Theory of Vagueness.” *Philosophical Topics* 28(1): 45-81.

Feldman, Richard. (1985). “Reliability and Justification.” *The Monist* 68: 159-174.

-----. (1993). “Review: Proper Functionalism.” *Noûs* 27(1): 34-50.

-----. (2005). “Review of William Alston’s *Beyond “Justification”: Dimensions of Epistemic Evaluation*.” August 10, 2010. <http://ndpr.nd.edu/review.cfm?id=3461>.

Feldman, Richard & Conee, Earl. (1985). “Evidentialism.” *Philosophical Studies* 48(1): 15-34. Reprinted in Conee & Feldman (2004): 83-101.

-----. (2001). “Internalism Defended.” *American Philosophical Quarterly* 38(1): 1-18. Reprinted in Conee & Feldman (2004): 53-80.

-----. (2002). “Typing Problems.” *Philosophy and Phenomenological Research* 65(1): 98-105.

-----. (2004). “Afterword to “Evidentialism”.” Reprinted in Conee & Feldman (2004): 101-7.

Fetzer, James H. (1972). “Statistical Explanations.” *PSA: Proceedings of the Biennial Meeting of the Philosophy of Science Association*: 337-347.

-----. (1977). “Reichenbach, Reference Classes, and Single Case ‘Probabilities’.” *Synthese* 34(2): 185-217.

-----. (ed.) (2001). *The Philosophy of Carl G. Hempel*: *Studies in Science, Explanation, and Rationality*. Oxford: Oxford University Press.

Firth, Roderick. (1978). “Are Epistemic Concepts Reducible to Ethical Concepts?.” In A. Goldman & J. Kim (eds.), *Values and Morals*. Dordrecht, Netherlands: Reidel: 215–29.

Fogelin, Robert J. (1994). *Pyrrhonian Reflections on Knowledge and Justification*. Oxford: Oxford University Press.

Fumerton, Richard A. (1995). *Metaepistemology and Skepticism*. Lanham, MD: Rowman & Littlefield.

-----. (2004). “Epistemic Probability.” *Philosophical Issues* 14: 149-164.

Gaifman, Haim. (1988). “A Theory of Higher Order Probabilities.” In B. Skyrms & W. Harper (eds.), *Causation, Chance, and Credence*. Dordrecht, Netherlands: Kluwer Academic: 275-292.

Gendler, Tamar S & Hawthorne, John. (2005). “The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies.” *Philosophical Studies* 124(3): 331-352.

Gettier, Edmund L. (1963). “Is Justified True Belief Knowledge?” *Analysis* 23(6): 121-123.

Giere, Ronald N. (1973). “Objective Single-Case Probabilities and the Foundations of Statistics.” In P. Suppes., L. Henkin., G. Moisil & A. Joja (eds.), *Logic, Methodology, and the Philosophy of Science*, Vol.4. Amsterdam, Netherlands: American Elsevier: 467-483.

Gillies, Donald. (2000). *Philosophical Theories of Probability*. London: Routledge Press.

Goldman, Alan. (1988). *Empirical Knolwedge*. Berkeley, CA: University of California Press.

Goldman, Alvin I. (1976). “Discrimination and Perceptual Knowledge.” *The Journal of Philosophy* 73(20): 771-791. Reprinted in Goldman (1992b): 85-103.

-----. (1979). “What is Justified Belief?” In G. Pappas (ed.), *Justification and Knowledge*. Dordrecht, Netherlands: Reidel: 1-23. Reprinted in Goldman (1992b): 105-126.

-----. (1983). “Review of Robert Nozick’s *Philosophical Explanations*.” *The Philosophical Review* 92(1): 81-88.

-----. (1986). *Epistemology and Cognition*. Cambridge, MA: Harvard University Press.

-----. (1988). “Strong and Weak Justification.” *Philosophical Perspectives* 2: 51-69. Reprinted in Goldman (1992b): 127-141.

-----. (1992a). “Epistemic Folkways and Scientific Epistemology.” In Goldman (1992b): 155-175.

-----. (1992b) *Liasons: Philosophy Meets the Cognitive and Social Sciences*, Cambridge, MA: MIT Press.

-----. (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.

-----. (2001). “Social Routes to Belief and Knowledge.” *The Monist* 84(3): 346-368.

-----. (2002). “What is Social Epistemology? A Smorgasbord of Projects.” In A. Goldman, *Pathways to Knowledge: Private and Public*. Oxford: Oxford University Press: 164–181.

-----. (2008). “*Reliabilism*.” *Stanford Encyclopedia of Philosophy*. August 10, 2010. <http://plato.stanford.edu/entries/reliabilism/>.

Goldstone, Robert L, Medin, Douglas L & Halberstadt, Jamin. (1997). “Similarity in Context.” *Memory & Cognition* 25(2): 237-255.

Good, Irving J. (1967). “On the Principle of Total Evidence.” *The British Journal for the Philosophy of Science* 17(4): 319-321.

Goodman, Nelson. (1955). *Fact, Fiction, & Forecast*. Cambridge, MA: Harvard University Press.

Greco, John. (1999). “Agent Reliabilism.” *Philosophical Perspectives* 13: 273-296.

-----. (2000). *Putting Skeptics in Their Place: The Nature of Skeptical Arguments and Their Role in Philosophical Inquiry*. New York, NY: Cambridge University Press.

-----. (2002). “Virtues in Epistemology.” In Moser (2002): 287-315.

-----. (2003). “Knowledge as Credit for True Belief.” In DePaul & Zagzebski (2003): 111-134.

-----. (2004a). “A Different Sort of Contextualism.” *Erkenntnis* 61(2-3): 383-400.

-----. (ed.) (2004b). *Ernest Sosa and His Critics*. Oxford: Blackwell.

-----. (2004c). “How to Preserve Your Virtue while Losing Your Perspective.” In Greco (2004b): 96-105.

-----. (2007a). “The Nature of Ability and the Purpose of Knowledge.” *Philosophical Issues* 17: 57-69.

-----. (2007b). “Worries about Pritchard’s Safety.” *Synthese* 158(3): 299-302.

-----. (2008). “What’s Wrong with Contextualism?” *The Philosophical Quarterly* 58(232): 416-436.

-----. (2009). “Knowledge and Success from Ability.” *Philosophical Studies* 142(1): 17-26.

-----. (2010). *Achieving Knowledge: A Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.

Greco, John & Sosa, Ernest. (eds.) (1999). *The Blackwell Guide to Epistemology*. Oxford: Blackwell.

Gundersen, Lars. (2010). “Tracking, Epistemic Dispositions and the Conditional Analysis.” *Erkenntnis* 72(3): 353-364.

Hacking, Ian. (1965). *Logic of Statistical Inference*. Cambridge: Cambridge University Press.

Hájek, Alan. (1996). ““Mises Redux” – Redux. Fifteen Arguments against Finite Frequentism.” *Erkenntnis* 45(2-3): 209-227.

-----. (2003). “Conditional Probability is the Very Guide of Life.” In Kyburg & Thalos (2003): 183-203.

-----. (2007a). “Interpretations of Probability.” *Stanford Encyclopedia of Philosophy*. August 10, 2010. <http://plato.stanford.edu/entries/probability-interpret/>.

-----. (2007b). “The Reference Class Problem is Your Problem too.” *Synthese* 156(3): 563-585.

Halliday, Daniel. (2007). “Contextualism, Comparatives and Gradability.” *Philosophical Studies* 132(2): 381-393.

Harman, Gilbert. (1968). “Knowledge, Inference, and Explanation.” *American Philosophical Quarterly* 5(3): 164-173.

-----. (1974). *Thought*. Princeton: Princeton University Press.

-----. (1980). “Reasoning and Evidence One Does not Possess.” *Midwest Studies in Philosophy* 5: 163-182.

-----. (1986). *Change in View: Principles of Reasoning*. Cambridge, MA: MIT Press.

Harper, William. (1996). “Knowledge and Luck.” *The Southern Journal of Philosophy* 34(3): 273-283.

Hawthorne, John. (2000). “Implicit Belief and A Priori Knowledge.” *The Southern Journal of Philosophy* Suppl. 38: 191-211.

-----. (2002). “Deeply Contingent A Priori Knowledge.” *Philosophy and Phenomenological Research* 65(2): 247-269.

-----. (2004a). *Knowledge and Lotteries*. Oxford: Oxford University Press.

-----. (2004b) “Replies.” *Philosophical Issues* 14: 510-523.

-----. (2005). “The Case for Closure.” In Sosa & Steup (2005): 26-43.

Heit, Evan. (1997). “Features of Similarity and Category-Based Induction.” *Proceedings of the Interdisciplinary Workshop on Similarity and Categorisation*: 115-121.

-----. (2000). “Properties of Inductive Reasoning.” *Psychonomic Bulletin & Review* 7(4): 569-592.

Heit, Evan & Rubinstein, Joshua. (1994). "Similarity and Property Effects in Inductive Reasoning." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 20(2): 411-422.

Heller, Mark. (1989). "Relevant Alternatives." *Philosophical Studies* 55(1): 23-40.

-----. (1995a). "Might-Counterfactuals and Gratuitous Differences." *Australasian Journal of Philosophy* 73(1): 91.

-----. (1995b). "The Simple Solution to the Problem of Generality." *Noûs* 29(4): 501-515.

-----. (1999). "The Proper Role for Contextualism in an Anti-Luck Epistemology." *Philosophical Perspectives* 13: 115.

Hempel, Carl G. (1960). "Inductive Inconsistencies." *Synthese* 12: 439-69.

-----. (1962). "Deductive-Nomological vs. Statistical Explanation." *Minnesota Studies in the Philosophy of Science* 3: 98-169. Reprinted in Fetzer (2001): 87-145.

-----. (1965). *Aspects of Scientific Explanation and Other Essays in the Philosophy of Science*. New York, NY: Free Press.

-----. (1968). "Maximal Specificity and Lawlikeness in Probabilistic Explanation." *Philosophy of Science* 35(2): 116-133. Reprinted in Fetzer (2001): 146-164.

-----. (1977). "Nachwort 1976: Neuere Ideen zu den Problemen der statistischen Erklarung." In *Aspekte wissenschaftlicher Erklarung*. Berlin & New York: de Gruyter. H. Maxian (trans.), "Postscript 1976: More Recent Ideas on the Problem of Statistical Explanation." in Fetzer (2001): 165-186.

Henderson, David & Horgan, Terence. (2006). "Transglobal Reliabilism." *Croatian Journal of Philosophy* 6: 171-195.

-----. (2007). "Some Ins and Outs of Transglobal Reliabilism." In S. Goldberg (ed.), *Internalism and Externalism in Semantics and Epistemology*. Oxford: Oxford University Press: 100-130.

Henderson, David., Horgan, Terence & Potrč, Matjaž. (2007). "Transglobal Evidentialism-Reliabilism." *Acta Analytica* 22(4): 281-300.

Hitchcock, Christopher. (2001). "Causal Generalizations and Good Advice." *The Monist* 84(2): 218-241.

Holland, John H., Holyoak, Keith J., Nisbett, Richard E & Thagard, Paul R. (1989). *Induction: Processes of Inference, Learning, and Discovery*. Cambridge, MA: MIT Press.

Holyoak, Keith J & Thagard, Paul R. (1996). *Mental Leaps: Analogy in Creative Thought*. Cambridge, MA: MIT Press.

Horwich, Paul. (1982). *Probability and Evidence*. Cambridge: Cambridge University Press.

-----. (2006). "The Value of Truth." *Noûs* 40(2): 347-360.

Howard-Synder, Daniel., Howard-Synder, Frances & Feil, Neil. (2003). "Infallibilism and Gettier's Legacy." *Philosophy and Phenomenological Research* 66(2): 304-327.

Howson, Colin & Urbach, Peter. (1993). *Scientific Reasoning: The Bayesian Approach*, 2nd Ed. Chicago, IL: Open Court.

-----. (2006). *Scientific Reasoning: The Bayesian Approach*, 3rd Ed. Chicago, IL: Open Court.

Huemer, Michael. (2005). "Logical Properties of Warrant." *Philosophical Studies* 122(2): 171-182.

Humberstone, Lloyd. (2000). "Parts and Partitions." *Theoria* 66(1): 41-82.

Jackson, Frank. (1998). *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Oxford University Press.

Jeffrey, Richard C. (1983). *The Logic of Decision*, 2nd Ed. Chicago, IL: University of Chicago Press.

Jenkins, Carrie. (2007). "Entitlement and Rationality." *Synthese* 157(1): 25-45.

Kahneman, Daniel & Tversky, Amos. (1979). "Prospect Theory: An Analysis of Decision under Risk." *Econometrica* 47(2): 263-291.

Kamp, Hans. (1975). "Two Theories about Adjectives." In E. Keenan (ed.), *Formal Semantics of Natural Languages*. Cambridge: Cambridge University Press: 123-155.

Kaplan, Mark. (1996). *Decision Theory as Philosophy*. Cambridge: Cambridge University Press.

Kappel, Klemens. (2006). "A Diagnosis and Resolution to the Generality Problem." *Philosophical Studies* 127(3): 525-560.

Kennedy, Christopher. (2007). "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Adjectives." *Linguistics and Philosophy* 30(1): 1-45.

Kitcher, Patricia. (1988). "Marr's Computational Theory of Vision." *Philosophy of Science* 55(1): 1-24.

Klein, Ewan. (1980). "A Semantics for Positive and Comparative Adjectives." *Linguistics and Philosophy* 4(1): 1–45.

Klein, Peter D. (2000). "Contextualism and the Real Nature of Academic Skepticism." *Philosophical Issues* 10: 108-116.

Korcz, Keith A. (1997). "Recent Work on the Basing Relation." *American Philosophical Quarterly* 34(2): 171-191.

Kornblith, Hilary. (2000). “The Contextualist Evasion of Epistemology.” *Philosophical Issues* 10: 24-32.

-----. (2002). *Knowledge and Its Place in Nature*. Oxford: Oxford University Press.

Kripke, Saul A. (1984). *Wittgenstein on Rules and Private Language: An Elementary Exposition*. Cambridge, MA: Harvard University Press.

Kvanvig, Jonathan. (2003). *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.

-----. (2005). “Is Truth the Primary Epistemic Goal: Truth is not the Primary Epistemic Goal.” In Sosa & Steup (2005): 285-296.

-----. (2008). “Critical Notices of Duncan Pritchard’s *Epistemic Luck*.” *Philosophy and Phenomenological Research* 77(1): 272-281.

Kyburg, Jr., Henry E. (1961). *Probability and the Logic of Rational Belief*. Middletown, CT: Wesleyan University Press.

-----. (1970a). “More on Maximal Specificity.” *Philosophy of Science* 37(2): 295-300.

-----. (1970b). *Probability and Inductive Logic*. New York, NY: MacMillan.

-----. (1974a). “Propensities and Probabilities.” *The British Journal for the Philosophy of Science* 25(4): 358-375.

-----. (1974b). *The Logical Foundations of Statistical Inference*. Dordrecht, Netherlands: Reidel.

-----. (1977). “Randomness and the Right Reference Class.” *The Journal of Philosophy* 74(9): 501-521.

-----. (1983). “The Reference Class.” *Philosophy of Science* 50(3): 374-397.

-----. (2003). “Don’t Take Unnecessary Chances!” In Kyburg & Thalos (2003): 277-294.

Kyburg, Jr., Henry E & Teng, Choh Man. (2001). *Uncertain Inference*. Cambridge: Cambridge University Press.

Kyburg, Jr., Henry E & Thalos, Mariam. (eds.) (2003). *Probability is the Very Guide of Life: The Philosophical Uses of Chance*. Chicago, IL: Open Court.

Lassaline, Mary E. (1996). "Structural Alignment in Induction and Similarity." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 22(3): 754-770.

Latus, Andrew. (2000). "Our Epistemic Goal." *Australasian Journal of Philosophy* 78(1): 28-39.

Lehrer, Keith. 1965. "Knowledge, Truth and Evidence." *Analysis* 25(5): 168-175.

-----. (2000). *Theory of Knowledge*, 2nd Ed. Boulder, CO: Westview Press.

-----. (2005). "Coherence and the Truth Connection." *Erkenntnis* 63(3): 413-423.

Leite, Adam. (2004). "Is Fallibility an Epistemological Shortcoming?" *The Philosophical Quarterly* 54(215): 232-251.

Lepock, Christopher. (2009). "How to Make the Generality Problem Work for You." *Acta Analytica* 24(4): 275-286.

Levi, Isaac. (1980). *The Enterprise of Knowledge: An Essay on Knowledge, Credal Probobility, and Chance*. Cambridge, MA: MIT Press.

-----. (1983). "Review of Richard Jeffrey (ed.), *Studies in Inductive Logic and Probability*." *The Philosophical Review* 92(1): 116-121.

Lewis, David K. (1973). *Counterfactuals*. Cambridge, MA: Harvard University Press.

-----. (1979). "Counterfactual Dependence and Time's Arrow." *Noûs* 13(4): 455-476.

-----. (1980). “A Subjectivist's Guide to Objective Chance.” In R. Jeffrey (ed.), *Studies in Inductive Logic and Probabilities*, Vol. II. Berkeley, CA: University of California Press: 83-133.

-----. (1996). “Elusive Knowledge.” *Australasian Journal of Philosophy* 74(4): 549-567.

Ludlow, Peter. (1989). “Implicit Comparison Classes.” *Linguistics and Philosophy* 12(4): 519-533.

-----. (2005). “Contextualism and the New Linguistic Turn in Epistemology.” In G. Preyer & G. Peter (eds.), *Contextualism in Philosophy: Knowledge, Meaning, and Truth.* Oxford: Oxford University Press: 11-50.

Luper-Foy, Steven. (1987a). “The Causal Indicator Analysis of Knowledge.” *Philosophy and Phenomenological Research* 47(4): 563-587.

-----. (ed.) (1987b). *The Possibility of Knowledge: Nozick and His Critics*. Totowa, NJ: Rowman & Littlefield.

-----. (2007). “The Easy Argument.” *Acta Analytica* 22(4): 321-331.

Manley, David & Wasserman, Ryan. (2007). “A Gradable Approach to Dispositions.” *The Philosophical Quarterly* 57(226): 68-75.

Markie, Peter J. (2005). “Easy Knowledge.” *Philosophy and Phenomenological Research* 70(2): 406-416.

Marr, David. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information*. San Francisco, CA: W. H. Freeman.

McDonald, John., Samuels, Mark & Rispoli, Janet. (1996). “A Hypothesis-Assessment Model of Categorical Argument Strength.” *Cognition* 59(2): 199-217.

McDowell, John H. (1982). “Criteria, Defeasibility, and Knowledge.” *Proceedings of the British Academy* 68:455-479. Reprinted in McDowell (1998): 369-394:

-----. (1986). “Singular Thought and the Extent of Inner Space.” In P. Pettit & J. McDowell (eds.), *Subject, Thought, and Context*. Oxford: Clarendon Press: 136-168. Reprinted in McDowell (1998): 228-259.

-----. (1995). “Knowledge and the Internal.” *Philosophy and Phenomenological Research* 55: 877–893. Reprinted in McDowell (1998): 395-413.

-----. (1998). *Meaning, Knowledge, and Reality*. Cambridge, MA: Harvard University Press.

-----. (2008). “The Disjunctive Conception of Experience as Material for a Transcendental Argument.” In A. Haddock & F. Macpherson (eds.), *Disjunctivism: Perception, Action, Knowledge*. Oxford: Oxford University Press: 376-389.

McEvoy, Mark. (2005). “Belief-Independent Processes and the Generality Problem for Reliabilism.” *Dialectica* 59(1): 19-35.

McGinn, Colin. (1984). “The Concept of Knowledge.” *Midwest Studies in Philosophy* 9: 529-54.

Medin, Douglas L., Coley, John D., Storms, Gert & Hayes, Brett K. (2003). “A Relevance Theory of Induction.” *Psychonomic Bulletin & Review* 10(3): 517-532.

Medin, Douglas L., Goldstone, Robert L & Gentner, Dedre (1993). “Respects for Similarity.” *Psychological Review* 100(2): 254-278.

Merricks, Trenton. (1995). “Warrant Entails Truth.” *Philosophy and Phenomenological Research* 55(4): 841-856.

-----. (1997). “More on Warrant’s Entailing Truth.” Philosophy and Phenomenological Research 57(3): 627-632.

Millar, Alan. (1991). *Reasons and Experience*. Oxford: Clarendon Press.

Millikan, Ruth G. (1984). “Naturalist Reflections on Knowledge.” *Pacific Philosophical Quarterly* 65(4): 315-334.

von Mises, Richard. (1981). *Probability, Statistics and Truth*, 2nd Revised Ed. New York, NY: Dover.

Morton, Peter. (1993). “Supervenience and Computational Explanation in Vision Theory.” *Philosophy of Science* 60(1): 86-99.

Moser, Paul K. (ed.) (2002). *Oxford Handbook of Epistemology*. New York, NY: Oxford University Press.

Myrvold, Wayne. (2003). “A Bayesian Account of the Virtue of Unification.” *Philosophy of Science* 70: 399–423.

Neta, Ram. (2007a). “Fixing the Transmission: the New Mooreans.” In Nuccetelli & Seay (2007): 62-83.

-----. (2007b). “Propositional Justification, Evidence, and the Cost of Error.” *Philosophical Issues* 17: 197-216.

-----. (2008). “What Evidence Do You Have?” *The British Journal for the Philosophy of Science* 59(1): 89-119.

-----. (forthcoming). “Liberalism and Conservatism in the Epistemology of Perceptual Belief.” *Australasian Journal of Philosophy*.

Neufeld, Eric & Goodwin, Scott. (1998). “The 6-49 Lottery Paradox.” *Computational Intelligence* 14(3): 273-286.

Nozick, Robert. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.

Nuccetelli, Susana & Seay, Gary. (eds.) (2007). *Themes from G. E. Moore: New Essays in Epistemology and Ethics*. Oxford: Oxford University Press.

Osherson, Daniel N., Smith, Edward E., Wilkie, Ormond., Lopez, Alejandro & Shafir, Eldar. (1990). "Category-Based Induction." *Psychological Review* 97(2): 185–200.

Pappas, George S. (1987). "Suddenly He Knows." In S. Luper-Foy (1987b): 152-162.

Peacocke, Christopher. (1999). Being Known. Oxford: Clarendon Press.

Pearl, Judea. (1990). "Jeffrey's Rule, Passage of Experience, and Neo-Bayesianism." In H. Kyburg., R. Loui & G. Carlson (eds.), *Knowledge Representation and Defeasible Reasoning*. Dordrecht, Netherlands: Kluwer Academic: 245-265.

Piccinini, Gualtiero. (2008). "Computation without Representation." *Philosophical Studies* 137(2): 205-241.

Plantinga, Alvin. (1993a). *Warrant and Proper Function*. New York, NY: Oxford University Press.

-----. (1993b). *Warrant: The Current Debate*. New York, NY: Oxford University Press.

Pollock, John L. (1984). "Reliability and Justified Belief." *Canadian Journal of Philosophy* 14(1): 103-114.

-----. (1990). Nomic Probability and the Foundations of Induction. New York, NY: Oxford University Press.

Popper, Karl R. (1957). "The Propensity Interpretation of the Calculus of Probability, and the Quantum Theory." In S. Körner (ed.) *Observation and Interpretation*. London: Butterworth Scientific: 65-70.

-----. (1959a). *The Logic of Scientific Discovery*, English Ed. New York, NY: Basic Books.

-----. (1959b). “The Propensity Interpretation of Probability.” *The British Journal for the Philosophy of Science* 10(37): 25-42.

Portner, Paul H. (2009). *Modality*. Oxford: Oxford University Press.

Pritchard, Duncan H. (2005a). *Epistemic Luck*. Oxford: Oxford University Press.

-----. (2005b). “The Structure of Sceptical Arguments.” *The Philosophical Quarterly* 55(218): 37-52.

-----. (2007). “Recent Work on Epistemic Value.” *American Philosophical Quarterly* 44(1): 85-110.

-----. (2008a). “Knowledge, Luck, and Lotteries.” In D. Pritchard & V. Hendricks (eds.), *New Waves in Epistemology*. London: Palgrave Macmillan: 28-51.

-----. (2008b). “Radical Scepticism, Epistemic Luck, and Epistemic Value.” *Proceedings of the Aristotelian Society* Suppl. 82: 19-41.

Pryor, James. (2000). “The Skeptic and the Dogmatist.” *Noûs* 34(4):517–549.

-----. (2004). “What’s Wrong with Moore’s Argument?” *Philosophical Issues* 14:349–378.

Putnam, Hilary. (1988). *Representation and Reality*. Cambridge, MA: MIT Press.

Pylyshyn, Zenon W. (1980). “Computation and Cognition: Issues in the Foundations of Cognitive Science.” *Behavioral and Brain Sciences* 3(1): 111-132.

-----. (1984). *Computation and Cognition: Toward a Foundation for Cognitive Science*. Cambridge, MA: MIT Press.

Rabinowicz, Wlodek & Rønnow-Rasmussen, Toni (2000). “A Distinction in Value: Intrinsic and for Its Own Sake.” *Proceedings of the Aristotelian Society* 100(1): 33-51.

-----. (2003). “Tropic of Value.” *Philosophy and Phenomenological Research* 66(2): 389-403.

Ramsey, Frank P. (1926). “Truth and Probability”, in R. B. Braithwaite (ed.) (1931), *Foundations of Mathematics and other Essays*. London: Routledge & P. Kegan: 156–198. Reprinted in D. H. Mellor (ed.) (1990), *Philosophical Papers*. Cambridge: Cambridge University Press: 52-94.

Reed, Baron. (2002). “How to Think about Fallibilism.” *Philosophical Studies* 107(2): 143-157.

Rehder, Bob. (2006). “When Similarity and Causality Compete in Category-Based Property Generalization.” *Memory & Cognition* 34(1): 3-16.

Reichenbach, Hans. (1949). *The Theory of Probability: An Inquiry into the Logical and Mathematical Foundations of the Calculus of Probability*. Berkeley, CA: University of California Press.

Roddy, Bernard. (1994). “Millar and the Problem of Generality.” *Ratio* 7(2): 183-190.

Rorty, Richard. (1979). *Philosophy and the Mirror of Nature*. Princeton, NJ: Princeton University Press.

-----. (1995). “Is Truth a Goal of Enquiry? Davidson vs. Wright.” *The Philosophical Quarterly* 45(180): 281-300.

Rosenberg, Jay F. (2002). *Thinking about Knowing*. Oxford: Clarendon Press.

Ross, Brian H & Murphy, Gregory L. (1999). “Food for Thought: Cross-Classification and Category Organization in a Complex Real-World Domain.” *Cognitive Psychology* 38(4): 495-553.

Sainsbury, Richard M. (1997). “Easy Possibilities.” *Philosophy and Phenomenological Research* 57(4): 907-919.

Salmon, Wesley C. (1971a). “Introduction.” In Salmon, Jeffrey & Greeno (1971): 3-17.

-----. (1971b). “Statistical Explanation.” In Salmon, Jeffrey & Greeno (1971): 29-87.

-----. (1984). *Scientific Explanation and the Causal Structure of the World*. Princeton, NJ: Princeton University Press.

-----. (1989). *Four Decades of Scientific Explanation*. Minneapolis, MN: University of Minnesota Press.

Salmon, Wesley C., Jeffrey, Richard C & Greeno, James G. (eds.) (1971). *Statistical Explanation and Statistical Relevance*. Pittsburgh, PA: University of Pittsburgh Press.

Sartwell, Crispin. (1991). “Knowledge is Merely True Belief.” *American Philosophical Quarterly* 28(2): 157-65.

-----. (1992). “Why Knowledge is Merely True Belief.” *The Journal of Philosophy* 89(4): 167-180.

Schaffer, Jonathan. (2004). “From Contextualism to Contrastivism.” *Philosophical Studies* 119(1): 73-103.

-----. (2005). “Contrastive Knowledge.” In T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Vol. 1. Oxford: Oxford University Press: 235-271.

-----. (2007a). “Closure, Contrast, and Answer.” *Philosophical Studies* 133(2): 233-255.

-----. (2007b). “Knowing the Answer.” *Philosophy and Phenomenological Research* 75(2): 383-403.

-----. (2008). “The Contrast-Sensitivity of Knowledge Ascriptions.” *Social Epistemology* 22(3): 235.

-----. (2009). “Knowing the Answer Redux: Replies to Brogaard and Kallestrup.” *Philosophy and Phenomenological Research* 78(2): 477-500.

Scheutz, Matthias. (2001). “Computational versus Causal Complexity.” *Minds and Machines* 11(4): 543-566.

Schmitt, Frederick F. (1992). *Knowledge and Belief*. London: Routledge.

Searle, John R. (1992). *The Rediscovery of the Mind*. Cambridge, MA: MIT Press.

Segal, Gabriel. (1989). “Seeing What is not There.” *The Philosophical Review* 98(2): 189-214.

-----. (1991). “Defence of a Reasonable Individualism.” *Mind* 100(4): 485-494.

Sellars, Wilfrid. (1956). “Empiricism and the Philosophy of Mind.” *Minnesota Studies in the Philosophy of Science* 1: 253-329.

-----. (1966). “The Refutation of Phenomenalism: Prolegomena to a Defense of Scientific Realism.” In P. Feyerabend & G. Maxwell (eds.), *Mind, Matter, and Method: Essays in Philosophy and Science in Honor of Herbert Feigl*. Minneapolis, MN: University of Minnesota Press: 198-214.

Shagrir, Oron. (1999). “What is Computer Science About?” *The Monist* 82(1): 131-149.

-----. (2001). “Content, Computation and Externalism.” *Mind* 110(438): 369-400.

Shapiro, Lawrence A. (1995). “Content, Kinds, and Individualism in Marr’s Theory of Vision.” *The Philosophical Review* 102(4): 489-513.

-----. (1997). “A Clearer Vision.” *Philosophy of Science* 64(1): 131-153.

Shatz, David. (1981). “Reliability and Relevant Alternatives.” *Philosophical Studies* 39(4): 393-408.

-----. (1987). “Nozick’s Conception of Skepticism.” In Luper-Foy (1987b): 242-266.

Shope, Robert. (2008). “Abnormality, Cognitive Virtues, and Knowledge.” *Synthese* 163(1): 99-118.

Silins, Nicholas. (2005). “Deception and Evidence.” *Philosophical Perspectives* 19: 375-404.

-----. (2008). “Basic Justification and the Moorean Response to the Skeptic.” In T. Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, Vol. 2. Oxford: Oxford University Press: 108-140.

Sloman, Steven A. (1993). “Feature-Based Induction.” *Cognitive Psychology* 25(2): 231-280.

-----. (1994). “When Explanations Compete: The Role of Explanatory Coherence on Judgements of Likelihood.” *Cognition* 52(1): 1-21.

-----. (1997). “Explanatory Coherence and the Induction of Properties.” *Thinking & Reasoning* 3(2): 81-110.

-----. (1998). “Categorical Inference is not a Tree: The Myth of Inheritance Hierarchies.” *Cognitive Psychology* 35(1): 1-33.

Smith, Edward E., Shafir, Eldar & Osherson, Daniel N. (1993). “Similarity, Plausibility, and Judgments of Probability.” *Cognition* 49(1-2): 67-96.

Smith, Martin. (2009). “Transmission Failure Explained.” *Philosophy and Phenomenological Research* 79(1): 164-189.

-----. (2010). “What Else Justification Could be.” *Noûs* 44(1): 10-31.

Sosa, Ernest. (1980). "The Raft and the Pyramid: Coherence versus Foundations in the Theory of Knowledge." *Midwest Studies in Philosophy* 5: 3-26. Reprinted in Sosa (1991b): 165-191.

-----. (1985). "Knowledge and Intellectual Virtue." *The Monist* 68(2): 224-45. Reprinted in Sosa (1991b): 225-244.

-----. (1991a). "Intellectual Virtue in Perspective." In Sosa (1991b): 270-293.

-----. (1991b). *Knowledge in Perspective: Selected Essays in Epistemology*. New York, NY: Cambridge University Press.

-----. (1991c). "Reliabilism and Intellectual Virtue." In Sosa (1991b): 131-145.

-----. (1993). "Review: Proper Functionalism and Virtue Epistemology." *Noûs* 27(1): 51-65.

-----. (1999). "How to Defeat Opposition to Moore." *Philosophical Perspectives* 13: 141-153.

-----. (2002). "Tracking, Competence, and Knowledge." In Moser (2002): 264-286.

-----. (2003). "The Place of Truth in Epistemology." In DePaul & Zagzebski (2003): 155-179.

-----. (2004a). "Relevant Alternatives, Contextualism Included." *Philosophical Studies* 119(1): 35-65.

-----. (2004b). "Reply to Keith DeRose." In Greco (2004b): 276-282.

-----. (2004c). "Reply to Linda Zagzebski." In Greco (2004b): 319-322.

-----. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Vol. 1. Oxford: Oxford University Press.

-----. (2009). *Reflective Knowledge: Apt Belief and Reflective Knowledge*, Vol. 2. Oxford: Oxford University Press.

Sosa, Ernest & Steup, Matthias. (eds.) (2005). *Contemporary Debates in Epistemology*. Oxford: Blackwell.

Sperber, Dan & Wilson, Deirdre. (1996). *Relevance: Communication and Cognition*, 2nd Ed. Oxford: Blackwell.

Sprevak, Mark D. (forthcoming) "Computation, Individuation, and the Received View on Representation." *Studies in History and Philosophy of Science*.

Stalnaker, Robert C. (1968). "A Theory of Conditionals." *Studies in Logical Theory, American Philosophical Quarterly Monograph Series* 2: 98-112.

-----. (1974). "Pragmatic Presuppositions." In M. Munitz & P. Unger (eds.), *Semantics in Philosophy*. New York, NY: New York University Press: 197-213. Reprinted in Stalnaker (1999): 47-62.

-----. (1978). "Assertion." In P. Cole (ed.), *Syntax and Semantics*. New York, NY: Academic Press: 315-332. Reprinted in Stalnaker (1999): 78-95.

-----. (1999). *Context and Content: Essays on Intentionality in Speech and Thought.* Oxford: Oxford University Press

-----. (2004). "Comments on "From Contextualism to Contrastivism"." *Philosophical Studies* 119(1): 105-117.

Stanley, Jason. (2000). "Context and Logical Form." *Linguistics and Philosophy* 23(4): 391–434. Reprinted in Stanley (2007): 30-68.

-----. (2002). "Nominal Restriction." In G. Preyer & G. Peter (eds.), *Logical Form and Language*. Oxford: Oxford University Press: 365–388. Reprinted in Stanley (2007): 111-132.

-----. (2004). "On the Linguistic Basis for Contextualism." *Philosophical Studies* 119(1-2): 119-146.

-----. (2005). *Knowledge and Practical Interests*. Oxford: Oxford University Press.

-----. (2007). *Language in Context: Selected Essays*. Oxford: Oxford University Press.

Stich, Stephen P. (1990). *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: MIT Press.

Sutton, Jonathan. (2007). *Without Justification*. Cambridge, MA: MIT Press.

Swain, Marshall. (1978). "Reasons, Causes, and Knowledge." *The Journal of Philosophy* 75(5): 229-249.

Swinburne, Richard G. (2001). *Epistemic Justification*. Oxford: Oxford University Press.

Thorn, Paul D. (2007). *Three Problems of Direct Inference*. Doctoral Dissertation. University of Arizona.

Tucker, Chris. (2009). "Perceptual Justification and Warrant by Default." *Australasian Journal of Philosophy* 87(3): 445.

Turri, John. (2010). "On the Relationship between Propositional and Doxastic Justification." *Philosophy and Phenomenological Research* 80(2): 312-326.

Tversky, Amos. (1977). "Features of Similarity." *Psychological Review* 84(4): 327-352.

Unger, Peter. (1968). "An Analysis of Factual Knowledge." *The Journal of Philosophy* 65(6): 157-170.

Vahid, Hamid. (2008). "The Puzzle of Fallible Knowledge." *Metaphilosophy* 39(3): 325-344.

-----. (2009). "Triangulation, Content and the Basing Relation." *Grazer Philosophische Studien* 78: 231-250.

Van Cleve, James. (2003). “Is Knowledge Easy – or Impossible? Externalism as the Only Alternative to Skepticism.” In S. Luper (ed.), *The Skeptics: Contemporary Essays*. Aldershot: Ashgate: 45-59.

van Fraassen, Bas C. (1980). *The Scientific Image*. Oxford: Clarendon Press.

-----. (1984). “Belief and the Will.” *The Journal of Philosophy* 81(5): 235-256.

-----. (1989). *Laws and Symmetry*. Oxford: Clarendon Press.

-----. (1995). “Belief and the Problem of Ulysses and the Sirens.” *Philosophical Studies* 77(1): 7-37.

Vogel, Jonathan M. (1990). “Are There Counterexamples to the Closure Principle?” In M. Roth & G. Ross (eds.), *Doubting: Contemporary Perspectives on Skepticism*. Dordrecht, Netherlands: Kluwer Academic: 13-27.

-----. (1999). “The New Relevant Alternatives Theory.” *Philosophical Perspectives* 13: 155.

-----. (2000). “Reliabilism Leveled.” *The Journal of Philosophy* 97(11): 602-623.

-----. (2004a). “Skeptical Arguments.” *Philosophical Issues* 14: 426-455.

-----. (2004b). “Speaking of Knowledge.” *Philosophical Issues* 14: 501-509.

-----. (2007). “Subjunctivitis.” *Philosophical Studies* 134(1): 73-88.

-----. (2008). “Epistemic Bootstrapping.” *The Journal of Philosophy* 105(9): 518-539.

von Mises, Richard. (1957). *Probability, Statistics and Truth*, revised English Ed. New York, NY: Macmillan.

Wallis, Charles. (1994). “Truth-Ratios, Process, Task, and Knowledge.” *Synthese* 98(2): 243-269.

White, Michael J. (1988). “The Unimportance of Being Random.” *Synthese* 76(1): 171-178.

White, Roger. (2006). “Problems for Dogmatism.” *Philosophical Studies* 131(3): 525-557.

Wildes, Richard. (1990). “Computational Vision with Reference to Binocular Stereo Vision.” In

K. Leibovic (ed.), *Science of Vision*. Berlin, Germany: Springer-Verlag: 332-364.

Williams, Michael. (1977). *Groundless Belief: An Essay on the Possibility of Epistemology*. New Haven, CT: Yale University Press.

-----. (1991). *Unnatural Doubts: Epistemological Realism and the Basis of Scepticism*. Oxford: Blackwell.

-----. (2001a). “Contextualism, Externalism and Epistemic Standards.” *Philosophical Studies* 103(1): 1-23.

-----. (2001b). *Problems of Knowledge: A Critical Introduction to Epistemology*. Oxford: Oxford University Press.

-----. (2004). “Knowledge, Reflection and Sceptical Hypotheses.” *Erkenntnis* 61(2-3): 315-343.

Williamson, Timothy. (2000a). *Knowledge and Its Limits*. Oxford: Oxford University Press.

-----. (2000b). “Scepticism and Evidence.” *Philosophy and Phenomenological Research* 60(3): 613-628.

-----. (2001). “Comments on Michael Williams’ Contextualism, Externalism and Epistemic Standards.” *Philosophical Studies* 103(1): 25-33.

-----. (2005). “Contextualism, Subject-Sensitive Invariantism, and Knowledge of Knowledge.” *The Philosophical Quarterly* 55(219): 213-235.

-----. (2005). “Knowledge and Scepticism.” In F. Jackson & M. Smith (eds.), *The Oxford Handbook of Contemporary Philosophy*. Oxford: Oxford University Press: 681-700.

-----. (2009a). “Probability and Danger.” *The Amherst Lecture in Philosophy* 4: 1-35. August 10, 2010. <http://www.amherstlecture.org/williamson2009/>

-----. (2009b). “Reply to Goldman.” In P. Greenough & D. Pritchard (eds.), *Williamson on Knowledge*. Oxford: Oxford University Press: 305-312.

Wittgenstein, Ludwig. (1953/2001). *Philosophical Investigations*, 3rd Ed. G. Anscombe & R. Rhees (eds.), G. Anscombe (trans.), Oxford: Blackwell.

-----. (1969) *On Certainty*. G. Anscombe & G. von Wright (eds.), G. Anscombe & D. Paul (trans.), Oxford: Blackwell.

Wright, Crispin. (1985). “Facts and Certainty.” *Proceedings of the British Academy* 71: 429-72.

-----. (1991). “Scepticism and Dreaming: Imploding the Demon.” *Mind* 100(397): 87-116.

-----. (2000). “Cogency and Question-Begging: Some Reflections on McKinsey’s Paradox and Putnam's Proof.” *Philosophical Issues* 10: 140-163.

-----. (2002). “(Anti-)Sceptics Simple and Subtle: G. E. Moore and John McDowell.” *Philosophy and Phenomenological Research* 65(2): 330-348.

-----. (2003). “Some Reflections on the Acquisition of Warrant by Inference.” In S. Nuccetelli (ed.), *New Essays on Semantic Externalism and Self-Knowledge*. Cambridge, MA: MIT Press: 57-78.

-----. (2004). “Warrant for Nothing (and Foundations for Free)?” *Proceedings of the Aristotelian Society* Suppl. 78(1): 167-212.

-----. (2007). “The Perils of Dogmatism.” In Nuccetelli & Seay (2007): 25-48.

Wunderlich, Mark E. (2003). “Vector Reliability: A New Approach to Epistemic Justification.” *Synthese* 136(2): 237-262.

-----. (2009). “Noncomparabilism in Epistemology.” *Philosophical Studies* 142(2), 133-151.

Zagzebski, Linda T. (1994). “The Inescapability of Gettier Problems.” *The Philosophical Quarterly* 44(174): 65–73.

-----. (1996). *Virtues of the Mind. An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge. Cambridge: Cambridge University Press.*

-----. (1999). “What is Knowledge?” In Greco & Sosa (1999): 92-116.

-----. (2003). “The Search for the Source of Epistemic Good.” *Metaphilosophy* 34(1-2): 12-28.

-----. (2004). “Epistemic Value Monism.” In Greco (2004b): 190-198.