

2017

Action, Knowledge Norms, and Moral Reasons

Delamont, Audrey

Delamont, A. (2017). Action, Knowledge Norms, and Moral Reasons (Master's thesis, University of Calgary, Calgary, Canada). Retrieved from <https://prism.ucalgary.ca>. doi:10.11575/PRISM/24688
<http://hdl.handle.net/11023/3748>

Downloaded from PRISM Repository, University of Calgary

UNIVERSITY OF CALGARY

Action, Knowledge Norms, and Moral Reasons

by

Audrey Delamont

A THESIS

SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR THE
DEGREE OF MASTER OF ARTS

GRADUATE PROGRAM IN PHILOSOPHY

CALGARY, ALBERTA

APRIL, 2017

© Audrey Delamont 2017

ABSTRACT

What is the relationship between knowledge and action? Many philosophers in recent years have proposed a very intimate connection between these two concepts. John Hawthorne and Jason Stanley—among others—have proposed that knowledge is both a necessary and sufficient condition for rational deliberate action, citing only that an agent must use reasons that are both known and relevant in their decision making. They also argue that an agent is blameworthy and deserving of criticism for failing to adhere to this knowledge norm of action. This thesis examines Hawthorne and Stanley’s norm, pulling apart concepts such as ‘relevance’ and ‘reasons’ to try and better understand what this knowledge norm of action entails. Ultimately this thesis proposes that Hawthorne and Stanley’s principles (and those like it) fail to adequately capture the scope of what a general norm of action should, and proposes tentative suggestions for what features a knowledge norm of action should have.

ACKNOWLEDGEMENTS

I have benefited greatly from the wisdom and knowledge of many colleagues over the years. Although it is in no way a worthy substitute, it is my pleasure to recognize the following individuals for the invaluable impact they have made on me and my work.

I owe Jeremy Fantl a profound debt of gratitude. He has helped me on my philosophical journey as my undergraduate honours advisor as well as the advisor for this undertaking. He has endured hours worth of skeptical arguments, but was always willing to listen, and never made me feel like my skeptical sympathies should be set aside. Jeremy helped broaden the scope of my epistemological interests, and is one of the best professors I've ever had the pleasure of working with and taking classes from. Despite the time lines on which we worked on this thesis together, Jeremy was always there when I needed him with insightful discussions, helpful suggestions, and (sometimes barely legible) margin-commentary on every page of every draft.

I also owe a debt of gratitude to Nicole Wyatt for her understanding and support in all things related to my time in her department, both professionally and personally. Without Nicole's honesty and willingness to encourage and help me while I was without an advisor, I likely would not be where I am today.

I want to thank Denise Retzlaff for *everything* that she has done for me. Denise has not only encouraged and pushed me to do my very best, but has gone over and above in everything I could have ever asked of her.

My work has also benefited greatly from an epistemology study group at the University of Calgary. Ryan Hebert and Caleb Lee served as sounding boards for

many early drafts of this document and the ideas presented throughout. Without their candid criticisms and constant willingness to put down their own work to talk about mine, the ideas articulated here would not be written with as much confidence as they are. Despite our unlikely friendship, Ryan's cleverness and humour has had a large impact on how I navigate the world. Thank you for all you have knowingly, and unknowingly, done for me.

Finally, I want to thank my husband Daniel Delamont, who has endured hours of my philosophical musings that I'm sure he would have been happy to do without. Thank you for learning this philosophical language so that you could talk to me about my ideas, and always being supportive of my process. I love you.

I dedicate this work to Daniel Delamont.

CONTENTS

ABSTRACT	i.
ACKNOWLEDGEMENTS	ii.
DEDICATION	iv.
TABLE OF CONTENTS	v.
1 <i>Introduction</i>	1
1.1 Formulating the Principle	7
1.2 What's Ahead	11
2 <i>Reasons</i>	13
2.1 Defining Reasons	14
2.2 Reasons You Know About and Ones You Don't	17
2.3 You Know Your Subjective Reasons	20
2.4 Moral Reasons	28
3 <i>The Case</i>	33
3.1 The Addition Method	37
3.2 The Background Condition Method	40
3.3 Objectively p-Dependent	43
3.4 Overriding Ethical Norms	48
3.5 Summary	53
4 <i>Objections & Responses</i>	54
4.1 Should Have Known Better	54
4.2 Other Objections	60
4.2 Summary	70
5 <i>Conclusion</i>	72
5.1 What a Knowledge Norm Might Look Like	74
 <i>Bibliography</i>	 82

I

Introduction

One of the central problems in epistemology is knowing what knowledge is and when an agent is in possession of it. But what good is knowledge once we have it? The contemporary epistemic literature is now not only concerned with when we might say an agent is in possession of knowledge, but now also concerns itself what that agent *does* with said knowledge. As Timothy Williamson (2005) notes,

Moreover, the concept of knowledge would lose some of its significance: one reason why it matters whether you know something is that, if you do, you are entitled to use it in ways in which you would not otherwise be so entitled. (228)

Regularly, we use knowledge to reason, make decisions, and carry out

actions. Questions that try to make sense of the relationship between knowledge, reasoning, and action, might help us understand agents who act with knowledge, those who act with less than knowledge, and to guide us in understanding and evaluating our criticisms of less-than-knowledge-level actions.

Consider the following case:

CHRISTMAS TREE: It is November 25th and Juliet would like to go and get a Christmas tree. She knows that it isn't yet December and that she is setting up her tree rather early, but she would like to extend the holiday season as long as she is able, which means setting her tree up early. Juliet drives to the store where she has purchased her tree in years previous—despite the cold temperatures and bad road conditions—only to find that the store will not begin selling trees until December 1.

What might we say of Juliet's actions? We might wonder why she didn't call the store first to see if they would have Christmas trees so early in the season. In effect, we would be wondering whether or not Juliet *knew* that they had Christmas trees, before acting as if they would. Surely if Juliet had known that she would not be able to purchase a Christmas tree until December 1st, she would not have driven to the store on November 25th. Instead, we would expect her to cite her knowledge that the store would *not* have Christmas trees until December 1st as her reason for staying home.

ORANGES: Juliet has run out of oranges. When asked, "Why are you going to the store?" Juliet replies that she is out of oranges, that she knows the store will have more, and that going to the

store will be the best way to re-stock her orange supply.

Because Juliet knows that the store will have oranges, when prompted, she cites this knowledge as her reason for going to the store.

BURGLARY: There has been a rash of garage-related break-ins in the neighbourhood where Juliet and Dan live. Having just left the house, Juliet asks Dan if he remembered to close the garage door. Dan replies that he did. Hours later upon their return, they find that the garage door was indeed left open, however nothing had been taken.

In **BURGLARY**, despite the fact that leaving the garage door open did not result in any bad consequences, we might nevertheless expect Juliet to ask Dan why he said the garage door was closed, when in fact it wasn't. In his defense, Dan might tell Juliet that he was, "pretty sure that the door was closed", to which she might point out that he shouldn't have acted as if it was closed (by saying that it was), if he didn't *know* that it was.

We encounter examples such as **CHRISTMAS TREE**, **ORANGES**, and **BURGLARY** often. When asked why we did something, or decided to refrain from doing something, we commonly cite knowledge (or a lack thereof) as our justification. When our actions are challenged or we need to provide an excuse, usually we offer knowledge-based answers to try and lessen our culpability: "I thought the store would have Christmas trees". One doesn't need to be familiar with epistemology to recognize the important role that knowledge plays in guiding and justifying our

actions. From these cursory but sufficiently common-place examples, we can see that there seems to be some sort of link between knowledge, reasons, and action. Having knowledge seems to play a crucial role in formulating our reasons for acting the way that we do. When we have knowledge it allows us to make decisions about how to navigate the world. When we do not possess knowledge, we are apt to make poor decisions. But it isn't always so straight forward. Being in possession of knowledge is not a free license to act, just as being without knowledge is not a reason to remain paralyzed by epistemic constraints.

The way some contemporary philosophers have attempted to explain the relationship between knowledge and action involves discussion of what they have called 'knowledge norms'.¹ These norms are meant to capture the correct state that an agent will find herself in when she is reasoning appropriately.

In their 2008 paper "Knowledge and Action", John Hawthorne and Jason Stanley offer a knowledge norm for action that they believe accurately captures the way we evaluate speakers in cases of action in our ordinary folk appraisals. As evidence for this view, they point to the way we criticize agents in cases of action based on anything less than knowledge.

1 See for instance Ball (2014), Brown (2008, 2008a), Carter and Gordon (2011), Cohen (2004), Davidson (1963), DeRose (2002), Douven (2006), Hawthorne (2004), Hawthorne and Stanley (2008), Hill and Schechter (2007), Hyman (1999), Kvanvig (2009, 2011), Lackey (2007, 2010), Littlejohn (2016), Maitra and Weatherson (2010), Fantl & Mcgrath (2009, 2012), McKinnon (2011, 2013), Raz (1999), Rescorla (2009), Shieber (2009), Stanley (2005), Turri (2010, 2010a, 2011), Weiner (2005), Williams (2002), Williamson (2000, 2005)

Consider,

DAN: Didn't you know if the store would have Christmas trees this early?

JULIET: No, I didn't.

DAN: You should have checked before we drove all the way over here.

Hawthorne and Stanley would agree with Dan's criticism of Juliet. Because she lacked the knowledge of whether or not the store was selling Christmas trees but acted as if she did, she opened herself up to criticism based on this lack of knowledge. In describing the sorts of cases Hawthorne and Stanley take to be straightforward examples of our ordinary folk appraisals of knowledge and action, they offer the following example:

Suppose, for example, that Hannah and Sarah are trying to find a restaurant, at which they have time-limited reservations. Instead of asking someone for directions, Hannah goes on her hunch that the restaurant is down a street on the left. After walking for some time, it becomes quite clear that they went down the wrong street. A natural way for Sarah to point out that Hannah made the wrong decision is to say, "You shouldn't have gone down this street, since you did not know that the restaurant was here." (571)

Similar to *CHRISTMAS TREE*, Hannah has acted on less than knowledge—specifically the belief that the restaurant was on the left—rather than knowledge that it was. It is because of this lack of knowledge that Hawthorne and Stanley

feel that Sarah's criticisms of Hannah are warranted. Intuitively it seems to be the case that when we act on less than knowledge, we are doing something wrong. Just as intuitively, in these cases we might be tempted to say that while there was something epistemically-off about Juliet and Hannah's actions, and that they might be deserving of a, "you shouldn't have done that" from Dan and Sarah respectively, their epistemic missteps were not monumental. Contrast that with a 'high stakes' type of case where a failure to know might have disastrous effects, such as the outcome of a murder trial. If a judge were to act on a hunch that the defendant was guilty rather than knowledge that they were, we would offer much stronger criticisms of the judge, than those given to Juliet and Hannah. In all three cases of the judge, Juliet, and Hannah, actions were taken despite a lack of knowledge. However despite the fact that it was the same *type* of offense committed, the extent of our criticisms are different.

So what's going on here?

If there is a knowledge norm relating knowledge with action as Hawthorne and Stanley think that our common folk appraisals seem to suggest that there is, then it seems to be one which is often violated. The examples sketched above are familiar to us precisely *because* of how often the knowledge norm is violated. Following Timothy Williamson (2000), Hawthorne and Stanley acknowledge that the knowledge norm is often broken, and the degree to which it is can be evaluated from the very excusable (as in the case of CHRISTMAS TREE, Hannah and Sarah, or even cases of Gettiered actions), to the very inexcusable (knowingly acting on a hunch, as with the judge).

FORMULATING THE PRINCIPLE

Hawthorne and Stanley take this picture to be relatively uncontentious, arguing that, "...our ordinary folk appraisals of the behaviour of others suggest that the concept of knowledge is intimately intertwined with the rationality of action." (571). More rigorously stated, Hawthorne and Stanley capture this link between knowledge and action in a knowledge norm comprised of two principles which they take to be jointly necessary and sufficient for rational action. Although these principles run contrary to views that suggest rational action ought to be guided by maximizing expected utility, Hawthorne and Stanley suggest that they accurately capture the way we ordinarily criticize speakers in cases of actions based on anything less than knowledge.² They argue that, "When someone acts on a belief that does not amount to knowledge, she violates the norm, and hence is subject to criticism" (577). It is important to note that what Hawthorne and Stanley are primarily concerned with is what is appropriate to treat as a reason for action. That is, they don't say anything about the conclusions that an agent ought to draw from those reasons, only what premises may permissibly appear in an agent's reasoning.

The necessity condition of Hawthorne and Stanley's knowledge norm states

² Although it may be thought that a theory which simply obligates agents to maximize expected utility does just as good of a job as Hawthorne and Stanley's knowledge norm despite not necessarily being as intuitive, Hawthorne and Stanley argue that their norm captures an important aspect of action that expected utilities theories do not. There are cases where an agent does what they ought to do, but they do it for the wrong reasons. A theory of expected utility does not have the framework to criticize an agent for negligently, luckily, or mistakenly doing what they should, whereas the knowledge norm does.

that only known propositions ought to be used as premises in practical reasoning. For instance, Hawthorne (2004) argues that the following is an “absurd” piece of reasoning:

1. I will lose the lottery
 2. If I keep the ticket I will get nothing
 3. If I sell the ticket, I will get a cent
- So I ought to sell the ticket. (572)

It is absurd because they argue that we never can know that we will lose a lottery and so can't know the first premise. It is precisely because premise 1 isn't known which explains why using it in practical reasoning is improper. Hawthorne and Stanley argue that this principle follows naturally from the way we criticize speakers in cases of reasoning that lack knowledge. They formulate the principle as follows:

THE ACTION-KNOWLEDGE PRINCIPLE (AKP): Treat the proposition that *p* as a reason for acting only if you know that *p*. (577)

This principle is not sufficient on its own; it does not stipulate specifically what knowledge can or should be used in rational deliberation. For example, if only known premises are needed, then why shouldn't the above lottery-ticket consideration include my knowledge of instruments or painting? The reason is that those pieces of knowledge are not *relevant* to my current deliberation of whether I ought to sell my lottery ticket. Thus Hawthorne and Stanley's second principle

provides a sufficiency condition:

THE REASON-KNOWLEDGE PRINCIPLE (RKP): Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting *iff* you know that p. (578)

This principle serves to narrow the scope of what might appropriately be treated as a reason for action. If my p-dependent choice is about whether or not I ought to sell my lottery ticket, then my knowledge of instruments and painting is irrelevant, and would be *prima facie* bizarre to treat as reasons for or against the sale of my lottery ticket. Importantly, this principle emphasizes the difference between there being a reason, and an agent having that reason. Knowledge is the crucial link that can turn brute facts about the world into reasons that an agent has, but this can only be case if they are known (580).

Hawthorne and Stanley are not the only ones in the literature to cite a knowledge-based norm of action. Consider a similar sentiment from Clayton Littlejohn (2016):

How are knowledge and reasons related? [...] It's arguably the case that knowledge is a way to be related to some facts such that they can guide you and serve as your reasons. Arguably, it's the only way that you can be guided by such facts as reasons. (18)

Or Michael Da Silva (2014) referring to Williamson (2000)

... [K]nowledge both helps explain actions in a manner that other theories cannot, and serves as a norm for action. It plays a role in explaining action that no non-factive notion can play. In Williamson's picture, one will act differently when acting on the basis of a knowledge-desire pair rather than a belief-desire pair. Normatively justified intentional action should be the result of knowledge-desire pairs. (324)

Or John Hyman (1999):

S knows that p if and only if there is some action such that S is able to perform that action for the reason that p.

Or Jeremy Fantl and Matthew McGrath (2002),

If you know that p, then it shouldn't be a problem to act as if p. If it is a problem to act as if p, you can explain why by saying that you don't know that p. (72)

Or finally, Timothy Williamson (2005)

There is an intuitively attractive link between knowledge and practical reasoning. If it is not true for me to say 'I know that my door is locked', something is wrong with practical reasoning in which I rely on the premise that my door is locked, for in some sense I am not entitled to that premise, however much it seems to me that I am. If I have good evidence that 'I know that my door is locked' is true, although it is in fact false for reasons

beyond my ken, then I have good evidence that I am entitled to the premise that my door is locked, and therefore a good excuse for relying on it, but a good excuse for doing something is not an entitlement to do it. I may have good evidence that I am entitled to take books out of the library, and therefore a *good* excuse for taking them out, even though I am not in fact entitled to do so, because the rules have recently been changed. Conversely, if it is true for me to say ‘I know that my door is locked’, so far nothing is wrong with practical reasoning in which I rely on the premise that my door is locked, for I am entitled to that premise. In other words, ‘I know that my door is locked’ is true in my context iff that my door is locked is an *appropriate premise* for my practical reasoning. (227)

As we can see, a knowledge norm of action is not uncommon in contemporary epistemic literature. Many authors link knowledge norms to how an agent is obligated to reason, and how they are obligated to act (or refrain from acting). In the following, I will attempt to problematize the formulation of Hawthorne and Stanley’s knowledge norm, and any knowledge norm of action that obliges an agent to reason and act in a specific way.

1.2

WHAT’S AHEAD

In this introduction I described John Hawthorne and Jason Stanley’s knowledge norm of action. Together, the **ACTION-KNOWLEDGE PRINCIPLE (AKP)** and the **REASON-KNOWLEDGE PRINCIPLE (RKP)** provide necessary and sufficient conditions for how an agent ought to reason. Using premises that are known and relevant in *p*-dependent reasoning, Hawthorne and Stanley argue that their norm captures

how we criticize agents in cases where they act on the basis of something they don't know, or reason from extraneous premises.

In chapter 2 I extrapolate what Hawthorne and Stanley might mean by 'reasons', and introduce the vocabulary of subjective and objective reasons to help clarify the type of reasons that Hawthorne and Stanley are concerned with. Briefly I consider moral reasons, and conclude that Hawthorne and Stanley seemingly cannot simultaneously hold moral realism, subjective reasons internalism, and their knowledge norm of action.

In chapter 3 I introduce a case—**BRIDGE CONSTRUCTION**—that I argue provides a genuine counterexample to Hawthorne and Stanley's norm. I provide several ways that a proponent of Hawthorne and Stanley's project might try and amend my case such that Hawthorne and Stanley's norm arrives at the proper conclusion, and argue that in virtue of the mechanics of their norm, none of these options are available. I also conclude Hawthorne and Stanley *do* seem as though they espouse moral realism and subjective reasons internalism simultaneously, and, given their assumed commitment to their own **AKP/RKP**, seem to have inconsistent principles at work.

In chapter 4 I consider major and minor objections to my case and show where they go wrong.

Finally, I conclude that whatever a knowledge norm of action might look like, those which obligate an agent to act or reason a certain way, such as Hawthorne and Stanley's and many others in the literature, will be susceptible to the arguments that I've outlined and argued for in this thesis.

2

Reasons

Both Hawthorne and Stanley's **RKP** and **AKP** make reference to reasons but there is no subsequent discussion of what reasons might be, or whether or not they play any important role in their knowledge norm. Are we meant to understand reasons as reducible to propositions, and when they are known to you, it is appropriate for you to act upon them? Are reasons meant to be understood normatively? Must they have some motivational bearing upon an agent's state of mind? Are reasons just the same as evidence? Are reasons facts?

Hawthorne and Stanley's account provides us with a knowledge norm of action. Though they have provided the epistemological mechanics for their norm, the link between propositions, reasons, and action has not been fully given. Without straying too far from the focus of this thesis, this section will serve to

elucidate a very small portion of the literature regarding reasons relevant to this discussion.

Although there is an entire literature that focuses on what reasons might be and how we are to understand them,³ for the purposes of this thesis I shall try to offer an account of reasons that is both broad enough so as not to commit Hawthorne and Stanley to any particular account, but gives enough clarity as to what reasons are and how they work so as to be useful for the following discussion.

2.1

DEFINING REASONS

Consider this example from Michael Schroeder (2007a):

Tonight there is going to be a party, and everyone is invited. There will be good food, drinks, friendly chat, music—and dancing. Ronnie and Bradley, like everyone else, have been invited to the party. But while Ronnie loves to dance, Bradley can't stand it. Not only does he not like dancing, he prefers to stay away from where it is going on, lest he come under pressure to be shown up in his awkward maneuvers by those with fewer left feet than he. So while the fact that there will be dancing at the party is a reason for Ronnie to go, it is not a reason for Bradley to go. (1)

³ See for instance Alvarez (2010), Audi (2004), Berker (2007), Broome (2004, 2007), Chang (1997), Dancy (2000, 2004), Darcy (2003), Davidson (1963), Kearns & Star (2008, 2009), Parfit (2001, 2011), Rax (1999), Raz (1978), Schroeder (2007, 2007a, 2008, 2012), Setiya (2004, 2007, 2014), Setiya & Paakkunainen (2012), Shafer-Landau (2003), Smith (1987, 1994, 1995, 2006), Sobel (1999), Williams (1981)

It seems that the same fact—that there will be dancing at the party—is only a reason for Ronnie to go to the party, but not Bradley. Schroeder notes that a *prima facie* explanation for the difference in Ronnie and Bradley’s reasons has something to do with what they are like, care about, or want (2007:1). This seems to be a fairly intuitive first gloss of how reasons might reduce to the psychological state of an agent. When asked repeatedly, “Why did you do that?” one’s answers might be thought to always bottom-out in a reason grounded in a want or desire.⁴ For example,

DAN: “Why did you go to the store to buy oranges?”
 JULIET: “Because I was out, and the store had oranges.”
 DAN: “But why did you want to replace the oranges in the first place?”
 JULIET: “Because I like to eat them.”
 DAN: “Why do you like to eat them?”
 JULIET: “Because I like the way they taste.”
 DAN: “Why do you like the way they taste.”
 JULIET: “Because I do.”
 DAN: “But why?”
 JULIET: “...”

Here, there is no response that Juliet can give to further support her reason for going to get more oranges; she simply likes the taste of them. Although not

⁴ For instance, David Hume believed that you could not have an action without an associated desire (T 2.3.3). This is, in effect, a psychological thesis; an agent cannot be said to be doing something that is going to count as an action, unless there is something that the agent desires such that the action is something that the agent desires—or is a means to, part of, or constitutive of something that the agent desires. Hume’s isn’t a theory about the nature of reasons, just of actions. Aristotle thought something similar, in that he believed that every activity had to have a final cause which *the good* was aimed, since there could not be an infinite regress of merely extrinsic goods. (Nic. Ethics I 2) Most theories of motivation internalism subscribe to this belief as well.

all reasons are so easily reducible, it at least seems that there is *some* relationship between the reasons we have, and the things that we want.

Is this psychological relationship between an agent, her desires, and her known propositions adequately captured by Hawthorne and Stanley's **RKP** and **AKP**? One might argue that it is. Remember that **RKP** is concerned with p-dependent choices. Presumably if you know what your desires are (and for the sake of argument we'll assume that you do)⁵, then those desires are known propositions and may appropriately be treated as reasons. In chapter 1's **ORANGES** Juliet cited her lack of oranges and the store's supply of them as her reason for going to the store. Now we might understand her reasoning to include something about her desires as well:

ORANGES (2): Juliet has run out of oranges. When deciding what to do, Juliet reasons that she knows that the store will have more oranges and moreover that she would like to have more oranges in the house, because she likes how they taste. Because going to the store will be the best way to re-stock her orange supply, she goes.

Though we likely do not always consciously articulate our desires as reasons, it is clear that they play an important role in our actions. After all, it would be bizarre if Juliet didn't much like oranges and didn't care whether there were any in the house, and nonetheless went out to buy more. Just as it would be bizarre

5 For the purposes of this thesis, I will assume luminosity about desire-states. That is, I will allow that agents not only have access to their beliefs about their desires, but that they also have access to their beliefs about their beliefs about their desires. Although luminosity is contentious when it comes to knowledge, for the remainder of our discussion of reasons and desires, given that Hawthorne and Stanley make no comment to this effect, I believe this assumption is the most charitable and the least limiting to their position.

if Bradley cited the fact that there will be dancing at the party as his reason for going, despite the fact that he hated being around dancing. Of course reasons are not infeasible; Bradley might know that his friends always turn up at parties where there is dancing, and so the fact that there will be dancing at the party might well turn out to serve as his reason for going. But even so, that reason still revolves around some psychological state of Bradley.

2.2

REASONS YOU KNOW ABOUT AND ONES YOU DON'T

Consider a continuation to Schroeder's example (2008):

Freddie, like Ronnie, likes to dance. But Freddie, unlike Ronnie, knows that there will be dancing at the party tonight. (59)

Importantly for Schroeder are the parallels between Ronnie and Freddie, and the one salient difference between the two: Both Ronnie and Freddie like to dance, and there will indeed be dancing at the party. So as before, we might say that there is a reason for both Ronnie and Freddie to go to the party. However, because Freddie *knows that there will be dancing*, there is a sense in which the reason for Freddie to go to the party is importantly different from the reason for Ronnie to go to the party. Ronnie, we might say, doesn't *have* the reason to go the party, whereas Freddie does.

But in what sense is this a difference?

Although propositions are not things we *have* in the same way we might have a cold or a new paintbrush, still we might say that Freddie *has* a reason to go to the party in a way that Ronnie does not (even though there *is* a reason for Ronnie to go).⁶ Broadly speaking, we might say that reasons fall into two camps: reasons that you know you have, and reasons that you're ignorant of, but that bear upon you.

Consider another of Schroeder's examples:

Take the case of Bernie, who is standing around at a cocktail party, holding a glass that he believes to contain the gin and tonic for which he has just asked his hostess, but which is in fact full of gasoline. (2008:60)

Here, there are clearly two sorts of 'reasons' at play. One is a reason for Bernie to take a sip of his drink, and one is a reason against. The reason in favour of taking a drink is that Bernie has just ordered a gin and tonic and so presumably he would like to drink it, and that moreover he believes that his glass contains gin and tonic.⁷ The reason against Bernie taking a drink is of course that the glass actually contains gasoline (presumably not a substance he wants to drink). One is a

6 Early in his discussion, Schroeder parses this difference in terms of 'has' and 'for'. For example, he would say that there is a reason *for* Ronnie to go to the party, but it isn't a reason that he *has*. Whereas Freddie *has* a reason to go to the party (which also implies that the reason is also a reason *for* him to go).

7 It has always struck me as rather odd that gasoline is what has been poured into Bernie's glass, and not say, vodka. Why would there be gasoline behind the bar at a party? Why wouldn't Bernie immediately smell the gasoline? Although the example can still be understood, it might be worth mentioning that we are assuming that Bernie does not smell the gasoline and that its appearance in this example is merely to fill the roll of 'colourless liquid'.

reason that Bernie takes himself to have, and one is an objective fact of the matter that he is unaware of. So does Bernie have a reason to take a drink from his glass? For the purposes of our discussion we will say *yes* (though what sort of reason this is will be clarified shortly). Remember that AKP states that only known propositions may appropriately be treated as reasons. Despite the fact that objectively we know that what is in the glass is gasoline, and that if Bernie wants to quench his thirst he should not drink it, AKP prevents Bernie from using this fact in his reasoning.

For clarity, during the remainder of this discussion we will be concerned only with what I will refer to as, *objective* and *subjective* reasons. *Subjective* reasons count as reasons just in case they are beliefs an agent has (or second-order beliefs about their beliefs), that the agent might reasonably be expected to offer in response to questions in the vein of *what reasons do you have to be like that, why did you decide to do that, what made you act like that*, etc., and are ones that an agent ought to use in their reasoning, when appropriate.⁸ *Objective* reasons, we will say, are ones that an agent has just in case were they to know about them, they would count them as subjective reasons. To be clear, subjective and objective, as I'm using them here, map precisely onto Schroeder's *reasons for* and reasons that one *has*.

8 It is worth noting that this use of 'subjective' is not meant to imply that they are in any way ultra-subjective. This is not the same type of subjectivity that is at play when I say that my beer tastes good, and after taking a sip, you disagree with me. For the purposes of this thesis, subjective reasons are simply those reasons that an agent is aware that they have. It should also be noted that this distinction is made to try and accommodate Hawthorne and Stanley's description of *reasons*, specifically what they refer to as *personal reasons*. After all, you can't just *believe* that you have a reason (which might be one interpretation of an ultra-subjective reason), you have to *know it*. So there is an objective sense to the subjective reasons that we are talking about, because it is in fact true that the agent knows that they have the reason. However, given the familiarity many of us have with the distinction between subjective and objective, I have opted to keep these names as they are, and hope that this footnote serves to avoid potential confusion.

So, out of the two broad camps that reasons-talk might fall into, it seems that Hawthorne and Stanley are committed to situating themselves in the camp that concerns itself with subjective ones. Because Hawthorne and Stanley's program is concerned with action based on known propositions, objective reasons that an agent doesn't know simply cannot figure into the equation. Imagine how strange it would be if, in deciding whether or not to take a sip from his glass, Bernie cited his subjective belief that his glass contained gin and tonic, as well as the objective fact that it actually contained gasoline. Notwithstanding the irrationality of believing that the glass contained gin if he was somehow also aware that it actually contained gasoline, it's absurd to imagine any scenario in which Bernie could possibly reason using "the glass contains gasoline" as one of his premises.

2.3

YOU KNOW YOUR SUBJECTIVE REASONS

Fundamentally for Hawthorne and Stanley, it seems you need to know your reasons, and this is what turns an objective reason, into a subjective one that you can use in your reasoning. This is exactly what Hawthorne and Stanley's **RKP** says. They explicitly state:

More generally we need to distinguish between the existence of a reason for acting and appreciating that reason in such a way as to make it *your* reason for action (between mere rationalizers and motivators). As we are thinking about things, it is knowledge that constitutes the relevant sort of appreciation that converts the mere

existence of a reason into a personal reason. (2008:580)

Interestingly, even this simplified portrayal of reasons as objective and subjective dovetails nicely with another facet of Hawthorne and Stanley's argument that also does not get an extended treatment in their paper. Repeatedly, Hawthorne and Stanley use examples of our willingness to criticize agents in cases where they act on less than knowledge, and we can see an example of this with Bernie.

Imagine that we are watching Bernie silently deliberate about whether or not to take a sip of his drink. *We* know what his glass actually contains, and moreover, we know that *he* does not. What possible reason could Bernie have to *not* take a sip from his glass? He has not lost interest in his gin and tonic, did not merely order it so that he would not look awkward at the party, and he believes that what is in his glass is precisely what he ordered. If he did not take a sip, we might reasonably ask him, "Bernie, why haven't you taken a sip of your drink?", and out of all of his possible responses, we would *not* expect him to say that his reason against taking a sip has anything to do with gasoline. So it seems that Hawthorne and Stanley are also in a position to say that not only are agents deserving of criticism when they *act* on less than knowledge, but they might also be deserving of criticism when they *fail* to act on what they know.

But, as the case of Bernie and our discussions of subjective and objective reasons shows, something interesting happens when we act on the basis of subjective reasons which are at odds with objective reasons. This applies for both actions that we take, and ones that we refrain from taking. For example, Jessica Brown (2008)

defends the criticizability of *non-action* based on Gettiered-knowledge. She offers an example wherein she asks her partner to pick up potatoes on his way home. The variation on the example includes her partner's Gettiered belief that they are out of potatoes:

For instance, suppose that it's a system between us for me to leave lists of urgently needed items on the kitchen table right where he'll see them. Knowing we're out of potatoes I put 'potatoes' on the list on the table. After I leave, he finds a shopping list on the table and forms the justified belief that we are out of potatoes. However, unbeknown to us, the kids have decided to make up mock shopping lists for fun. What my partner picked up was one of their lists. As he has no reason to suspect that the kids are putting fake shopping lists around the house, he has a justified belief that we are out of potatoes. Further, the belief is true: we are out of potatoes. Discovery of the kids' prank in no way mitigates the negative assessment. For instance, it would be ridiculous for my partner to try to defend his action by pointing out that he didn't know that we lacked potatoes, but only had a justified belief that we lacked them. Further, it seems that after the kids' tricks have been discovered, I can still criticize his action by saying, 'You should have got potatoes. You thought they were on the shopping list'. (173)

Brown seems to be suggesting that even though her partner's justified, true, belief was *Gettiered* by the list that their kids made, because he had a subjective reason for stopping and getting potatoes but did not, his actions were deserving of criticism. Despite the fact that he lacked knowledge that they were out of potatoes, this sort of subjective reasoning seems as though it is just as well-suited for criticism (primarily *because* his belief was justified). If Brown's partner *believed* that they were out of potatoes, he seems just as deserving of criticism for not stopping as he would

have been had his knowledge not been foiled by his kid's list. Brown's case reads *prima facie* as a potentially problematic counterexample to Hawthorne and Stanley's knowledge norm of action. After all, it is a case where an agent is in compliance with RKP and AKP, and yet is deserving of criticism. In response, two things might be said.

The first, is that Hawthorne and Stanley are concerned with cases of deliberative reasoning.⁹ Their arguments do not discuss unconscious decision making, as presumably is the type of decision making employed when we are forgetful, or simply do something without thinking (as Brown's partner might have, if he simply auto-piloted home and didn't even think about stopping for potatoes). Further, Hawthorne and Stanley even admit that what is an appropriate way to reason might not yield what one pragmatically ought to do. In a response to an objection, Hawthorne and Stanley consider reasoning based on partial beliefs.

Suppose I am driving to a restaurant, when I come upon a fork in the road. I think that it is somewhat more likely that the restaurant is to the left than to the right. Given that these are my only options, and (say) I do not have the opportunity to make a phone call or check a map, it is practically rational for me to take the left fork. Yet I do not know that the restaurant is on the left. (2008:581)

9 It is worth noting that many of our cases of reasoning are done unconsciously, or at least in such a way that we don't actually consciously reason from individual premises to a conclusion. Although Hawthorne and Stanley do not elaborate on the constraints of deliberate reasoning, I take it that primarily the type of cases which are meant to be excluded are cases where no reasoning happened whatsoever, as in cases of reflexes, habitual behaviour, or simply acting without thinking.

This type of objection once again underscores standard counterexamples to Hawthorne and Stanley's norm.¹⁰ There seem to be many instances wherein one acts on less than knowledge in a way that seems entirely rational. Hawthorne and Stanley respond to this sort of objection by arguing that knowledge of chances constitutes knowledge. Therefore if the agent in the above case reasoned the following way: I know there is a higher chance that the restaurant is to the left rather than the right, therefore I ought to go to the left, they would be reasoning in accordance with the knowledge norm. The probability of the restaurant being on the left is both known, and relevant, thus is in accordance with both **AKP** and **RKP**.¹¹ So one way for Hawthorne and Stanley to respond to Brown's example would be to say that her partner reasoned from knowledge of chances (or should have so reasoned): seeing potatoes on the list makes it likely that they are out of potatoes, and thus makes it rational for him to reason in such a way that he should stop and get more.

However there is another way that Hawthorne and Stanley might respond to Brown. To avoid Gettier-related worries, consider an un-Gettiered variation of Brown's example:

10 For instance, see Ram Neta's (2009) paper that argues precisely for the success of these sorts of counterexamples.

11 They also respond to these sorts of objections by pointing out that what it is rational to do can be distinguished from what is appropriate to treat as a reason (581). Therefore in cases such as this (and other similar counterexamples) Hawthorne and Stanley suggest that they are simply very excusable violations of the knowledge norm rather than counterexamples to it.

POTATOES IN THE HOUSE: Suppose that it's a system between us for me to leave lists of urgently needed items on the kitchen table right where he'll see them. The kids, wanting to practice their 'adult' printing so that they could write notes to their teachers to get themselves out of class, write 'potatoes' as neatly as possible, on one of my shopping lists. My partner, finding the list, forms the belief that we are out of potatoes. However, there are plenty of potatoes in the house.

As in Brown's original case, her partner does not stop and get potatoes. However in this example, despite the fact that her partner had the same level of evidence that they were out of potatoes as in the previous case, would he still be deserving of criticism? Many I think, would have the intuition that he is. After all, assuming the kid's hand-writing was sufficiently legible (and we'll say that it was), he had no reason to think that they were *not* out of potatoes and every reason to think that they *were*. Despite the fact that Brown didn't require potatoes it seems perfectly reasonable of her to criticize her partner for failing to act on what he believed that he knew. After all, it would be strange of her partner to use as an excuse—when he found out about the forgery—the fact that they *weren't* out of potatoes. As in the Gettiered case, Brown could still criticize his action by saying, 'You should have got potatoes. You thought they were on the shopping list'.

What might Hawthorne and Stanley say of this example? Brown's partner couldn't have used '*we are out of potatoes*' as one of the premises in his reasoning, because he didn't know that they were. So on the face of it, it seems that Hawthorne and Stanley's **AKP** and **RKP** are unable to locate the source of the partner's criticizability. But what if we substituted, for '*we are out of potatoes*', the following:

1. When the list says that we are out of something, I should stop and get more of that item.
2. The list says 'potatoes'.
Therefore, I should stop and get more potatoes.

Each of these premises are known and relevant to Brown's partner's decision making. On this reconstruction, there *is* a way for both Hawthorne and Stanley—as well as Brown—to criticize Brown's partner's actions. The point to press is that the presence of an objective reason does not necessarily have to map onto what an agent ought to do (*qua* RKP and AKP), if a suitable substitute for their reasons can be made.

Consider,

Raj and Nik are starting a small business. The business is near a river that hasn't flooded in recent memory, but around which there isn't much flood protection. They could buy flood insurance which would be useful in a flood, naturally, but would be costly in the much more likely event that there is not a flood. Raj has done the calculations of the likelihood of a flood, the amount this would damage the business, the utility loss of not having this damage insured, and the utility loss of paying flood insurance premiums. He has concluded that buying flood insurance is the thing to do. As it happens, this was a good conclusion to draw: it does, in fact, maximise his (and Nik's) expected utility over time. (It doesn't maximise their actual utility, as there actually won't be a flood over the next twelve months. So, the insurance premium is an expense they could have avoided. But that doesn't seem particularly relevant for prudential evaluation. Prudential buyers of insurance should maximise expected utility, not actual utility. Or so we must say unless we want to be committed to the view that everyone who buys an insurance policy and doesn't make a claim on it is imprudent). (Maitra & Weatherston, 2010, p3)

In this case, I don't think we would find fault in Raj's reasoning, despite the fact that objectively there was no need to buy flood insurance. Although Raj's subjective reasons differed greatly from the objective reason, he is not criticizable in the same way that the partner in Jessica Brown's example is. On Hawthorne and Stanley's view, one might want to say that Raj is criticizable because he didn't *know* that it would flood, envisioning that he reasoned this way:

1. It know it is going to flood.
So, I should buy flood insurance.

But, as with Browns' partner, there is an appropriate substitute for Raj's reasoning. Instead, Hawthorne and Stanley might say that Raj reasoned this way:

1. Buying flood insurance maximizes expected utility.
2. I want to maximize our expected utility.
So, I should buy flood insurance.

Using this substitution, we see that Raj's reasoning complies with **RKP** and **AKP**.

These two examples are meant to show how objective and subjective reasons can come apart. It might be tempting—especially for the skeptically-inclined—to argue that anytime there is an objective reason that differs from an agent's subjective reasons, the agent will always violate **RKP**. However, using appropriate substitutions, we can still use **RKP** and **AKP** to explain away examples where it seems

that agents commonly act on less than knowledge. Returning to Hawthorne's previous example, despite the fact that many would argue that you cannot know whether your lottery ticket will be a loser, if, instead of using "I know this ticket is a loser" and deciding to use it as scrap paper for your shopping list, you instead use, "the expected value of using this ticket as scrap paper is higher than the chance of it being a winner", suddenly your reasoning no longer violates **RKP** or **AKP**.

2.4

MORAL REASONS

If only known reasons should be used in rational deliberation then this set of reasons will also cover moral ones. Although there are a number of ways in which we might delineate ethical theories, given that the proceeding (and upcoming) discussion revolves around moral facts and *reasons*, this section aims to give a high level overview of moral realism and reasons internalism.¹²

For the purposes of this thesis we'll simply say that moral realism affirms the following:

MORAL REALISM: There exist moral facts that exist independently of an agent's knowledge, beliefs, or desires.

¹² Of course there are views other than these two in the literature, however in the following chapter I will present a case that any of these other theories—whatever they are—will have to take a stand on. In order to be plausible, a moral theory will have to conclude that the agents in the case are not immaculate in their behaviour. I will be arguing that regardless of which moral view is applied to the agents, once that view concludes that they are not immaculate in their behaviour, the knowledge norm becomes incompatible with the theory. I will later argue in section 4.2.5 that only the most extreme relativist theories have the theoretical framework to avoid the criticisms that I will offer shortly.

Reasons internalist theories try to answer the question of what normative reasons there are for an agent to do—or refrain from doing—a certain act. This type of thesis is a completely general theory about reasons for action, so if it turns out that moral facts exist, then reason internalism will accommodate those as well as any other sort of fact that can serve as a reason for an agent. Although it is not primarily a thesis about motivating reasons, from the point of view of these sorts of theories, there is usually an important connection between normative reasons and motivation.¹³ However in Hawthorne and Stanley’s context, **AKP** and **RKP**—and knowledge norms of action generally—are not concerned with motivating reasons. Instead, these norms try to provide answers to questions of, “What reasons would count in favour of an agent ϕ -ing?”, rather than questions of, “Why did she do that?” This latter question targets an agent’s motivation, and is not relevant to our discussion of reasons. For the remainder of this discussion, all references to ‘reasons internalism’ will refer to reasons internalism as it applies in a moral context (that is, ‘moral reasons internalism’), and will exclude considerations of motivation.

Given the introduction of our subjective and objective reasons-language, I think it will be helpful to apply this classification to reasons internalism as well, to

13 A version of reason internalism that incorporates motivational internalism might be:

Reasons internalism affirms that if M is a moral fact and M calls for an act x in circumstances C then necessarily:

(a) $\forall A$: if A is an agent, then M gives A a *ceteris paribus* reason to comply with M in C

AND

(b) $\forall A$: if A is an agent, the M will *ceteris paribus* motivate A (in the sense that it would *ceteris paribus* **cause** A) to comply with M in C

allow these terms to map onto our previous distinctions and upcoming discussion.

REASONS INTERNALISM (OBJECTIVE): If M is a moral fact then necessary M gives an agent A a *ceteris paribus* objective reason to comply with M.

REASONS INTERNALISM (SUBJECTIVE): If M is a moral fact then necessary M gives an agent A a *ceteris paribus* subjective reason to comply with M, that is, a reason that they ought to use in their reasoning.¹⁴

Importantly for Hawthorne and Stanley is going to be the distinction between objective and subjective reasons internalism. In effect, the combination of moral realism and reasons internalism makes it the case that moral judgments are true if true, and false if false independently of the beliefs or attitudes of agents about the judgement or what is being judged, *and* give an agent (a subjective or objective) reason to comply with them. So if moral realism is true, it is wrong to poke someone in the eye, and doesn't matter whether anyone believes or knows that this is true, or even whether or not anyone likes that it is. It is still wrong. And—using our subjective/objective reasons-language—if objective reasons internalism is true, these facts will serve as reasons for an agent to behave, act, or reason a certain way, regardless of whether or not the agent knows about them (so if it's morally wrong to poke someone in the eye, you will have an objective reason not to poke someone in the eye, even if you are unaware of this reason). If subjective reasons internalism is

¹⁴ It is important to remember that in our terms a 'subjective reason' just *is* one that an agent ought to use in their reasoning.

true, then an agent is obligated to use moral facts in her reasoning (so if it's morally wrong to poke someone in the eye, this is a fact that you should use in your reasoning when deciding whether or not to poke someone in the eye).

So, if reasons internalism and moral realism are true, agents can be obligated to act in certain ways regardless of what that agent knows or believes. Further, they will be obligated to use moral facts in their reasoning, no matter what anyone believes about those moral facts. If this is the case, then Hawthorne and Stanley's **AKP** and **RKP** would almost always be in constant violation. Agents would either be reasoning from premises that they didn't know, or would fail to use premises that were relevant in their reasoning (assuming that agents are fallible when it comes to their knowledge of moral facts).

From this brief overview, it should be clear that Hawthorne and Stanley cannot hold moral realism *and* subjective reasons internalism. Remember that moral realism holds that a moral fact can be true even if an agent is unaware of it, and subjective reasons internalism *obligates* an agent to use subjective moral facts in her reasoning, when that moral fact is true. However **AKP** stipulates that only known propositions can be treated as reasons, which is in direct opposition to subjective reasons internalism.

We can assume that Hawthorne and Stanley are not giving up on **AKP**, and

there is evidence to suggest that they also affirm a version of moral realism.¹⁵ If **AKP** and **RKP** are true, then it can't be the case that an agent could be obligated to be, reason, or act according to a moral fact, if the agent was unaware of it. Recall that for Hawthorne and Stanley, if an agent is deserving of criticism for her actions, it is because she has violated a norm. In a case where an unknown moral fact obligates an agent to be, reason, or act in a certain way and she failed to do so, barring any other failures of reasoning, Hawthorne and Stanley would likely conclude that she was *not* deserving of criticism whatsoever, as far as the knowledge norm is considered. That is, her premises were known and p-dependent. Thus they will have to say that although moral facts might exist objectively of whether an agent knows, believes, or affirms them, and can serve as *objective reasons* (that is, reasons for an agent to do—or refrain from doing—a certain act), that these should not appear in an agent's reasoning unless known.

In the following section I will try and provide good reasons for thinking that in fact Hawthorne and Stanley's norm *commits* them to a subjective reasons internalists view. If this is the case, then their view turns out to be logically inconsistent when dealing with cases of moral reasoning.

15 For instance, consider what Hawthorne (2004) says in one of his footnotes: For comparison, most of us think that people in the past made many moral claims that were mistaken, but we don't think they did not know what 'good' meant. The fact that they thought that it was morally acceptable to have slaves, for example, or morally unacceptable for women to have equal civil status to men shows that they made moral errors (with alarming regularity)—but it does not show that they failed to understand the meanings of moral terms....This involves semantic ignorance *of a sort*. They did not know that 'bad' was true of slavery. And so on. But this hardly meant that they did not know what 'good' meant good and so on. (115)

3

The case

Standardly, criticisms of Hawthorne and Stanley's knowledge norm proceed by way of counterexample. These counterexamples articulate cases wherein an agent is proper to act on the basis of p , but does not know that p . These examples are rejected in the literature as simply being very excusable violations of the principle, but not counterexamples to it. Consider one of Hawthorne and Stanley's own examples :

Suppose Hannah's husband Mordechai has gone off to war, and goes missing for many years. Hannah remarries after waiting five years, reasonably assuming her husband to be dead. (2008:573)

As it turns out, Mordechai has not died, and returns home to find Hannah

remarried. Despite the fact that Hannah might want to defend herself against Mordechai's criticisms by saying, "But I waited for 5 years. There was nothing more I could do," Hawthorne and Stanley still hold that Hannah has violated the knowledge norm, as "Mordechai is dead" would have been one of the premises she used in her reasoning to decide whether or not to remarry, and happens to be false. That we might be understanding of Hannah's situation only moves her blameworthiness to the very excusable end of the scale, rather than absolving her of having violated the norm. Remember that violations, for Hawthorne and Stanley, exist on a spectrum from the very excusable to the very inexcusable. This is how they account for many of our folk appraisals and intuitions in instances where an agent has failed to satisfy the knowledge norm, but has seemingly 'done nothing wrong'. Such cases merely exist on the 'very excusable' end of the spectrum.

In order for a counterexample to be successful, it must be a case wherein an agent reasons and acts in accordance with Hawthorne and Stanley's **AKP** and **RKP**, and yet is still deserving of criticism.

Consider,

BRIDGE CONSTRUCTION: Two small, unique, and autonomous indigenous groups of people live on an isolated island. Each group is aware of the other, but they never make contact. One group, the **BUILDERS**, are constructing a bridge over a gully in the centre of the island. The other group, the **DWELLERS**, live at the bottom of the gully. In order to build the base of the bridge on one side of the gully, the **BUILDERS** will have to clear a path, which involves pushing a number of large boulders over the cliff into the gully. Because of the proximity and size of the population below, it is nearly certain that all of the **DWELLERS** will be killed, and this fact

is known to the **BUILDERS**. Having no knowledge of the moral fact ‘murder is wrong’, the **BUILDERS** reason this way:

1. We need to clear a path to build the bridge.
2. In order to clear a path, we must push boulders into the gully.
So we ought to push the boulders into the gully.

Structurally there is nothing wrong with the **BUILDERS**’ reasoning according to **RKP** and **AKP**. They have reasoned from premises that are known and relevant to their decision to clear the path, and yet something about it seems defective, given that their decision will cause the death of all of the **DWELLERS**. This seems similar to the way we might criticize historical figures such as Hitler for his decisions despite the fact that he would have, allegedly, thought he was doing the good. Likewise in **BRIDGE CONSTRUCTION**, despite the adherence to **RKP** and **AKP**, committing mass genocide simply to build a bridge seems improper. Regardless of the level of culpability we would ascribe to the **BUILDERS** for being ignorant of the relevant moral fact that ‘murder is wrong’ (perhaps we would ascribe very little), plausibly they are still blameworthy for causing the genocide of the **DWELLERS**. When reading this case, it might be tempting to say that the **BUILDERS** have reasoned incorrectly somehow, and yet we have seen that based on **AKP** and **RKP**, they haven’t. Despite the fact that the case explicitly states that the **BUILDERS**’ have no knowledge of the moral fact that ‘murder is wrong’, the disagreement with the **BUILDERS**’ conclusion seems to stem from a misfire between their premises and their conclusion. This misfire seems to be most easily accounted for by the fact that the **BUILDERS** have failed to take into account the well-being of the **DWELLERS** below, who will be obviously impacted by their decision to push the boulders into the gully. This suggests that on some

level we take rational reasoning to necessarily include certain facts, namely moral ones, when they are relevant to our reasoning. This most clearly accounts for the culpability we ascribe to the **BUILDERS**.

But what if someone were to deny this? In the case of Bernie and his glass of gasoline, when he takes a sip, we don't think that he is deserving of criticism. He is wrong (in a sense), to take a sip, because his glass does not contain the gin that he thinks it does (and that he wants to drink), but instead contains gasoline. Similarly one might think in **BRIDGE CONSTRUCTION**, although there is an objective reason that we are aware of for the **BUILDERS** to refrain from pushing the boulders into the gully, they are unaware of this. And so, like Bernie, they are not deserving of criticism. Let us quickly deal with this objection.

If we accept that *something* about what the **BUILDERS** have done is deserving of criticism, then it must be the case that they have violated the knowledge norm. Recall that Hawthorne and Stanley argue that when an agent violates the knowledge norm, they are necessarily deserving of criticism (2004:577). So if the impropriety of the **BUILDER'S** actions can't be found in their reasoning—given that Hawthorne and Stanley's norm is meant to govern action generally—then it can't be found anywhere.¹⁶

In **BRIDGE CONSTRUCTION**, the **BUILDERS** are deserving of criticism for their

¹⁶ A second and more anecdotal response is simply to point to the difference between these two facts: the glass contains gasoline, and murder is wrong. These certainly seem very different. Certainly there seems to be something worse about the **BUILDERS** committing genocide to build a bridge than for Bernie to take a sip of a drink that isn't his gin (barring all the nasty effects that a mouthful of gasoline are sure to cause).

reasoning, even if it's hard to say why, at least on first glance. However in its most crude form, Hawthorne and Stanley's knowledge norm simply states that one ought to use only known and relevant premises in one's reasoning. And indeed, the **BUILDERS** *have* reasoned from known and relevant premises (and haven't left any out). Given that there is no obvious violation of the norm, it is difficult to see why the **BUILDERS** are deserving of criticism and yet Hawthorne and Stanley insist that the fact that they are criticizable, means that they must have.

In the following sections I will examine what we might think of as the most natural ways to supplement the **BUILDERS'** reasoning to explain why they are criticizable in a way that is available in Hawthorne and Stanley's framework. In each instance I will show that Hawthorne and Stanley's account is unable to capture what is defective about the **BUILDERS'** reasoning.

3.1

THE ADDITION METHOD

A *prima facie* way of resolving **BRIDGE CONSTRUCTION** is by adding into the **BUILDERS'** reasoning what seems to be a very obvious (and absent) premise. Specifically, that the **DWELLERS** below will all be killed should the **BUILDERS** decide to push the boulders into the gully. As the description of the case states, this is a fact that the **BUILDERS** *are* aware of, and so it seems perfectly innocuous to simply add it into their reasoning to fix the problem.

Thus, we might think the **BUILDERS** ought to reason this way:

1. We need to clear a path to build the bridge.
2. In order to clear a path, we must push boulders into the gully.
3. If the boulders fall into the gully, the **DWELLERS** will all be killed.
So we ought not clear the path.

This reasoning seems far more appropriate, and intuitively the **BUILDERS** have arrived at the appropriate conclusion (that is, genocide is too high of a price to pay for the convenience of building a bridge). If this fix seems suitable, let us move on to see whether or not it can be accepted by Hawthorne and Stanley's knowledge norm.

This formulation of the reasoning in **BRIDGE CONSTRUCTION** now includes the fact that should the **BUILDERS** push the boulders into the gully, the **DWELLERS** will all be killed. Clearly the addition of premise 3 does not violate Hawthorne and Stanley's Action-Knowledge Principle (**AKP**), as it is indeed a fact that the **BUILDERS** know. What about the Reason-Knowledge Principle (**RKP**)? Recall that Hawthorne and Stanley describe **RKP** this way:

Where one's choice is p-dependent, it is appropriate to treat the proposition that p as a reason for acting iff you know that p. Our principle concerns what is appropriate to treat as a reason for action ... It would be overly demanding to require someone to treat all of their relevant knowledge as reasons for each action undertaken. (2008:578)

Does the addition of premise 3 violate **RKP**? The answer is yes. In order to treat p as a reason it must be relevant to the decision at hand. I presume that the

underlying motivation for thinking that the addition of premise 3 is so crucial is because *prima facie* it seems relevant. However **BRIDGE CONSTRUCTION** offers one other piece of information regarding the **BUILDERS'** epistemic states: they are unaware of the moral fact that 'murder is wrong'. Bearing this in mind it should become clear to see why premise 3 is not relevant to their decision. That the **DWELLERS** will all be killed is as peripheral to the **BUILDERS'** decision-making as is the current temperature on the island, or whether or not the lead **BUILDER** owns a dog. Without the associated moral fact there is simply nothing to tie the death of the **DWELLERS** to the construction of the bridge in such a way as to make it relevant for the **BUILDERS**. So while it is known (and thus allowable by **AKP**), it is not relevant and so is precluded by **RKP**. Even if premise 3 were to be improperly added to the **BUILDERS'** reasoning, it would still yield a defective conclusion:

1. We need to clear a path to build the bridge.
 2. In order to clear a path, we must push boulders into the gully.
 3. If the boulders fall into the gully, the **DWELLERS** will all be killed.
- So we ought to clear the path.

Premise 3 is an irrelevant fact. It would only become relevant in a way which could affect the **BUILDERS'** reasoning, if they were to become aware of the moral fact 'murder is wrong'. Thus were one tempted to simply violate **RKP** and insert premise 3 into the **BUILDERS'** reasoning, they would have offered no further aid to Hawthorne and Stanley, nor to the **BUILDERS'** level of culpability.

Although simply adding in the intuitive third premise has failed, there are perhaps still other ways to explain the culpability of the **BUILDERS** in **BRIDGE CONSTRUCTION**.

3.2

THE BACKGROUND CONDITION METHOD

In the previous section I stated that there was nothing to tie the death of the **DWELLERS** to the **BUILDERS**' reasoning regarding the construction of the bridge and the moving of boulders. Let us take up that idea in more detail. For many there will be a fairly obvious tie between premise 3 and the wrongness of the **BUILDERS**' reasoning. That the fact 'murder is wrong' is a moral fact is *precisely* what explains the impropriety of the **BUILDERS**' reasoning, regardless of whether or not they know it. The problem with the reasoning employed in the Addition Method, is that without knowing the relevant moral fact there was nothing to promote premise 3 to being a relevant fact for the **BUILDERS**, in the sense required by **AKP**. So perhaps what we ought to say about the **BUILDERS** is that even though they are ignorant of the moral fact 'murder is wrong', the mere fact alone has the ability to promote premise 3 to being relevant, whether or not it is known.

This should not seem like an outrageous suggestion. For instance we think that people are capable of violating and breaking laws even if they are unaware of their existence or jurisdiction. We might hold agents who violate them to a lower standard of culpability based on the extent and reasonableness of their ignorance

(or we may not), but the fact that certain legal facts were unknown to them only has the ability to diminish culpability, not exonerate it all together. This is precisely what the conjunction of moral realism and subjective reasons internalism affirms. Because ‘murder is wrong’ is true, it applies to the **BUILDERS** regardless of their knowledge, beliefs, or acceptance of it, and when relevant, ought to be used in their reasoning.

We might be tempted to apply this framework to **BRIDGE CONSTRUCTION**. Although the **BUILDERS** were ignorant of the moral fact ‘murder is wrong’, nonetheless that moral fact makes premise 3 relevant to their reasoning. Because ‘murder is wrong’ serves as a background premise—or inference rule—it has the ability to promote premise 3 above other mundane and innocuous facts. That the **BUILDERS** failed to take this into account in their reasoning regarding the construction of the bridge is precisely what explains our criticisms of them.

We think they ought to reason the following way:

1. We need to clear a path to build the bridge.
 2. In order to clear a path, we must push boulders into the gully.
 3. If the boulders fall into the gully, the **DWELLERS** will all be killed.
- ** (Background Condition) Murder is wrong.
So we ought not clear the path.

The **BACKGROUND CONDITION METHOD (BCM)** supplements what was lacking in the Addition Method, specifically a background condition that serves to promote premise 3 to a status wherein it becomes relevant and contentful. Should the

BUILDERS reason this way and nevertheless decide to clear the path, we would have a meaningful way of criticizing their actions.

Is such a method open to Hawthorne and Stanley?

Again the answer is no.

BCM takes advantage of a distinction between there being a reason that there is and there being a reason that an agent has (what I have been calling subjective and objective reasons). However exploiting this connection is not a move that is available to Hawthorne and Stanley, who argue that the *only* method of turning reasons that there are (objective reasons) into reasons that an agent has (subjective reasons), is knowledge. This should not come as a surprise, given the intimate connection that they believe exists between knowledge and action. They argue,

More generally we need to distinguish between the existence of a reason for acting and appreciating that reason in such a way as to make it your reason for action (between mere rationalizers and motivators). As we are thinking about things, it is knowledge that constitutes the relevant sort of appreciation that converts the mere existence of a reason into a personal reason. (2008:580)

As such, because the **BUILDERS** do not know the moral fact ‘murder is wrong’, it cannot convert the mere existence of the moral fact not to murder, into a reason that the **BUILDERS** have and can use as a premise in their reasoning. Thus for Hawthorne and Stanley, **BCM** cannot serve as a way of solving the impropriety of the **BUILDERS**’ decision to push the boulders into the gully thereby killing all of the

DWELLERS. Further, as mentioned in section 2.4 MORAL REASONS, it is inconsistent for Hawthorne and Stanley to hold simultaneously the following three views: moral realism, subjective reasons internalism, and AKP/RKP.

3.3

OBJECTIVELY *P*-DEPENDENT

Given that the *objective reasons* vocabulary does not make an appearance in Hawthorne and Stanley's characterization of their knowledge norm, perhaps they might update their formulation of RKP and AKP to involve objective reasons. This would give them the ability to criticize the BUILDERS as well as have their norm yield satisfactory results in similar cases where an objective—but relevant—reason is present. Presumably what Hawthorne and Stanley want is a way for the moral fact 'murder is wrong' to somehow appear in a way that is relevant to the BUILDERS' reasoning. The Addition Method shows that 'murder is wrong' cannot simply be inserted into the BUILDERS' reasoning as a premise, because it isn't known, and the BCM shows that it can't serve as a background condition, because even that would need to be known in order to be relevant. Hawthorne and Stanley might argue that 'murder is wrong' doesn't actually need to be a premise *or* a background condition: it's not a reason that the BUILDERS need to have (and so not a fact that they need to know), because it is a *relevance-maker*. Relevance makers, using our objective-reasons vocabulary, are what turn objectively relevant facts, such as moral facts, into things that we ought to take into account in our reasoning. Because they aren't

reasons themselves, we don't need to know them. Remember, Hawthorne and Stanley's principles only say that you need to know your *reasons*, not the things that make your reasons *relevant*.

Perhaps in light of the introduction of objective and subjective reasons, Hawthorne and Stanley might update RKP as follows:

THE OBJECTIVE REASON-KNOWLEDGE PRINCIPLE (ORKP): Where one's choice is **objectively** p-dependent, it is appropriate to treat the proposition that p as a reason for acting *iff* you know that p.

What this revised principle allows Hawthorne and Stanley to say is that the BUILDERS ought to include premise 3 (pushing the boulders into the gully will kill the DWELLERS) in their practical reasoning. Why must they? Because the moral fact 'murder is wrong' is now a *relevance-maker*. What does it make relevant? The fact that pushing the boulders into the gully will kill all the DWELLERS. Because the BUILDERS' choice to push the boulders is objectively p-dependant, the BUILDERS must treat 'pushing the boulders into the gully will kill the DWELLERS' as a reason. This treatment of Hawthorne and Stanley's principle ties moral facts to the BUILDERS' conclusion in a way that I think most find intuitively lacking from the original description of the BUILDERS' reasoning. Because 'murder is wrong' is now a relevance-maker, choosing whether or not to commit genocide simply to build a bridge suddenly becomes relevant. And because it does, the BUILDERS are required to treat the proposition that 'pushing the boulders into the gully will kill the DWELLERS' as a reason for acting. And of course this is something that the

BUILDERS do know, and so it isn't in violation of **RKP**. Thus, with **ORKP**, we have a way of criticizing the **BUILDERS** for drawing the conclusion to push the boulders into the gully without forcing them to use facts that are unknown or irrelevant in their decision making.

If this sounds like a tempting solution, take a moment to think about how absurd the implications of such a norm would be. What **ORKP** crudely says, is if something is objectively p-dependent, then it's relevant to your reasoning, regardless of what you think, know, believe, or want.

Consider the following case:

WHERE TO SIT: It is Juliet's first day of university, and as she enters her Introduction to Epistemology class, she tries to decide where to sit. She considers a number of reasons for sitting at the front (she wants to pay attention, and knows that she will be distracted if she sits at the back), but also considers reasons for sitting in the middle (she doesn't want to look like a keener, and she sees someone she might like to be friends with already in the middle). Seats are filling up quickly, and there is only one seat left in the middle row. Unbeknownst to Juliet and everyone else, a bomb will go off if the middle row of chairs in the classroom is filled.¹⁷

If **ORKP** is correct, then the fact that the middle row must have at least one empty seat to prevent a bomb from going off is relevant to Juliet's reasoning about where to sit. Juliet doesn't know this of course (but **ORKP** doesn't require that she does, since it's a relevance-maker). This means that she must reason the following way:

¹⁷ Thanks to Jeremy Fantl for suggesting this type of case to me.

1. If I sit in the middle row, the row will be filled
So, I shouldn't sit in the middle row.

AKP says that reasons must be known, and Juliet does indeed know that she will fill the row if she sits in the middle; a reason that she is now required to use in her reasoning because of ORKP. This should seem very odd. Despite the fact that Juliet *doesn't know* that a bomb will go off if she sits in the middle row of the classroom—because it is objectively relevant to her decision—she is required to use “If I sit in the middle row, the row will be filled”, as a reason. If you don't see the strangeness, consider Juliet's answer if asked, “So why did you decide to sit in the front row?” She might say, “oh, because I get distracted if I sit anywhere else”, but will also say, “and because if I sat in the middle row, the row would have been filled.” Puzzled by this, someone might say, “why does that matter?” To which Juliet would have no reply. Remember, neither ORKP or RKP say that you need to *know* what makes something relevant, your reasons simply need to *be* relevant. And, because of the presence of the bomb, ORKP says that because it is objectively relevant to Juliet's decision, it must be included in her reasoning despite the fact that *she has no idea that it is*. This strangeness does not occur with RKP because when you are not required to use objectively-relevant facts (that is, facts that you have no knowledge of) in your reasoning, you *are* aware of all the relevance-makers. They are just *subjective* relevance-makers, as they are known to you.

As Brian Ball (2014) puts it,

But reasons, one might think, are just propositions; one has them just in case one knows them. Accordingly [...] normative facts about what one ought to do are not independent of one's knowledge of them; it is impossible for propositions concerning what an agent ought to do to be true without being known by the agent in question. (555)

Despite how attractive **ORKP** might have seemed at first blush, the idea that an unknown fact with no epistemic relation could create reasons for an agent to do something unbeknown to them, is very odd indeed. Even if we allow **ORKP**, surely it must be the case that you need to know *that* there is a relevance-maker at play.¹⁸ Were that the case, Juliet could reason like this:

1. If I sit in the middle row, the row will be filled
 2. There is something that makes premise 1 relevant to my decision
- So, I shouldn't sit in the middle row.

But of course *that* wouldn't be allowed in virtue of *either* **ORKP** or **RKP**. As Michael Schroeder (2008) says,

If there are reasons for one to do something of which one is not aware, one can't be reasonably expected to do these things, reasonably criticized for not doing them, or do these things for those reasons. This much is true. On the face of it, this is because to have a reason to do something, something else must happen, over and above there being a reason for one to do it. One must also *have* this reason. (60)

18 For worries related to Lewis Carroll style of regress, see section 4.2.3

3.4 OVERRIDING ETHICAL NORMS

A *prima facie* way to respond to the arguments sketched above is to simply categorize them as defeating Hawthorne and Stanley's knowledge norm of action, such that they merely *override* action-norms, rather than constitute counterexamples. For instance many would be tempted to say that because knowledge norms *should* be defeasible, in cases of tension between ethical norms and knowledge norms, the knowledge norms necessarily yield to the ethical ones. This is the manoeuvre that Timothy Williamson (2000) uses in his discussion of the knowledge norm of assertion. In the face of counterexamples Williamson argues that,

Such cases do not show that the knowledge rule is not the rule of assertion. They merely show that it can be overridden by other norms not specific to assertion. (256)

However Hawthorne and Stanley explicitly stipulate that they will not be following Williamson in this move.

Here we see a strong analogy with the knowledge-assertion connection. Where one has every reason to think that one knows that *p* but does not, an assertion that *p* is quite excusable. But that is no objection to the norm that one ought to assert that *p* only if one knows that *p*. On the contrary, the need for an excuse in the

case is explained by that norm. *Mutatis mutandis* in the action-theoretic setting. (586)

But,

The Action-Knowledge Principle is a norm governing action generally. So it is not an option that it can be trumped by norms governing action in some more general sense. (587)

Some of the difficulties raised above would be removed if Hawthorne and Stanley were to simply allow the knowledge norm to be defeasible. Although they have stated that their norm is meant to govern action generally, they don't necessarily say why this is the case. In the following, I will try and provide motivation for why their knowledge norm can't (and shouldn't) be overridden.¹⁹

If violations of Hawthorne and Stanley's norm exist on a spectrum from the very excusable, to the very inexcusable, consider cases which occupy the lower end of the spectrum (very excusable cases): cases involving brains-in-vats. In ordinary situations we reason properly when we reason from premises that are known and relevant. Yet in brain-in-vat cases our envatted-counterparts fail to do so because all of their known premises turn out to be false, however excusably and unavoidably they have reasoned. If our envatted-counterparts became aware of their situation, they might offer an excuse to try and exonerate their reasoning ("There is no way

¹⁹ In section 4.2.4 I discuss what would happen if Hawthorne and Stanley's norm was overridden by ethical norms, or if was on par with them.

I could have known I was a brain-in-a-vat. All of my reasoning would have been acceptable otherwise.”). This is of the form often cited in counterexamples to Hawthorne and Stanley’s norm. There seems to be nothing more that we could expect our envatted-counterparts to do to ensure they were reasoning correctly. Thus, although they have reasoned using unknown premises, they have done so in a way seemingly undeserving of blame or criticism. However Hawthorne and Stanley maintain that precisely *because* they would offer an excuse in these cases, that demonstrates that the knowledge norm has been violated (albeit excusably), as excuses are only offered when a norm has been violated (580). However if knowledge norms can be overridden, brain-in-vat cases would no longer be classified as very excusable violations of the norm, but in compliance with it.

So why shouldn’t this be the case?²⁰

Hawthorne and Stanley’s principles exemplify the ideal and proper state that practical reasoners should seek to be in; this is precisely the state a norm should reflect. Therefore, anything which falls short—even in an excusable way—*should* violate the norm. In ethical situations (such as **BRIDGE CONSTRUCTION**), it is better to know all of the relevant facts (including relevant moral facts) than to fail to do so. This state of omniscience would constitute the ideal state for an agent to reason

20 It’s worth noting that most epistemologists *do* want a way of delineating normal reasoning based on knowledge from the reasoning employed in BIV situations. If we say that those reasoning from BIV situations are not even on the very excusable end of culpability, and are deserving of no criticism, then it is unclear why we should prefer to know things, as opposed to fail to know them in a way that isn’t deserving of criticism and is very excusable (as is the case with our envatted-counterparts). Ideal-norm states aside, it should be clear that Hawthorne and Stanley must maintain that BIV reasoning is deserving of criticism. To do otherwise would entail that an agent needn’t know their premises to use them in their reasoning, which is in direct opposition to what **AKP** says.

from. Thus if the knowledge norm is meant to capture ideal and proper reasoning and general action, this standard will permeate across all extraneous considerations (including ethical ones), and therefore cannot be overridden.

Consider,

WATER DRINKING: Stella the scientist is deciding whether or not to drink the liquid in her cup. It is clear and looks like water, but she is skeptical because her coworkers have a habit of playing tricks using glasses that are left unattended. So, Stella spends the next hour examining the liquid and running it through various tests in her lab. Confident that the liquid is in fact water (and it is) based on her practical reasoning (using the test results as her premises) she drinks it.

Berkeley the scientist is in an identical situation to Stella, and proceeds exactly as she does. The only difference is that Berkeley, regrettably, is a brain-in-a-vat.

As mentioned, this is a common way of formulating a counterexample to Hawthorne and Stanley's norm. Stella has reasoned properly, while Berkeley has failed to do so, however excusably. Yet there is literally nothing more that Berkeley could have done, or that we could have expected him to do, given his situation. As Hawthorne and Stanley note, this 'excuse' demonstrates that Berkeley has indeed violated the norm, but in a way that we find very excusable. Stella's situation exemplifies the knowledge norm in action, while Berkeley's falls slightly short. Those who argue against a knowledge norm try to show that Berkeley and other similar cases are genuine counterexamples: action based upon premises which aren't known in a way that doesn't warrant blame or criticism. In these instances,

Hawthorne and Stanley maintain that precisely because Berkeley would offer an excuse were his actions criticized (“There is no way I could have known I was a brain-in-a-vat. All of my reasoning would have been acceptable otherwise”), demonstrates that his reasoning was deserving of criticism for having violated the norm.

Another way these cases might be approached, is simply to consider which one is better. Given the choice, would you opt to be in Stella’s position, or Berkeley’s? Would you rather rely on premises that are known? Or ones that aren’t? In **BRIDGE CONSTRUCTION**, it certainly would be better if the **BUILDERS** were aware of the moral fact that ‘murder is wrong’. It is precisely because they don’t know (however excusably), that they’ve ended up in the situation that they are in. Thus, in ethical situations, it is better to know all of the relevant facts and how they are relevant to you (including moral facts) than to fail to do so.

Hawthorne and Stanley’s principles are meant to exemplify the ideal and correct state that practical reasoners should seek to be in; this is the state a norm should reflect. Anything which falls short, even in an excusable way, violates the norm. This is why we shouldn’t expect the knowledge norm to be overridden by extraneous considerations.

SUMMARY

Above I have sketched out a case that I believe provides a genuine counterexample to Hawthorne and Stanley's knowledge norm: a case wherein an agent is deserving of criticism, but nonetheless has reasoned in accordance with **RKP** and **AKP**. I have offered four potential ways that Hawthorne and Stanley might try to explain the **BUILDERS'** impropriety and have shown why they are unsuccessful within Hawthorne and Stanley's own framework. If it is the case that when an agent's reasoning is deserving of criticism it is because they have violated the knowledge norm (as Hawthorne and Stanley argue that it is), then it is unclear in **BRIDGE CONSTRUCTION** how to explain this impropriety.

If **BRIDGE CONSTRUCTION** provides a genuine counterexample to Hawthorne and Stanley's knowledge norm of action, then either one or both of their principles fail to capture some fundamental aspect of an action norm.

4

Objections and Responses

Among the various objections to the case that I have presented here is one that I believe stands out on its own as being the most pressing. In the following section I will try and flesh out this objection as I understand it, and provide reasons to think the objection is misguided and inert. Other more minor objections and their associated replies will also be given.

4.1

SHOULD HAVE KNOWN BETTER

When considering the BUILDERS and their actions in relation to the DWELLERS—regardless of how I have set up the case—it seems fairly natural to

want to offer the following defense for the BUILDERS: *They should have known better.* Despite the fact that BRIDGE CONSTRUCTION stipulates that the BUILDERS have no knowledge of the moral fact that murder is wrong, it still might seem to many that they *ought* to have known better, and that their perceived culpability derives from that fact.

Of course the brief and uncharitable answer to this challenge is to merely say that Hawthorne and Stanley's principles are concerned only with what an agent *does* in fact know and says nothing about what they *should* have known. RKP and AKP together crudely say that if you know a proposition and it's relevant to your reasoning, then it can be used. Therefore, whatever we might think about what the BUILDERS should or shouldn't have known about moral facts or the death of the DWELLERS, is completely irrelevant to Hawthorne and Stanley's knowledge-action scheme. However I am sure that this answer will strike proponents of this criticism as being too quick. So let us look at this suggestion more carefully.

Prima facie it seems to me that there are two ways that the BUILDERS could have 'known better'.

1. The BUILDERS should have known better in a way that is culpably ignorant
2. The BUILDERS should have known better in a way that is blamelessly ignorant

For the purposes of this discussion, let us say that culpable ignorance consists

in some sort of epistemic negligence on the part of the BUILDERS. An example of this might be the following:

TIGERS IN THE VILLAGE: The island where the BUILDERS and DWELLERS live is inhabited by a number of dangerous creatures, one of which are tigers. One night, one of the BUILDERS goes to bed without properly closing the door to his hut. In the morning he finds that his hut has been ransacked, all his food has been eaten, and his clothing has been torn. One of the elders scolds the BUILDER for being so careless. The BUILDER replies that, “I didn’t know there would be tigers. I’m not psychic.” To which the elder replies, “There were paw prints on the ground around the village, and tigers are often seen lurking in the shadows.”

In effect, the elder is citing the BUILDER’S easy epistemic access to the facts he had overlooked, noting that the BUILDER *should* have used them as reasons which could have guided his actions. This is an example of culpable ignorance: the BUILDER’S actions are deserving of criticism despite his claims that he didn’t know there would be tigers, because of the epistemic ease with which he *could* have known.

Blameless ignorance we will say, are cases wherein an agent is ignorant of a relative epistemic fact in a way that deserves no blame. An example of this might be Gettier cases, or brain-in-a-vat cases. In such instances, agents arrive at epistemically incorrect conclusions, and yet the methods which with the correct conclusions could have been arrived at are impossible to obtain.

So, either the BUILDERS are culpably ignorant of the moral fact ‘murder is wrong’, or they are blamelessly ignorant of it. If they are blamelessly ignorant of it,

then it would seem that the ‘should have known better’ criticism immediately falls short. We do not say of brain-in-vat or Gettier cases that the agents *should have known better*, because there is no sense in which the agent *could* have known better. Thus if it was impossible for the BUILDERS to have known the moral fact ‘murder is wrong’, then this does little to assuage them of the culpability that they seem owed.

If the BUILDERS fall subject to the ‘should have known better’ criticism, then it is because they are culpably ignorant of the moral fact that ‘murder is wrong’. Let us see if they are.

How would they know better?

BRIDGE CONSTRUCTION specifically states that the BUILDERS had no knowledge of the moral fact that ‘murder is wrong’. For them to be culpably ignorant of it in the way that we’ve defined, there must have been *some way* that they could have known. Were there perhaps signs of the moral fact, like signs of the tigers in TIGERS IN THE VILLAGE ? It doesn’t seem so. Although it is up for debate whether moral facts occur objectively of us or not, one thing we can agree on is that they aren’t literally strewn on the ground the way that the paw-prints of large cats are. They typically do not roar either. So if the BUILDERS could not literally stumble over the moral fact, how else might they have encountered it?

BRIDGE CONSTRUCTION also stipulates that the BUILDERS and DWELLERS are autonomous groups and that the island does not interact with the outside world.

Given that the BUILDERS see no moral wrong-doing with causing the genocide of the DWELLERS by building their bridge, we might even speculate their system of morality espouses that ‘murder is good’. Or perhaps they associate the killing of their brethren with the killing of animals for food, and see it only as a necessity and not a moral line in the sand. Taken this way, it is hard to see where or how the BUILDERS have been epistemically negligent. There were no obvious indicators of the moral fact, no sense or skill that they possessed which could have led them to it, nor anyone on the island with knowledge of it. That it is an objective moral fact might be granted, but without knowledge of it, it’s hard to see how we might expect it to come to bear on their decision making in a way that makes them culpably ignorant.

However there might be a way in which it could: *golden rule reasoning*.

Despite the fact that BRIDGE CONSTRUCTION explicitly states that the BUILDERS have no knowledge of the moral fact ‘murder is wrong’, some may find this stipulation untenable. Even though the BUILDERS are very primitive, one might argue that they *must* have some notion of murder and its *badness*. After all, murder is such a basic sort of wrong that many will find it hard to believe that even an exceptionally primitive and isolated group will have developed absolutely no conception of something resembling the moral fact that ‘murder is wrong’.

Imagine the BUILDERS as some extreme group of consequentialists. They do not view murder the same way as you or I (assuming you view murder to be wrong). Instead, they view murder as a morally unconstrained sort of choice; it’s simple consequentialist’ calculus. For the BUILDERS, these two scenarios would be precisely the same:

ICE CREAM: Bob and Marie are **BUILDERS** and are going for ice cream. Marie asks Bob what kind of ice cream he would like, and Bob shrugs saying that it's up to Marie. Marie comes out with two scoops of mint chocolate chip, a flavour that Bob doesn't really like. Still, it isn't like Marie has done anything wrong, it just would have been preferable to Bob if she had done something else. Bob says thank you and takes his ice cream.

BRIDGE BUILDING: A bridge is being built to cross a ravine. As construction is happening, an important tool—a saw made of teeth and a skull—suddenly breaks. It is impossible to continue the construction without this tool, and the **BUILDERS** have no more saws in their inventory. In order to make a new one, they require a full set of adult teeth and a skull. Unfortunately, no one has died recently. Because Bob is the newest **BUILDER** on the job and the one with the least amount of bridge-building skills, he is selected to be killed so that his teeth and skull can be used to create a new saw. Although the **BUILDERS** have done nothing wrong by selecting Bob to refresh their inventory, he would prefer that he not be killed.

With this sort of consequentialist response, we can see that 1. The **BUILDERS** do not view murder as a moral act, but merely the outcome of consequentialist reasoning, and 2. That although Bob would certainly prefer if he did not have to die, he does not view this decision as any more unmoral or unjust than he viewed Marie's ice cream selection. What is important to note is that the **BUILDERS** are not making any factual errors when it comes to their considerations of whether or not to murder Bob. They just have a fundamentally erroneous moral belief. Therefore, what we might think of as *golden rule reasoning* does not provide a solution for the **BUILDERS**, as precisely 'doing onto others what you would have done to you', is what they are doing to the **DWELLERS**.

OTHER OBJECTIONS

This section will serve to give a brief treatment to other objections which might immediately be raised against **BRIDGE CONSTRUCTION** as a counterexample to Hawthorne and Stanley's proposed **RKP** and **AKP** knowledge norm principles.

4.2.1 Substitution

OBJECTION: In response to my variation of Jessica Brown's potato case, **POTATOES IN THE HOUSE**, I suggested that Hawthorne and Stanley could find a way to criticize Brown's partner for violating **RKP** on the basis of something like proposition-substitution: instead of using "I know that we are out of potatoes", Brown's partner instead was using the propositions, "When the list says that we are out of something, I should stop and get more of that item" and "The list said we are out of potatoes" both of which were known by Brown's partner and relevant. Is such a move available in **BRIDGE CONSTRUCTION**?

RESPONSE: No. There is no appropriate substitution for the moral fact 'murder is wrong' that the **BUILDERS** know, or that is relevant to their decision whether or not to push the boulders into the gully. As is shown in my response to **4.2 SHOULD HAVE KNOWN BETTER** it isn't the case that the **BUILDERS** are unfamiliar with death, or that they're ignorant of the fact that their decision to build their bridge (by pushing

boulders into the gully) will cause the death of all the **DWELLERS** below. As I see it, there is no suitable substitution available to the **BUILDERS**, such that we could point to some other fact that they knew and could correctly reason from.²¹

4.2.2 Probability can do all of the work

OBJECTION: Why can't probability do all of the work? In the previous section (4.2.1 **SUBSTITUTION**), as well as in response to section 2.3 **YOU KNOW YOUR SUBJECTIVE REASONS**, mentions were made both of reasoning based on partial knowledge, and the acceptability of substituting reasons. Why not then say that while it's true that the **BUILDERS** don't know that 'murder is wrong', instead they're acting on the basis of partial knowledge (that is, of the *probability* that 'murder is wrong')?

RESPONSE: In **BRIDGE CONSTRUCTION**, it isn't as though the **BUILDERS** have some credence to think that 'murder is wrong', but have other reasons which swamp those reasons. Nor is it the case that the **BUILDERS** are acquainted with the fact that 'murder is wrong', but choose to disregard it, or are skeptical of it. Instead, they have absolutely *no credence whatsoever* that murder is wrong. Even if Hawthorne and Stanley's norm was a weak justification norm of action, it would still fail, as the **BUILDERS** aren't even *weakly* justified. In fact, we could just as easily construct

²¹ Two previous substitutions have already been considered; "The **DWELLERS** will all be killed" in **3.1 THE ADDITION METHOD**, and "murder is wrong" in **3.2 THE BACKGROUND CONDITION METHOD**.

a case where the **BUILDERS** had every reason to think that not only was pushing the boulders into the gully—causing the genocide of all of the **DWELLERS**—acceptable, but that it was actually encouraged. Perhaps they reasoned that by eliminating all of the **DWELLERS** there would be more food on the island for them, or perhaps they were taught from a very young age that the only way for **DWELLERS** to reach Heaven was to die at the result of boulders falling on their heads. However silly these examples might seem, I think they reinforce the idea that the **BUILDERS** need not have any credence at all in thinking that murder is wrong.

4.2.3 It's a regress

OBJECTION: The argument put forward in the section 3.2 **BACKGROUND CONDITIONS METHOD** seems to suggest that one ought to know the principles that implicitly link one's premises to one's conclusion. After all, what made it impossible to merely insert 'murder is wrong' into the **BUILDERS**' reasoning in **BCM** and what made Juliet's reasoning so bizarre in **WHERE TO SIT**, was a lack of relevance. Part of the reason why there was a lack of relevance had to do with the unknown linking principles working silently in the background. I've articulated them in previous sections to differentiate the way in which they might be employed: objectively-p dependent reasons, background conditions, and so on. One objection to this sort

of argument might be that it facilitates a Lewisian-style regress.²² After all, as Lewis quickly demonstrates, without knowing certain principles show that a given inference is valid, one can easily generate an infinite series of logical leaps which need to be made in order to arrive at a conclusion:

If you accept *A* and *B* and *C*, you *must* accept *Z*.”
 “And why *must* I ?”
 “Because it follows *logically* from them. If *A* and *B* and *C* are true, *Z must* be true. You don’t dispute *that*, I imagine ?”
 “If *A* and *B* and *C* are true, *Z must* be true,” the Tortoise thoughtfully repeated. “ That’s *another* Hypothetical, isn’t it? And, if I failed to see its truth, I might accept *A* and *B* and *C*, and *still* not accept *Z*, mightn’t I ?”
 “You might,” the candid hero admitted; “though such obtuseness would certainly be phenomenal. Still, the event is *possible*. So I must ask you to grant *one* more Hypothetical.”
 “Very good. I’m quite willing to grant it, as soon as you’ve written it down. We will call it (*D*) If *A* and *B* and *C* are true, *Z* must be true. Have you entered that in your note-book?”
 “I *have!*” Achilles joyfully exclaimed, as he ran the pencil into its sheath. “And at last we’ve got to the end of this ideal race-course! Now that you accept *A* and *B* and *C* and *D*, *of course* you accept *Z*.” “Do I “ said the Tortoise innocently. “ Let’s make that quite clear. I accept *A* and *B* and *C* and *D*. Suppose I *still* refused to accept *Z*?” (1995: 692)

And so it goes. Carroll’s paradox underscores an objection that might be raised against one of my proposed (and rejected) solutions for Hawthorne and Stanley (BCM), but also one of my assumptions in **WHERE TO SIT**; specifically that part of the ridiculousness of Juliet’s reasoning based on **ORKP**, was that there was

22 See Lewis Carroll (1995), What the Tortoise Said to Achilles.

an objective fact (what I called a relevance-maker) bearing on her rational decision making that she was unaware of. In both cases my criticisms were centered around the fact that an agent cannot be held responsible (that is, obligated) to act or reason in a certain way, based on principles that they do not know.

Instead of insisting that an agent *must know* the principles which link premises to a conclusion (and thus opening the door to a regress), why is it not enough that the linking principles/background conditions simply *be* connected to the conclusion, as a matter of fact?

RESPONSE: I will offer two suggestions.²³ One is specific to Hawthorne and Stanley's project, and the second is a broad commentary on how these cases bear on ethical theories at large.

As this objection pertains to Hawthorne and Stanley, I think it is appropriate to reiterate why the **BCM** failed to be a potential solution for them as a way of addressing **BRIDGE CONSTRUCTION**. Recall that the reason why 'murder is wrong' could not simply be added as a background condition to the **BUILDERS'** reasoning, was that without the appropriate linking principles between it, and the fact that pushing the boulders into the gully would kill all the **DWELLERS**, the background condition wasn't able to do any of the philosophical heavy lifting that it would

23 There is a third type of response that I might give. It is to delineate between cases like **BRIDGE CONSTRUCTION** and cases like **WHERE TO SIT**. The important difference between these two cases is that it seems one deals with an objective moral fact (**BRIDGE CONSTRUCTION**), and the other deals with a complicated logical entailment. In epistemology, we commonly refer to closure principles—that is, knowledge is either closed under known entailments, or it isn't. Ethics does not make this distinction, and so there seems to be a very real difference between our judgments in moral cases about what is entailed by an agent's set of beliefs or justification. The same does not hold in a logical setting.

need to. So while I think it an appropriate Lewisian response to simply say that—just as the laws of chess and logic are what they are, so too are the laws of basic rationalization and deduction—it isn't the case that such reasoning is only one level deep. I leave it unanswered whether or not one needs to know the linking principle between the moral fact 'murder is wrong', and a grander moral schemata, but certainly for Hawthorne and Stanley and their insistence on an agent using only known and relevant reasons, it seems clear from my rejection of a background condition such as 'murder is wrong' in section 3.2, that at *least* first-level linking conditions do need to be known.

The second sort of response that I'll offer has a wider net. This regress problem does not strike me as an epistemic problem, but a moral one. It is the obligation of an ethical theory to provide a suitable answer to questions of, "Why ought I ϕ ?", and perhaps just as importantly, to provide equally suitable answers to follow up questions of, "But why ought I do *that*?" If regress worries seem compatible with the arguments given here, that is simply a function of whatever ethical theory has been implicitly offered as the BUILDERS' and Juliet's ethical theory.

4.2.4 Ethical norms vs. knowledge norms

OBJECTION: One of the most tempting ways to alleviate the bulk of the criticisms I've levied at Hawthorne and Stanley's knowledge norm of action, is either to say that ethical norms can trump knowledge norms, or if not, that they are

at least on par with them. I'll consider each of the implications for Hawthorne and Stanley should knowledge norms be on *par* with ethical norms, or if they can be *trumped* by them. In both cases I'll argue that Hawthorne and Stanley are forced to accept some unhappy consequences as a result.

RESPONSE 1: IF ETHICAL NORMS TRUMP KNOWLEDGE NORMS. Hawthorne and Stanley's principles are meant to reflect the ideal and correct state that agents want to reason from. Even in ethical situations it is better to know all of the relevant facts (including potential background conditions and/or relevance-makers) than to fail to do so. Thus the ideal state across *all* considerations and subsequent norms, is going to be one where an agent acts on what they know.

Recall also that Hawthorne and Stanley explicitly argue that their knowledge norm of action could not be defeated:

The Action-Knowledge Principle is a norm governing action generally. So it is not an option that it can be trumped by norms governing action in some more general sense. (587)

Thus, by Hawthorne and Stanley's own hand, they will not allow their norm to be overridden. In section 3.4 I gave reasons to motivate why Hawthorne and Stanley are correct to say that their norm should not be overridden. If it is the case that this norm is meant to govern action *generally*—and it is also the case that a norm should reflect the *ideal state* that an agent reasons (or acts) from—then there should be no situations where another norm somehow offers a *more general* or *more*

ideal way for an agent to be.

RESPONSE 2: IF ETHICAL NORMS ARE ON PAR WITH KNOWLEDGE NORMS. If ethical norms are on par with knowledge norms, dilemmas are sure to follow. Consider the case at hand: the knowledge norm of action requires that the BUILDERS act on only known premises—thus they are permitted to push the boulders into the gully. However moral norms require that they do not do this. If moral norms of action and knowledge norms of action are on a par, the BUILDERS are normatively irredeemable in some way. They both are required to use as a premise something they don't know (and/or don't believe), but if they did, it would be entirely bizarre. They can't do the morally right thing because that would be to act in accordance with a principle they don't know or believe, and so they would end up acting irrationally. But as soon as they limit themselves to the reasons that they're rational to use (those that they know and know to be relevant), they end up doing something morally wrong. They cannot help but go astray.

If this is the case—and this knowledge norm of action is meant to cover *all* spectrums of action—then one would assume it could handle cases of dilemmas, where one is epistemically damned if you do, but morally damned if you don't. That Hawthorne and Stanley's norm seems unable to handle cases such as this seems to suggest that it fails insofar as it purports to be a norm governing action generally.

As Kvanvig (2011) argues,

What we do not want and cannot tolerate is multiple answers to the questions of what to do and what to think. If a theorist says,

“well, if you do A you will be justified in so doing, and if you refrain, it will be excusable,” the appropriate reply is simply to repeat the request: tell me what to do. (243)

In addition to the fact that Hawthorne and Stanley’s ‘general’ norm fails to provide an answer to the question of “what should I do?”, in cases such as **BRIDGE CONSTRUCTION** their norm is forced to provide an immensely unsatisfactory and deeply unintuitive characterization of the case. *If* we allow that there was no impropriety in how the **BUILDERS** reasoned, and *if* it is the case that Hawthorne and Stanley’s is a norm that governs action generally, then we are left with the conclusion that despite the fact that the **BUILDERS** have committed genocide simply to build a bridge, no violations—ethical or otherwise—have transpired.

4.2.5 Can the relativist escape?

OBJECTION: *Prima facie* it seems clear that the arguments I’ve provided apply to accepting both moral realism and subjective reasons internalism. But what about theories of moral relativism? For the moral relativist, moral truth (crudely) is just a matter of being entailed by your values, and if something is entailed by your values, one would think that it’s knowable by you. If ‘murder is wrong’ isn’t known by the **BUILDERS**, then it likely isn’t entailed by their beliefs. And if it isn’t entailed by their beliefs, it isn’t morally obligatory for them. Prior to my response in 4.1 **SHOULD HAVE KNOWN BETTER** I tacitly argued that the **BUILDERS** had absolutely no conception

of the moral fact 'murder is wrong'. If that's the case and there's no inference that the BUILDERS can make from any of their attitudes to the conclusion that 'murder is wrong', then it seems as though moral relativist accounts can escape the previous criticisms quite easily.

RESPONSE: I think there are two responses to this sort of objection. The first and quickest response is to say simply that neither Hawthorne nor Stanley are moral relativists, and so this particular response is not open to them. Although it is open to a relativist to adopt a knowledge norm of action, the type of relativist that espouses the type of theory described above strikes me as a very rare breed. After all, if it is the case that relativistic moral truths are entailed by the beliefs of an agent, that suggests that an agent can never be in moral error. They might commit a morally bad act, but they cannot reason in such a way that their *conclusion* is in error, if it really is the case that moral truths are *entailed* by an agent's beliefs.

Most relativists do not think that moral truths are closed, that is, just because a belief of yours entails a claim about moral truth, does not necessarily mean you know that it does. Theories of this second type instead say that despite the relativistic nature of morality, it can still be difficult to figure out what is right and wrong for you to do. Other types of views which adopt an 'anything goes' stance of morality will of course be immune to my arguments here. However I take it given the many arguments to the contrary found in the literature, the ability for such a theory to take advantage of a Hawthorne and Stanley-type norm does not likely add much appeal.

SUMMARY

In the above sections I have tried to respond to the bulk of major and/or obvious objections that might be offered in response to my criticisms of Hawthorne and Stanley. Although initially I think many will believe that there are a few quick fixes that could be applied to their norm such that these fixes would easily render it functioning and free of my allegations, I hope that this section has shown that the majority of these—if not all of them—fail to adequately defend Hawthorne and Stanley’s account against my arguments and counterexample.

One type of ‘fix’ that was not mentioned above has to do with the special-making features of moral reasons (whatever they might be).²⁴ This is because I don’t think it’s actually an objection, but simply an ‘update’ that someone might argue is open to Hawthorne and Stanley. The update might go something like this: Realist theory or not, there seems to be something important about moral facts that distinguish them from regular sorts of facts. If the crux of Hawthorne and Stanley’s difficulty in responding to **BRIDGE CONSTRUCTION** lies in their norm’s

24 One reason that moral facts might be different in kind from other sorts of facts, is that they have normative power built into them. If this is true, then the number of reasons that an agent *has* (and reasons *for* that agent) to act morally, diminish. If you have a moral duty of easy rescue and see a child drowning, then you only have one reason to save the child: that moral duty of easy rescue. On this sort of picture, it would be nonsensical for an agent in this situation to say, “I know I have a duty of easy rescue, but why *ought* I rescue the child?” as if they were looking for some sort of further motivating reason to actually carry out their moral duty. This would explain why *prima facie* there seems to be something different between these two facts: *it is morally wrong to murder* and *the distance from Kamloops, British Columbia and Calgary, Alberta is 619 km*. The first fact does seem to suggest that an agent ought to be, reason, or act in a certain way, whereas the second fact suggests no particular epistemic or moral state of mind that the agent ought to take.

inability to accommodate moral reasons, then why not simply load the ‘special making feature’ into p-dependence, making some sort of tacit connection between the moral reason and an agent’s conclusion? As attractive as this might sound, I am skeptical that it is even open to Hawthorne and Stanley. For one thing, I don’t see how they could insert this ‘special making feature’ into p-dependence in such a way that my criticisms of—and the weirdness that resulted from—**ORKP** wouldn’t result once again. *Prima facie* someone might want to say that because moral reasons are slippery—in that they are theoretically hard to pin down—it isn’t surprising that when they subjectively (or objectively) enter into our reasoning, the premises and conclusions which are yielded are similarly slippery.

Unfortunately I don’t think that this sort of reasoning holds water for a few reasons. The first being that ethicists do not normally throw up their hands when asked, “How do moral facts get their normative statuses?” Ethical theories provide answers to precisely these sorts of questions. The second reason is that if Hawthorne and Stanley’s norm is really meant to be a general action-guiding norm, then it ought to be able to answer the same sorts of questions that ethical theories purport to, namely, “What ought I do to?” and “What ought I to do when I don’t know what I ought to do?” If Hawthorne and Stanley simply agree that cases involving moral reasons are going to be complicated and difficult to decipher, and so we should expect that how an agent ought to reason and act to be similarly complicated and difficult to decipher, it seems to me that their norm shouldn’t suitably be called general at all, if it only yields clear answers in mundane epistemic cases of restaurant directions and lottery tickets.

5

Conclusion

If I'm correct, the arguments that I've proposed in this thesis will apply to any knowledge norm of action that obliges an agent to reason and act in a specific way.²⁵ I have demonstrated how these arguments apply to Hawthorne and Stanley's principles, but there are others in the literature who formulate a similar style of knowledge norm of action. For example, consider Fantl and McGrath's (2009) norm:

25 Early on, I cite John Hyman in the camp of philosophers with similar knowledge norms of action. It should be mentioned that Hyman's account is a motivational account of knowledge and reasons. For Hyman, an agent must be *motivated* to act on the basis of their reasons. As such, many of the arguments levied above are not as applicable to his view. Without knowledge of the moral fact that 'murder is wrong' it is hard to see how the BUILDERS could possibly be motivated by it. That being said, whatever special-making features moral facts possess, it is possible that Hyman's view will be similarly affected by my arguments if it turns out that agents *ought* to be motivated by certain kinds of moral facts. See for instance Derek Parfit's agony argument in his 2011 *On What Matters*.

JJ: You are justified in believing p iff p is warranted enough to justify you in ϕ -ing, for any ϕ . (123)

Despite the fact that Fantl and McGrath's norm focuses on justification rather than knowledge (as in Hawthorne and Stanley's case), I believe that the arguments I've put forth are still applicable. I take what Fantl and McGrath's JJ is saying, is that if p is warranted enough to justify you to do something, then you are knowledge-level justified in believing it. And knowledge-level justification, for Fantl and McGrath is an obliging notion: if p obliges you to do something, then you have knowledge-level justification for it.²⁶ In **BRIDGE CONSTRUCTION**, the **BUILDERS** are obligated to use in their reasoning the fact that pushing the boulders into the gully will kill all of the **DWELLERS** below, and thus also the moral fact that 'murder is wrong' (for arguments given earlier). Given Fantl and McGrath's bi-conditional, if the **BUILDERS** are obligated to use 'murder is wrong' in their reasoning, then they have knowledge-level justification of it. But in **BRIDGE CONSTRUCTION** it isn't just that the **BUILDERS** fail to have *knowledge* level justification, they don't have *any* level of justification for it.

Similarly, Williamson offers his own norm:

KPR: A first-person present-tense ascription of 'know' with respect to a proposition is true in a context iff that proposition is an appropriate premise for practical reasoning in that context. (2005:227)

26 For a more detailed discussion, see Fantl and McGrath's *Knowledge in an Uncertain World*, pp. 110-111

Williamson's bi-conditional states that if something is an appropriate premise for practical reasoning in a given context, then that proposition is true for the agent in that context. It should be clear now why there is good reason to think that this isn't the case. Although it is appropriate (and even obligatory) for the **BUILDERS** in their context to use the complete demise of the **DWELLERS** as one of their premises in practical reasoning with the moral fact that 'murder is wrong', they cannot, even though it is true for them. Thus Williamson's bi-conditional fails: it isn't the case that what is an appropriate premise is true for the **BUILDERS**. Although Williamson does not go into a related defense of **KPR**, he does acknowledge that if it fails to hold, then so too might his arguments for 'shifting standards' related to sensitive invariantism.²⁷ Thus, I argue that because his conception does indeed fail for the precisely the same reasons as those that apply to Hawthorne and Stanley and Fantl and McGrath, then his arguments against the contextualist and the subjective sensitive invariantist are likely on less stable grounds.

WHAT A KNOWLEDGE NORM MIGHT LOOK LIKE

Although it is not the purpose of this thesis to provide a positive account of

27 Williamson (2005) writes "Although appealing, KPR is not self-evidently correct. If it fails, so does the envisaged strengthening of the argument for shifting standards, and the insensitive invariantist's task is correspondingly easier. More generally, the more the practical is separated from the epistemic, the easier it is to reply to arguments from practical differences to shifting semantic standards for epistemic terms. A full evaluation of KPR will not be undertaken here." (228)

what a knowledge norm of action might be, given the proceeding discussion there seem to be a few key elements that should be taken into account, if it is to avoid similar criticisms.

A norm should govern action generally, or be defeasible

As was argued in section 3.4 OVERRIDING ETHICAL NORMS and objection 4.2.4 ETHICAL NORMS VS. KNOWLEDGE NORMS, Hawthorne and Stanley's norm fails to be a norm that adequately governs action generally. There are two routes I could see an appropriate knowledge norm of action following: be defeasible, or govern action generally.

Timothy Williamson in *Knowledge and Its Limits* (2000) argues for a knowledge norm of assertion. However unlike Hawthorne and Stanley's norm, Williamson readily admits that his norm *can* be overridden by norms not general to assertion:

Sometimes one knows that one does not know *p*, but the urgency of the situation requires one to assert *p* anyway. I shout 'That is your train', knowing that I do not know that it is, because it probably is and you have only moments to catch it. Such cases do not show that the knowledge rule is not the rule of assertion. They merely show that it can be overridden by other norms not specific to assertion. The other norms do not give me warrant to assert *p*, for to have such warrant is to satisfy the rule of assertion. Similarly, when I am speaking a foreign language, the urgency of the situation may require me to speak ungrammatically, because

it would take me too long to work out the correct grammatical form for what I want to say; it does not follow that my utterance satisfied the rules of grammar in that context. (257)

Such an attitude, I think, is likely to be the most pragmatic. And, given that epistemologists are generally concerned by the appropriate/inappropriateness, right/wrongness, good/badness, obligatory/permissibility of epistemic-decision making and action, having a knowledge norm of action that is defeasible allows for other norms to step in and provide explanatory power when necessary, but continue to allow epistemologists to evaluate actions and reasoning in their own terms.

If a knowledge norm is meant to govern action generally however, such is a tall order. As this thesis hopefully shows, there are a myriad of factors that purport to determine the rightness/wrongness, goodness/badness, appropriateness/inappropriateness, obligatory status, etc., of a given deliberation and/or corresponding action. To have a norm that captures them all, and does so appropriately and fairly to all other disciplines involved, is surely to be a more complex task than two necessary and sufficient principles can account for.

In another paper, Williamson (2005) captures this sentiment precisely:

Suppose that an agent relies on the premise q in otherwise impeccable practical reasoning, and (in her context) q is indeed appropriate, but she is in no position to know that q is appropriate. Then our assessment of her reasoning should be ambivalent: it is good that she used an appropriate premise, bad that she was in no position to know that it was appropriate. No simple verdict captures the complexity of the position. (230)

A norm shouldn't confuse agent appraisals and act appraisals

Although it was not the scope of this thesis to cover a discussion of agent appraisals and act appraisals, it strikes me as incorrect that Hawthorne and Stanley— without argument—pair an agent's criticizability with a violation of their norm. Recall that they argue that “When someone acts on a belief that does not amount to knowledge, she violates the norm, and hence is subject to criticism” (2008:577). Hawthorne and Stanley are connecting an agent's action (a deontic appraisal) with their evaluation of the agent (a hypological appraisal).²⁸ As Ryan Hebert (2016) argues,

The apparent goodness or badness of reasons in virtue of which an agent deserves credit or criticism might be merely apparent. It need not track their reality. In point of fact, hypological appraisals do not track deontic appraisals for fallible agents: it is possible for an agent to deserve credit for an objectively impermissible performance and it is possible for an agent to deserve criticism for an objectively permissible performance. This is because fallible agents do not have privileged access to the actual quality of reasons. Rather, they must make due with mere appearances, which but imperfectly track reality. (4)

In effect, Hebert argues that there need not be any necessary link between what an agent does, and what sort of criticism or credit that agent deserves. For

²⁸ For an in-depth understanding of these sorts of appraisals and a discussion of the impropriety of linking agent and action appraisals, see Ryan Hebert (2016).

instance, I might fail to stop at a stop sign, but I might do so because otherwise I would have run over a puppy. Rather than incorrectly confusing my action with my reasoning as an agent, we can pull these appraisals apart, and praise me as an agent, while still condemning my action as being in violation of a law.

Understanding that there is no necessary link between agent and act appraisals also makes understanding norms far simpler. As Boyd (2014) argues,

A practice that involves a rule linking assertion to knowledge but allows minor violations of the rule is more complex than a practice that is based on a less demanding rule...In a practice of the former sort there will have to be a convention allowing excuses together with one or more rules specifying the gravity of various forms of infraction. Moreover, each participant in such a practice will have to keep track of the various psychological and epistemic factors that determine whether particular infractions should be excused. (110)

Douven (2006) argues in the same vein:

A first remark is that it seems simpler, and thus methodologically preferable, to explain our intuitions about false but reasonable assertions without having to appeal to an extra story about how one can breach rules blamelessly, or about there being two senses of proper assertability. In fairness, however, it should be said that we will need a story about breaching rules blamelessly anyway, regardless of what we are going to say is the rule of assertion. To give an example already pointed to, someone who is knowingly asserting a falsity in order to save her life has breached the rule of assertion both on the knowledge account and on the rational credibility account-but no one will think she deserves any blame for her act. (478)

If agent appraisals and act appraisals are pulled apart, not only is the resulting norm a much simpler one for agents to adhere to, but it would also provide a much more clear and concise framework for philosophers to evaluate.

An agent should be able to violate a norm, even if she is unaware of it.

When it comes to a knowledge norm it might seem rather bizarre that despite the fact that reasons, relevance makers, background conditions, etc., must all be known by an agent, the norm itself need not. Recall that nowhere in our preceding discussion were any of the agents involved criticized for not reasoning from the premise, “I ought to use (a principle like [RKP]) and (a principle like [AKP]) in my reasoning.” Despite any initial discomfort with this concept, I don’t think we should find it that strange at all. Norms, like laws, apply in a given context regardless of whether or not an agent is aware of them.

As Kvanvig (2011) argues, many do not think that it’s the case that norms are luminous:

According to Stanley and Hawthorne, whenever we have a norm, there will be cases on which a person unwittingly violates the norm, leading to an appearance of objection to the claim in question, an objection that is claimed to dissolve once we notice that every norm is subject to such a problem and that can be explained away by bringing a secondary notion into the picture. The reason such apparent problems are unavoidable is because norms are not

themselves “luminous”: we are not infallible about whether the conditions specified in the norm have themselves been satisfied... The argument that such a distinction is unavoidable begins with the denial of luminosity. The denial of luminosity is a reasonable assumption: there is very little about which we cannot reasonably be mistaken, and I will assume here that such fallibility is universal. (240)

Further, the idea that we ought to be more forgiving of agents who unknowingly violate a norm confuses the purpose and scope of the norm. If agent appraisals and act appraisals are appropriately distinguished (as in the proceeding section I suggest that they should be), then there is no need to entertain notions of what Keith DeRose (2009:94) and others have referred to as ‘primary’ and ‘secondary’ propriety. As Engels (2008) argues,

[T]here is only one way of violating a norm: simply by not doing or believing what it prescribes. The fact that I am unaware that I crossed the yellow line does not make me less liable to receive the fine from the policeman who notices my bad driving behaviour. That may excuse me, but that does not change the fact that I have violated the rule. The fact that one is unaware of violating the norm does not change one’s status with respect to the norm. (11)

Despite the initial attractiveness of primary and secondary propriety, it is simply no longer required when it comes to adjudication of a norm violation. It is not necessary if we delineate our appraisals of agents and actions, rather than forcing them to rise and fall together, nor does it seem entirely appropriate. After all, when I plead my case in traffic court for running a red light, there is no ‘secondary’

sense in which I am free of culpability despite the fact that I have still broken the law. Either I am guilty, or I am not, and whether or not I am depends on if I have broken a rule of law. My sentencing might be lowered based on other features, but whether or not I have broken the law (or in our case, violated a norm), focuses solely on my actions.

Bibliography

- Alvarez, Maria. (2010). *Kinds of reasons*. Oxford: Oxford University Press.
- Audi, Robert. (2004). *The right in the good*. Princeton, NJ: Princeton University Press.
- Ball, Brian. (2014). The Knowledge Rule and the Action Rule. *Southern Journal of Philosophy* 52 (4):552-574.
- Berker, Selim. (2007). Particular reasons. *Ethics*, 108(1), 109–139.
- Broome, John. (2003). Reasons In R. J. Wallace, M. Smith, S. Scheffler, & P. Pettit (Eds.), *Reason and value: Themes from the moral philosophy of Joseph Raz*
- (2004). Reasons. In R. J. Wallace, M. Smith, S. Scheffler, & P. Pettit (Eds.), *Reason and value: Themes from the moral philosophy of Joseph Raz* (pp. 28–55). Oxford: Oxford University Press.
- (2007). Does rationality consist in responding correctly to reasons? *Journal of Moral Philosophy*, 4(3), pp. 349–374.
- Boyd, Kenneth. (2015). Assertion, practical reasoning, and epistemic separabilism *Philosophical Studies* July 2015, Volume 172, Issue 7, pp. 1907–1927
- Brown, Jessica. (2008). Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning. *Noûs* 42: 167-89.
- (2008a). The Knowledge Norm for Assertion, *Philosophical Issues*, vol. 18, no.

1, pp. 89-103.

Carroll, Lewis. (1995). What the Tortoise Said to Achilles. *Mind*, 104(416), new series, 691-693

Carter, J. Adam & Emma Gordon. (2011). Norms of Assertion: The Quantity and Quality of Epistemic Support, *Philosophia*, vol. 39, no. 4, pp. 615-635.

Chang, Ruth. (1997). *Incommensurability, incomparability, and practical reason*. Cambridge, MA: Harvard University Press.

Cohen, Stewart. (2004). Knowledge, Assertion, and Practical Reasoning, in *Ernest Sosa and Enrique Villanueva (eds.), Philosophical Issues* 14: 482–91

Dancy, Jonathan. (2000). *Practical reality*. Oxford: Oxford University Press.

——— (2003). What Do Reasons Do? In *The Southern Journal of Philosophy* (2003) Vol. XLI. pp. 95 - 113

——— (2004). *Ethics without principles*. Oxford: Oxford University Press.

Davidson, Donald. (1963). Actions, reasons, and causes. *Reprinted in Essays on actions and events* (pp. 3–19). Oxford: Oxford University Press, 1980.

DeRose, Keith. (2002). Assertion, Knowledge, and Context, *Philosophical Review*, vol. Ill, no. 2, pp. 167-203.

——— (2009) *The Case for Contextualism* Oxford: Oxford University Press.

Douven, Igor. (2006). Assertion, Knowledge, and Rational Credibility, *Philosophical Review*, vol. 115, no. 4, pp. 449-485.

Engel, Pascal. (2008). In what sense is knowledge the norm of assertion, *unpublished manuscript*, Université de Geneve

Hawthorne, John. (2004). *Knowledge and Lotteries*. Oxford : Oxford University Press.

Hawthorne, John & Stanley, Jason. (2008). Knowledge and Action. *Journal of Philosophy* 105 (10):571-590.

Hebert, Ryan. (2016). *Teleologism Full Stop: A General Theory of Ability, Agency, Obligation, and Justification* unpublished manuscript, University of Calgary, Calgary, AB

Hill, Christopher, & Schechter, Joshua. (2007). Hawthorne's Lottery Puzzle and the Nature of Belief, *Philosophical Issues*, vol. 17, pp. 102-122.

Huemer, Michael. (2007). Moore's Paradox and the Norm of Belief. In *Susana Nuccetelli and Gary Seay (eds.), Themes From G. E. Moore: New Essays in Epistemology and Ethics*. Clarendon Press.

Hyman, John. (1999). How Knowledge Works. *Philosophical Quarterly* 49: pp. 433–51.

Kearns, Stephen, & Star, Daniel. (2008). Reasons: Explanations or evidence? *Ethics*, 119(1), 31–56.

—— (2009). Reasons as evidence. *Oxford Studies in Metaethics*, 4, 215–242.

Kvanvig, Jonathan. (2009). Assertion, Knowledge, and Lotteries, in *Williamson on Knowledge*, ed. Patrick Greenough and Duncan Pritchard (Oxford: Oxford University Press), pp. 140-160.

—— (2011). "Norms of Assertion." In *Jessica Brown and Herman Cappellen (eds.). Assertion: New Philosophical Essays*. Oxford: Oxford University Press.

Lackey, Jennifer. (1999). Testimonial Knowledge and Transmission, *Philosophical Quarterly*, vol. 49, pp. 471-490.

—— (2007) Norms of Assertion, *Noûs*, vol. 41, no. 4, pp. 594-626.

—— (2010) Assertion and Isolated Second-hand Knowledge, in *Assertion: New Philosophical Essays*, ed. Jessica Brown and Herman Cappellen. Oxford: Oxford University Press, pp. 251-276.

Littlejohn, Clayton. (2016). Reasons and Theoretical Rationality. In *Daniel Star (ed.), Oxford Handbook of Reasons and Normativity*. Oxford University Press.

Maitra, Ishani & Weatherson, Brian. (2010). Assertion, knowledge, and action. *Philosophical Studies* 149 (1):99-118.

Fantl, Jeremy & McGrath, Matthew. (2009). *Knowledge in an Uncertain World*. Oxford University Press

- (2012). Pragmatic Encroachment: It's not just about knowledge. *Episteme*, 9(1), 27-42
- McKinnon, Rachel. (2011). Lotteries, Knowledge, and Practical Reasoning. *Logos and Episteme* (2): pp. 225-231.
- (2013). The Supportive Reasons Norm of Assertion. *American Philosophical Quarterly*, 50(2), pp. 121-135.
- Neta, Ram. (2009). Treating Something as a Reason for Action. *Noûs* 43:4. pp 684-699
- Parfit, Derek. (2001). Reasons and rationality. In D. Egonsson, J. Josefsson, B. Petersson, & T. Rønnow-Rasmussen (Eds.), *Exploring practical philosophy* (pp. 17–39). Aldershot: Ashgate.
- (2011). *On What Matters* Oxford University Press.
- Putnam, Hilary. (1981). *Brains In a Vats, in Reason, Truth, and History*. Cambridge University Press
- Raz, Joseph. (1978). Introduction. In J. Raz (Ed.), *Practical reasoning* (pp. 1–17). Oxford: Oxford University Press.
- (1999). *Practical Reason and Norms*. Oxford: Oxford University Press.
- Rescorla, Michael. (2009). "Assertion and Its Constitutive Norms," *Philosophy and Phenomenological Research*, vol. 79, no. 1, pp. 98-130.
- Schroeder, Mark. (2007). Reasons and Agent-neutrality. *Philosophical Studies*, doi: 10.1007/s11098-007-9087-y.
- (2007a). *Reason as Slaves of the Passions*, Oxford University Press
- (2008). Having reasons. *Philosophical Studies* 139 (1):57 - 71.
- (2012). Precs of Slaves of the Passions, *Philosophical Studies* 157, Issue 3, pp. 431-434
- Setiya, Kieran. (2004). Against internalism. *Noûs*, 38(2), 266–298.
- (2009). Reply to Bratman and Smith. *Analysis*, 69(3), 531–540.

- (2014). What is a reason to act? *Philosophical Studies*, 167. pp. 221-235
- Setiya, Kieran., & Paakkunaninen, Hille. (2012). *Internal reasons: Contemporary readings*. Cambridge, MA: MIT Press.
- Shafer-Landau, Russ. (2003). *Moral realism*. Oxford: Oxford University Press.
- Shieber, Joseph. (2009). Epistemological Contextualism and the Knowledge Account of Assertion, *Philosophia*, vol. 37, no. 1, pp. 169-181.
- Smith, Michael. (1987). The Humean theory of motivation. *Mind*, 96(381), 36–61.
- (1994). *The moral problem*. Oxford: Blackwell.
- (1995). Internal reasons. *Philosophy and Phenomenological Research*, 55(1), 109–131.
- (2006). Moore on the right, the good, and uncertainty. In T. Horgan & M. Timmons (Eds.), *Metaethics after Moore* (pp. 133–148). Oxford: Oxford University Press.
- Sobel, David. (1999). Do the desires of rational agents converge? *Analysis*, 59(3), 137–147.
- Sosa, Ernest. (2011). *Knowing Full Well* Princeton University Press
- Stanley, Jason. (2005). *Knowledge and Practical Interests*. Oxford : Oxford University Press.
- Turri, John. (2010). Epistemic Invariantism and Speech Act Contextualism, *Philosophical Review*, vol. 119, no. 1, pp. 77-95.
- (2010a). Prompting Challenges, *The Analysis of Trust*, Vol 70, Num 3. Oxford University Press. pp 456-462
- (2011). The Express Knowledge Account of Assertion, *Australasian Journal of Philosophy*, vol. 89, no. 1, pp. 37-53
- Unger, Peter. (1975). *Ignorance*. Oxford University Press.

Weiner, Matthew. (2005). Must We Know What We Say? *The Philosophical Review* 114: pp. 227–51.

Williams, Bernard. (1981). Internal and external reasons. In *B. Williams (Ed.), Moral luck* (pp.101–113). Oxford: Oxford University Press

—— (2002). *Truth and Truthfulness*. Princeton : Princeton University Press.

Williamson, Timothy. (1996). Knowing and Asserting. *The Philosophical Review* 105: pp. 489–523.

—— (2000). *Knowledge and its Limits*. Oxford : Oxford University Press.

—— (2005). Contextualism, subject-sensitive invariantism and knowledge of knowledge. *Philosophical Quarterly*, 55 (219):213–235.