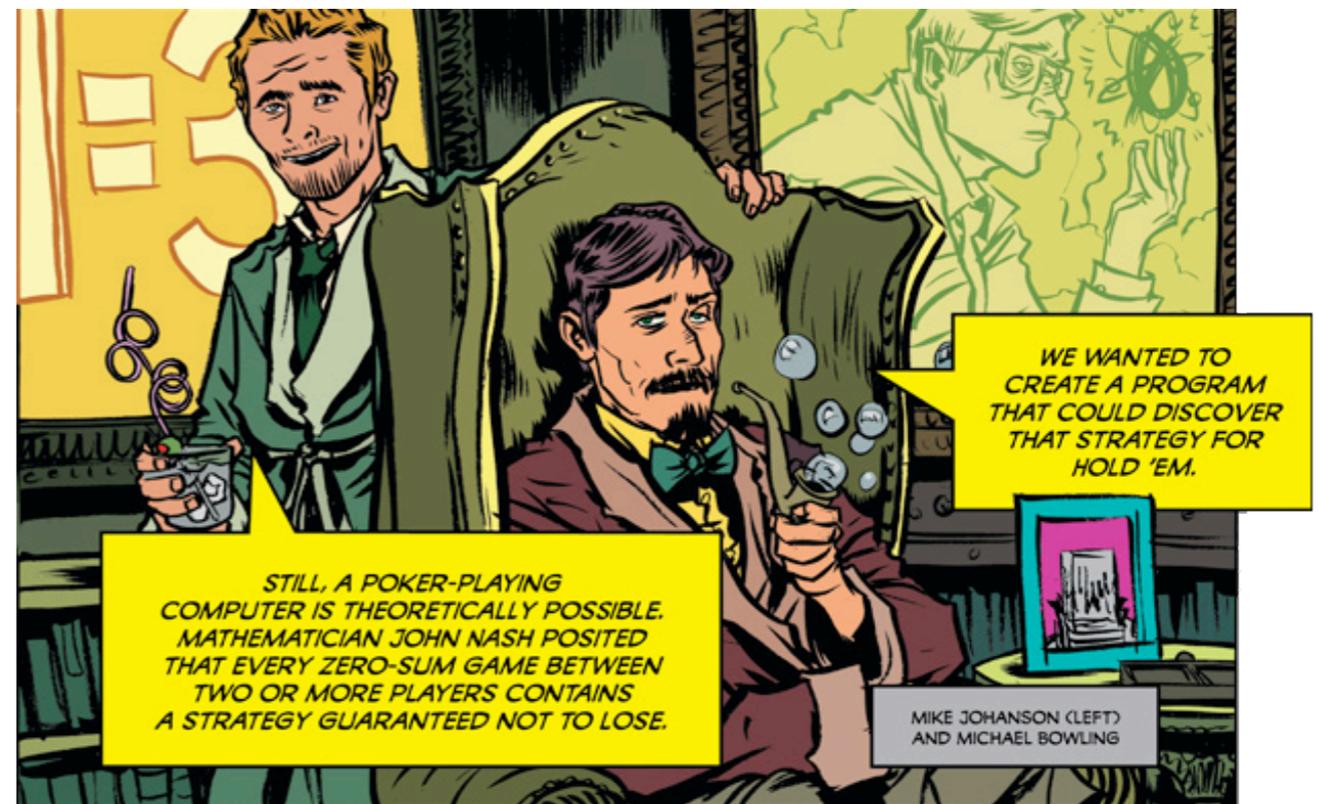


# The Polaris Computer Poker Program

---

Dr. Michael Bowling  
University of Alberta

*Wired*, December 2008



# The Polaris Computer Poker Program

Dr. Michael Bowling  
University of Alberta

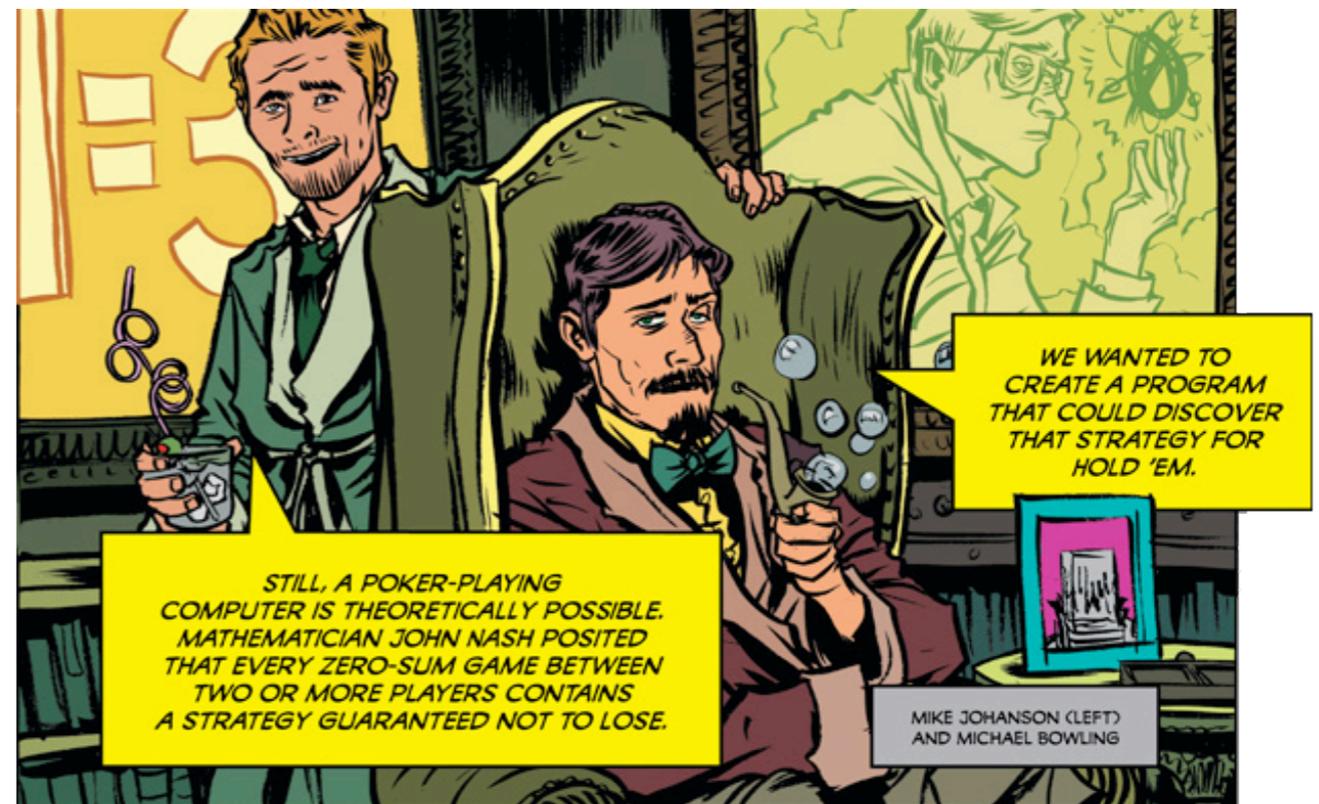
*Wired*, December 2008



# The Polaris Computer Poker Program

Dr. Michael Bowling  
University of Alberta

*Wired*, December 2008



# Why am I here?

---

- I am not a sociologist, economist, or political scientist
- I am a computer scientist
- I study artificial intelligence: how do we make computers make intelligent decisions?
- I have a passion for games
- My five year obsession: computers beating humans at **poker**

# The Computer Poker Research Group

---

- Started in 1997 (by J. Schaeffer and D. Szafron)
- Poker research at Alberta started in 1992 with Darse Billings (M.Sc. student)
- 9 M.Sc. and 1 Ph.D. dissertations later...
  - 4 Faculty
  - 6 Students
  - 2 Programmers
  - 2 Consultants

# The Computer Poker Research Group

---

- Started in 1997 (by J. Schaeffer and D. Szafron)
- Poker research at Alberta started in 1992 with Darse Billings (M.Sc. student)
- 9 M.Sc. and 1 Ph.D. dissertations later...
  - 4 Faculty
  - 6 Students
  - 2 Programmers
  - 2 Consultants



# The Game

---

# The Game

---



# Why Poker?

---

# Why Poker?

---

“The real world is much more like a poker game...”

David Moschella in Computerworld, 1999

# Why Poker?

---

“The real world is much more like a poker game...”

David Moschella in Computerworld, 1999

- Imperfect information

# Why Poker?

---

“The real world is much more like a poker game...”

David Moschella in Computerworld, 1999

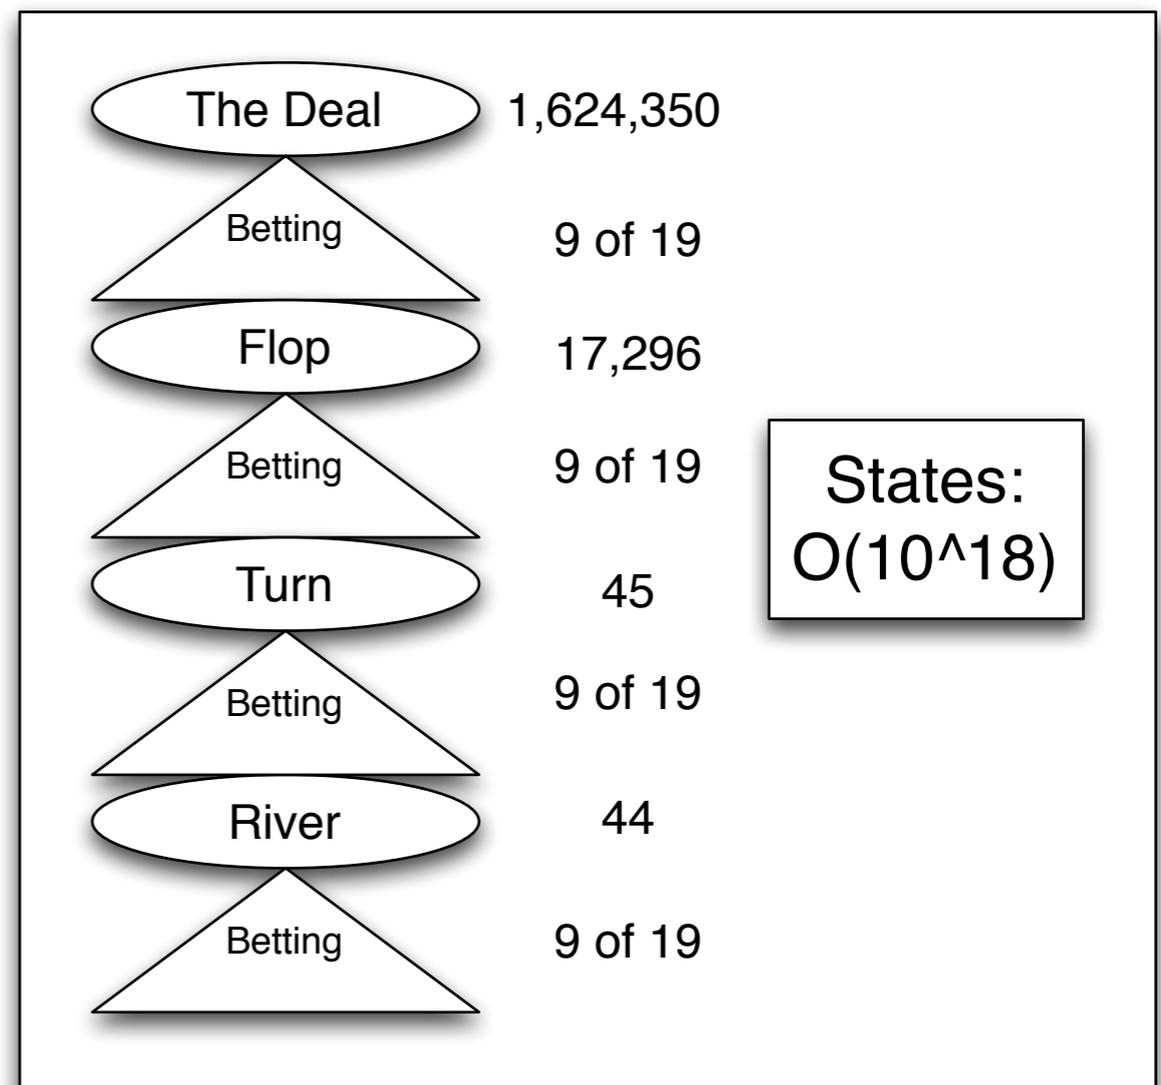
- Imperfect information
- Large game tree

# Why Poker?

---

“The real world is much more like a poker game...”  
David Moschella in Computerworld, 1999

- Imperfect information
- Large game tree



# Why Poker?

---

“The real world is much more like a poker game...”

David Moschella in Computerworld, 1999

- Imperfect information
- Large game tree
- Variable number of players (2-10, or more)

# Why Poker?

---

“The real world is much more like a poker game...”

David Moschella in Computerworld, 1999

- Imperfect information
- Large game tree
- Variable number of players (2-10, or more)
- Stochastic element

# Why Poker?

---

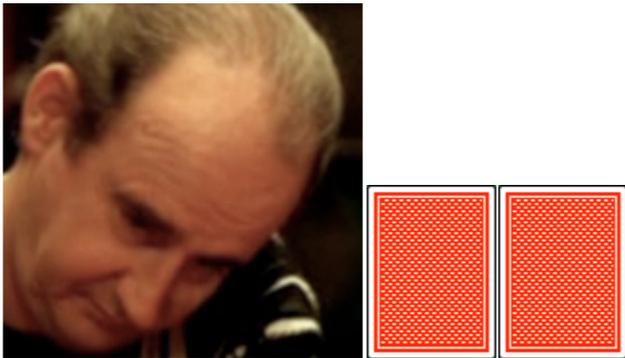
“The real world is much more like a poker game...”

David Moschella in Computerworld, 1999

- Imperfect information
- Large game tree
- Variable number of players (2-10, or more)
- Stochastic element
- Amount of winnings matters

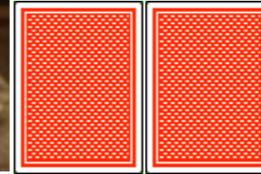
# The Challenges

---

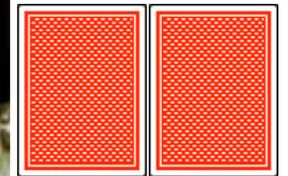
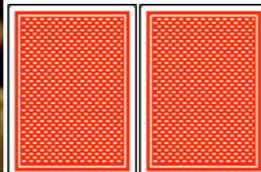


# The Challenges

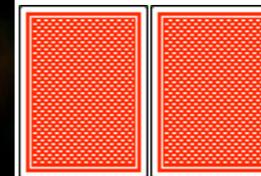
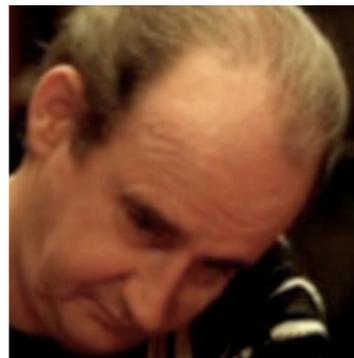
---



Small Blind



Big Blind



# The Challenges

---



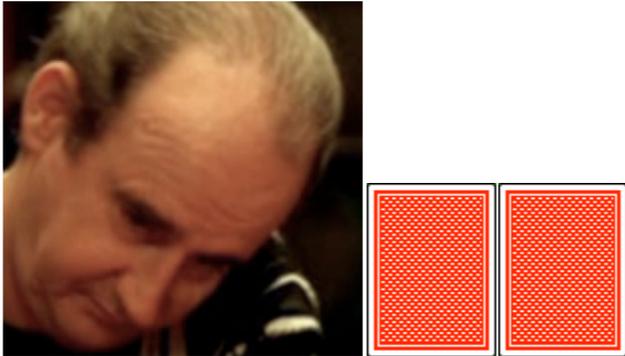
Small Blind



Call



Big Blind



# The Challenges

---



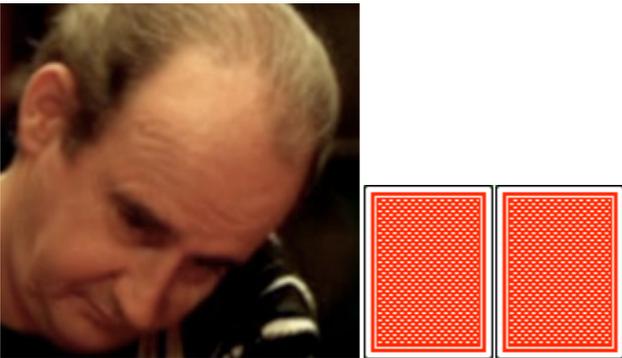
Small Blind



Raise



Big Blind



# The Challenges

---



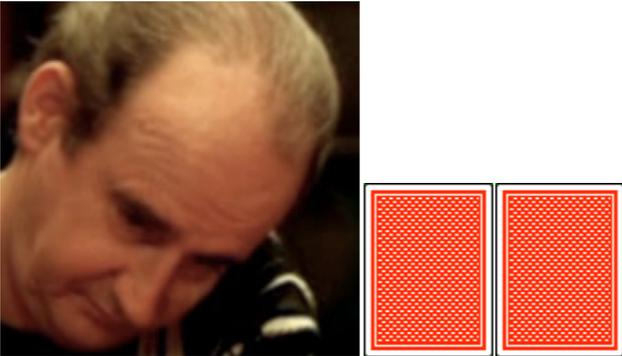
Small Blind  
Fold



Raise



Big Blind



Call

# The Challenges

---



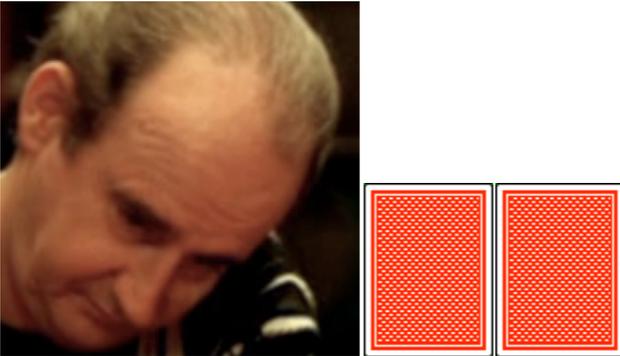
Small Blind  
Fold



Raise



Big Blind  
Call



# The Challenges

---



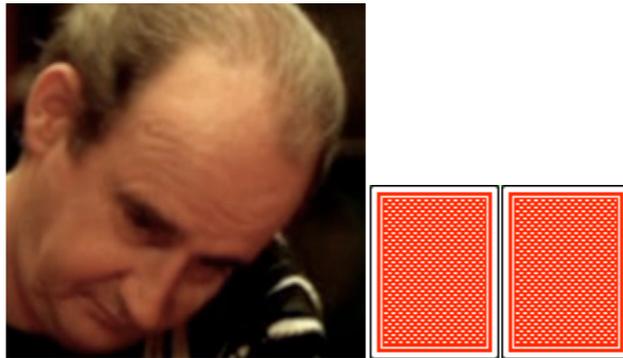
Small Blind  
Fold



Raise



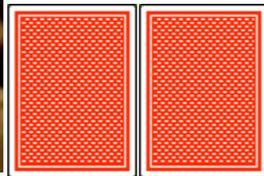
Big Blind  
Call



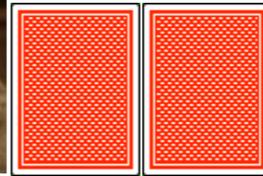
Call  
Fold

# The Challenges

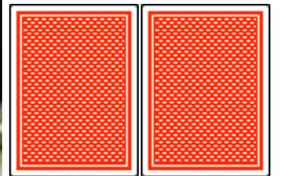
---



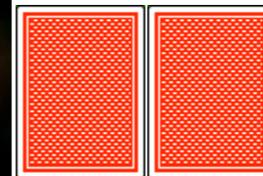
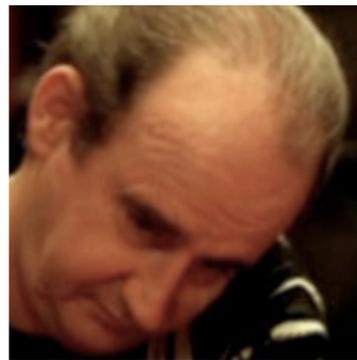
Raise



Small Blind  
Fold



Big Blind  
Call

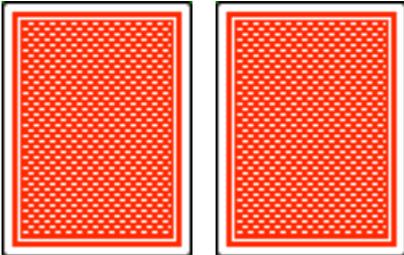
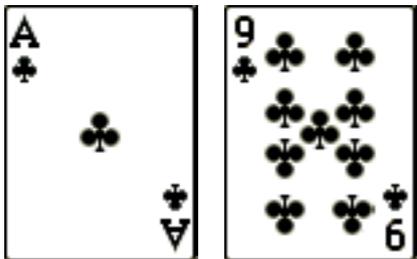


Call  
Fold

“It’s a position raise.” (> 2 player)

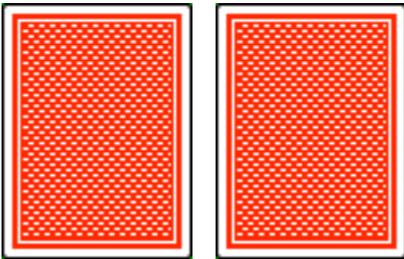
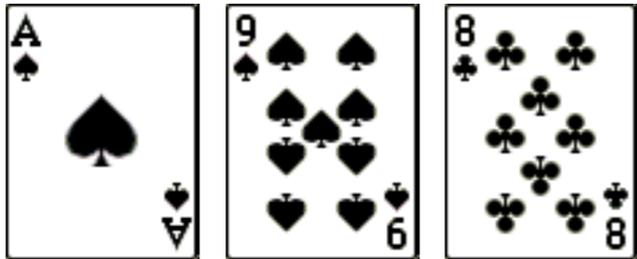
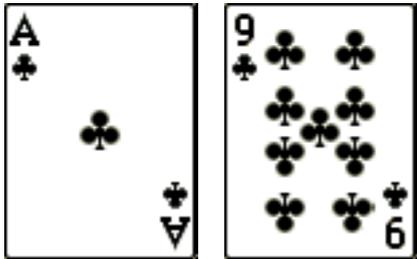
# The Game of Poker

---



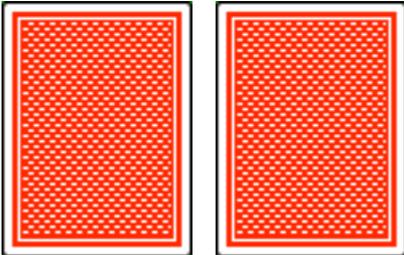
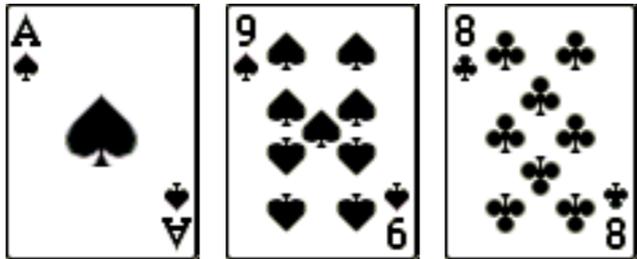
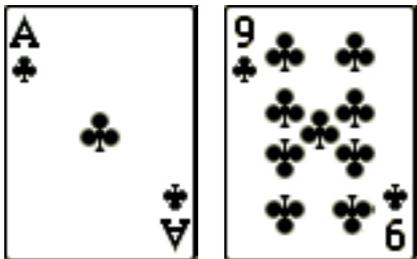
# The Game of Poker

---



# The Game of Poker

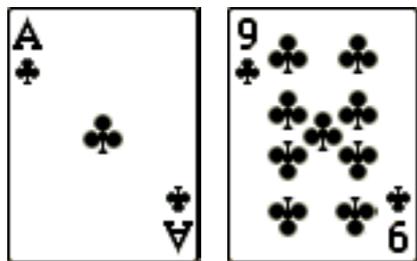
---



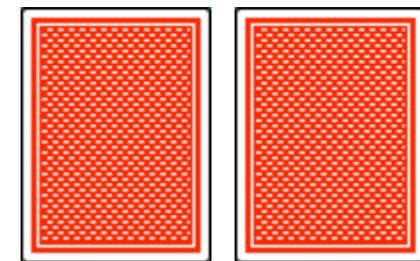
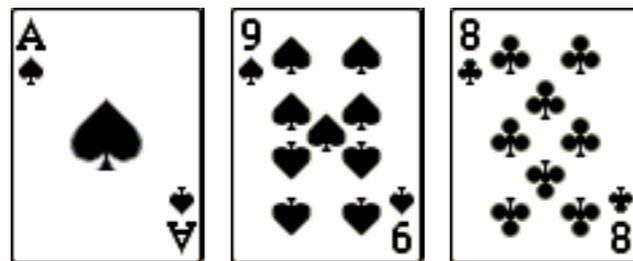
Check

# The Game of Poker

---



Raise

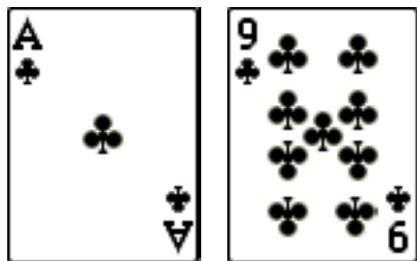


Check

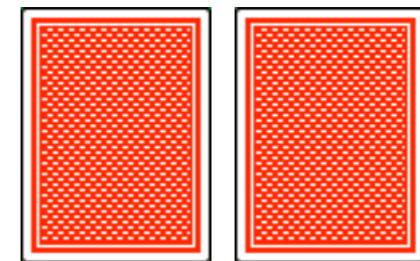
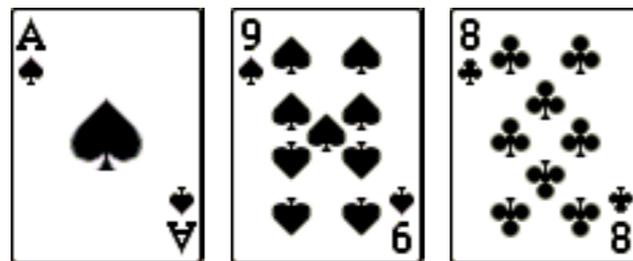
“Overbet to fake weakness” (winnings matter)

# The Game of Poker

---



Raise



Check

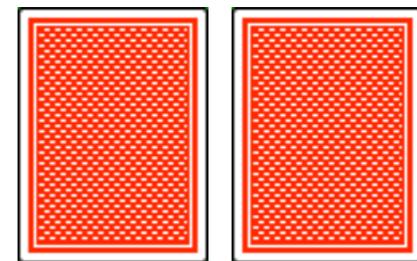
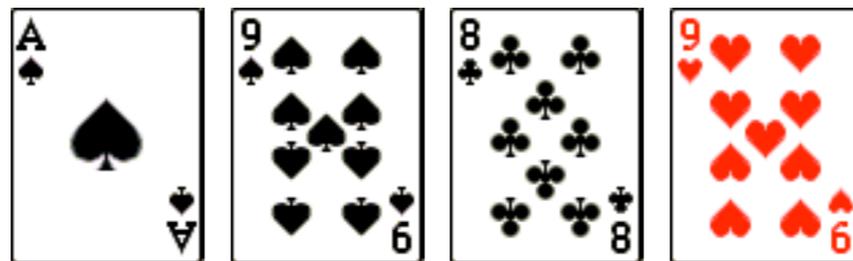
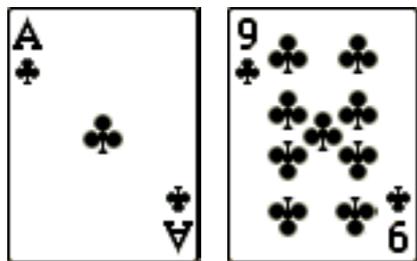
Call

“Overbet to fake weakness” (winnings matter)

“Teddy has a flush draw” (imperfect information)

# The Game of Poker

---

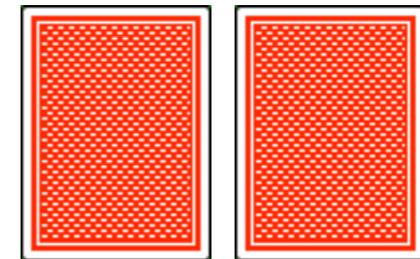
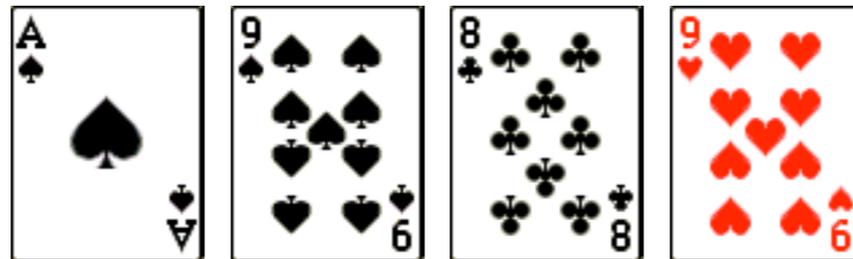
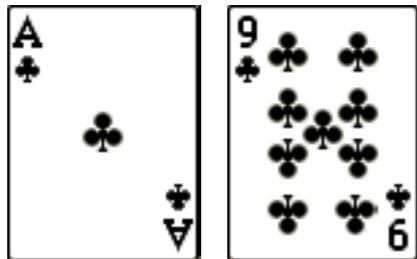


“Overbet to fake weakness” (winnings matter)

“Teddy has a flush draw” (imperfect information)

# The Game of Poker

---



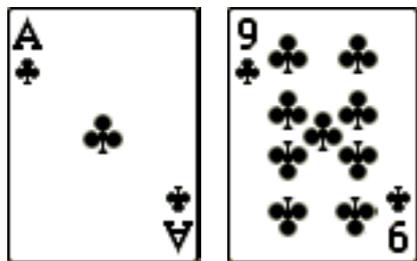
Check

“Overbet to fake weakness” (winnings matter)

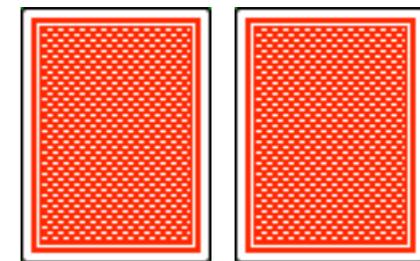
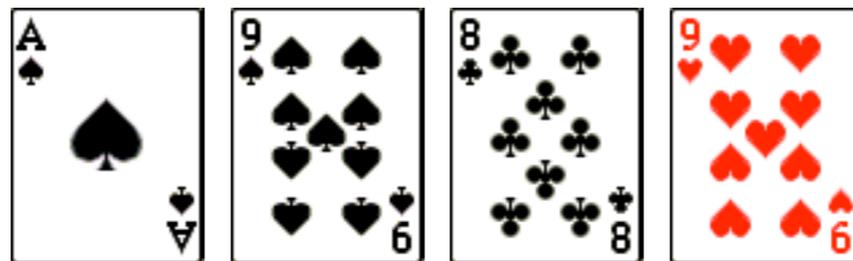
“Teddy has a flush draw” (imperfect information)

# The Game of Poker

---



Check



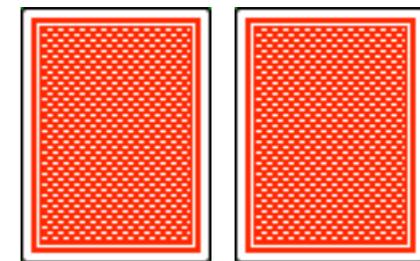
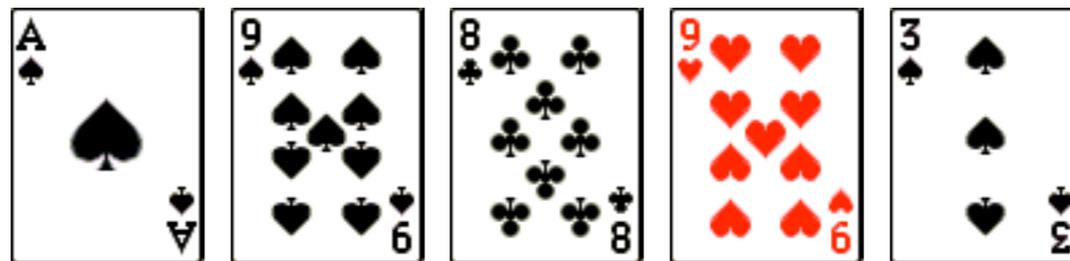
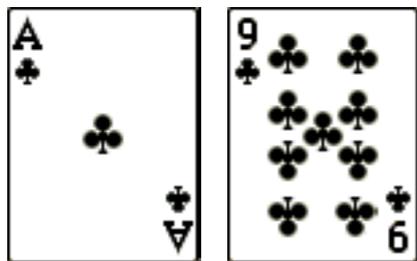
Check

“Overbet to fake weakness” (winnings matter)

“Teddy has a flush draw” (imperfect information)

# The Game of Poker

---

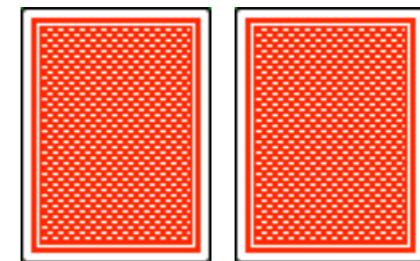
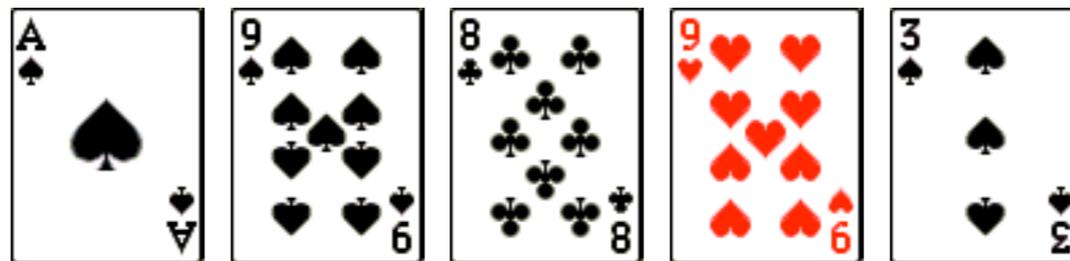
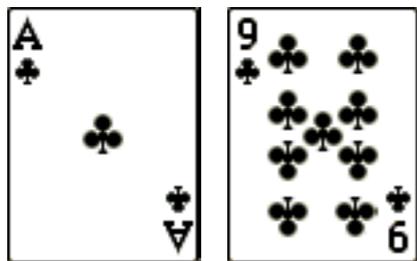


“Overbet to fake weakness” (winnings matter)

“Teddy has a flush draw” (imperfect information)

# The Game of Poker

---



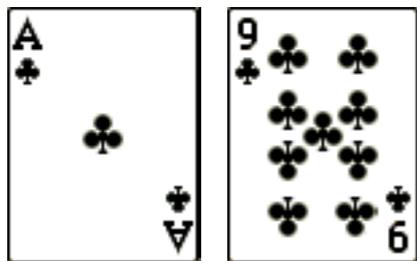
Raise

“Overbet to fake weakness” (winnings matter)

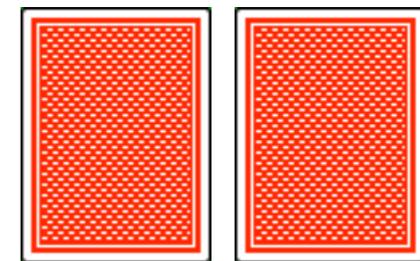
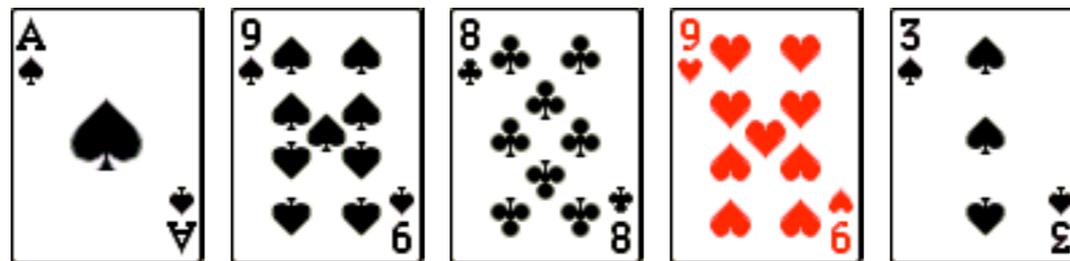
“Teddy has a flush draw” (imperfect information)

# The Game of Poker

---



All-In



Raise

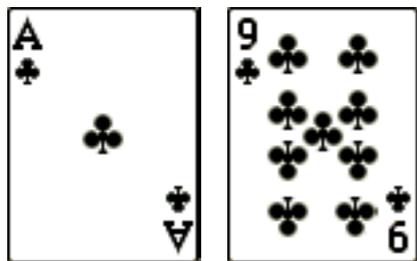
“Overbet to fake weakness” (winnings matter)

“Teddy has a flush draw” (imperfect information)

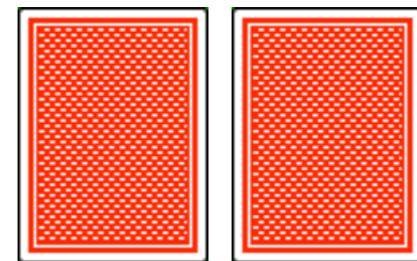
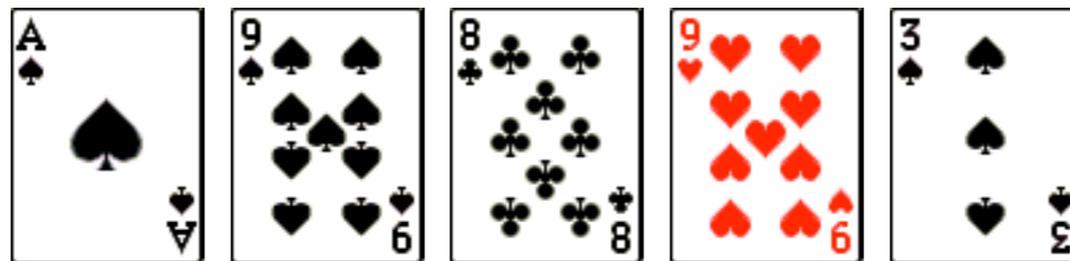
“Got'em... All-In” (opponent modelling)

# The Game of Poker

---



All-In



Raise

Call

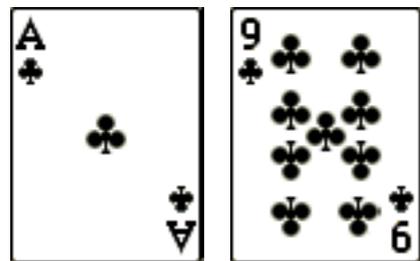
“Overbet to fake weakness” (winnings matter)

“Teddy has a flush draw” (imperfect information)

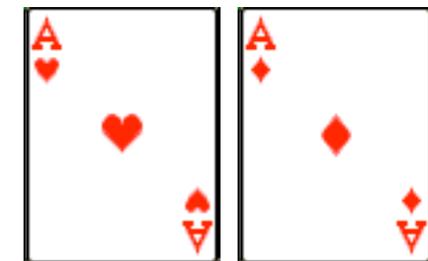
“Got'em... All-In” (opponent modelling)

# The Game of Poker

---

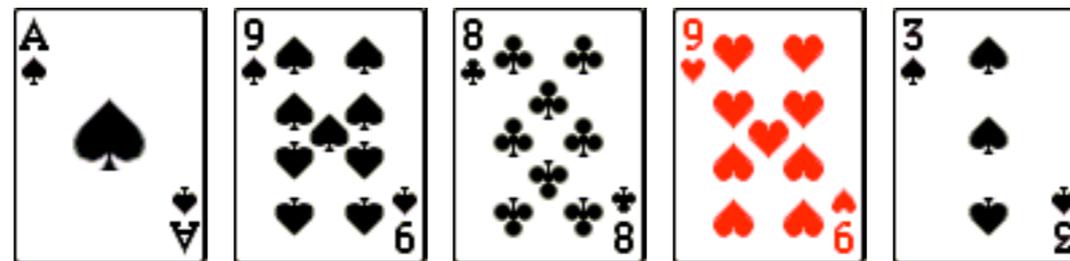


All-In



Raise

Call



“Overbet to fake weakness” (winnings matter)

“Teddy has a flush draw” (imperfect information)

“Got'em... All-In” (opponent modelling)

“Ouch!” (stochastic element)

# The Challenges

---

“The real world is much more like a poker game...”

David Moschella in Computerworld, 1999

- Imperfect information
- Large game tree
- Variable number of players (2-10, or more)
- Stochastic element
- Amount of winnings matters

# The Challenges

---

“The real world is much more like a poker game...”  
David Moschella in Computerworld, 1999

- Imperfect information
- Large game tree
- Variable number of players (2-10, or more)
- Stochastic element
- Amount of winnings matters



# The Challenges

---

“The real world is much more like a poker game...”  
David Moschella in Computerworld, 1999

- Imperfect information 
- Large game tree 
- Variable number of players (2-10, or more)
- Stochastic element 
- Amount of winnings matters 

Four Years of Research (2003-2007)... One Slide

---

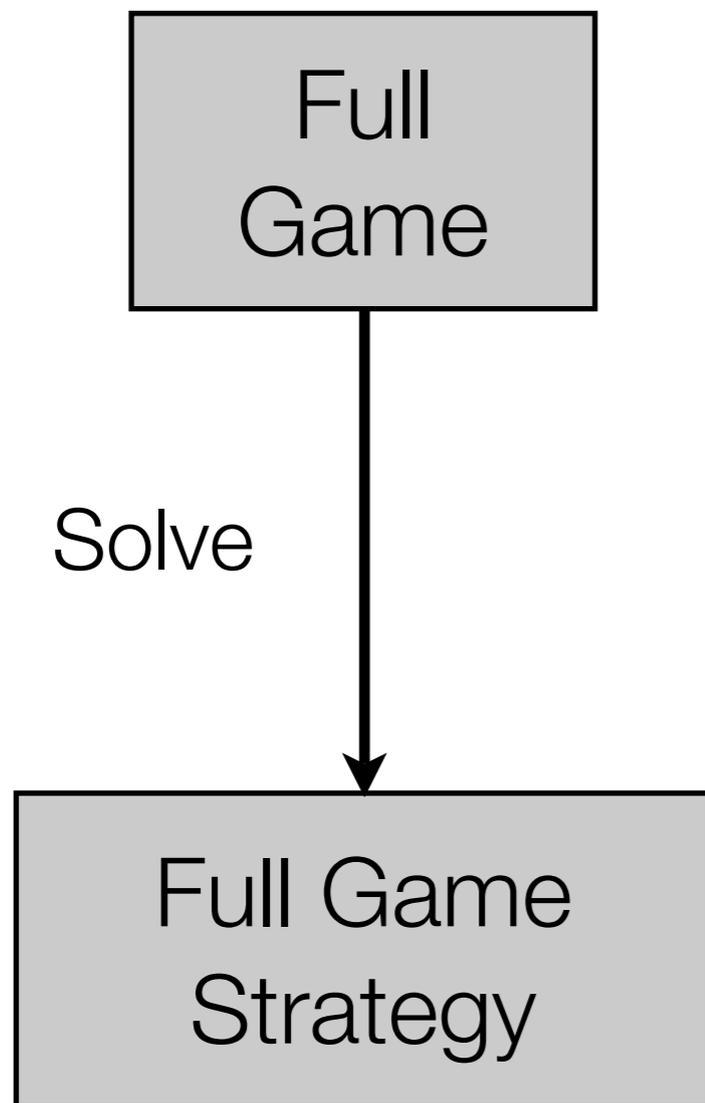
# Four Years of Research (2003-2007)... One Slide

---

- Game theoretic optimal strategies!
  - Static strategies, but may employ randomness
  - Guaranteed not to lose money in the long run

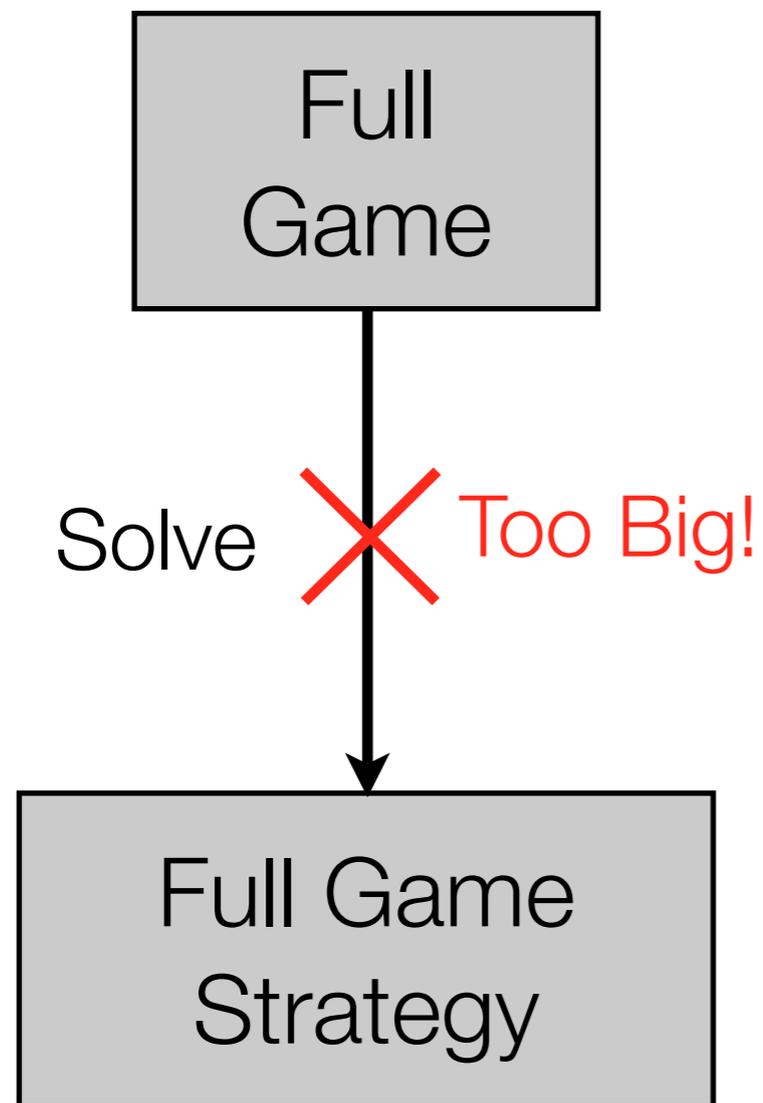
# Four Years of Research (2003-2007)... One Slide

---



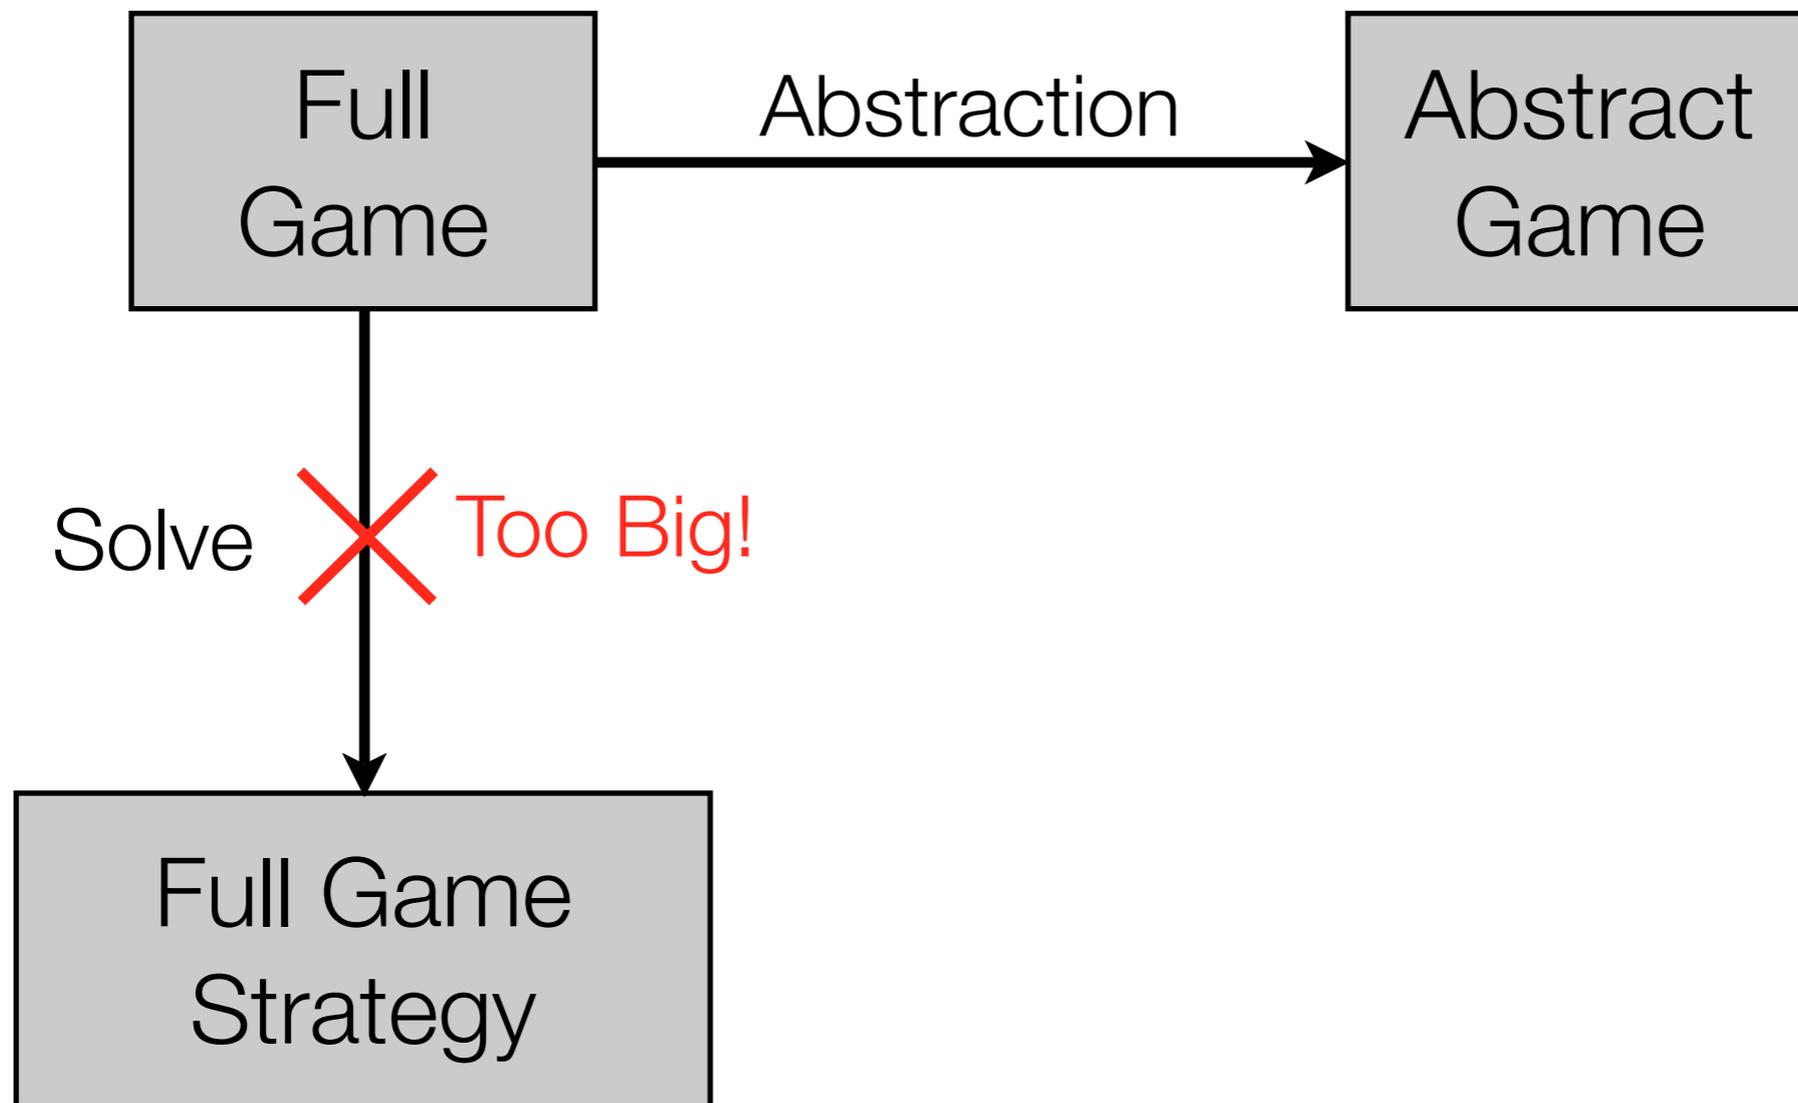
# Four Years of Research (2003-2007)... One Slide

---



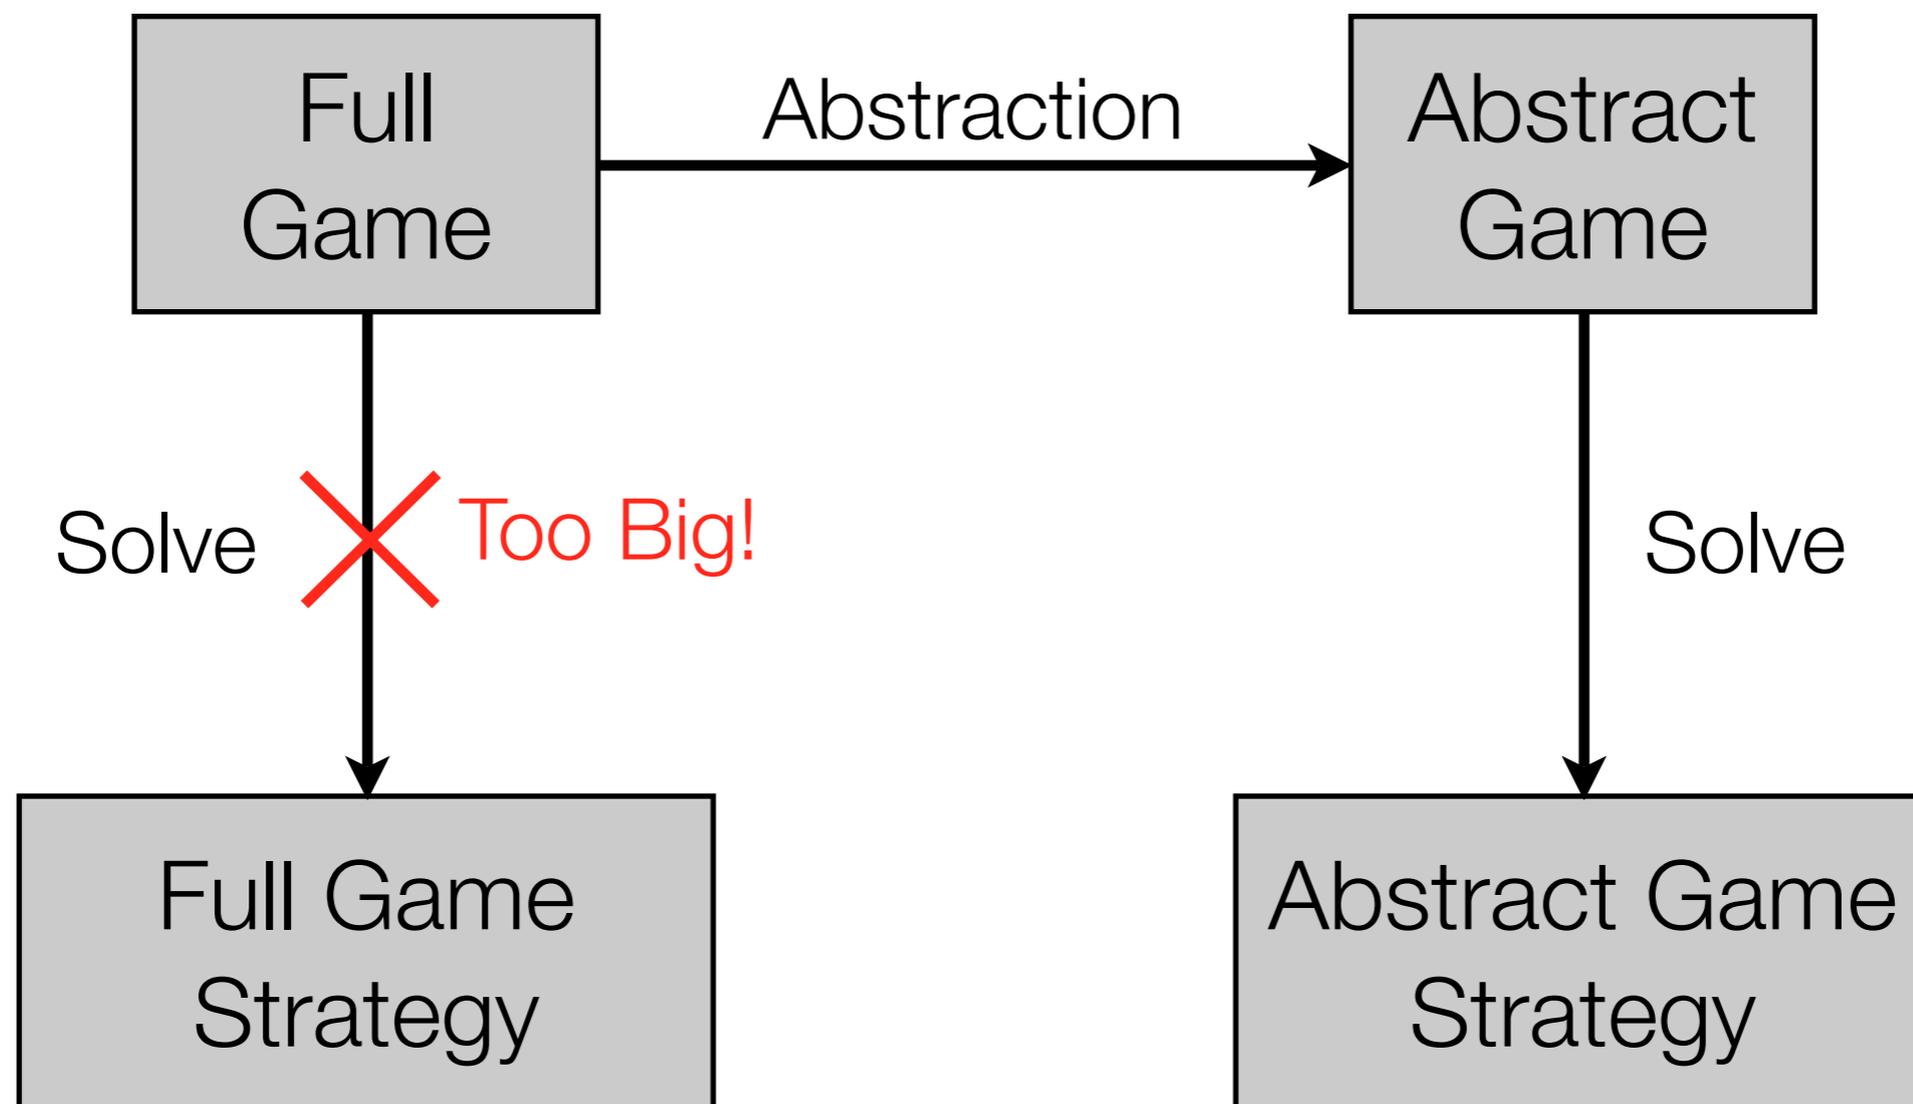
# Four Years of Research (2003-2007)... One Slide

---



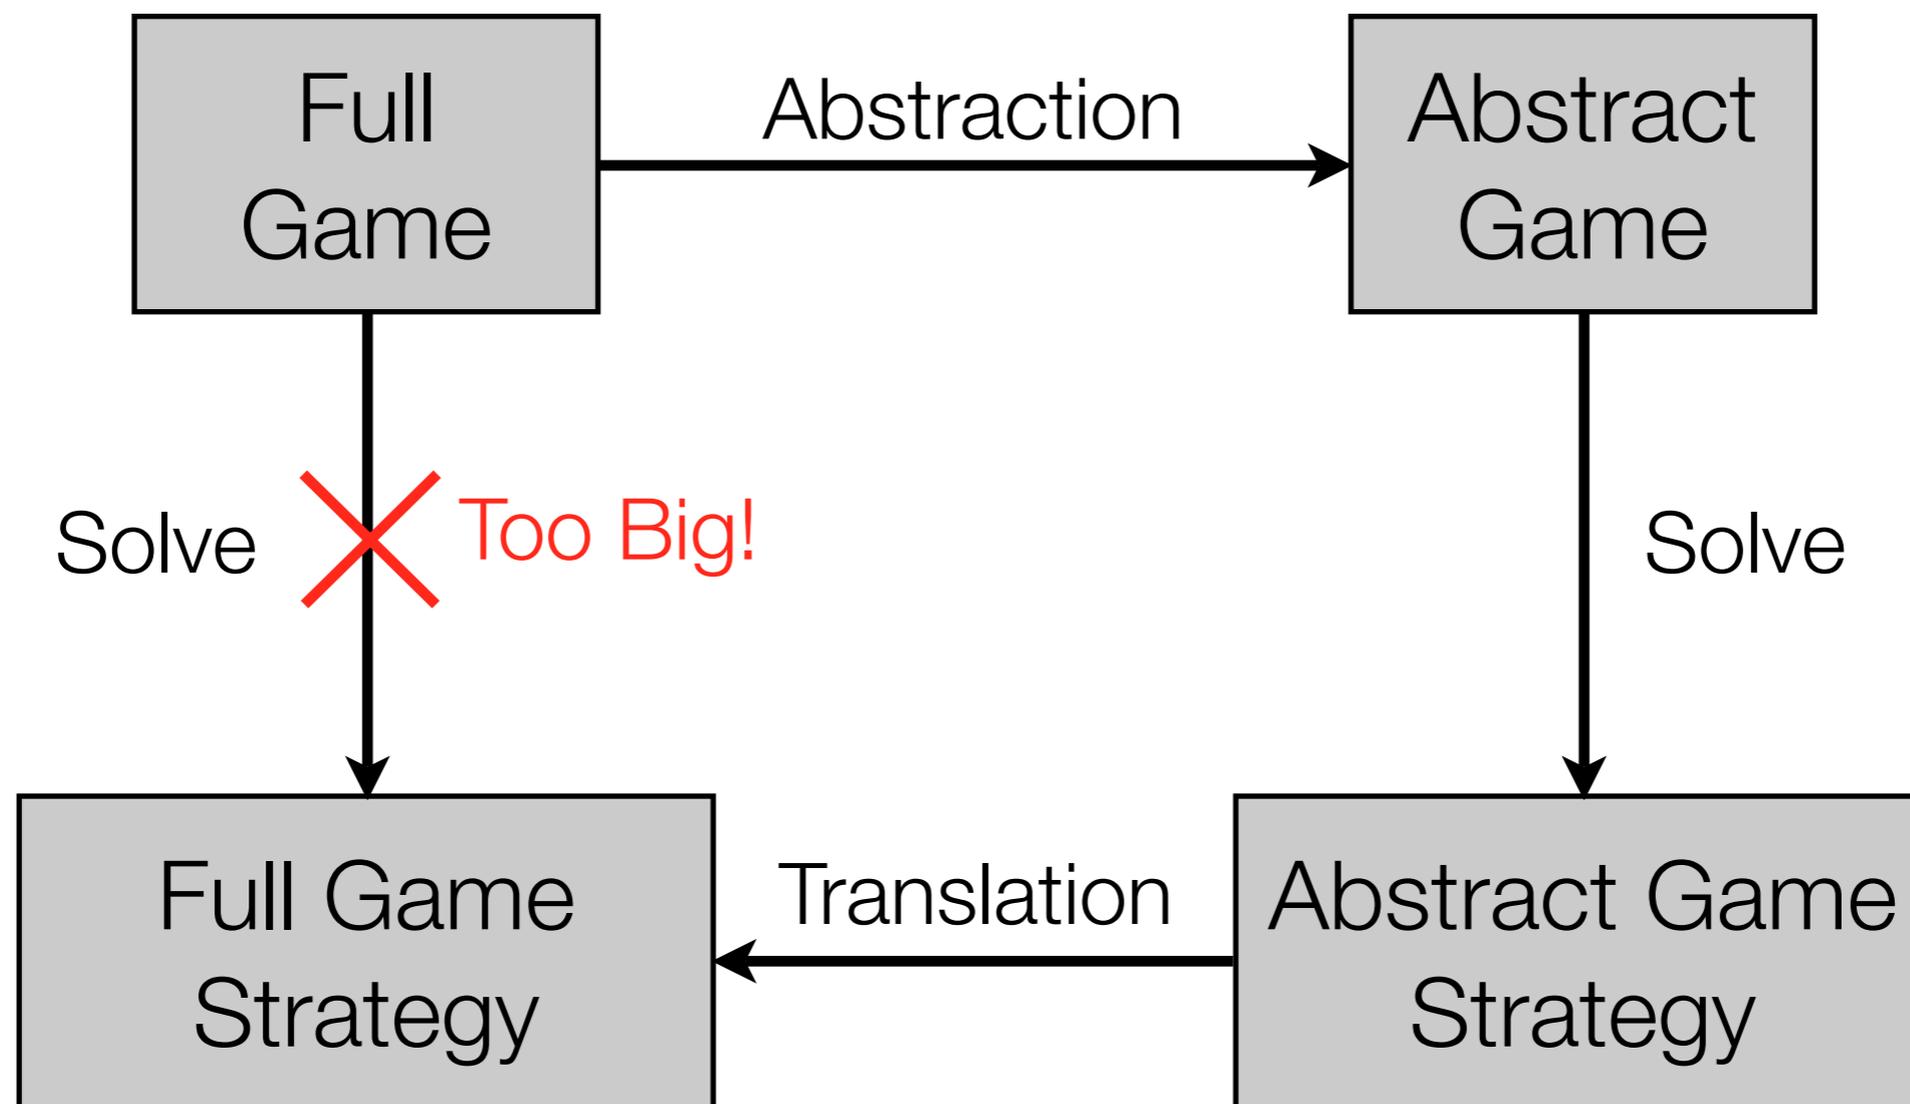
# Four Years of Research (2003-2007)... One Slide

---



# Four Years of Research (2003-2007)... One Slide

---



# First Man-Machine Poker Championship

---

- AAI, Vancouver 2007



Ali Eslami

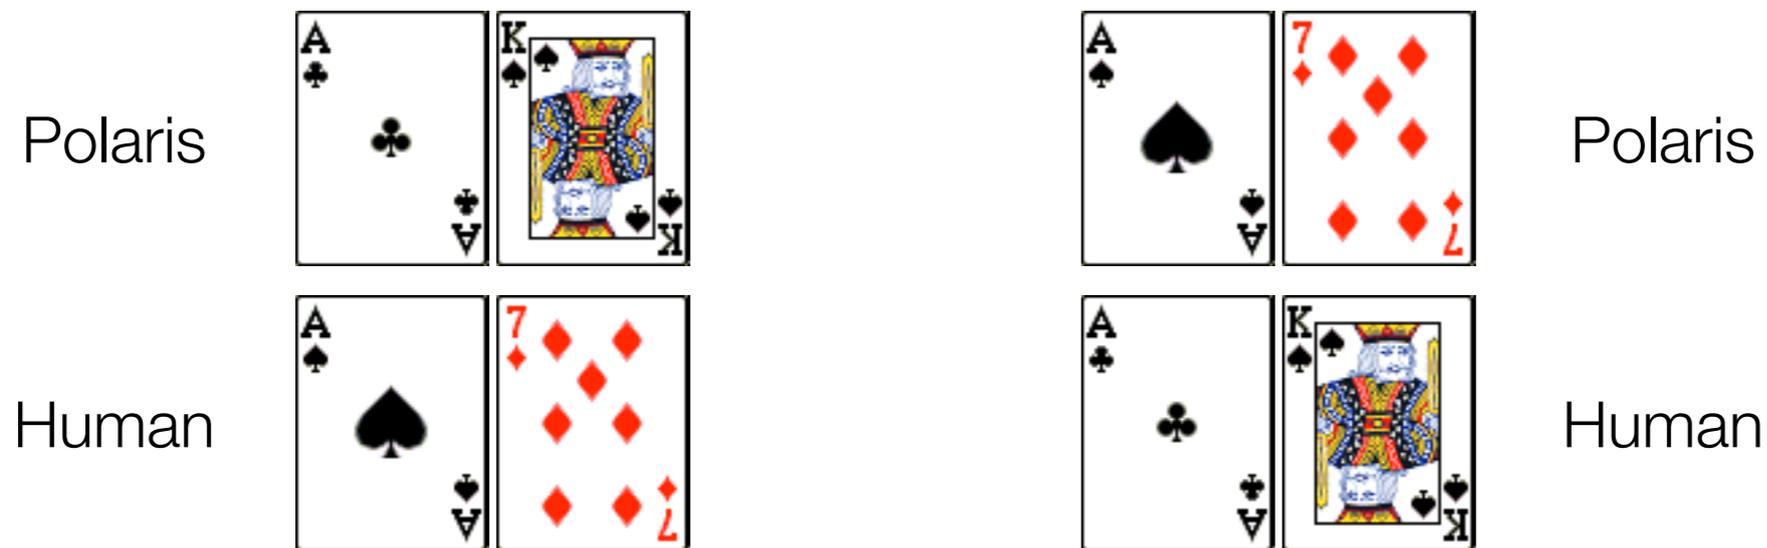


Phil Laak  
“The Unabomber”

# First Man-Machine Poker Championship

---

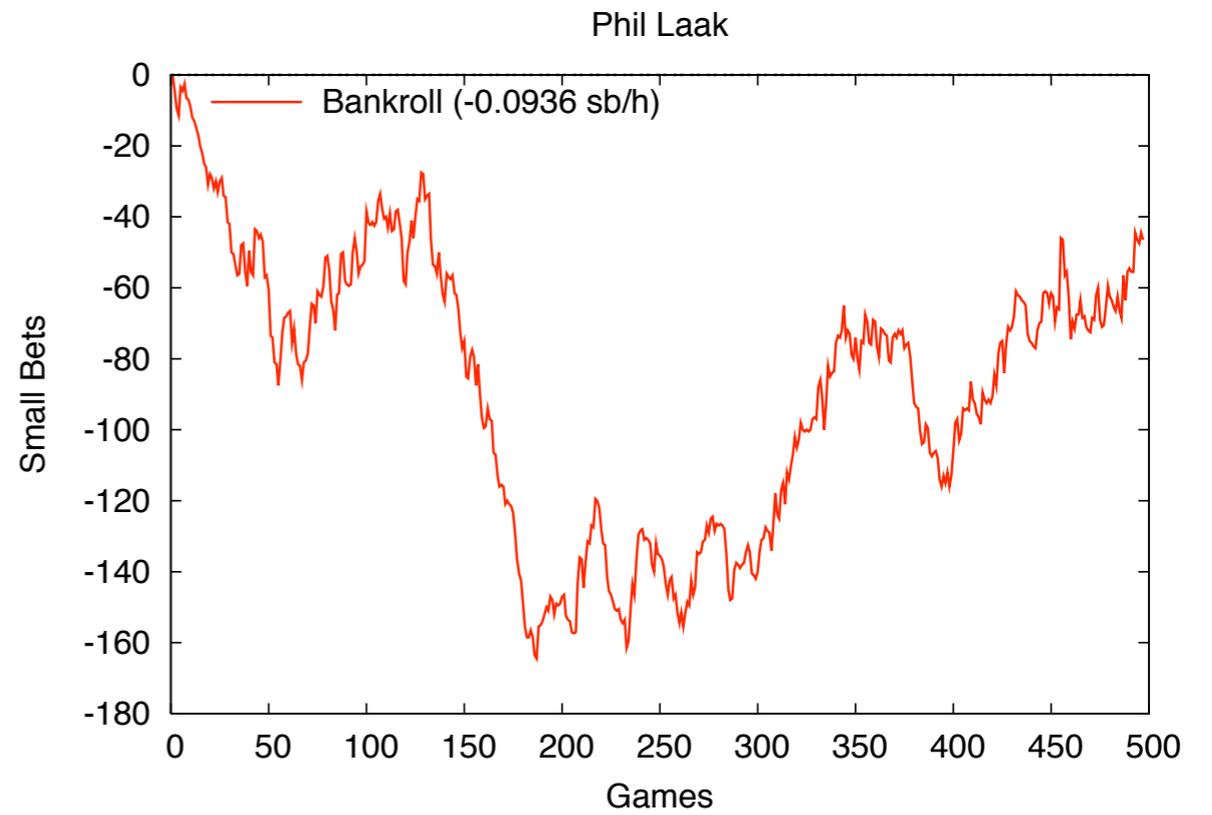
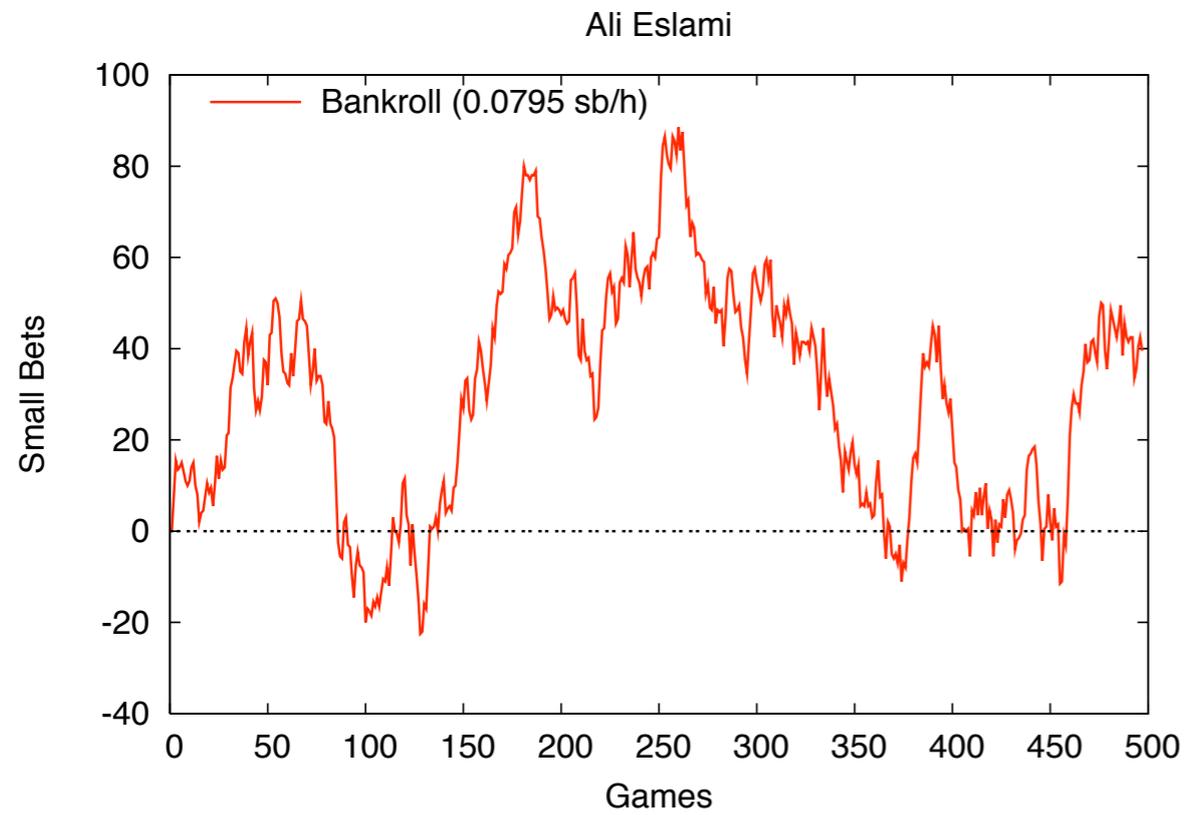
- 500 hand duplicate matches



- One match played in public
- Other played simultaneously in private
- Team's total winnings must exceed 25sb to win

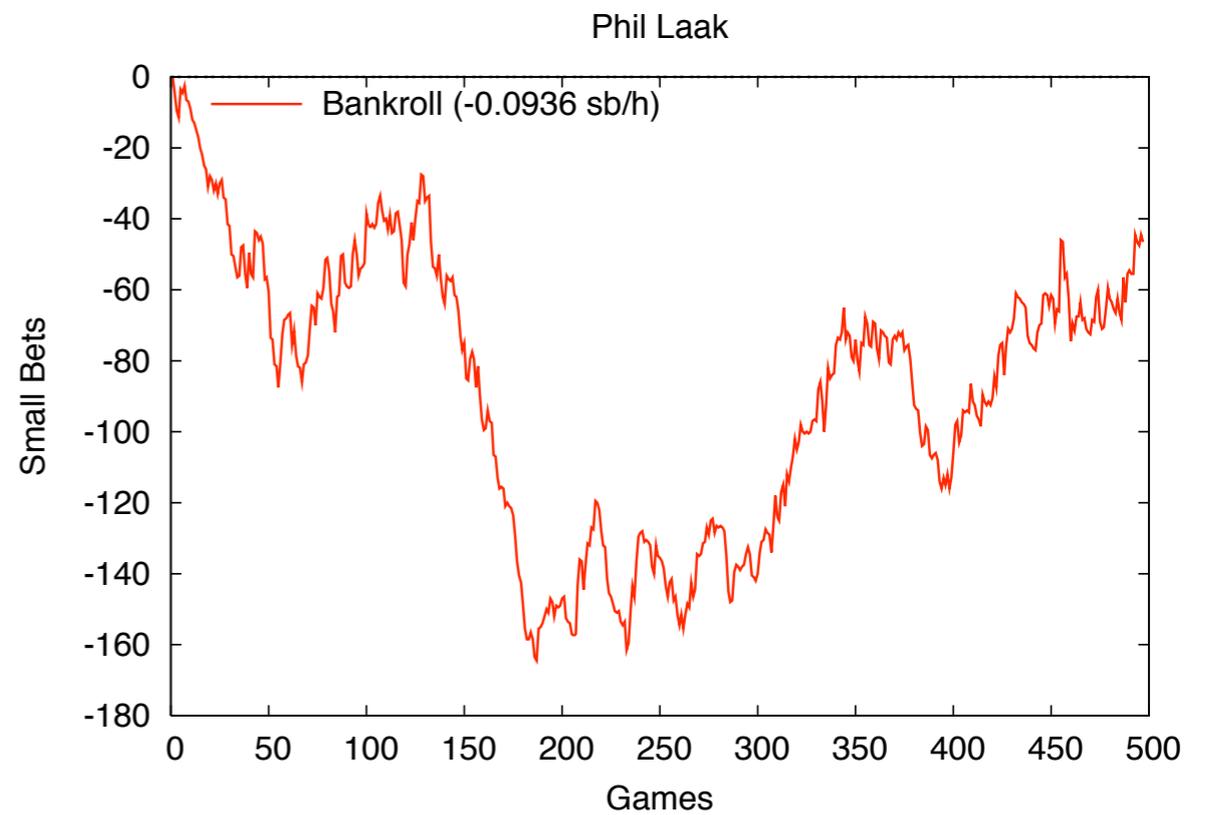
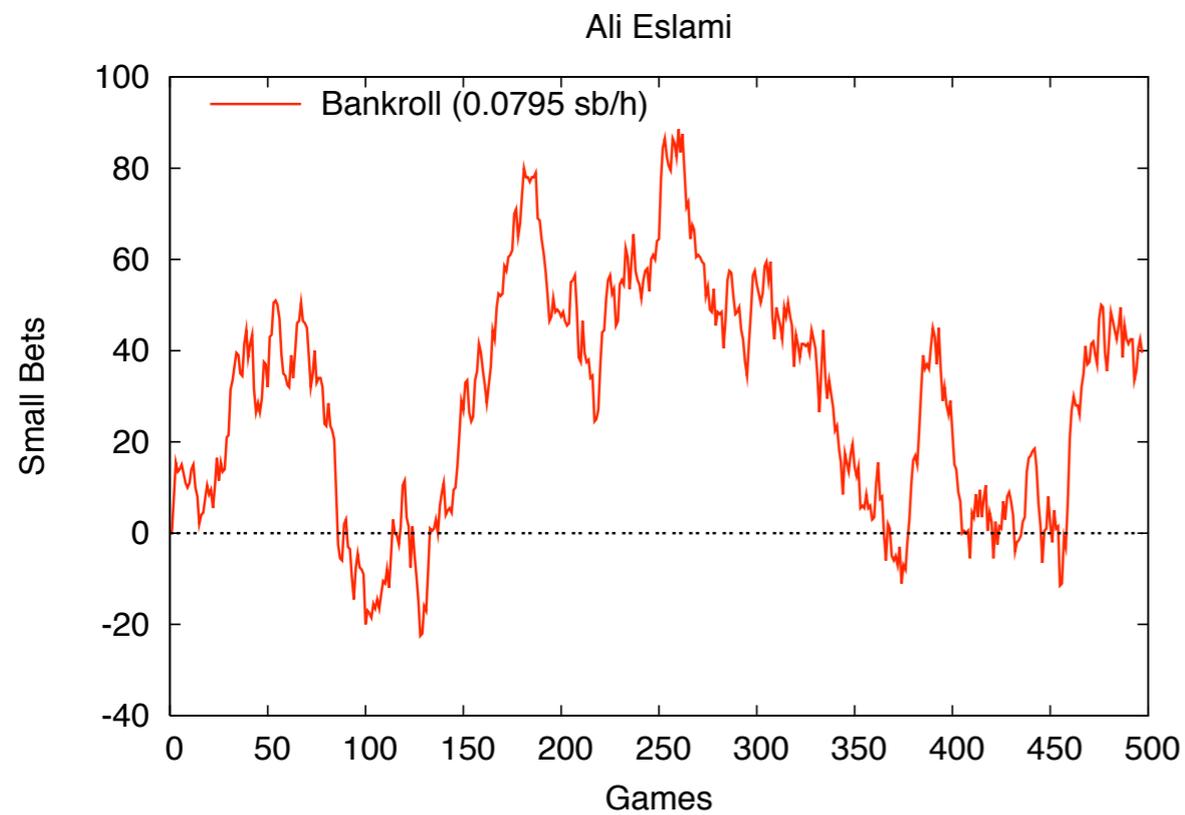
# The Results

---



# The Results

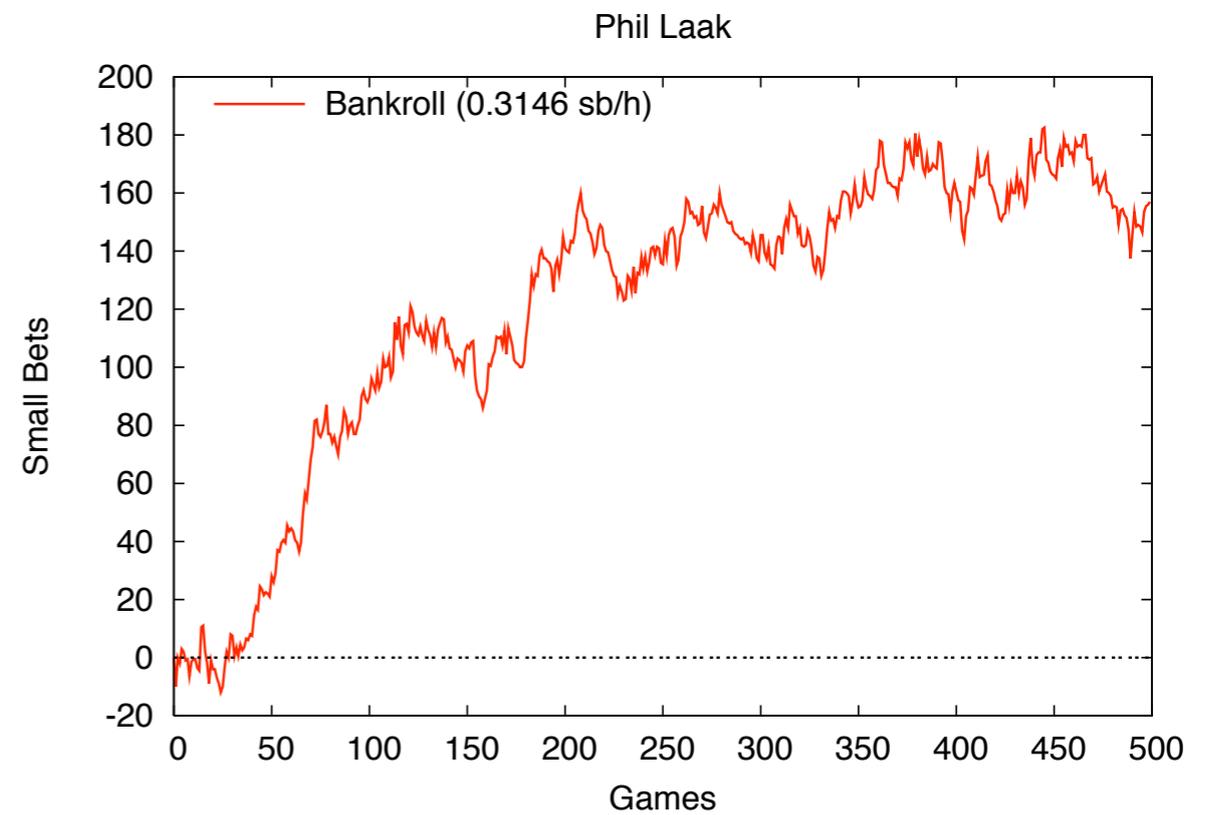
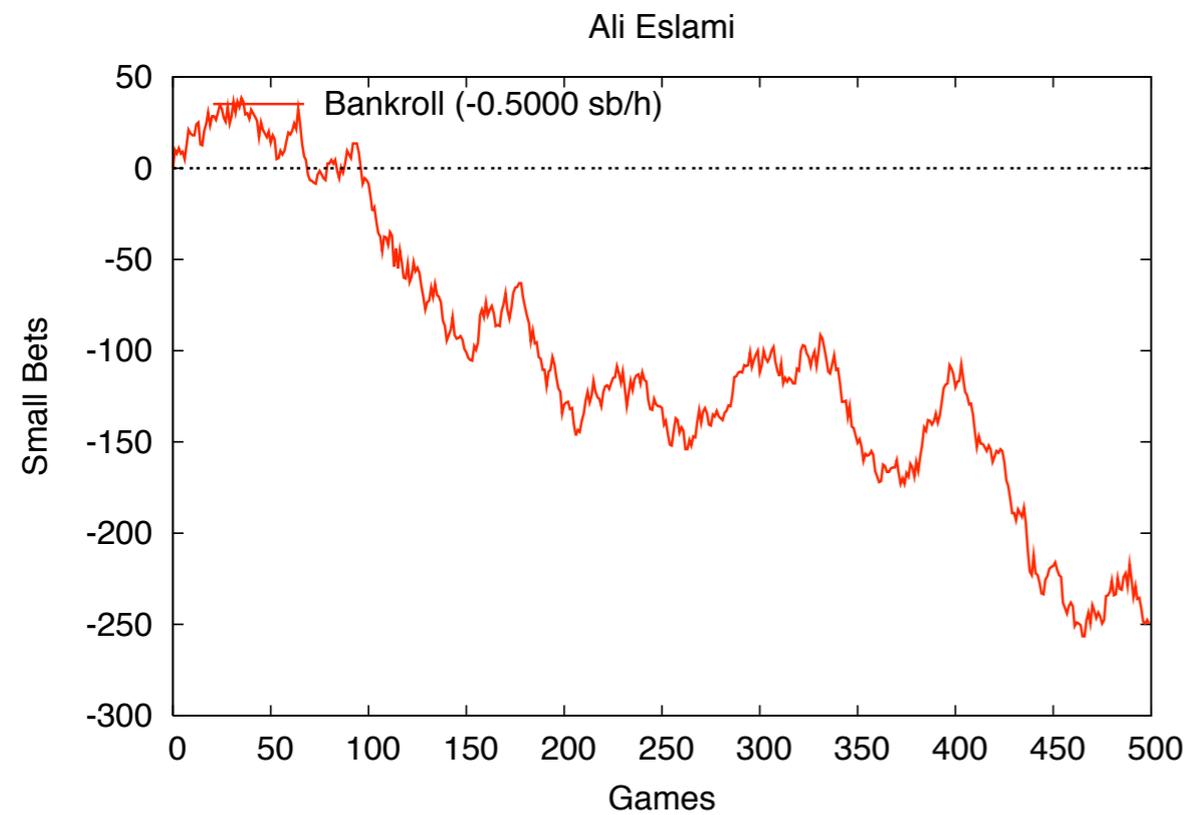
---



- Match 1: Polaris up by 7 bets – Draw

# The Results

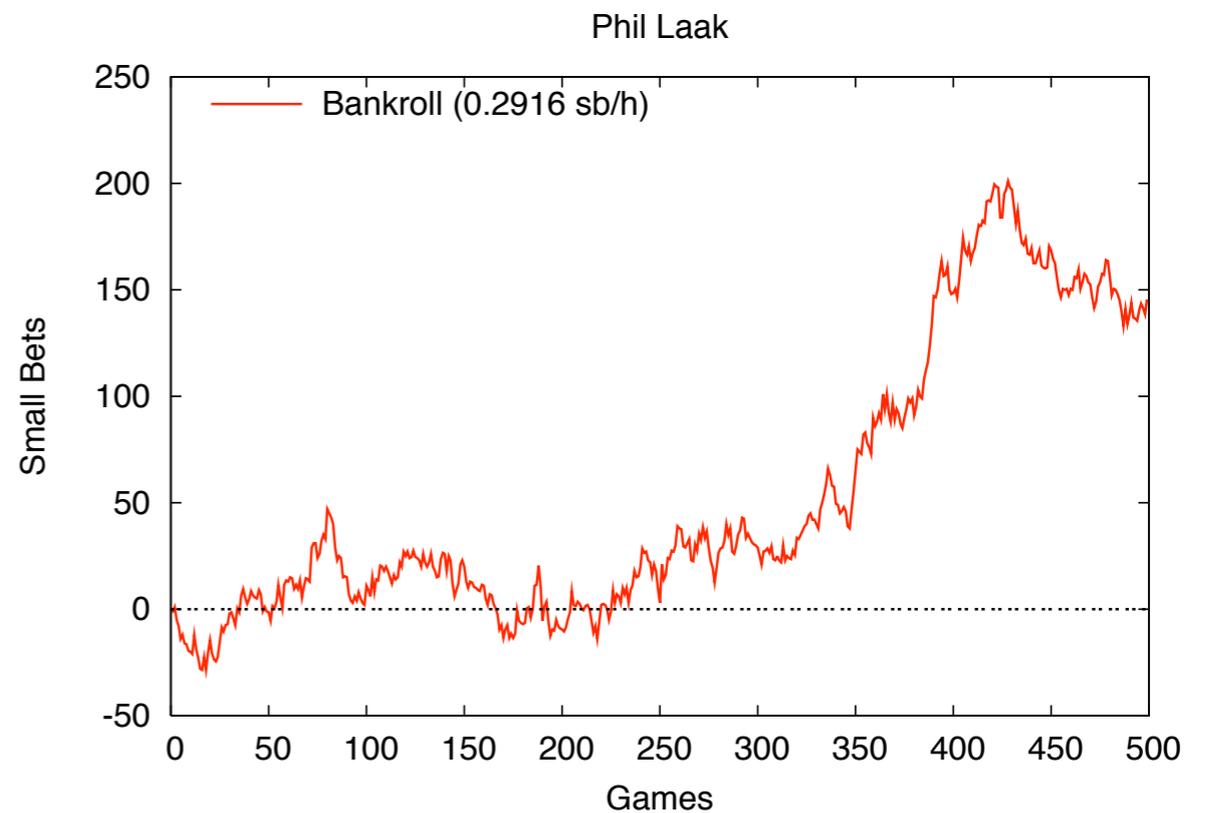
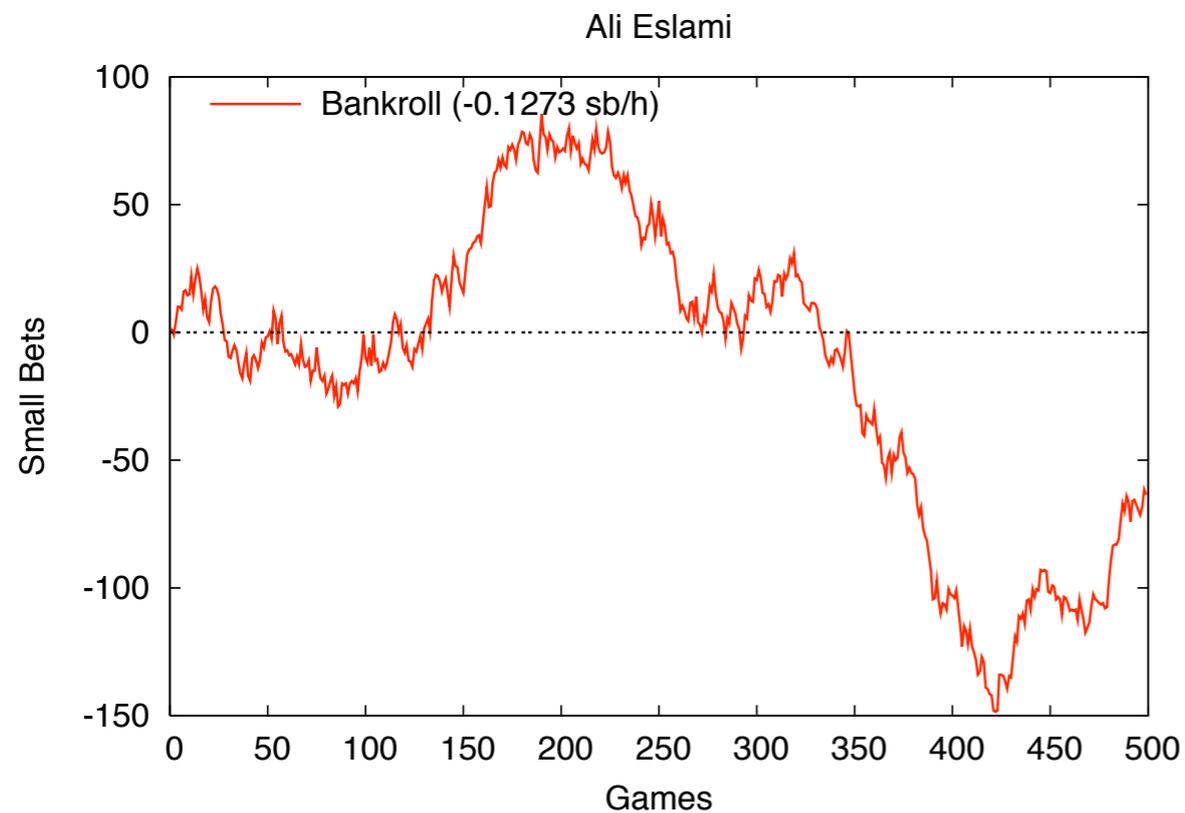
---



- Match 1: Polaris up by 7 bets – Draw
- Match 2: Polaris up by 93 bets – Win

# The Results

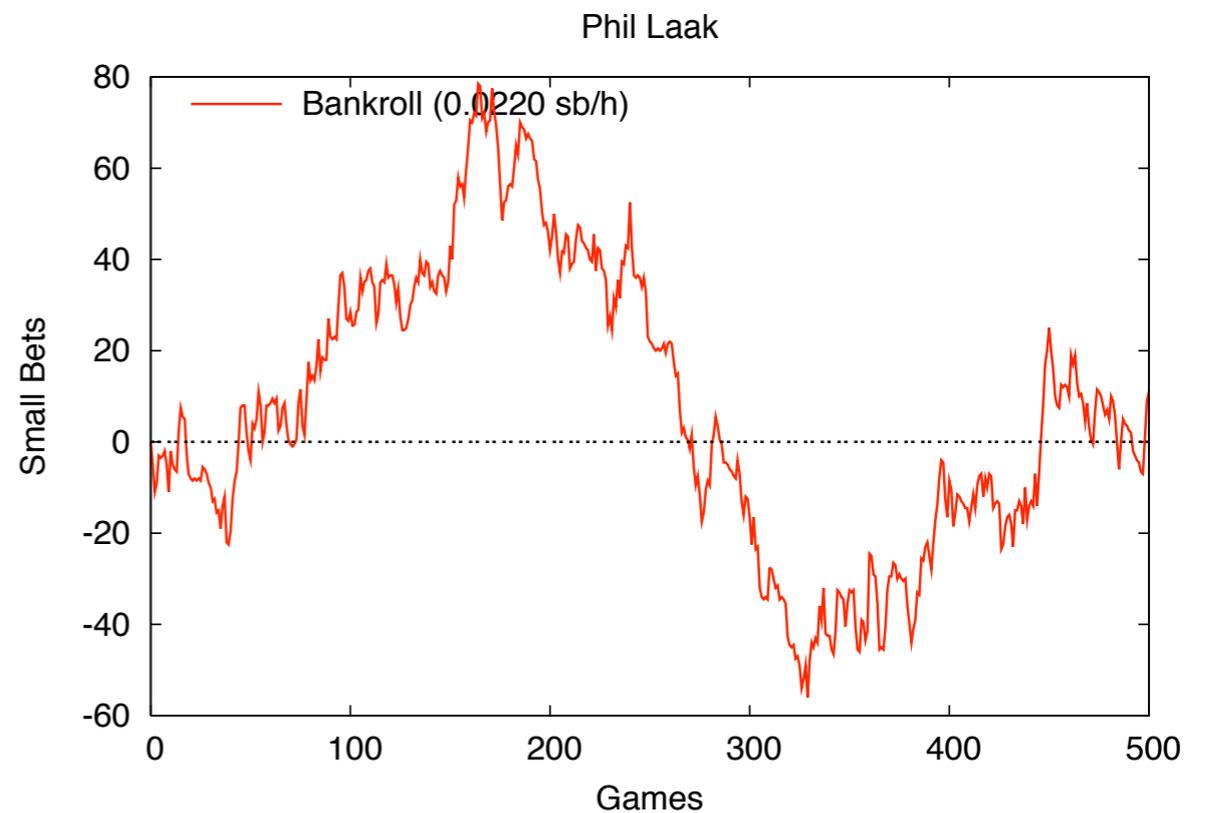
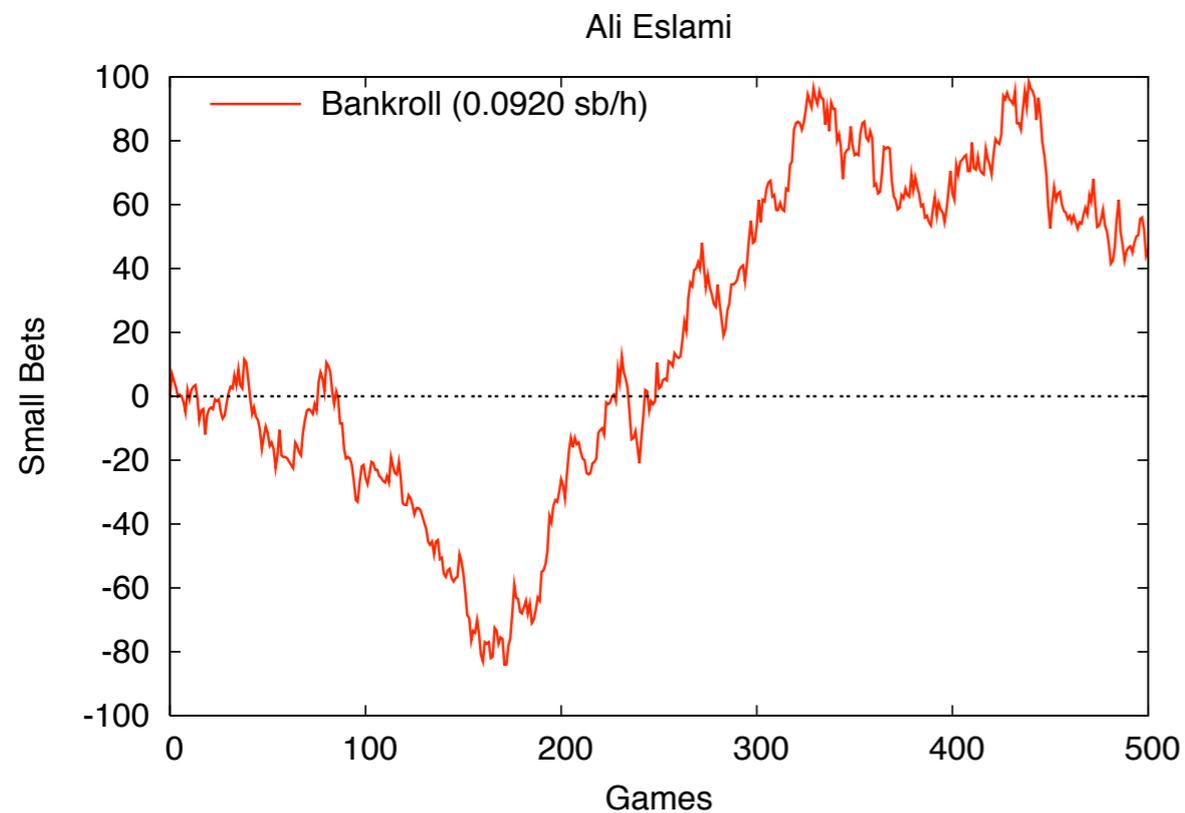
---



- Match 1: Polaris up by 7 bets – Draw
- Match 2: Polaris up by 93 bets – Win
- Match 3: Polaris down by 82 bets – Loss

# The Results

---



- Match 1: Polaris up by 7 bets – Draw
- Match 2: Polaris up by 93 bets – Win
- Match 3: Polaris down by 82 bets – Loss
- Match 4: Polaris down by 57 bets – Loss

# The Results

---

- 1 Win, 2 Losses, 1 Draw
- First Law of Poker

# The Results

---

- 1 Win, 2 Losses, 1 Draw
- First Law of Poker

*All players emphasize the skill in poker when they win and the luck when they lose.*

# The Results

---

- 1 Win, 2 Losses, 1 Draw
- First Law of Poker

*All players emphasize the skill in poker when they win and the luck when they lose.*

- The humans got lucky!

# The Results

---

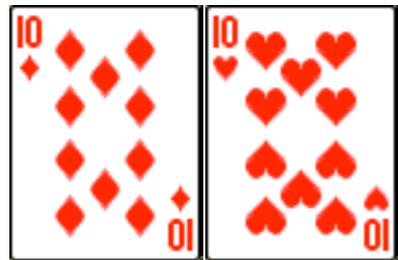
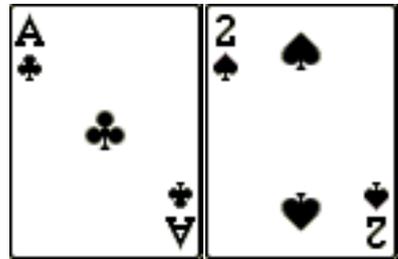
- 1 Win, 2 Losses, 1 Draw
- First Law of Poker

*All players emphasize the skill in poker when they win and the luck when they lose.*

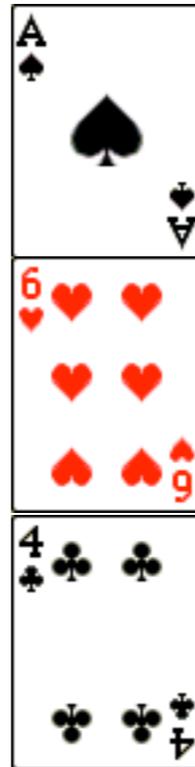
- The humans got lucky!
  - None of the match results were statistically significant
  - Need 20K hands without duplicate!
  - Need 2K pairs of hands with duplicate!

# Removing More Luck: DIVAT

---



check  
bet  
raise  
call



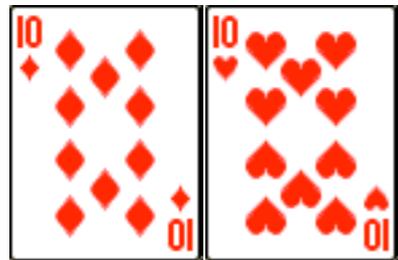
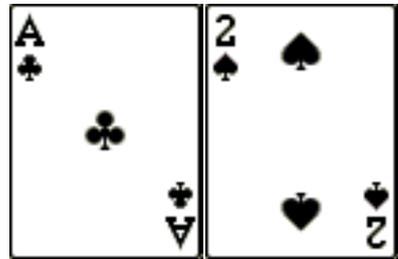
bet  
call



bet  
fold

# Removing More Luck: DIVAT

---



+3 sb

check  
bet  
raise  
call



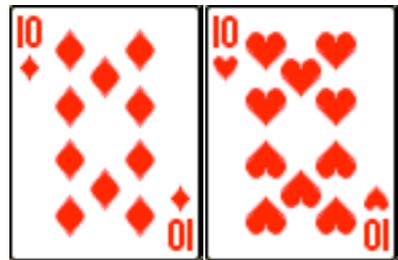
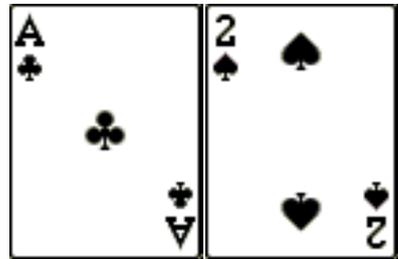
bet  
call



bet  
fold

# Removing More Luck: DIVAT

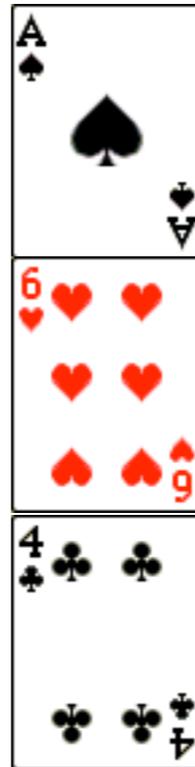
---



+3 sb

check  
bet  
raise  
call

+4 sb



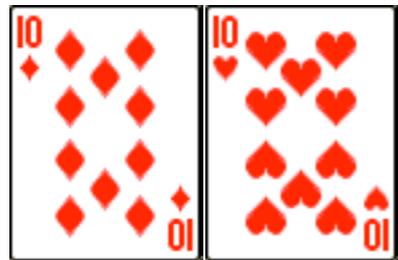
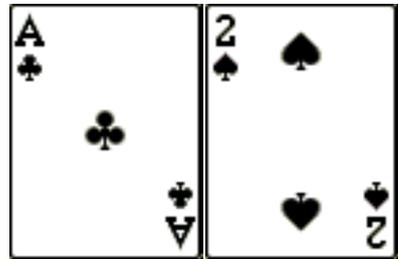
bet  
call



bet  
fold

# Removing More Luck: DIVAT

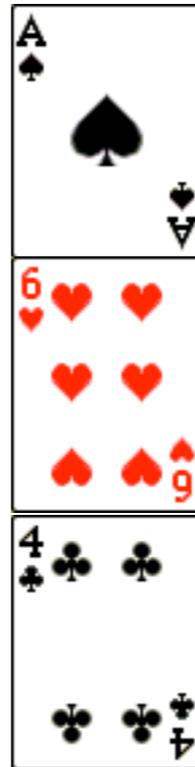
---



+3 sb

check  
bet  
raise  
call

+4 sb



-8 sb

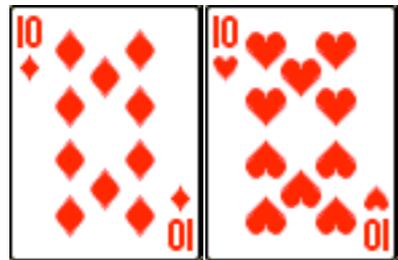
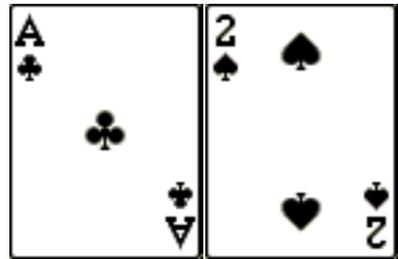
bet  
call



bet  
fold

# Removing More Luck: DIVAT

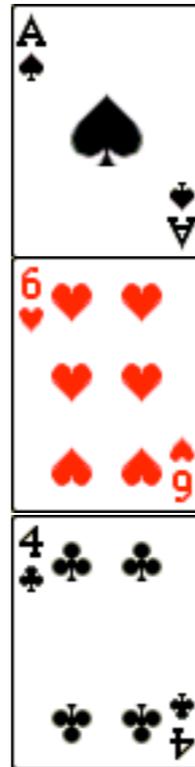
---



+3 sb

check  
bet  
raise  
call

+4 sb



-8 sb

bet  
call

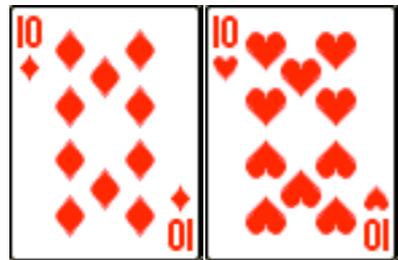
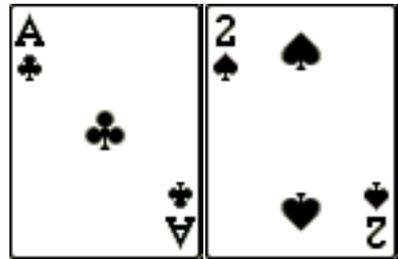
-7 sb



bet  
fold

# Removing More Luck: DIVAT

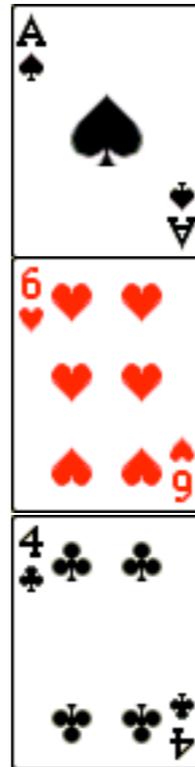
---



+3 sb

check  
bet  
raise  
call

+4 sb



-8 sb

bet  
call

-7 sb

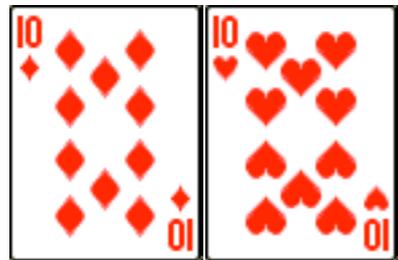
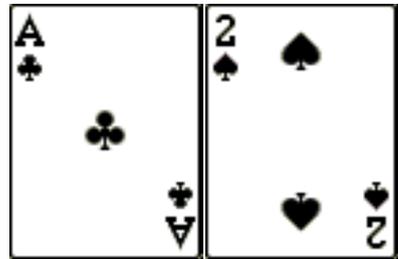


-6 sb

bet  
fold

# Removing More Luck: DIVAT

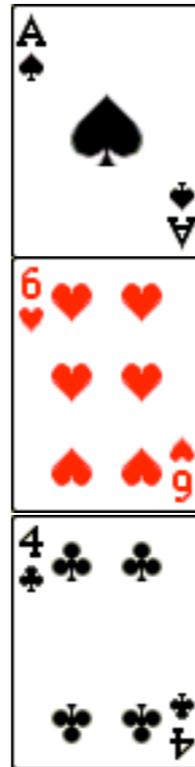
---



+3 sb

check  
bet  
raise  
call

+4 sb



-8 sb

bet  
call

-7 sb

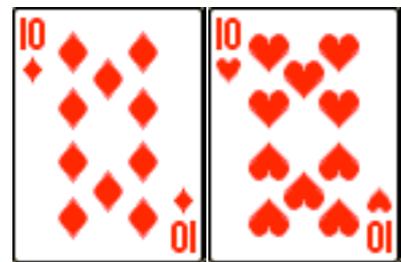
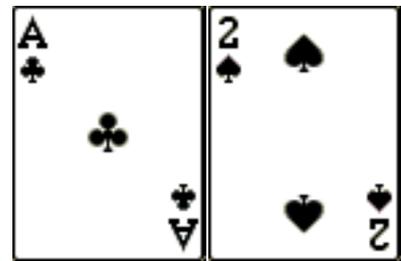


-6 sb

bet  
fold

-4 sb

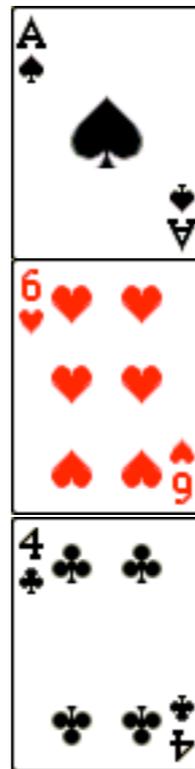
# Removing More Luck: DIVAT



+3 sb

check  
bet  
raise  
call

+4 sb



-8 sb

bet  
call

-7 sb



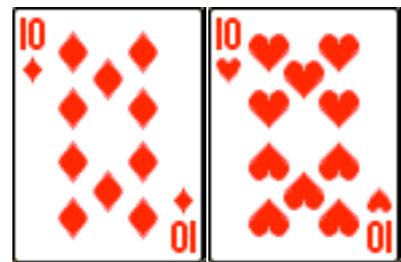
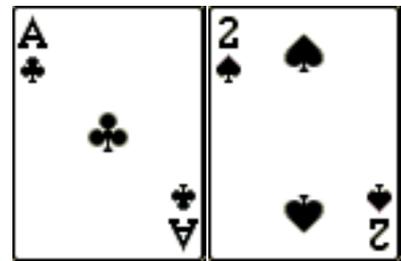
-6 sb

bet  
fold

-4 sb

LUCK

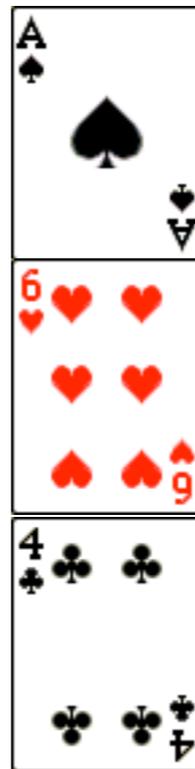
# Removing More Luck: DIVAT



check  
bet  
raise  
call

+3 sb

+4 sb



bet  
call

-8 sb

-7 sb



bet  
fold

-6 sb

-4 sb

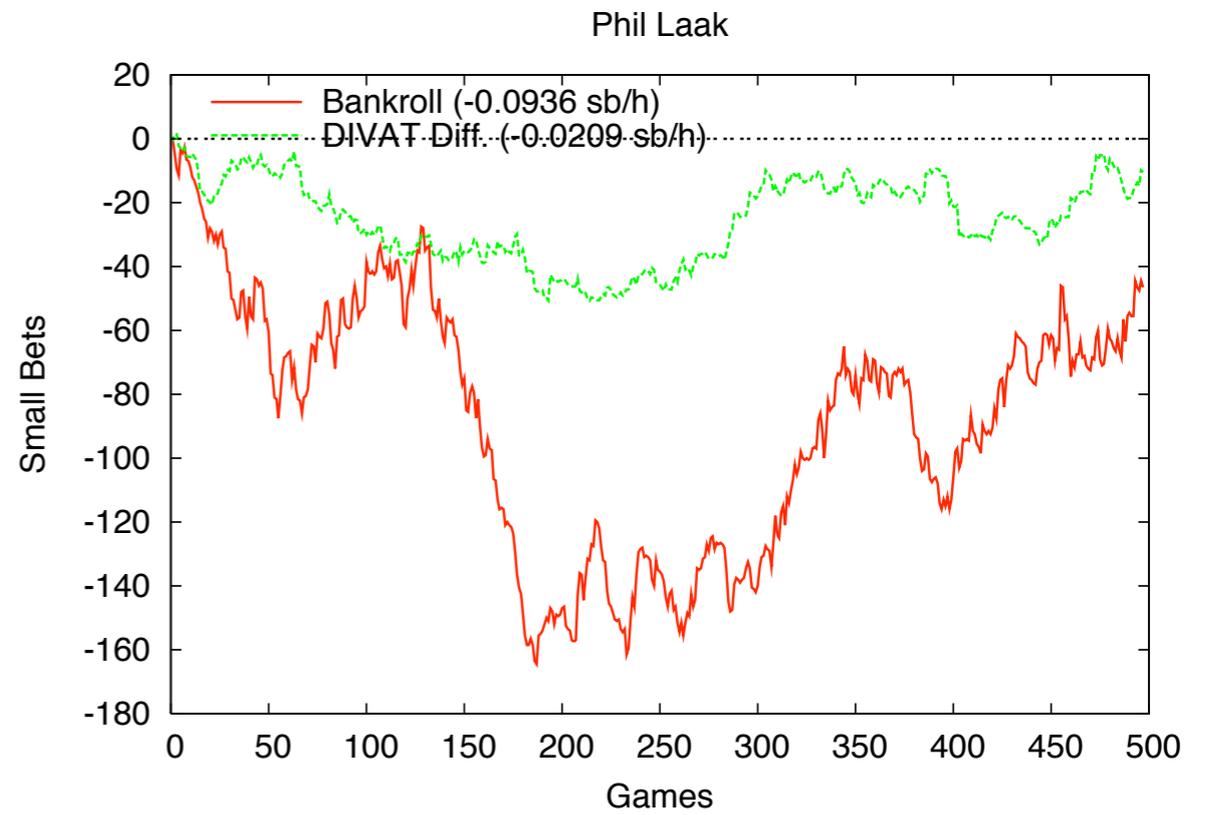
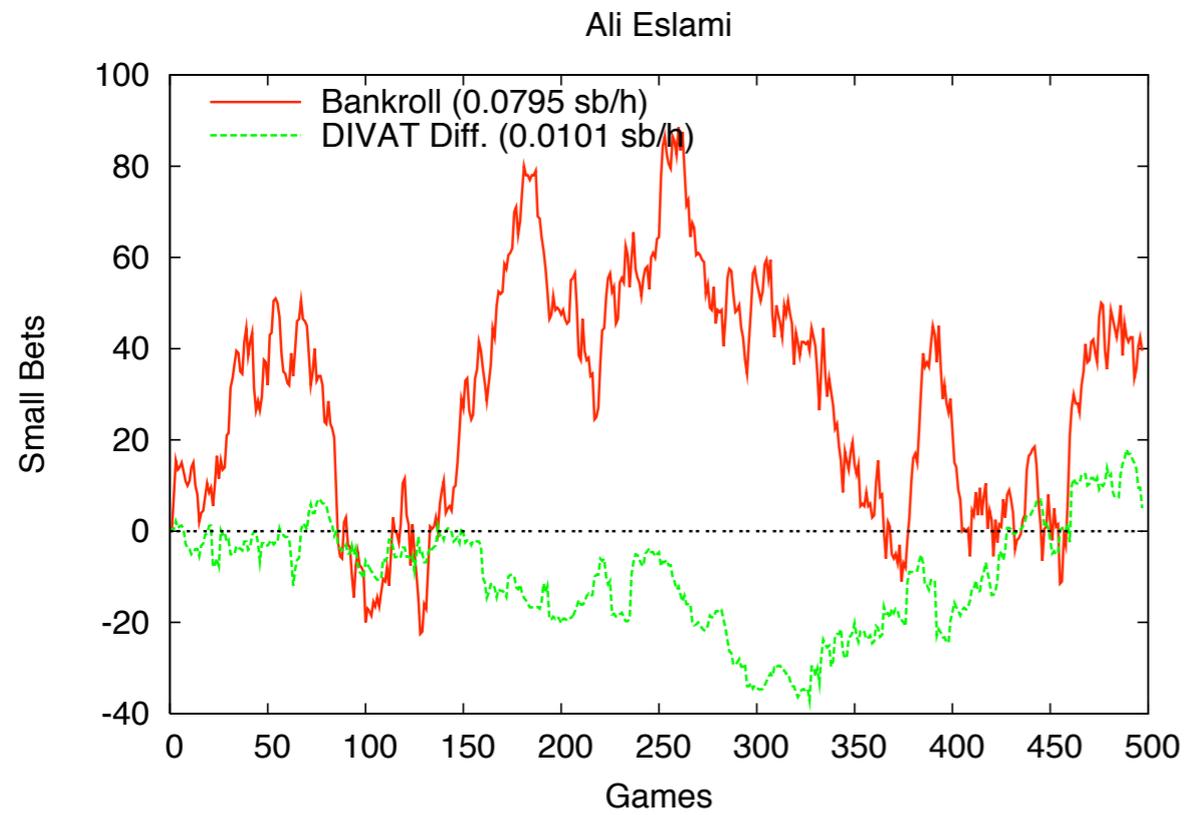
SKILL

# Evaluating Poker Players

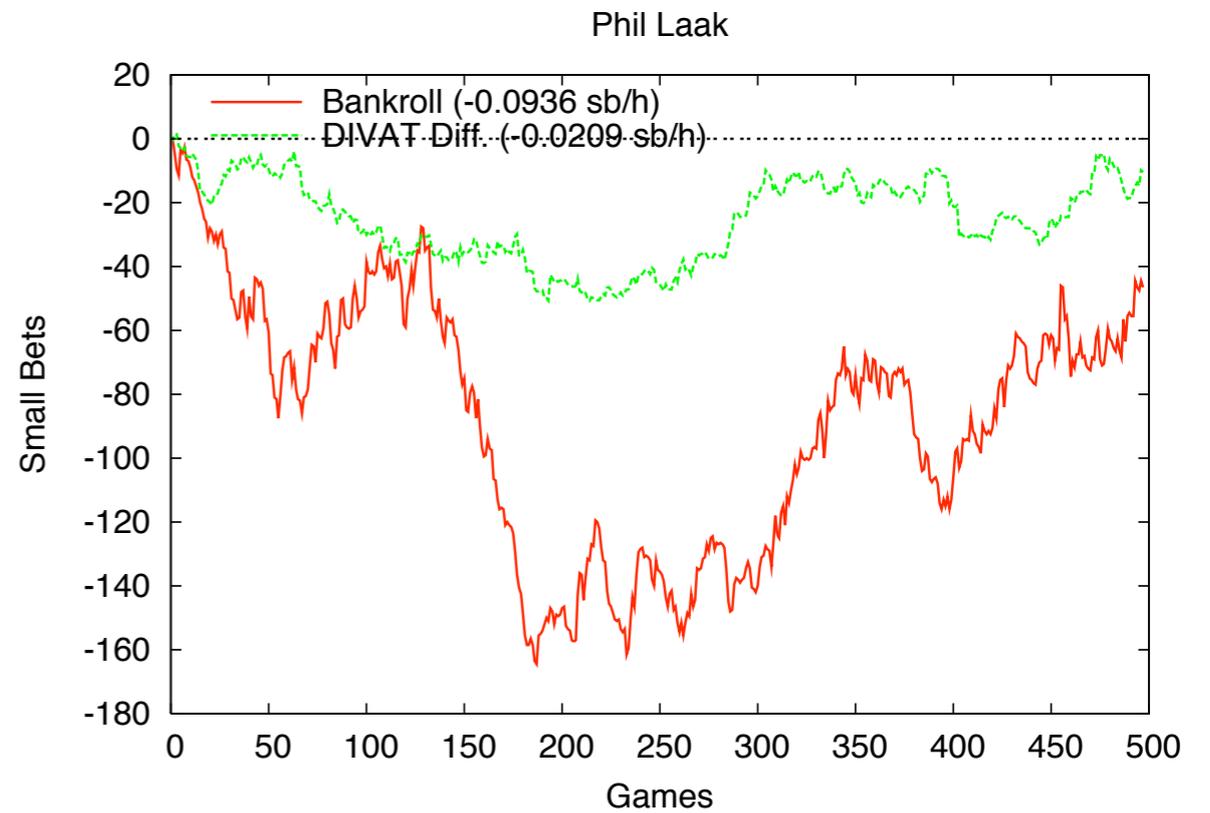
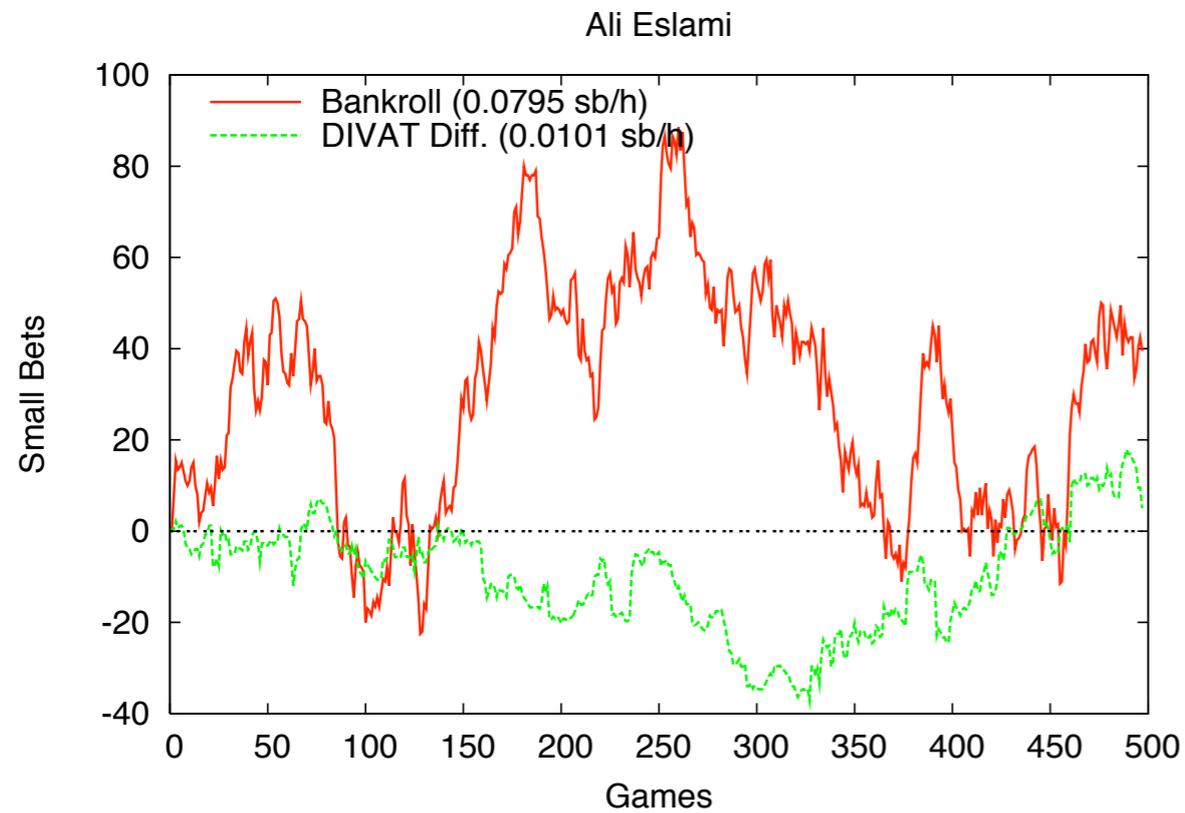
---

- Money is a terrible estimator!
  - ~20K hands needed for statistical significance
- Duplicate helps
  - ~2000 pairs of hands
- DIVAT+Duplicate is almost enough
  - ~1200 pairs of hands

# The Results

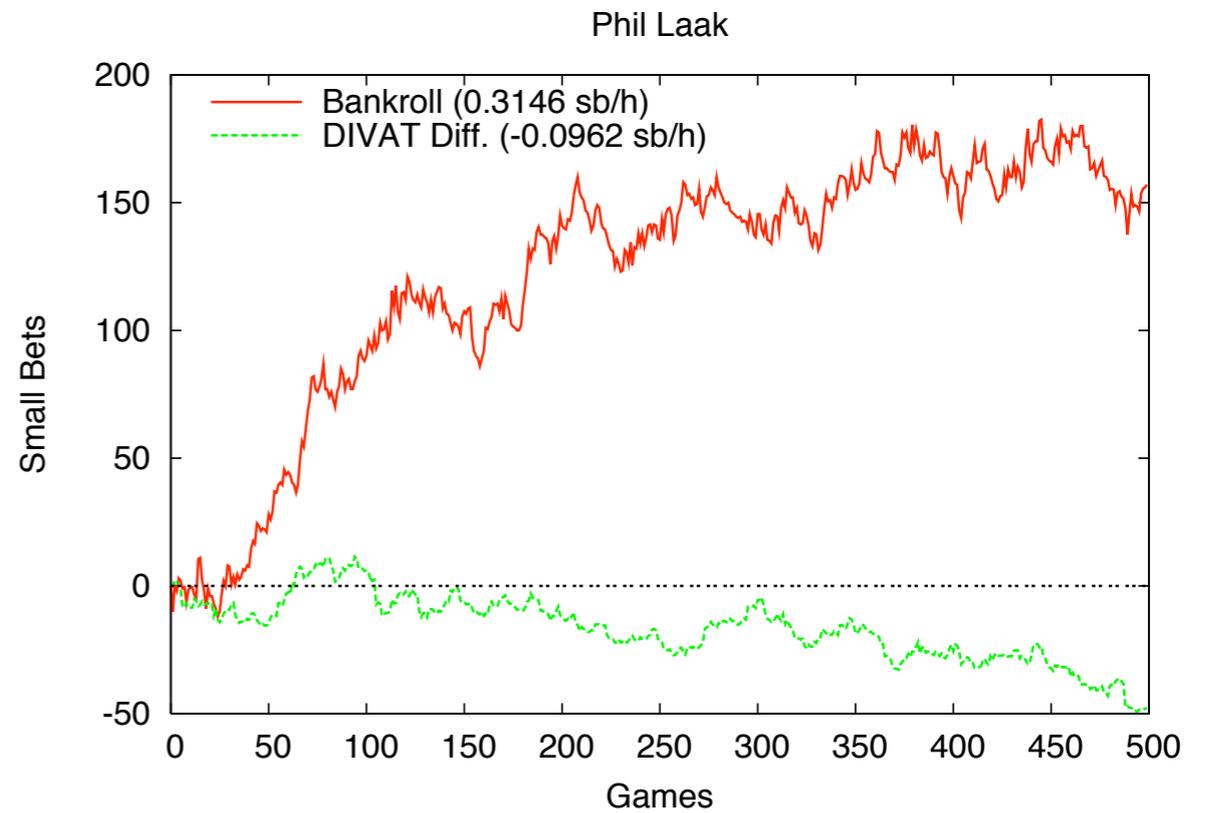
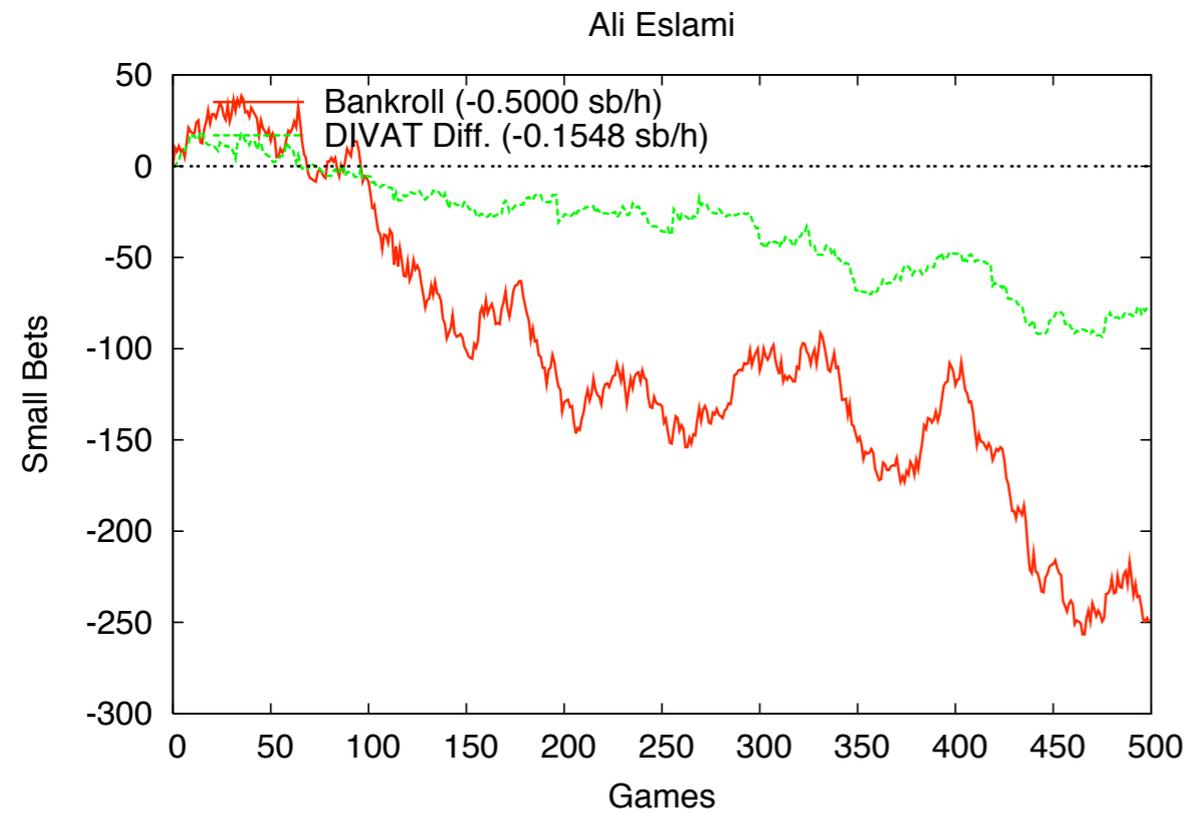


# The Results



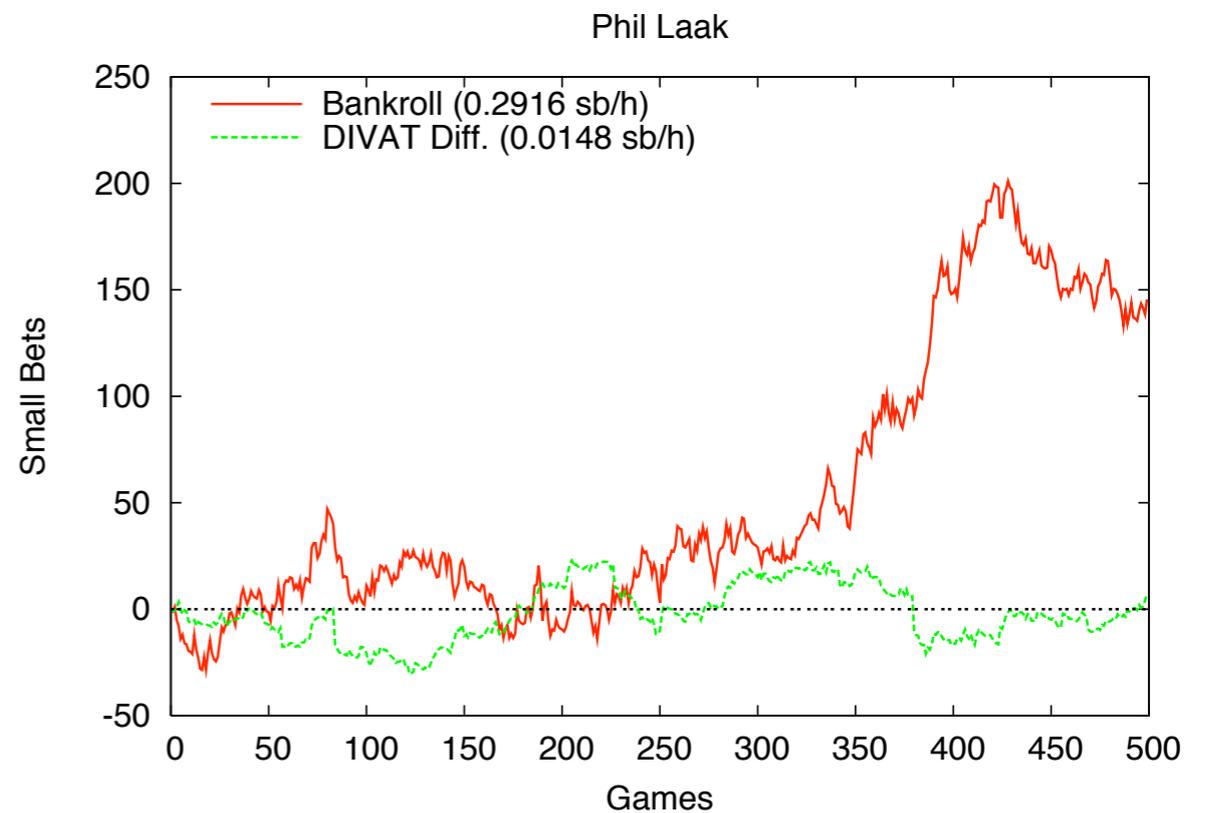
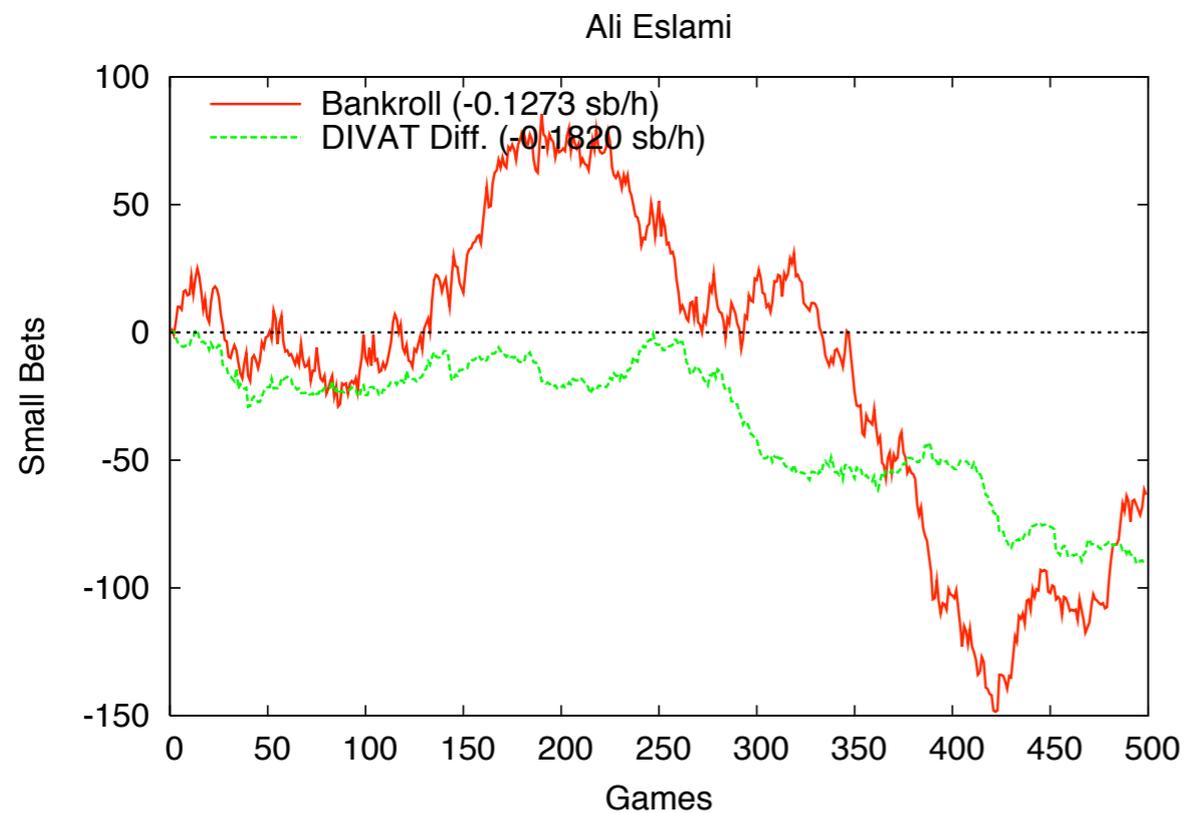
- Match 1: Polaris “up” by 6 bets

# The Results



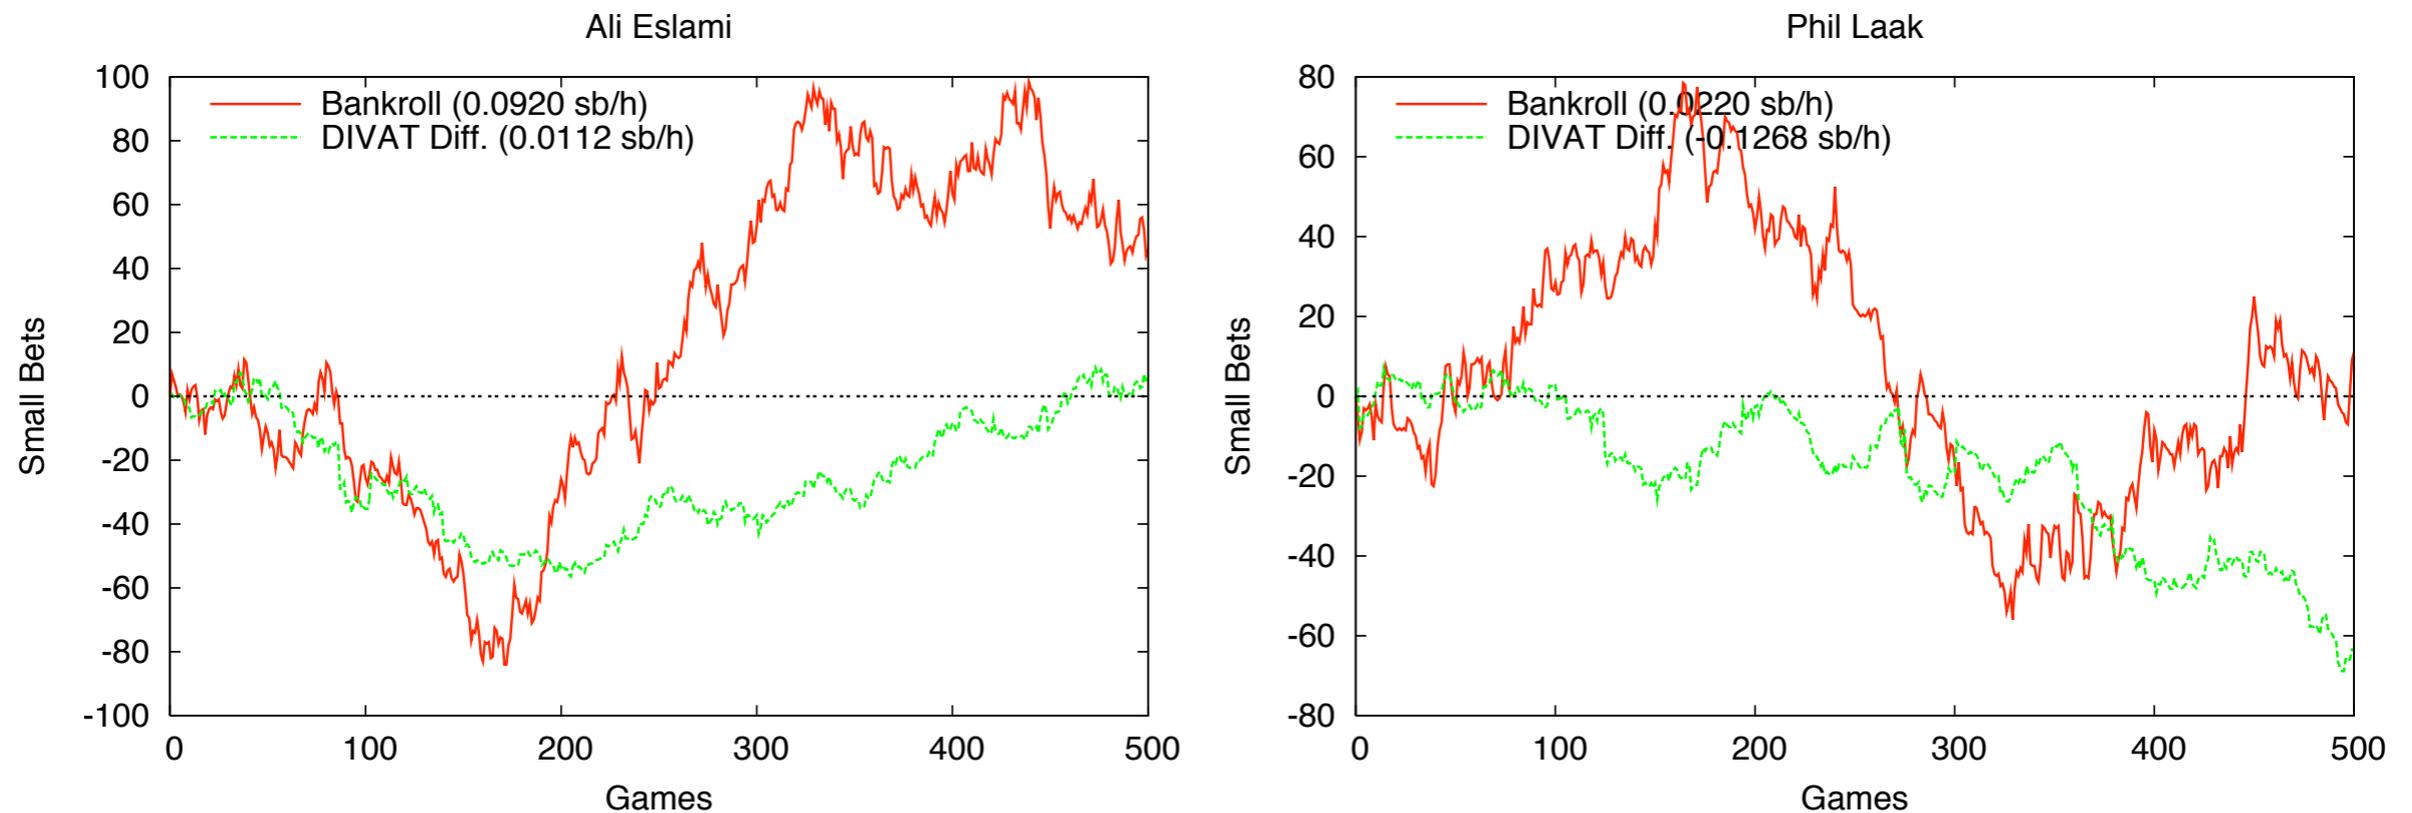
- Match 1: Polaris “up” by 6 bets
- Match 2: Polaris “up” by 126 bets — p-value of 0.0018

# The Results



- Match 1: Polaris “up” by 6 bets
- Match 2: Polaris “up” by 126 bets — p-value of 0.0018
- Match 3: Polaris “up” by 84 bets — p-value of 0.07

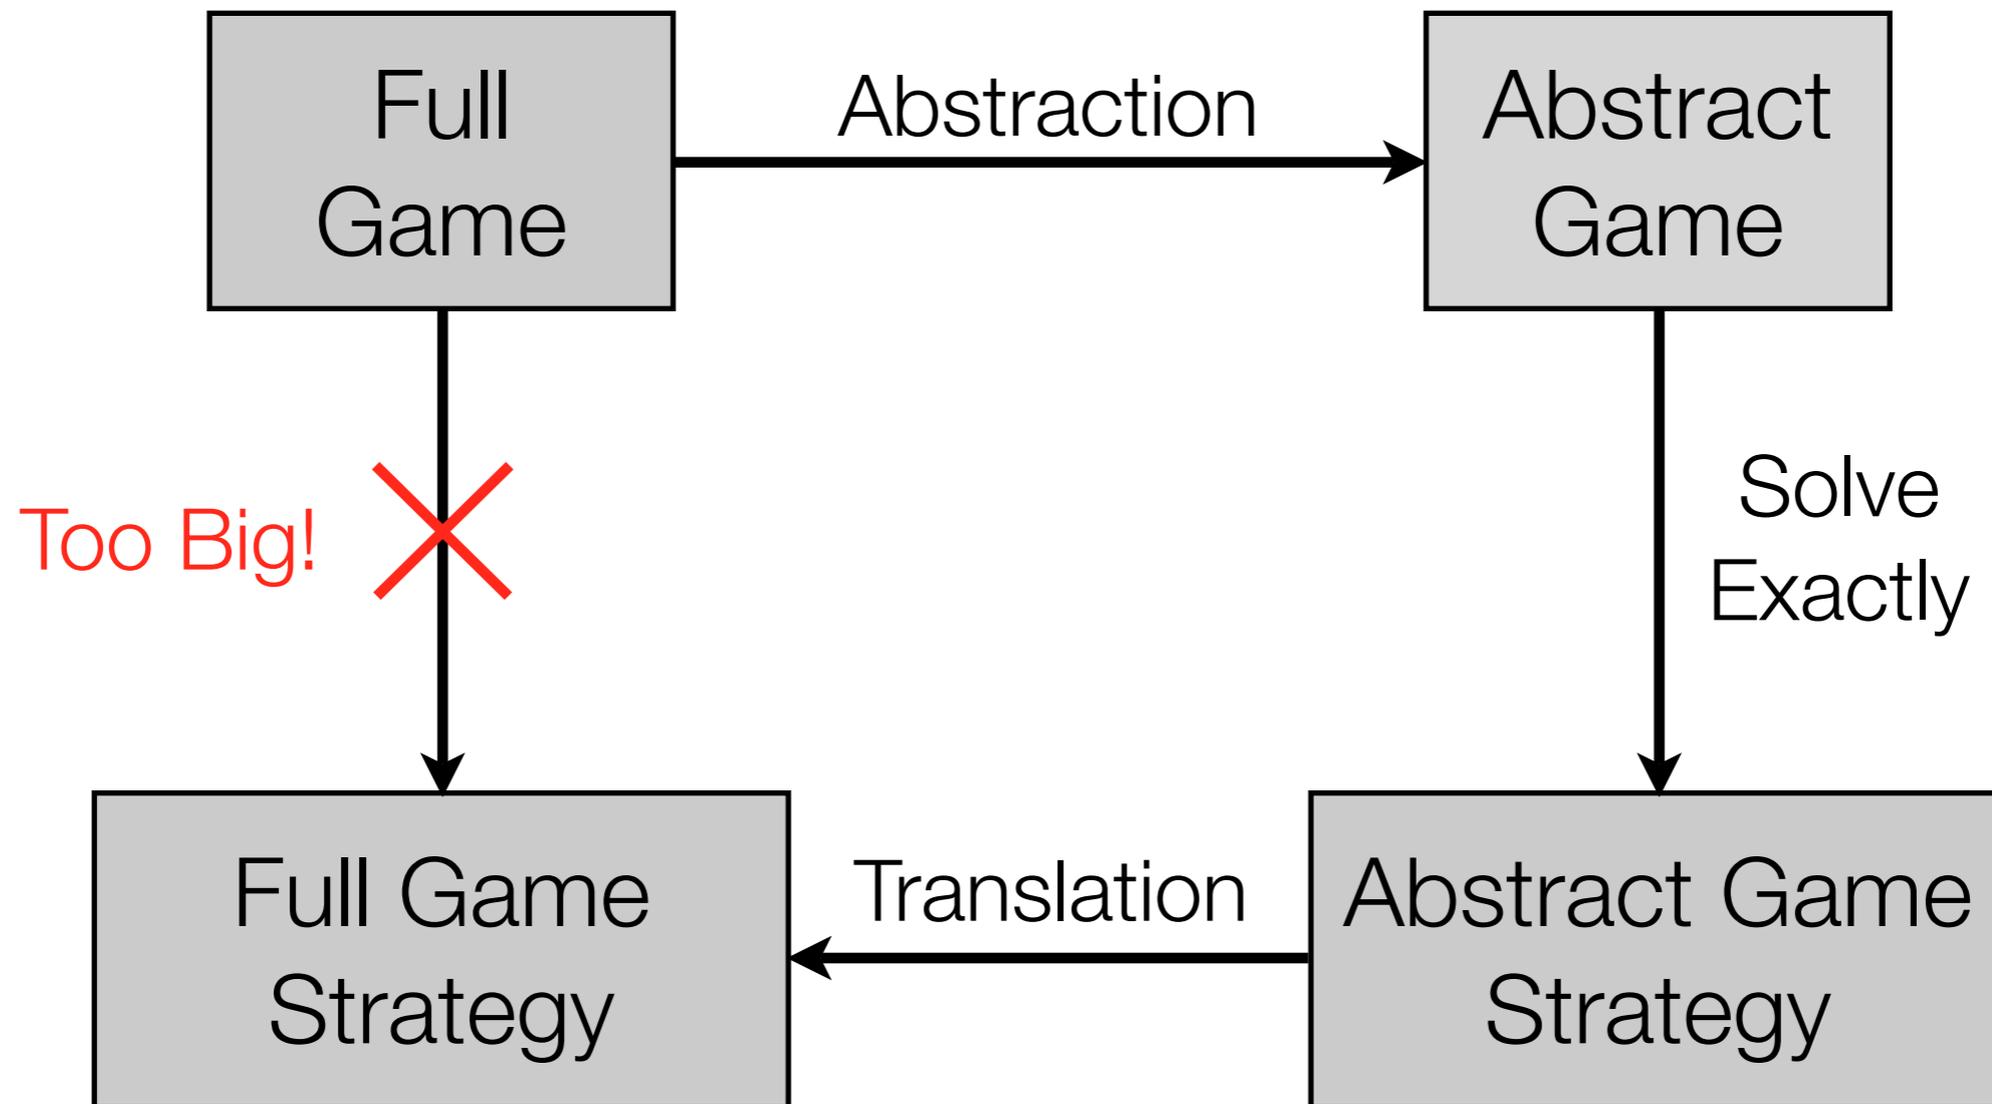
# The Results



- Match 1: Polaris “up” by 6 bets
- Match 2: Polaris “up” by 126 bets — p-value of 0.0018
- Match 3: Polaris “up” by 84 bets — p-value of 0.07
- Match 4: Polaris “up” by 58 bets — p-value of 0.14

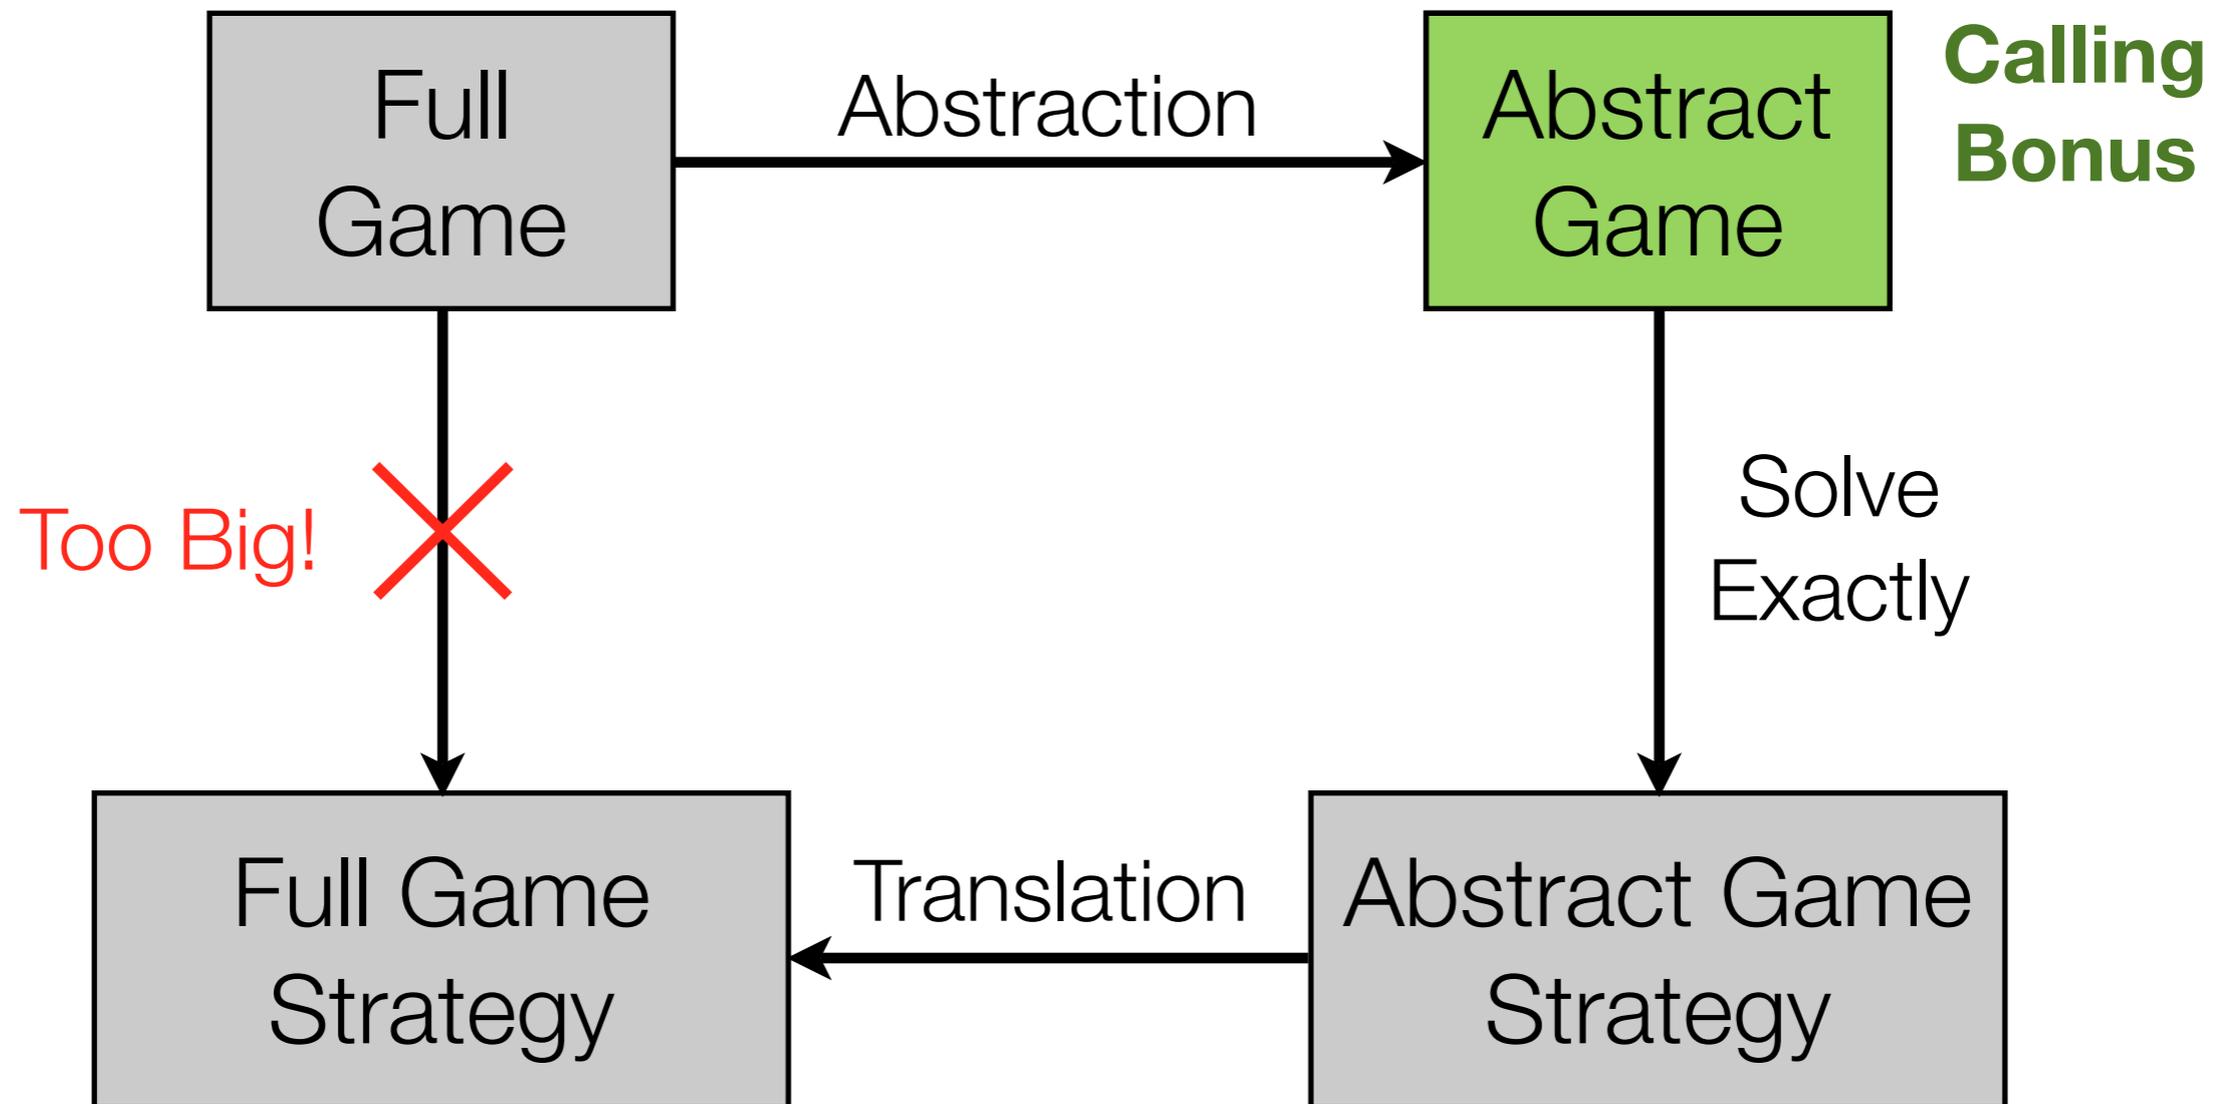
# One More Year of Research... In Two More Slides

---



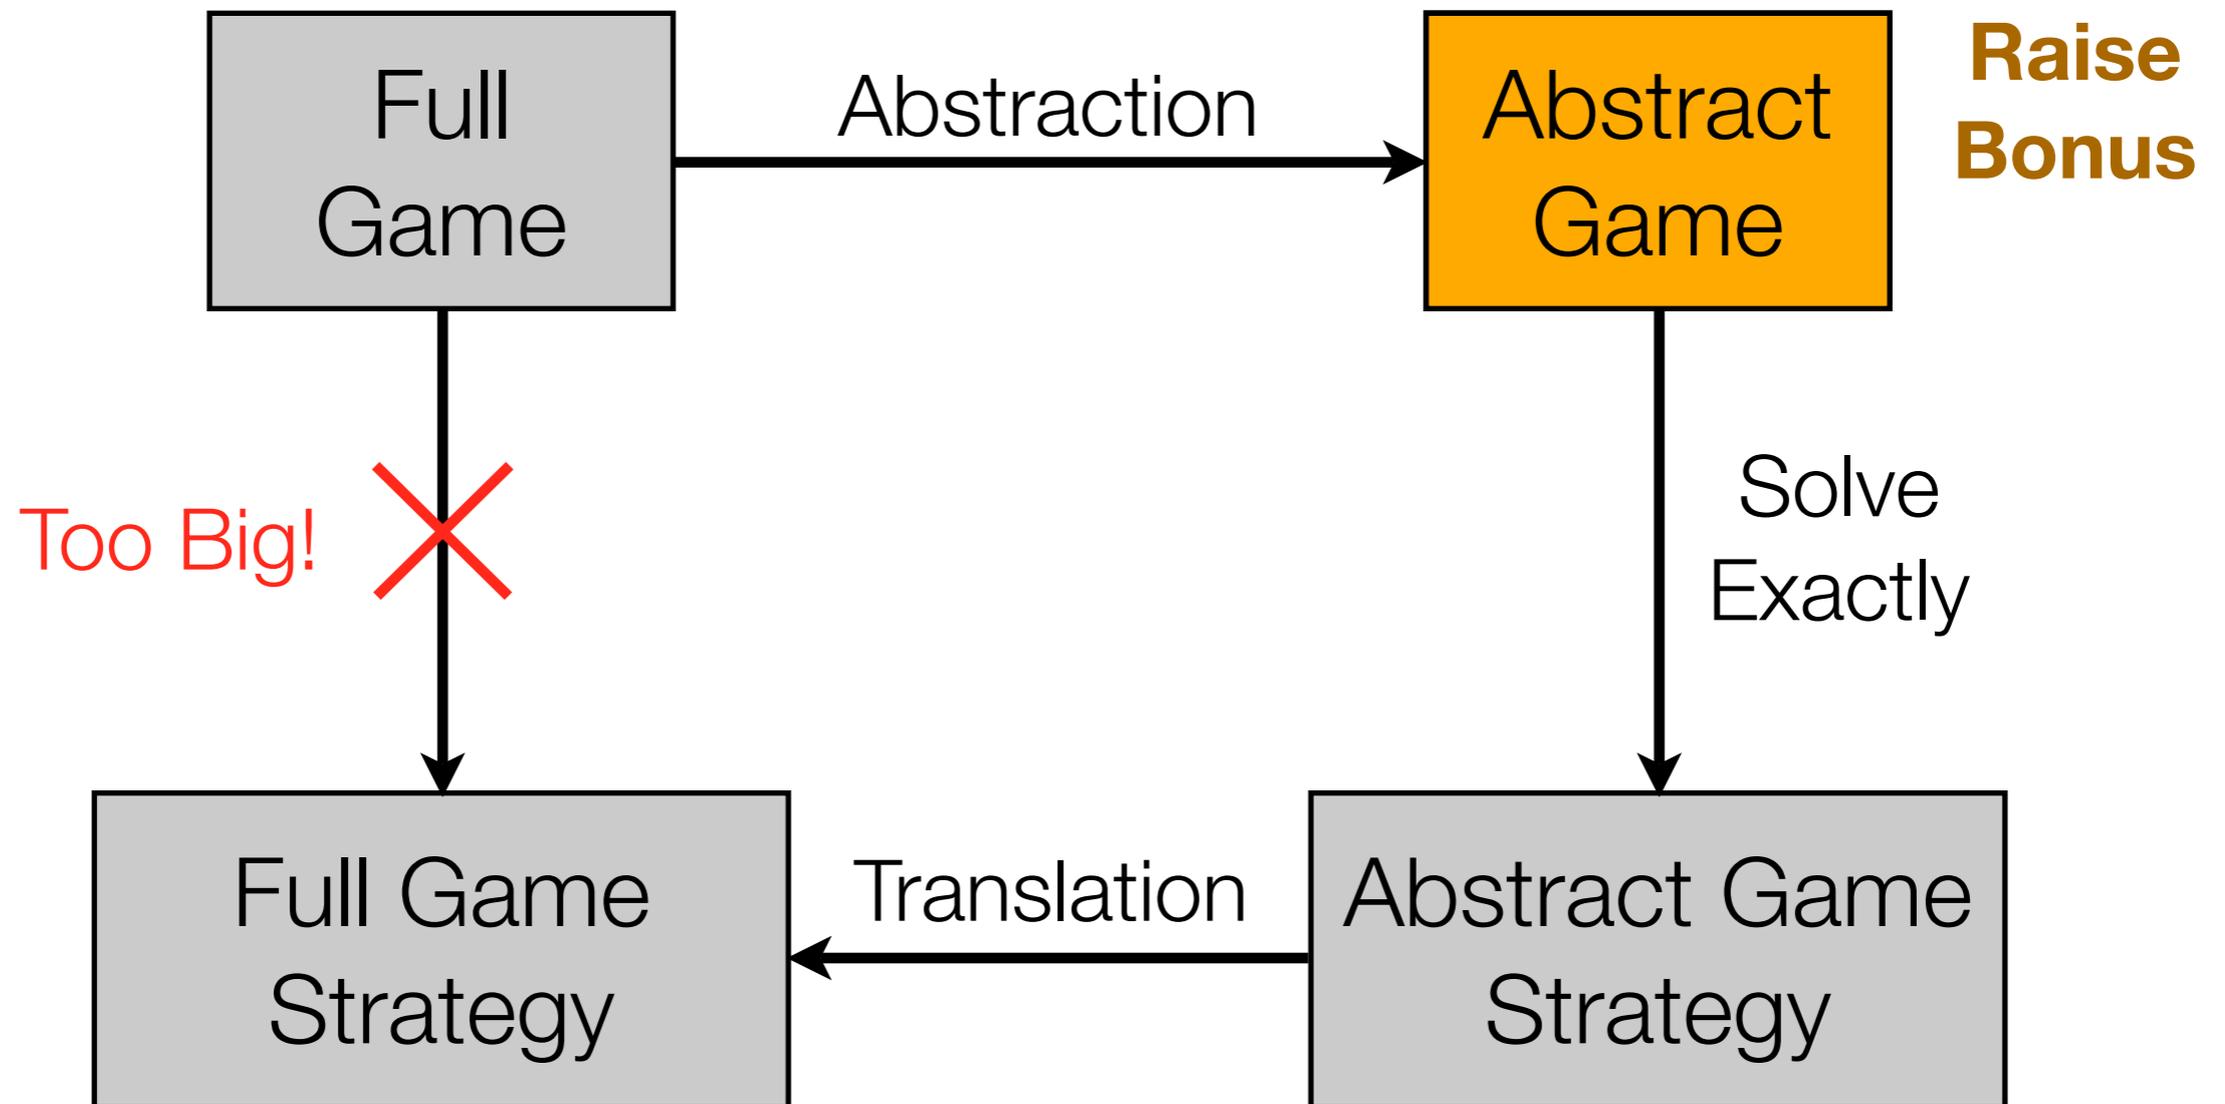
# One More Year of Research... In Two More Slides

---



# One More Year of Research... In Two More Slides

---



# Polaris 2.0

---

- 5 personalities: varying degrees of aggression

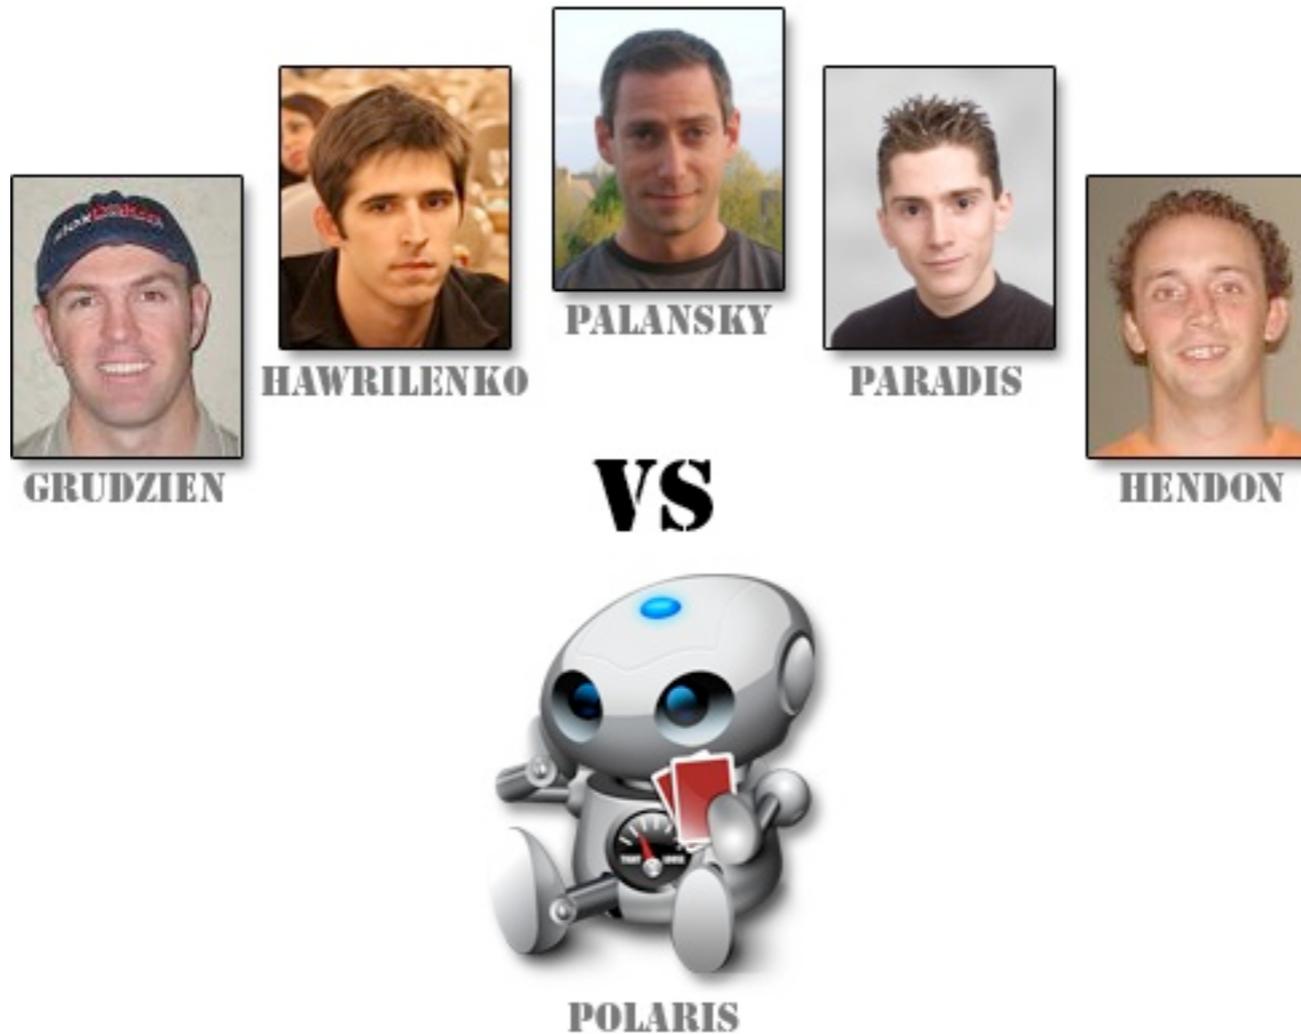


- Each personality is a solution to a game with  $10^{12}$  states
- Adapts online to the personality most effective against its opponent's style

# The Second Man-Machine Poker Championship

---

- Las Vegas, July 2008



# The Competitors

---

- 7 Professionals: Hawrilenko, Grudzien, Palanski, Hendon, Newhouse, McRoberts, Acosta
- Coaches with Stoxpoker.com
- Matt “Hoss\_TBF” Hawrilenko
  - World’s best current heads-up, limit player
  - Plays at \$1K/\$2K games online



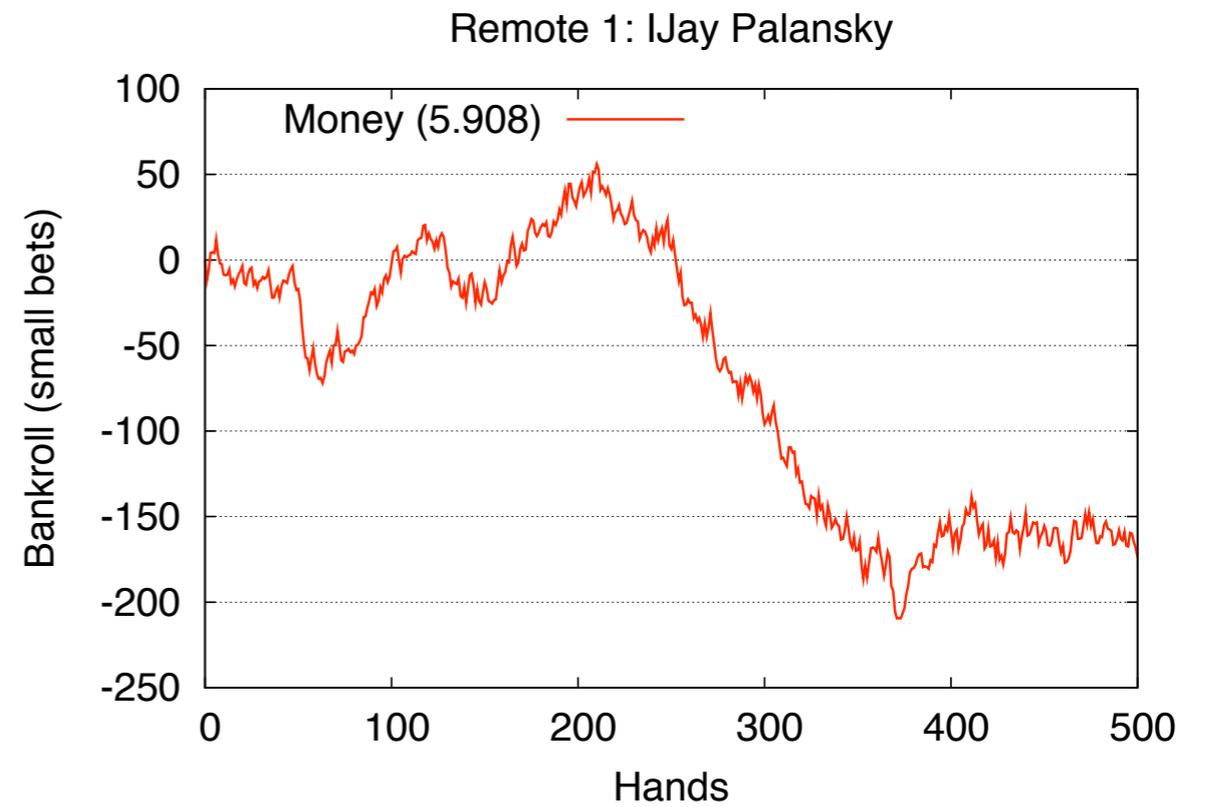
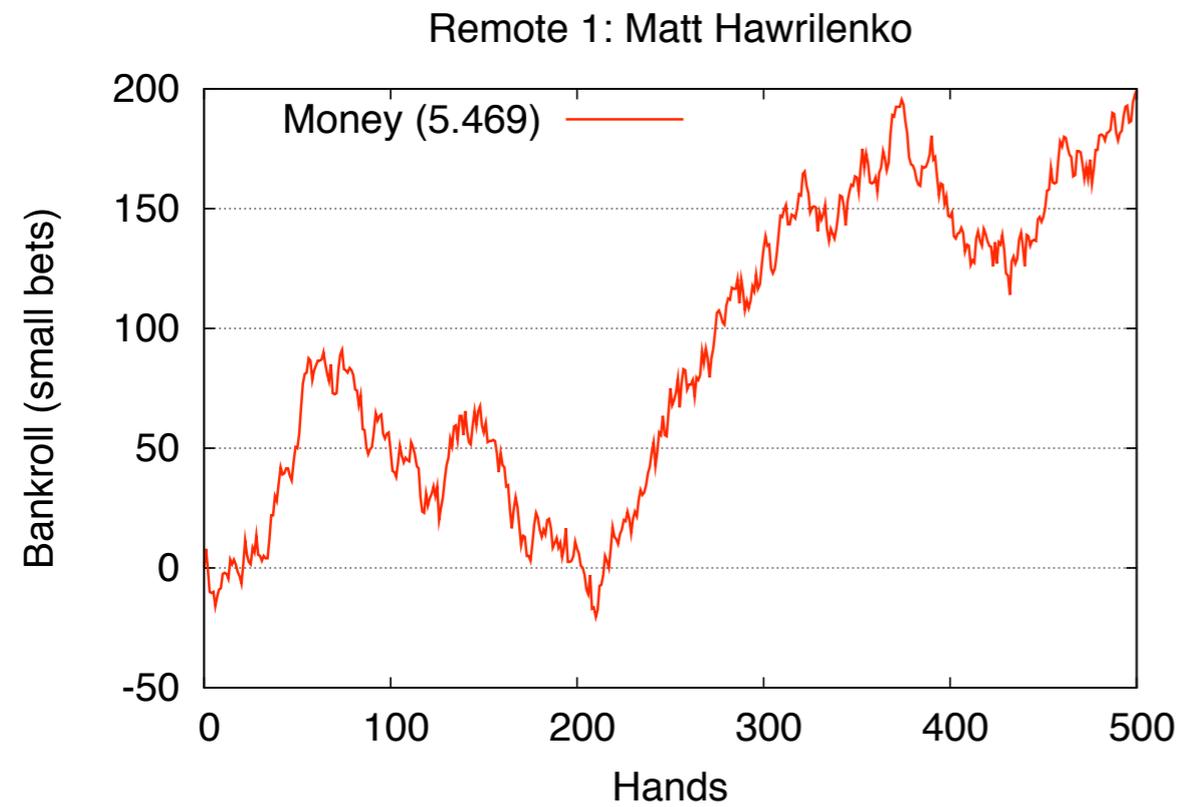
# The Rules

---

- 500 hand duplicate matches
- 2 remote, 4 live (1 each day)
- Must win by \$25K (\$1K/\$2K Limits) to win the match

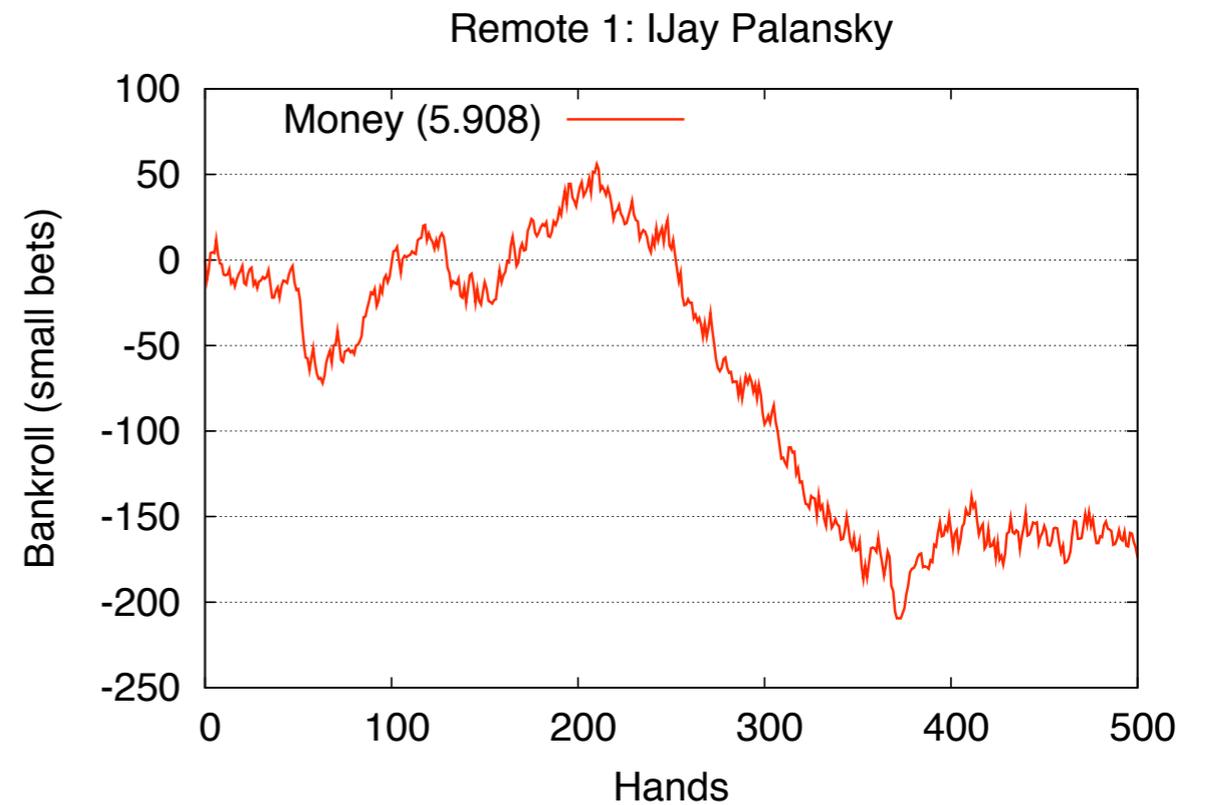
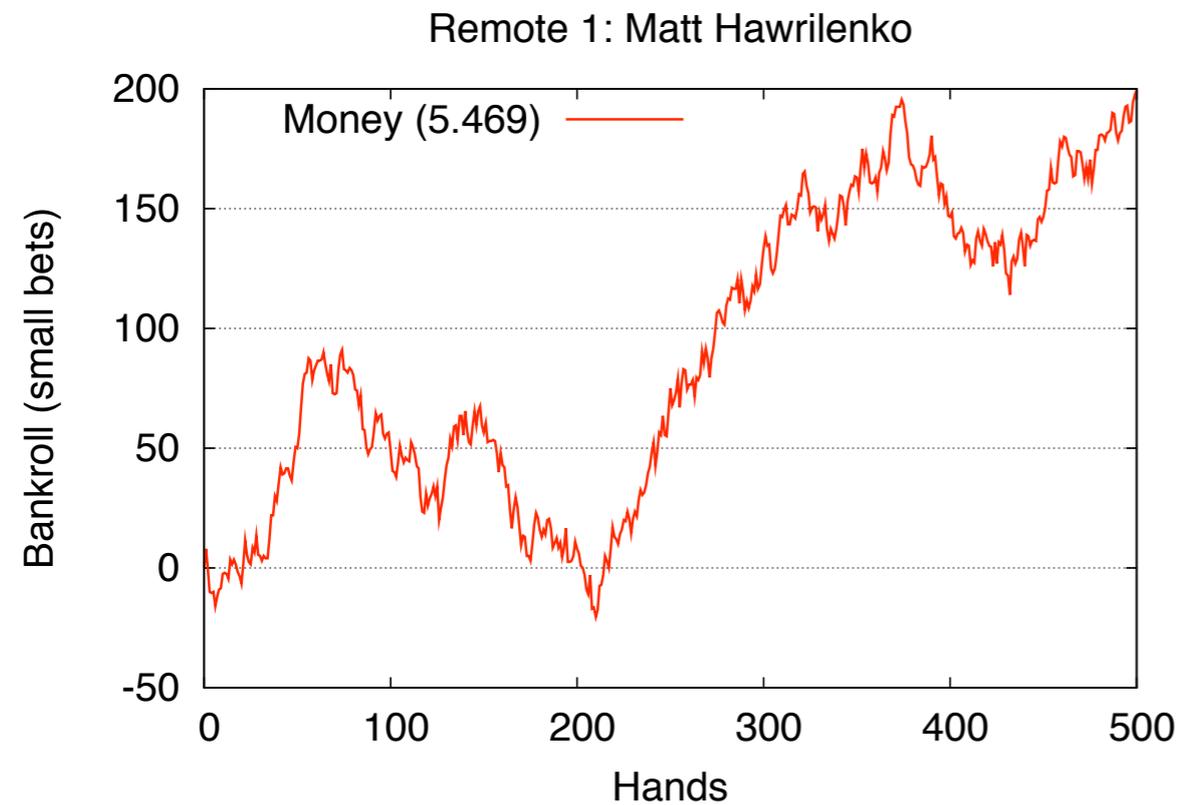
# Remote Match 1: Hoss-IJay

---



# Remote Match 1: Hoss-IJay

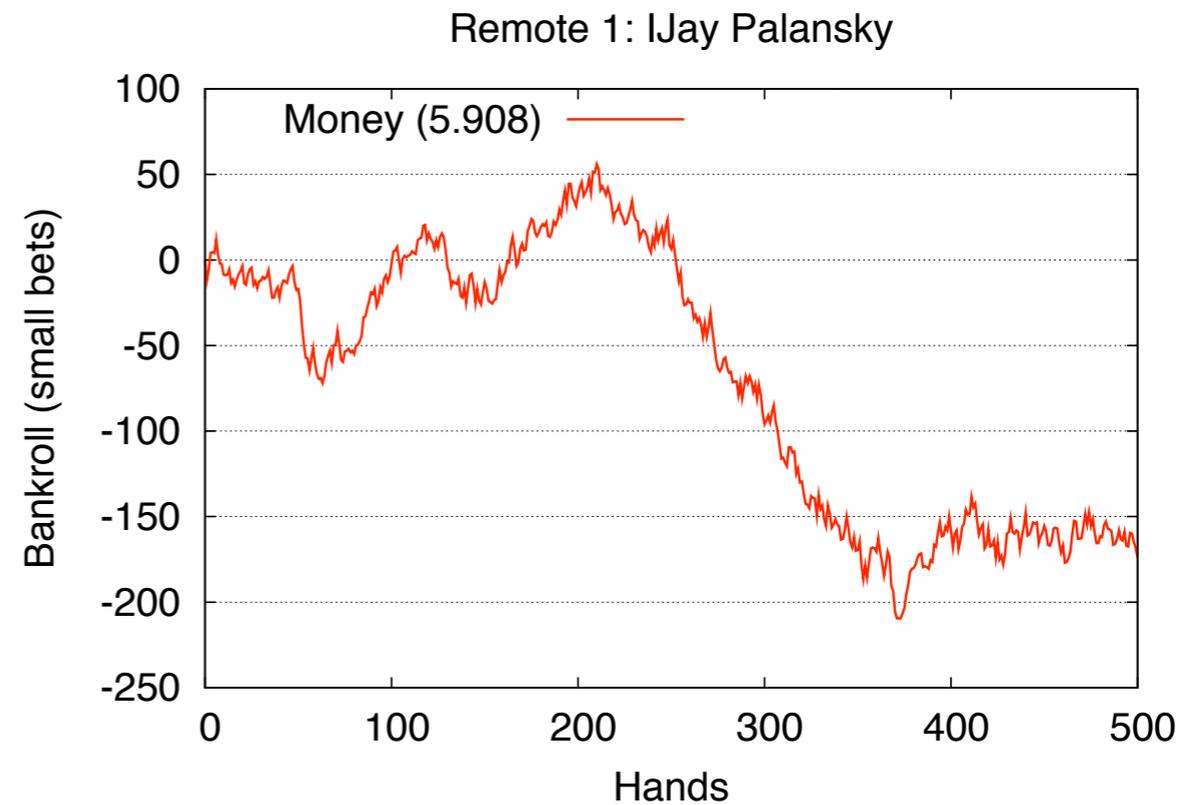
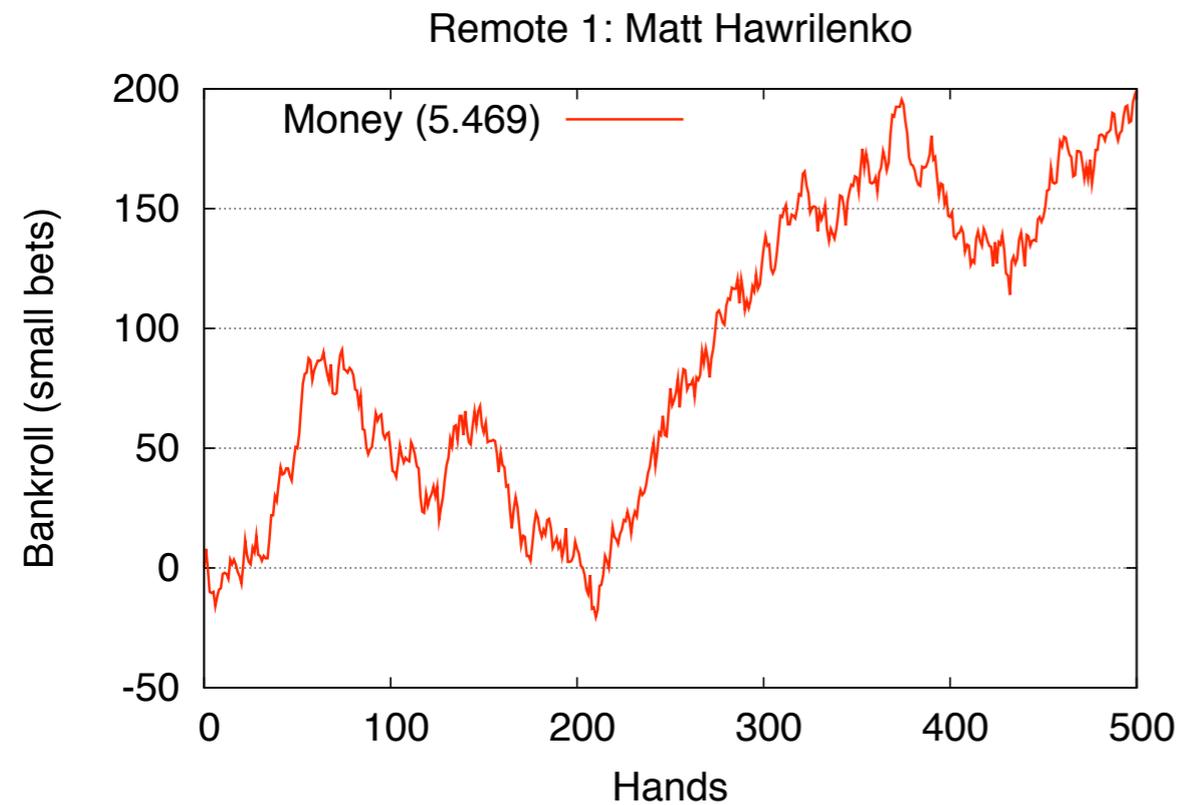
---



- Humans up by 25.5 bets — Polaris loses!

# Remote Match 1: Hoss-IJay

---

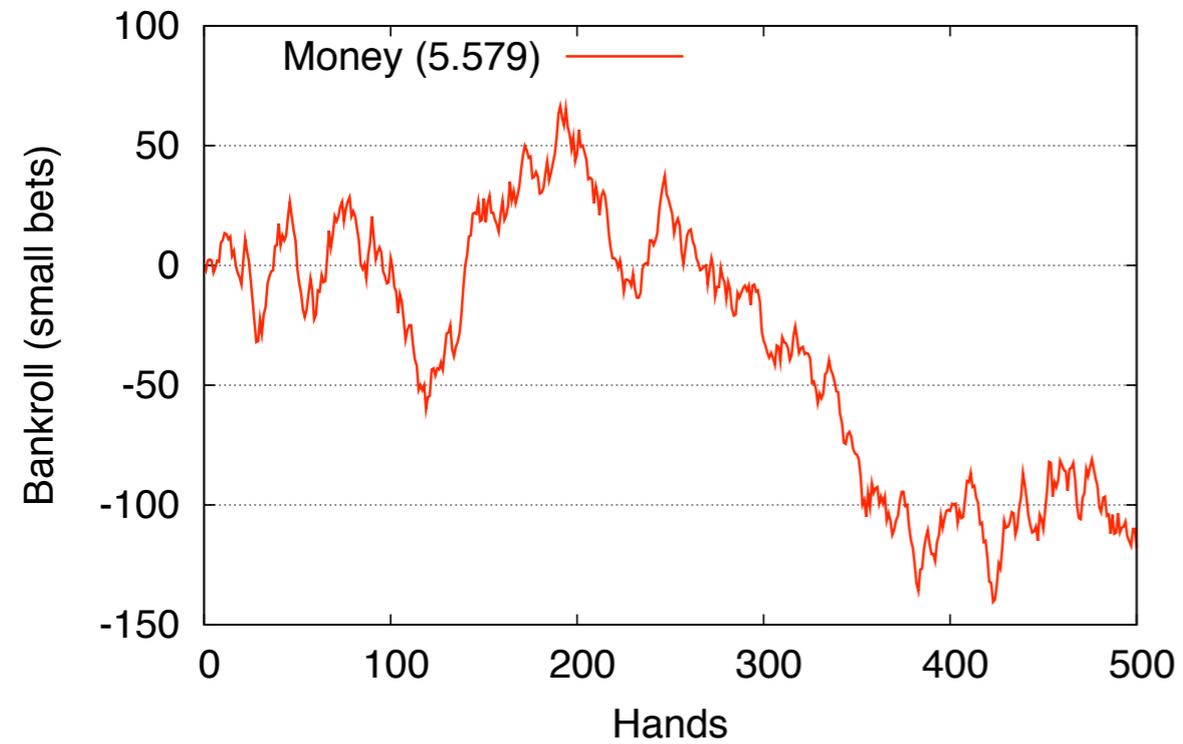


- Humans up by 25.5 bets — Polaris loses!
- The smallest possible margin of victory!

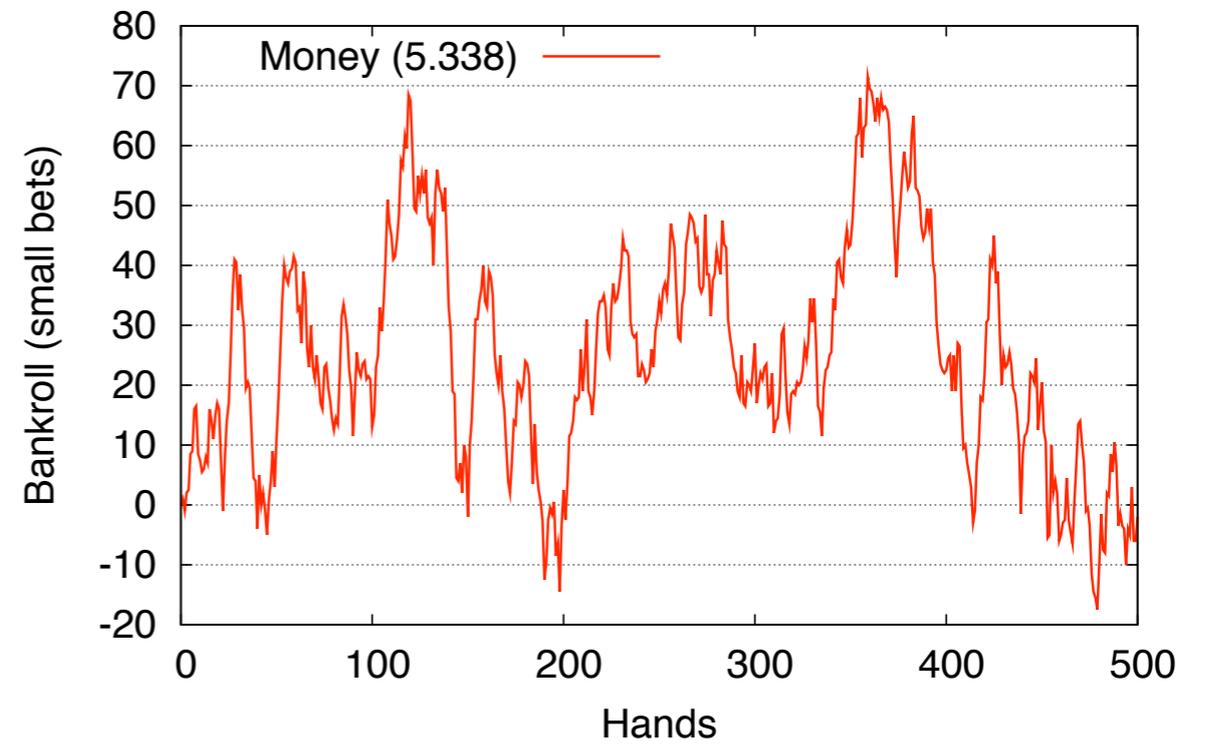
# Remote Match 2: Kyle-Nick

---

Remote 2: Kyle Hendon

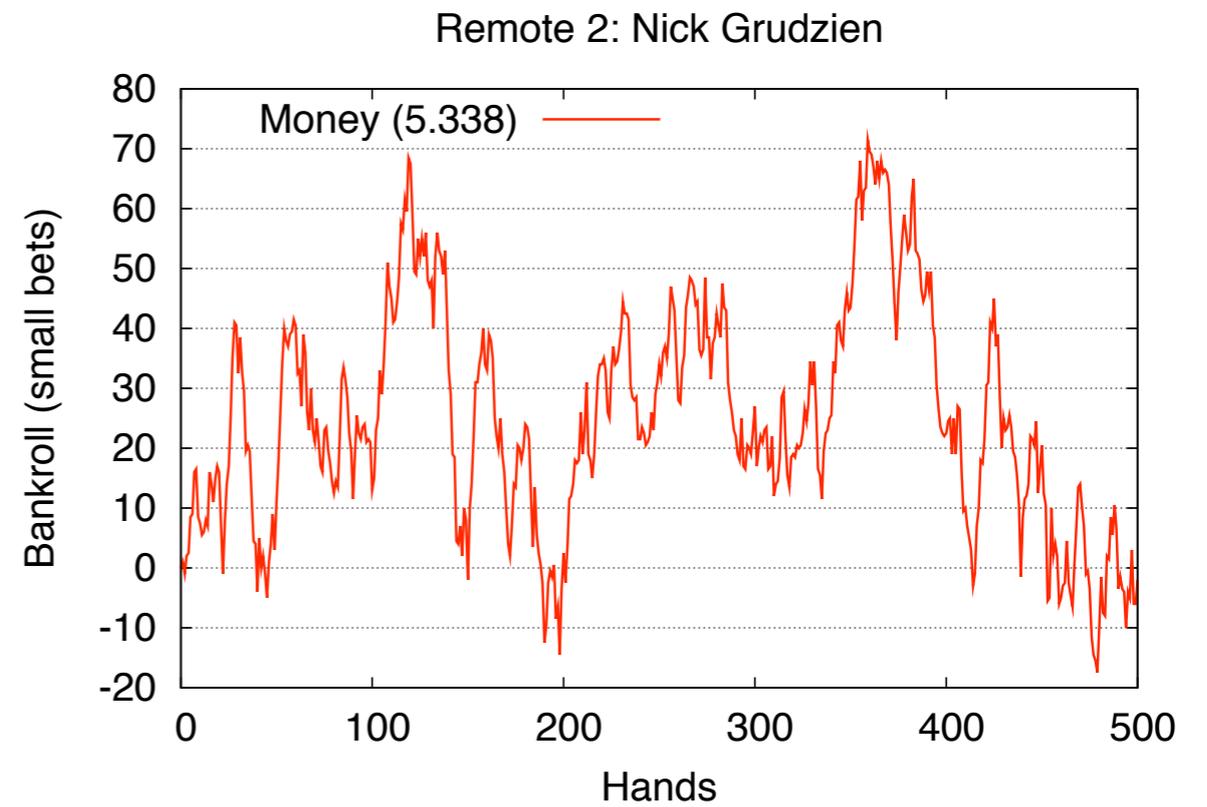
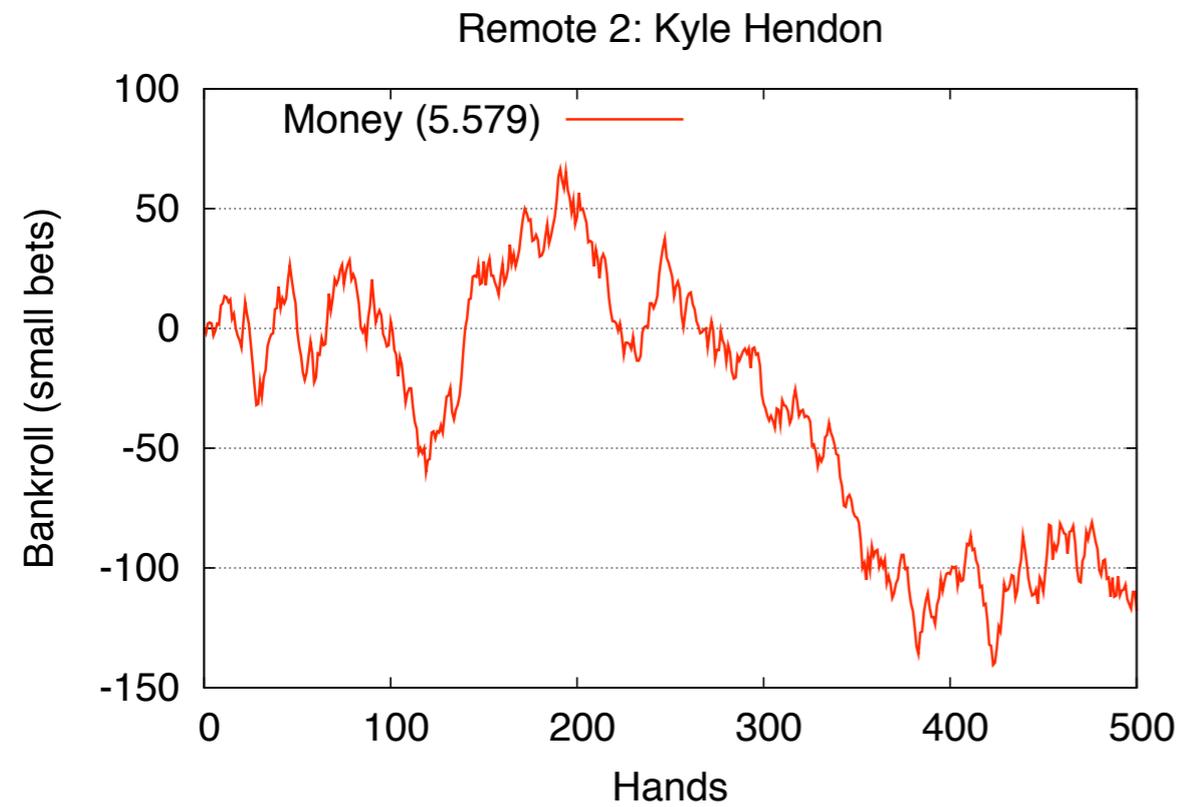


Remote 2: Nick Grudzien



# Remote Match 2: Kyle-Nick

---

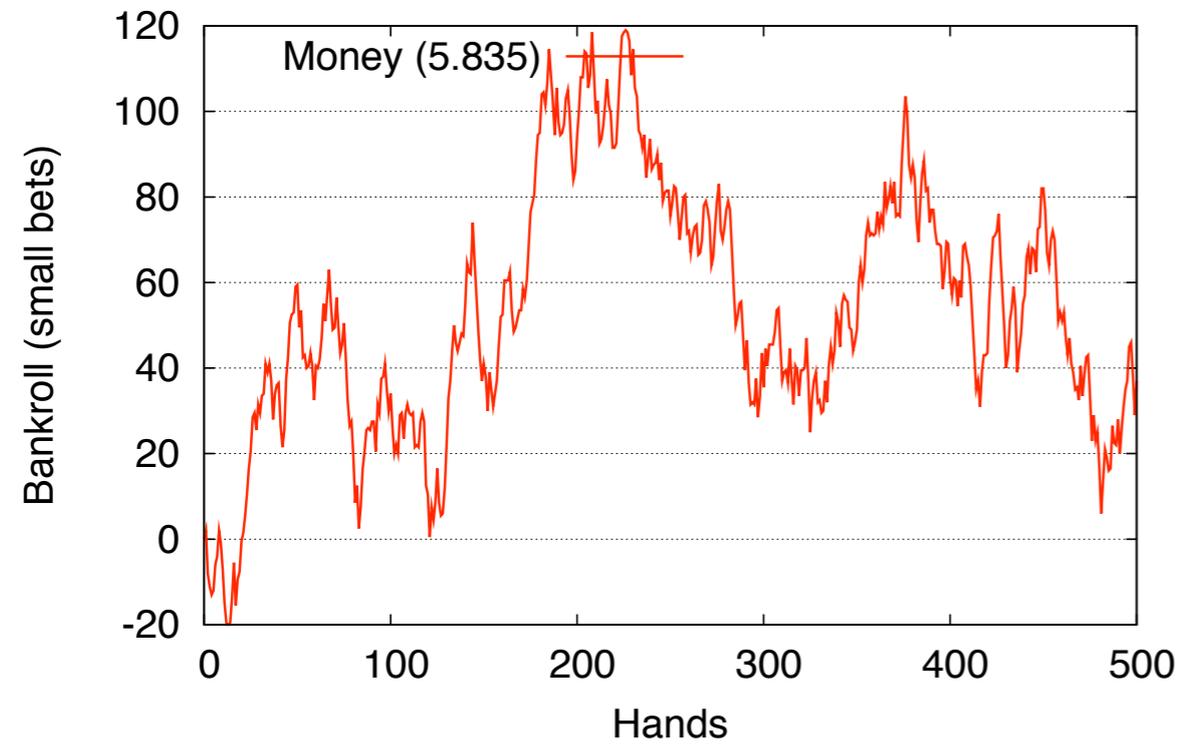


- Polaris up by 120 bets — Polaris wins!

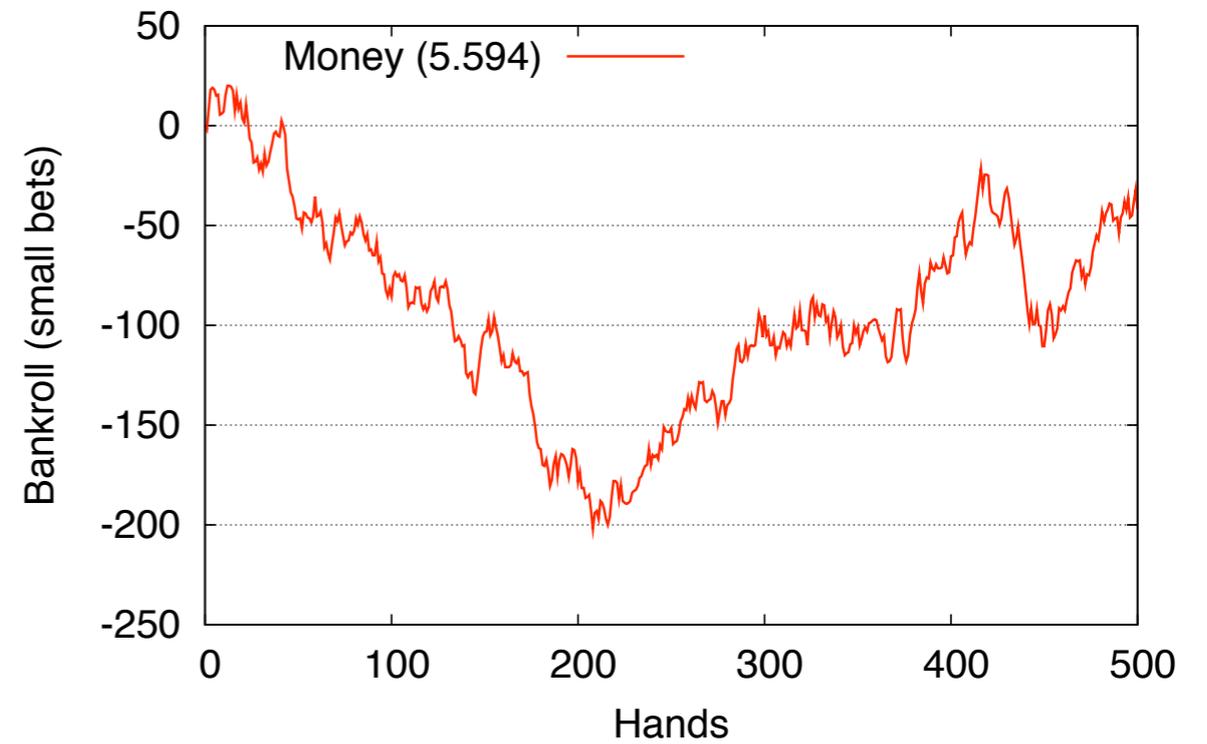
# Live Match 1: Kyle-Nick

---

Live 1: Kyle Hendon

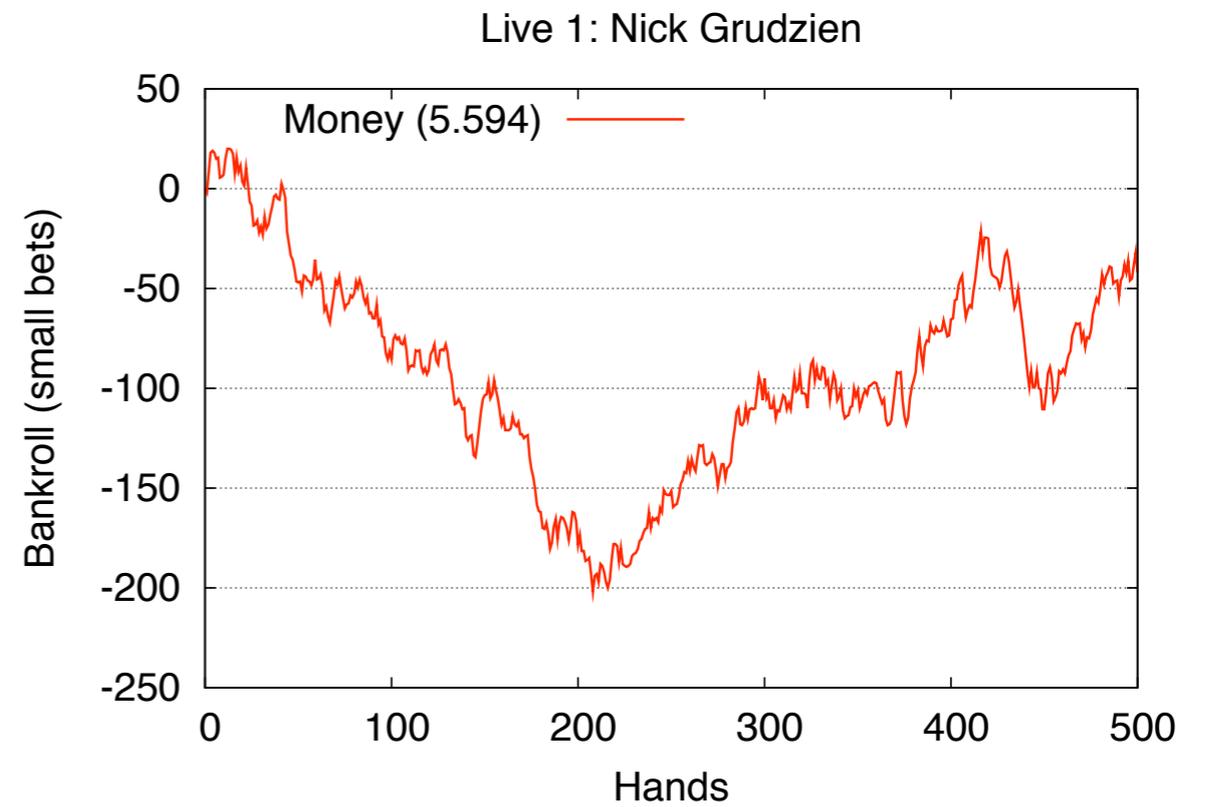
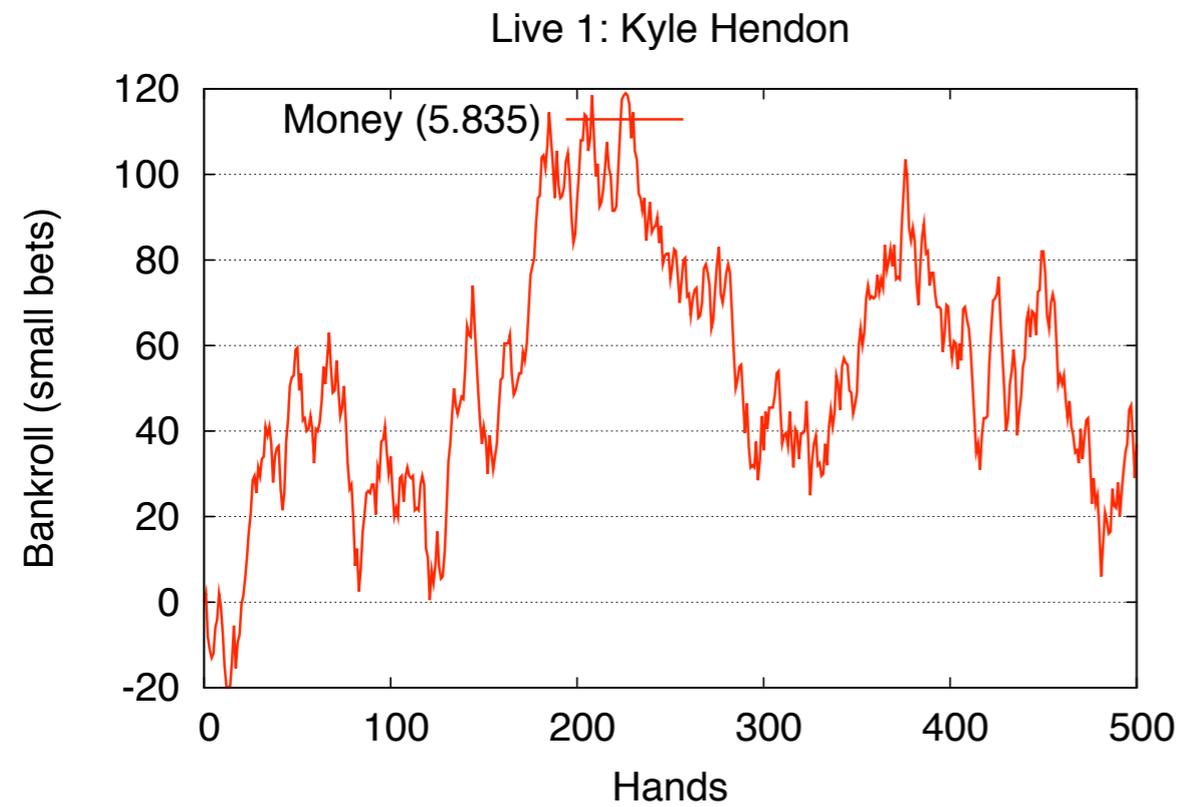


Live 1: Nick Grudzien



# Live Match 1: Kyle-Nick

---

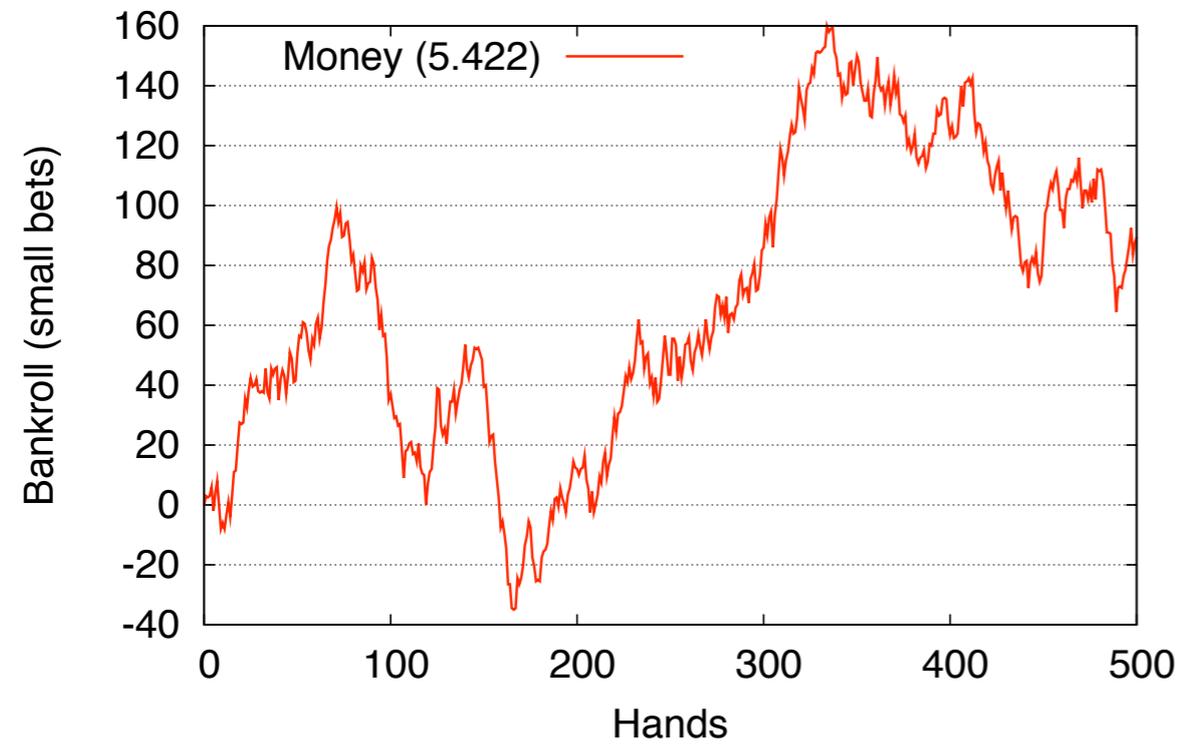


- Polaris up by 5 bets — Draw

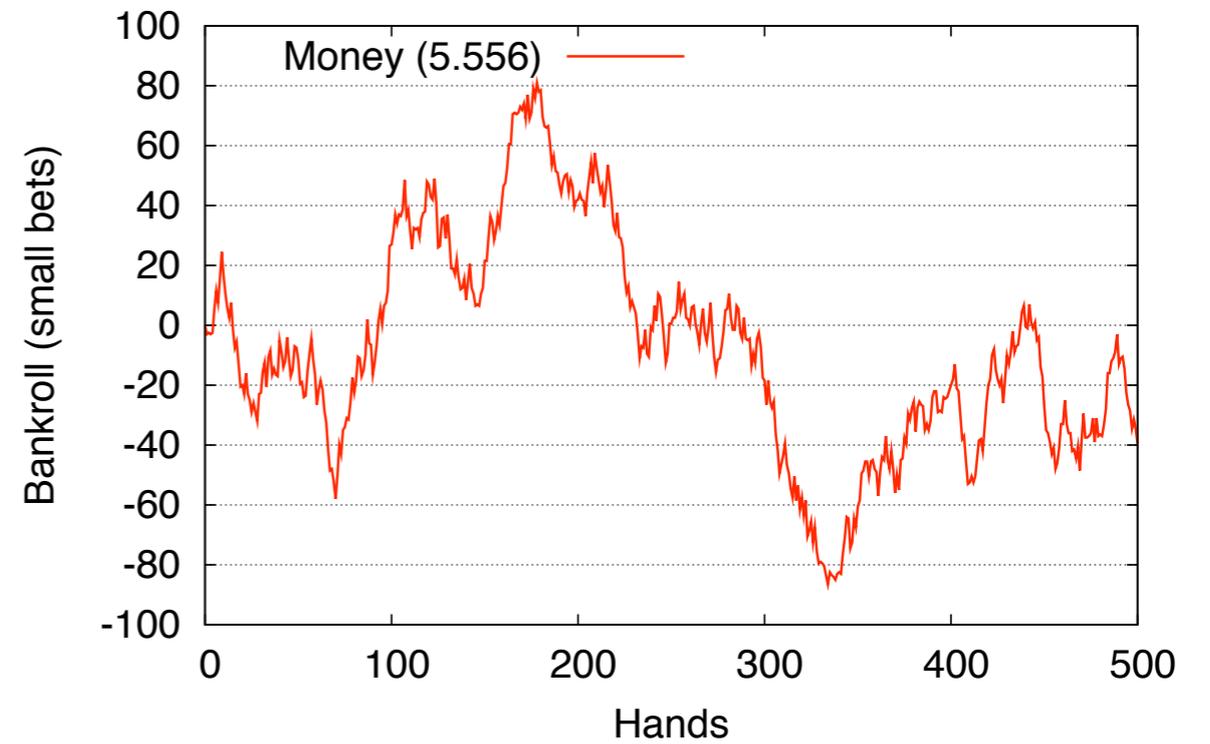
# Live Match 2: Rich-Victor

---

Live 2: Rich McRoberts

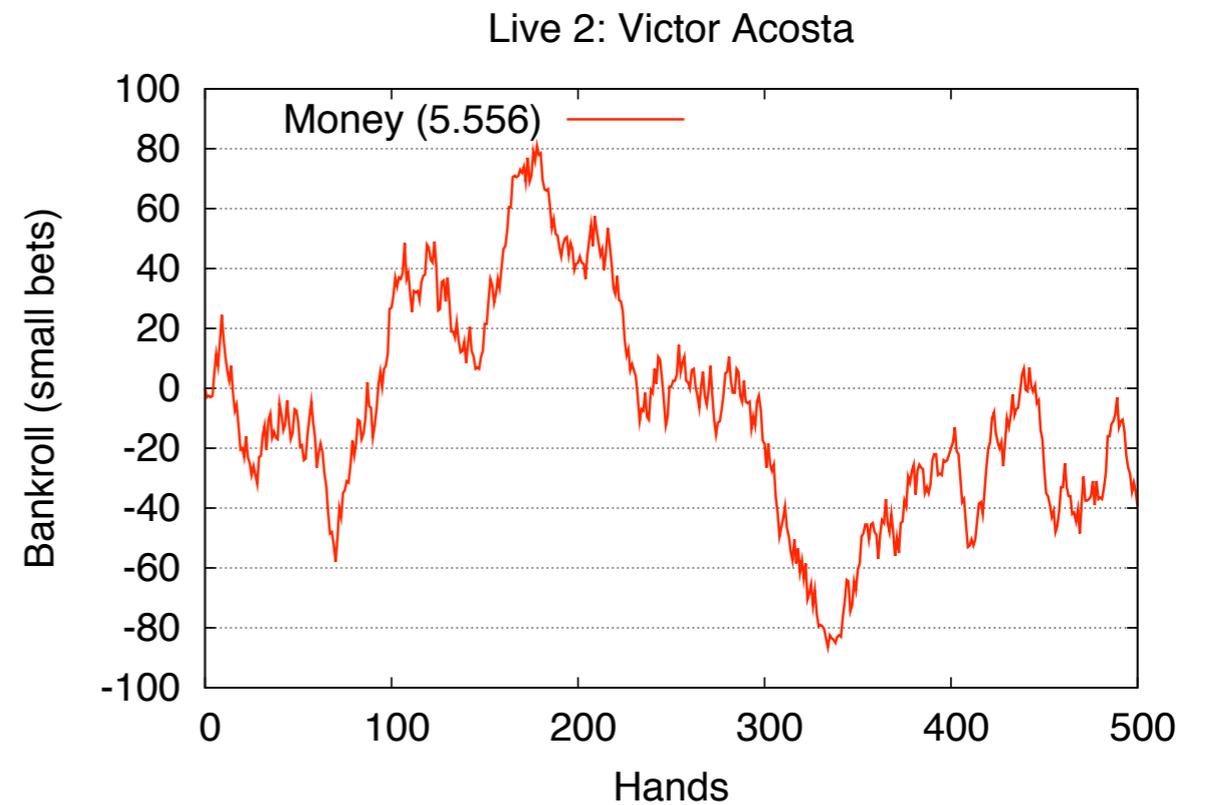
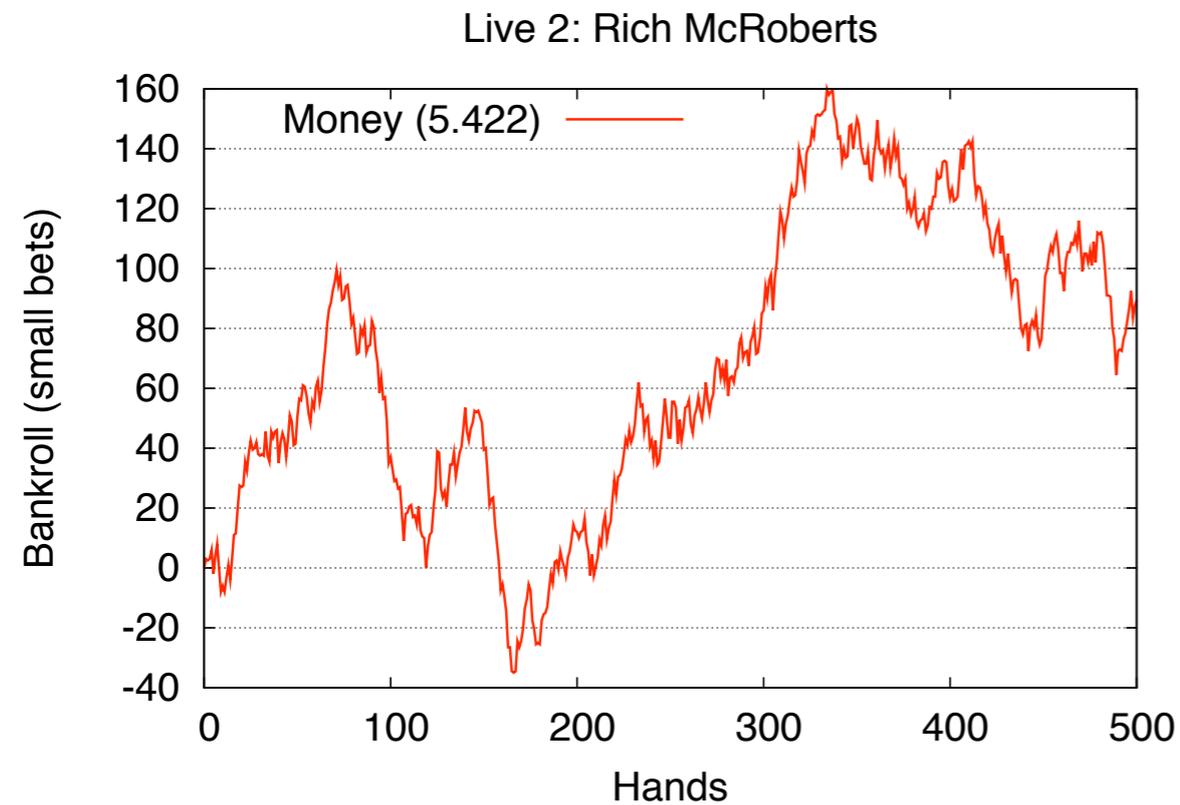


Live 2: Victor Acosta



# Live Match 2: Rich-Victor

---

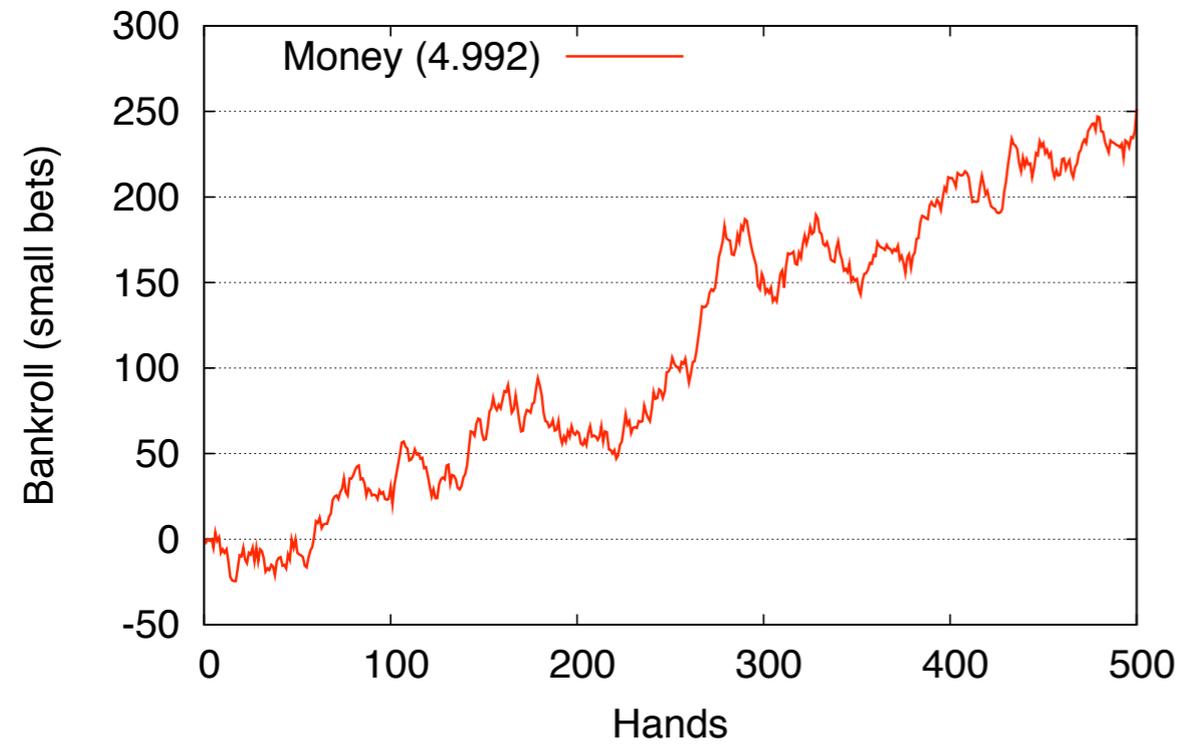


- Humans up by 50 bets — Polaris loses!

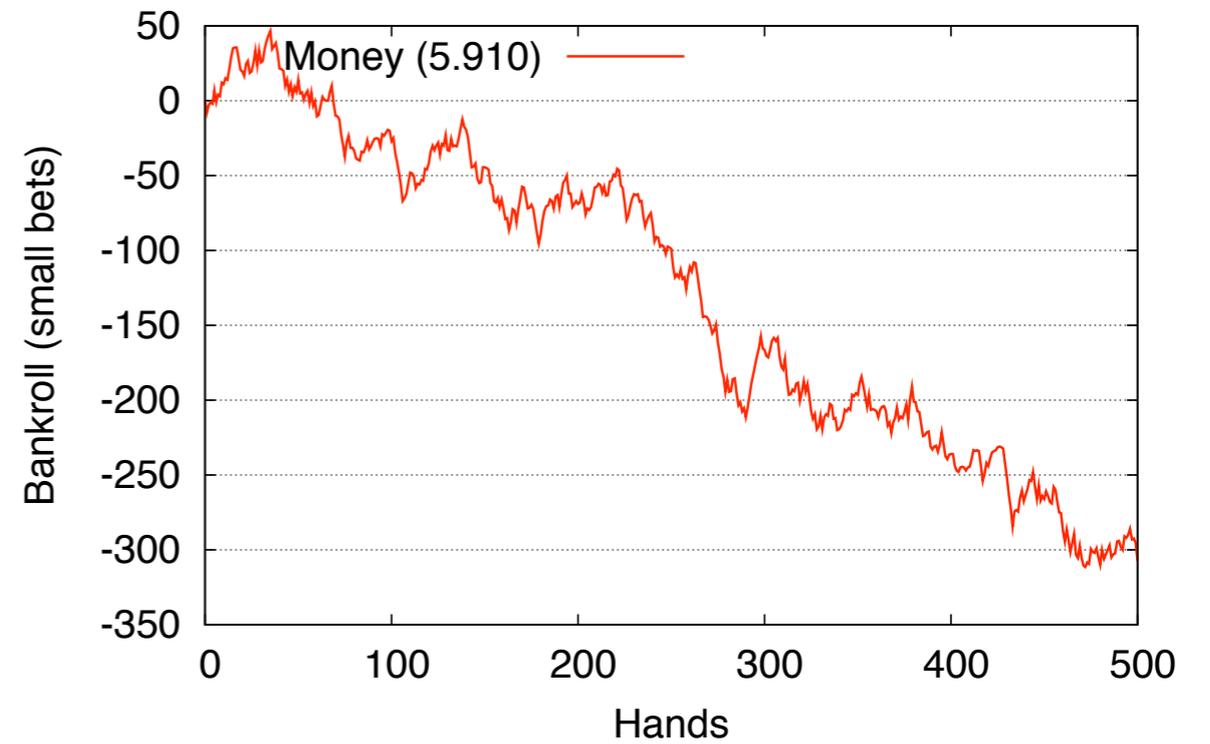
# Live Match 3: Mark-IJay

---

Live 3: Mark Newhouse

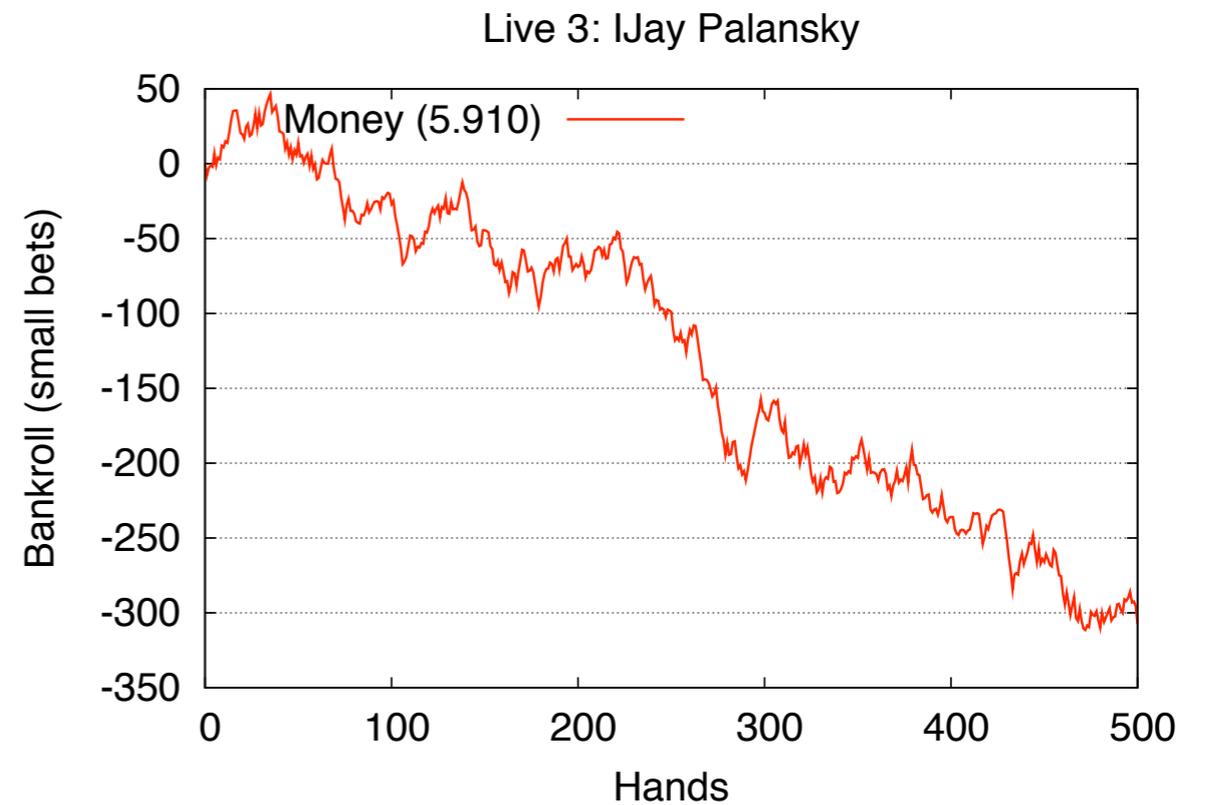
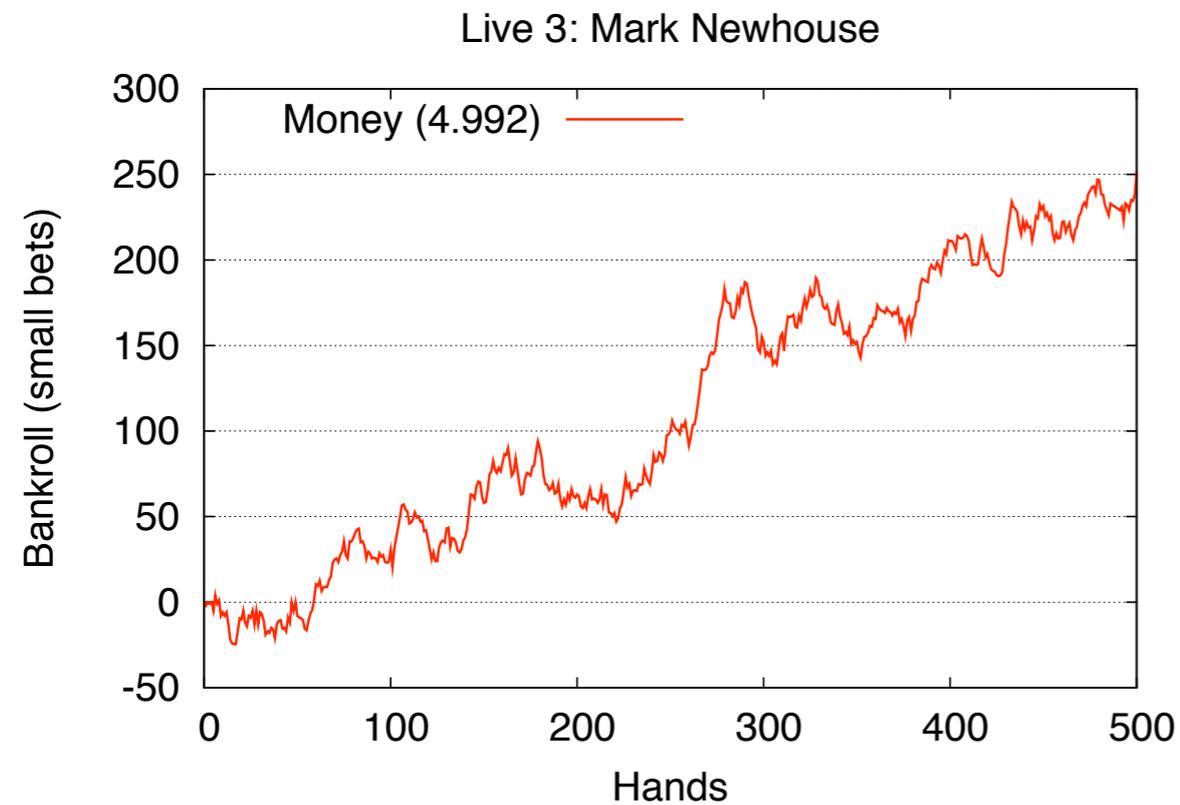


Live 3: IJay Palansky



# Live Match 3: Mark-IJay

---

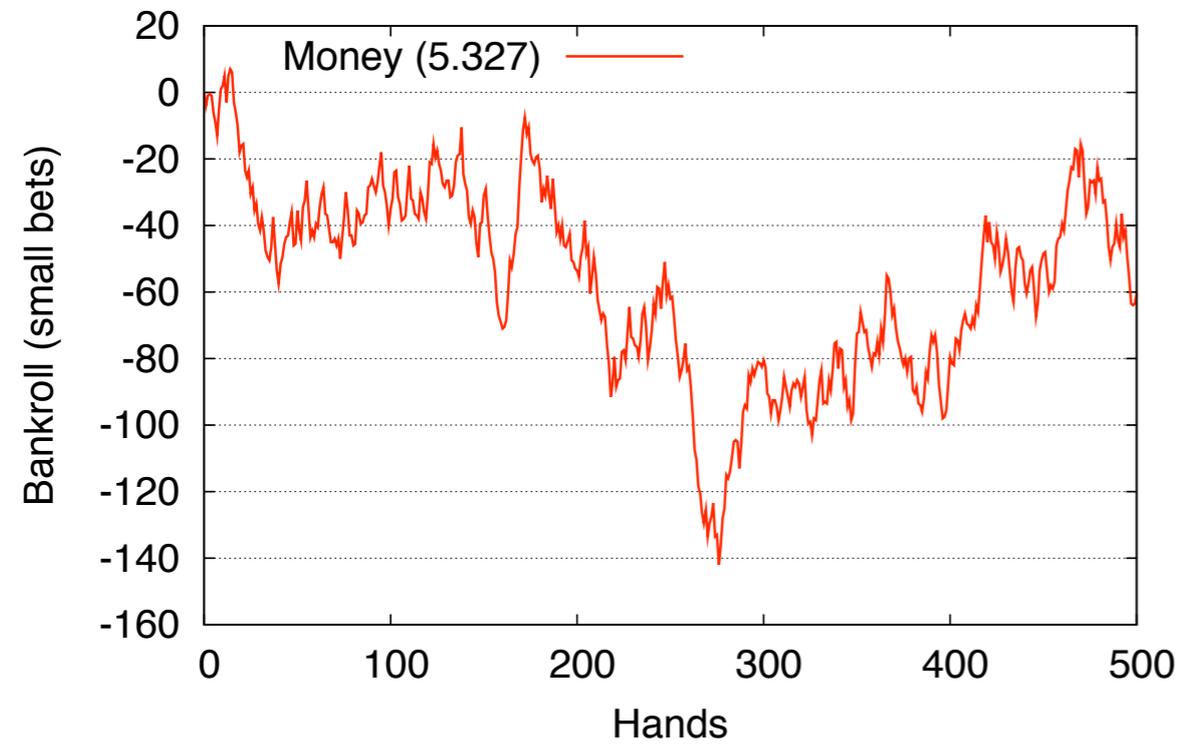


- Polaris up by 56 bets — Polaris wins!

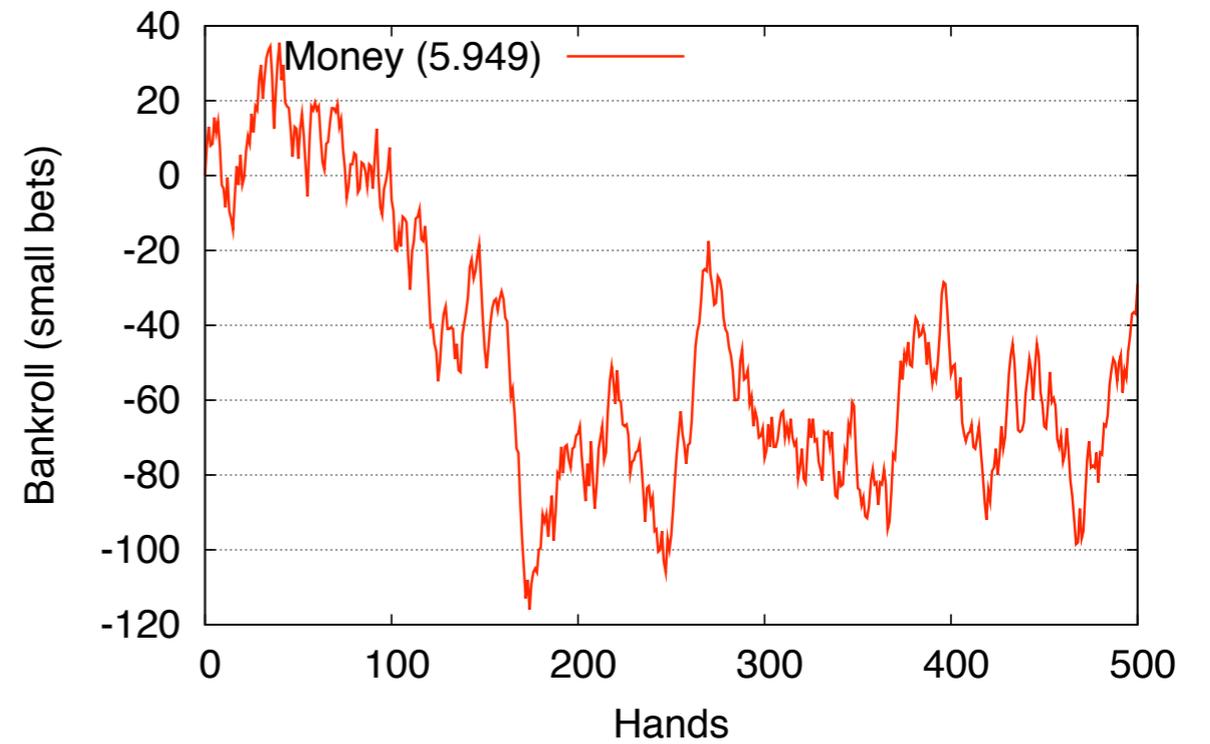
# Live Match 4: Hoss-IJay

---

Live 4: Matt Hawrilenko



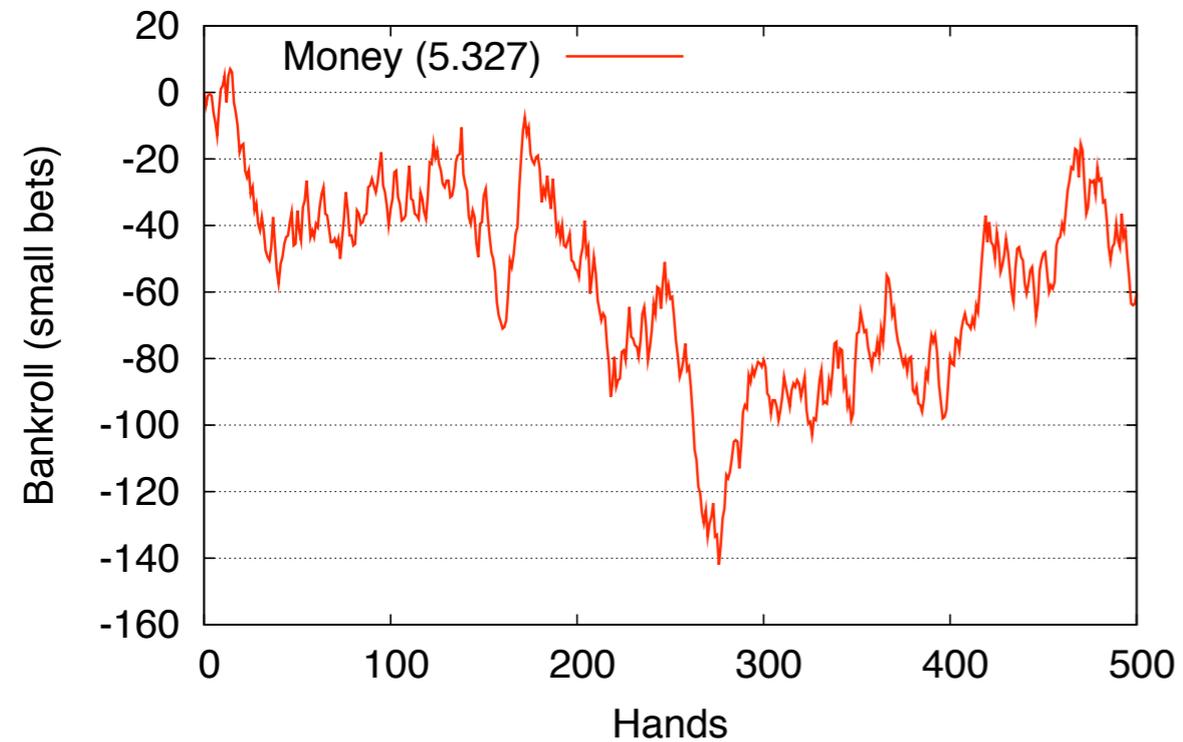
Live 4: IJay Palansky



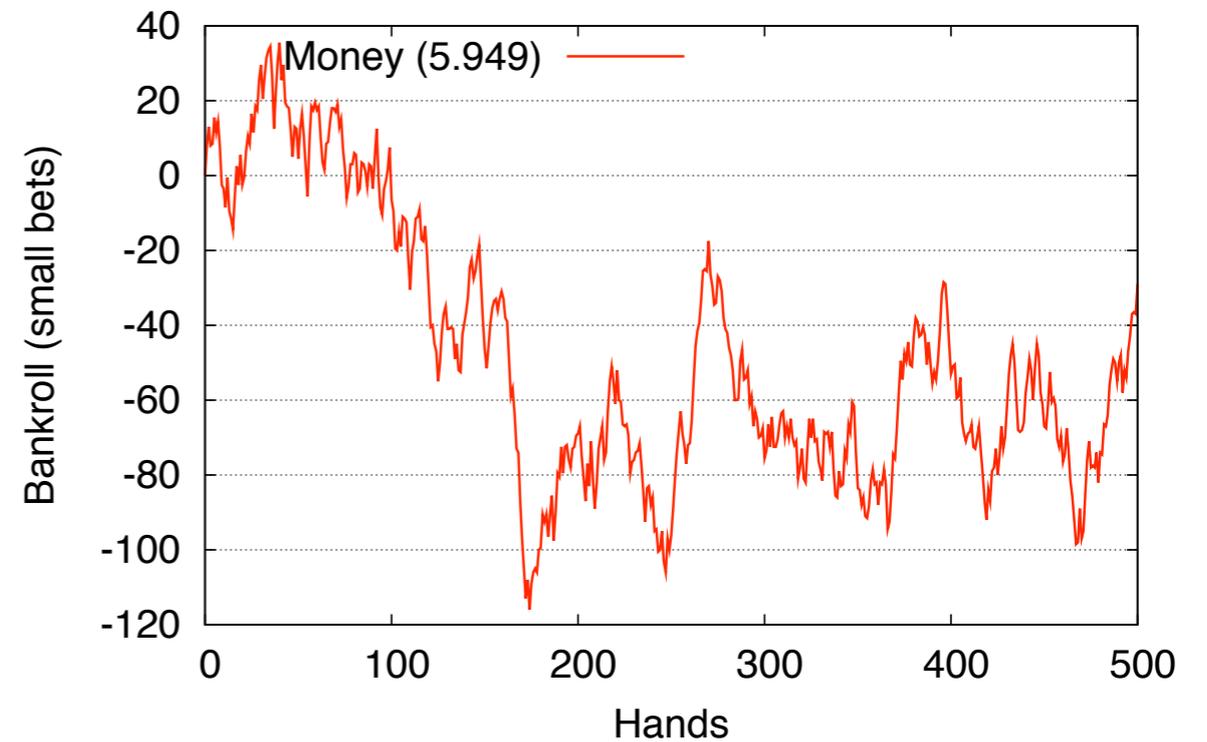
# Live Match 4: Hoss-IJay

---

Live 4: Matt Hawrilenko



Live 4: IJay Palansky



- Polaris up by 89.5 bets — Polaris wins!

# The Results

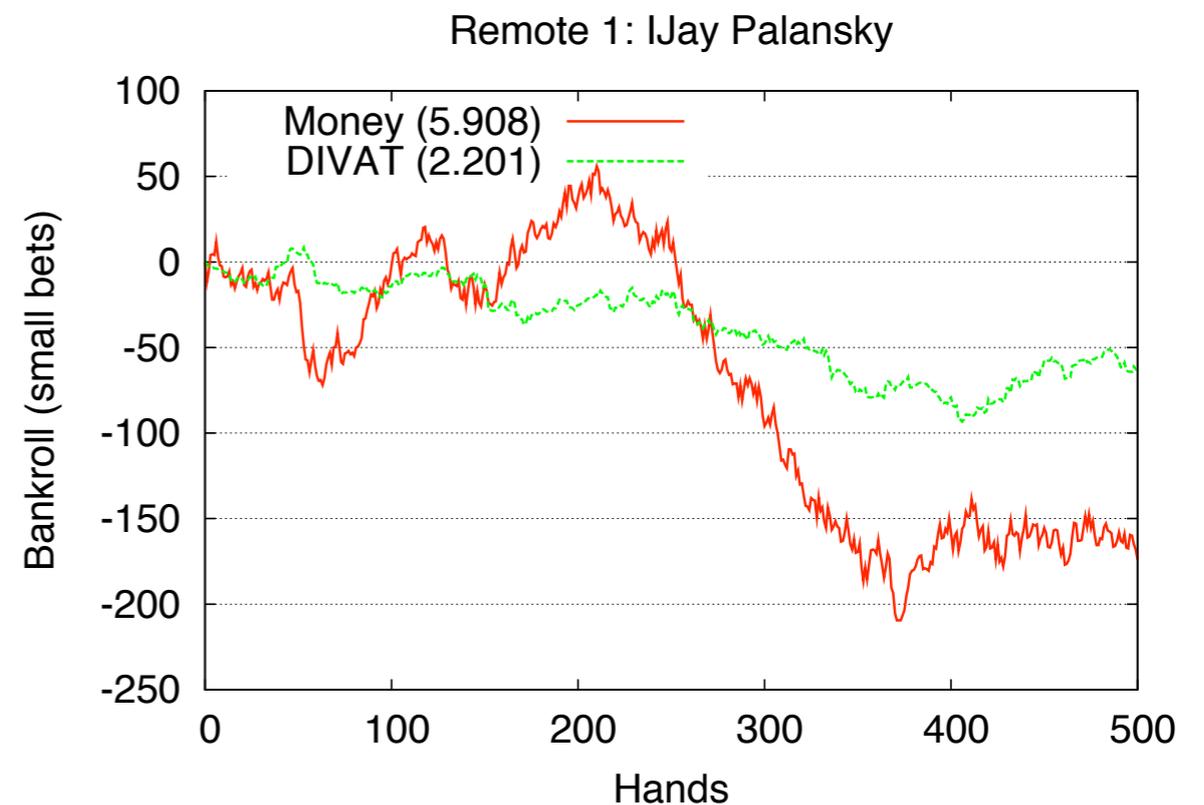
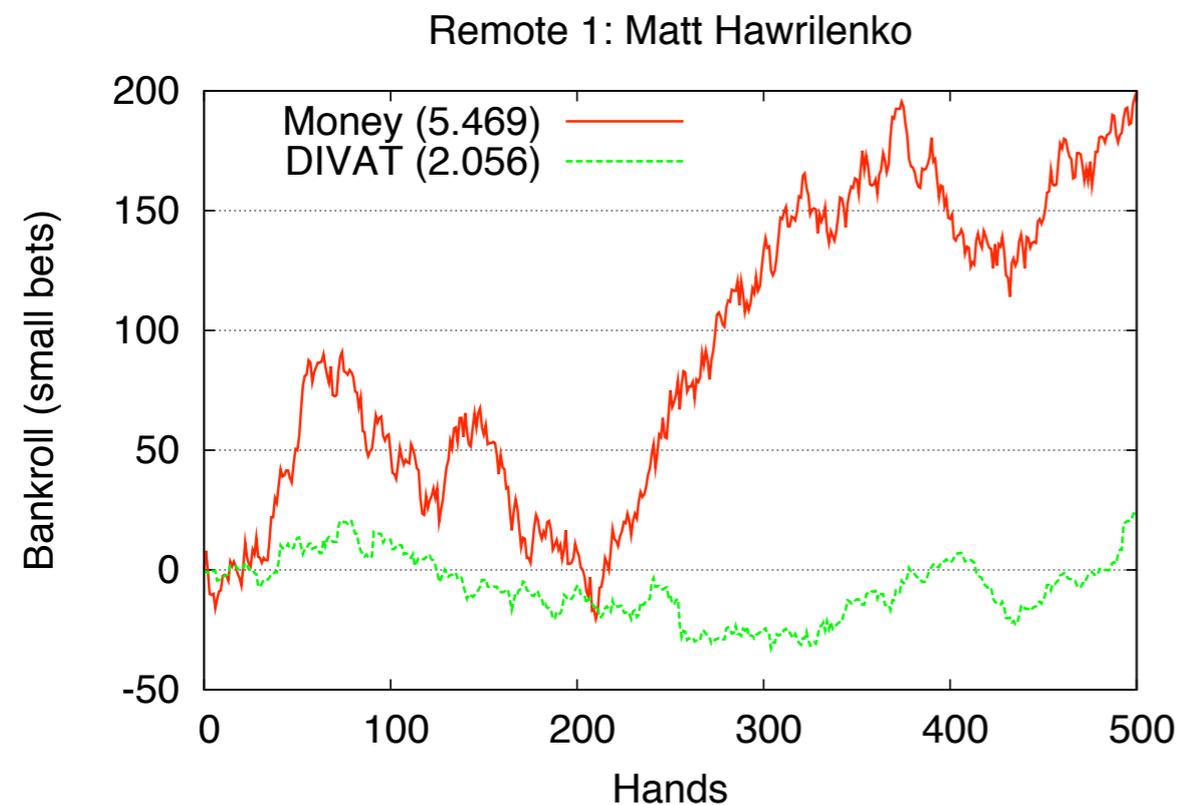
---

- Final Result: 3 Wins, 2 Losses, 1 Draw
- First time a computer has beaten professionals in a meaningful poker competition!

# First Law of Poker?

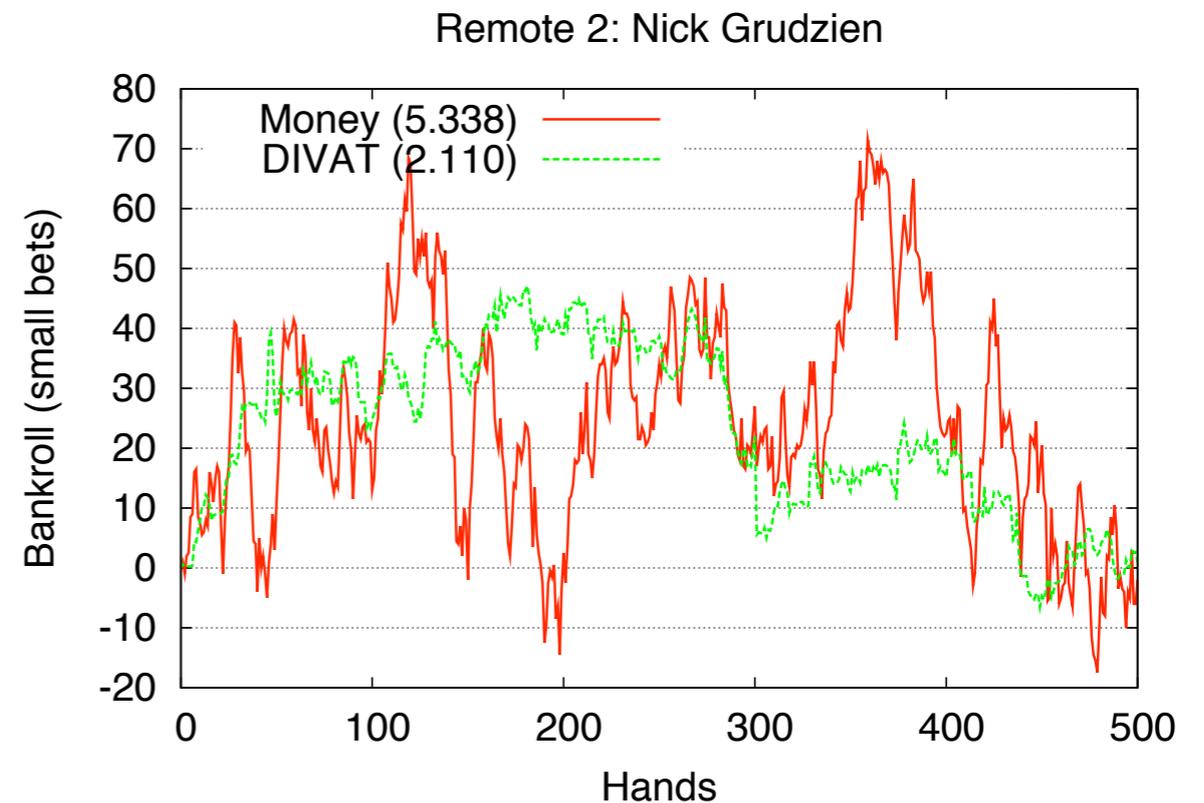
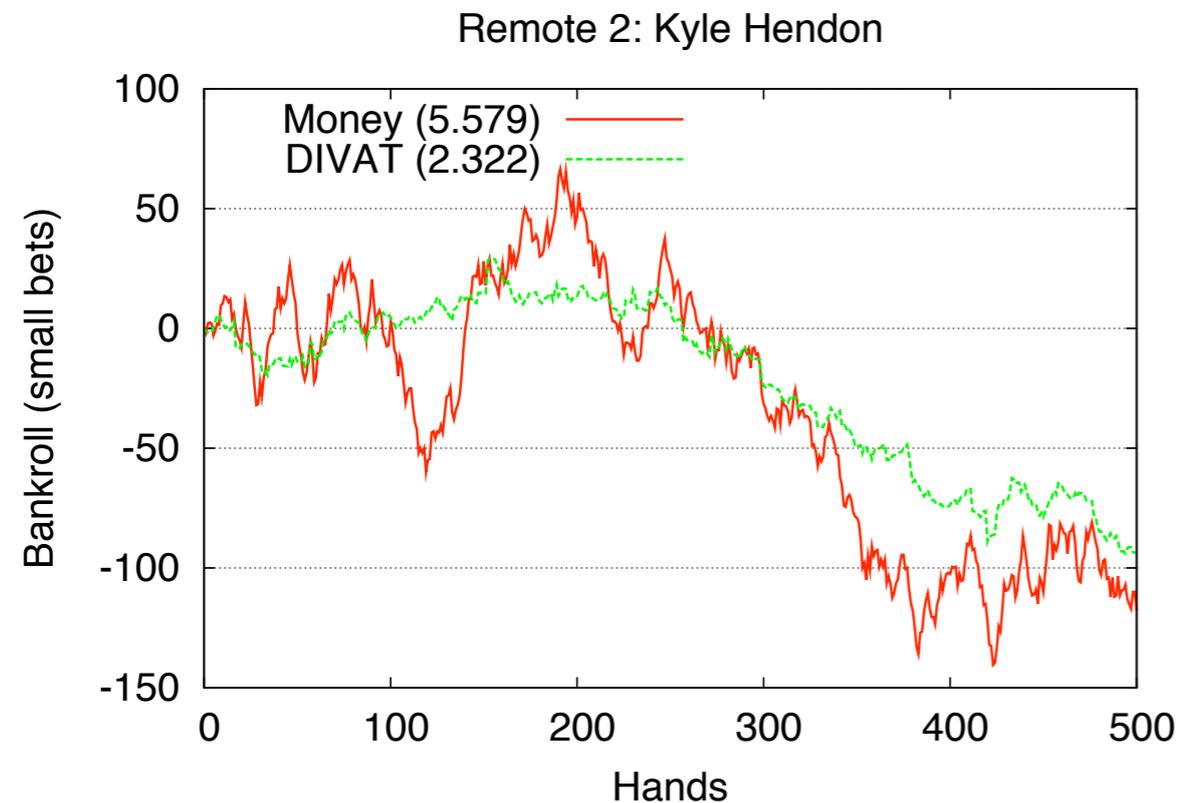
---

# Remote Match 1: Hoss-IJay



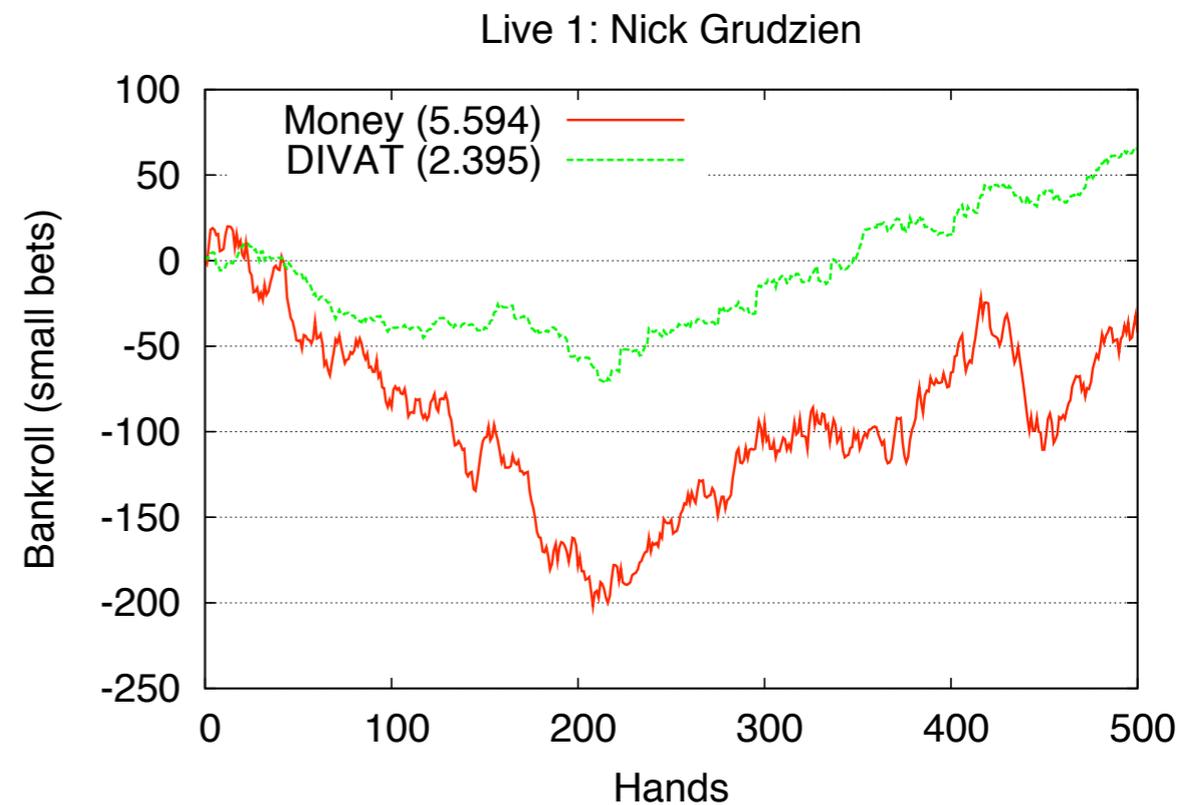
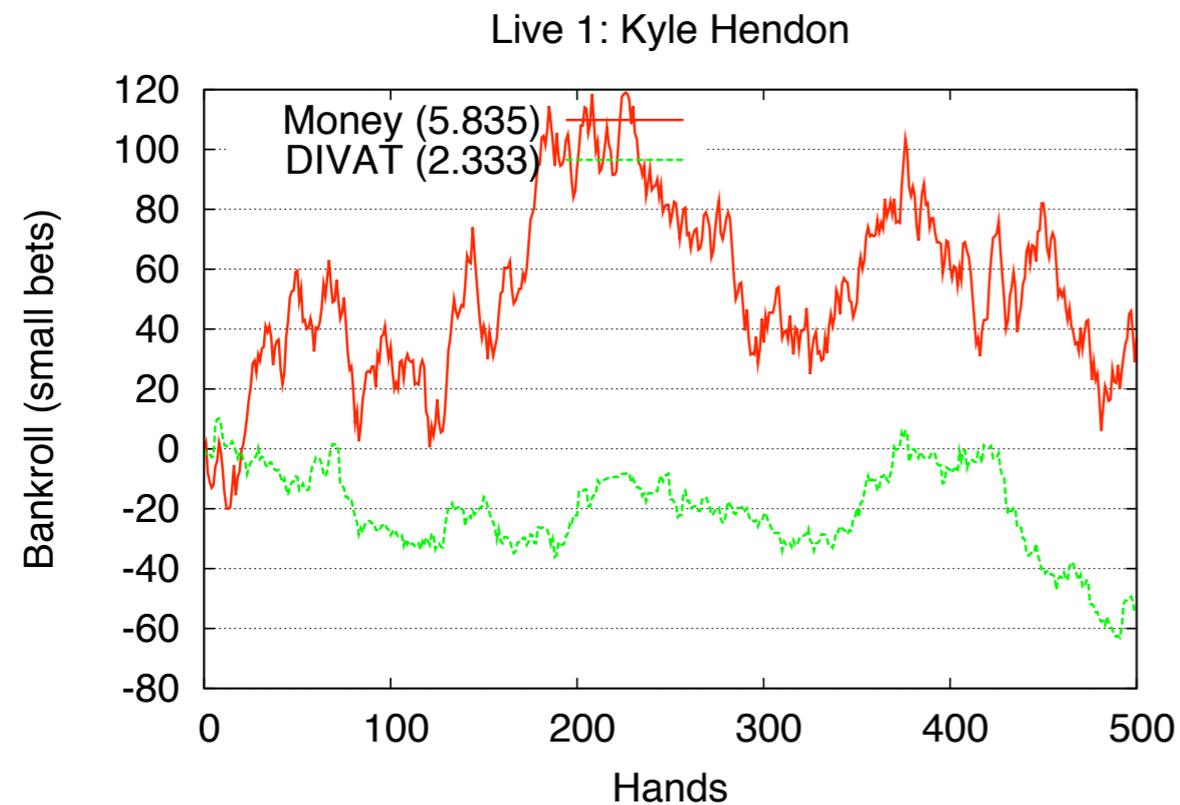
- Humans up by 25.5 bets — Polaris loses!
  - Not statistically significant, p-value = 0.38
- DIVAT showed Polaris up by 41 bets
  - Not statistically significant, p-value = 0.26

# Remote Match 2: Kyle-Nick



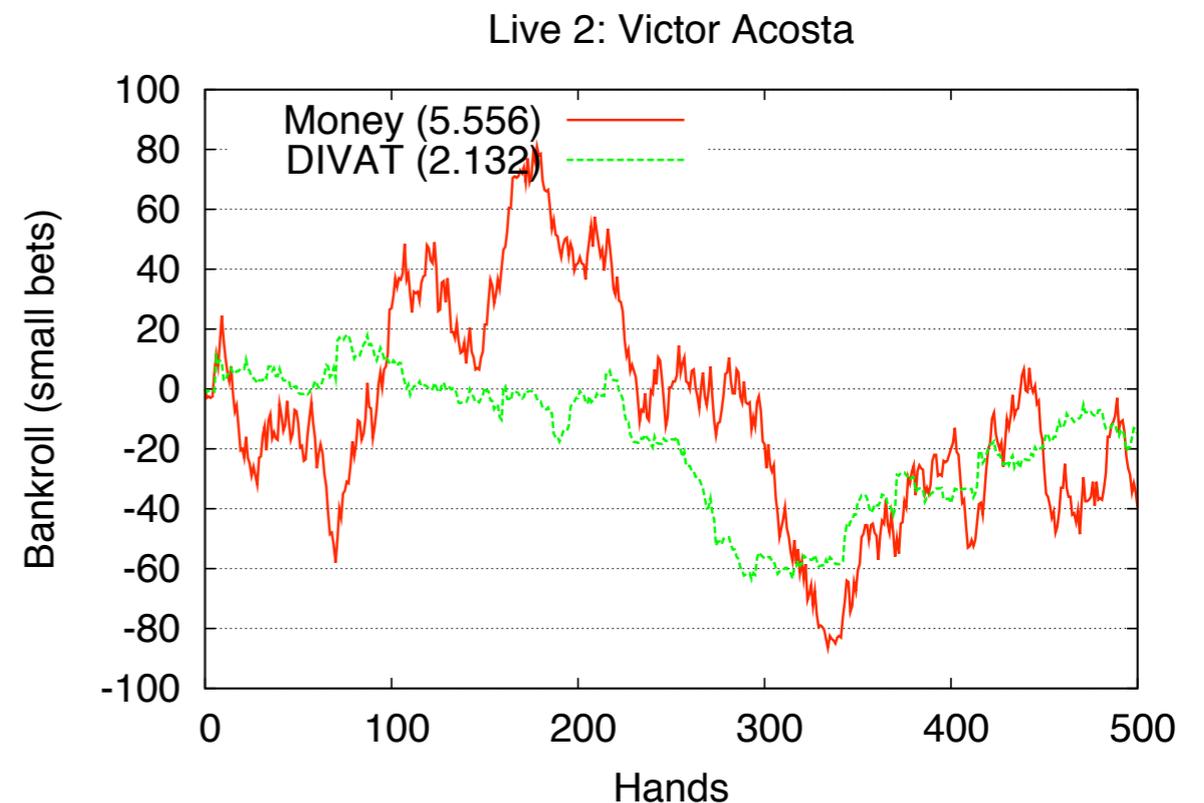
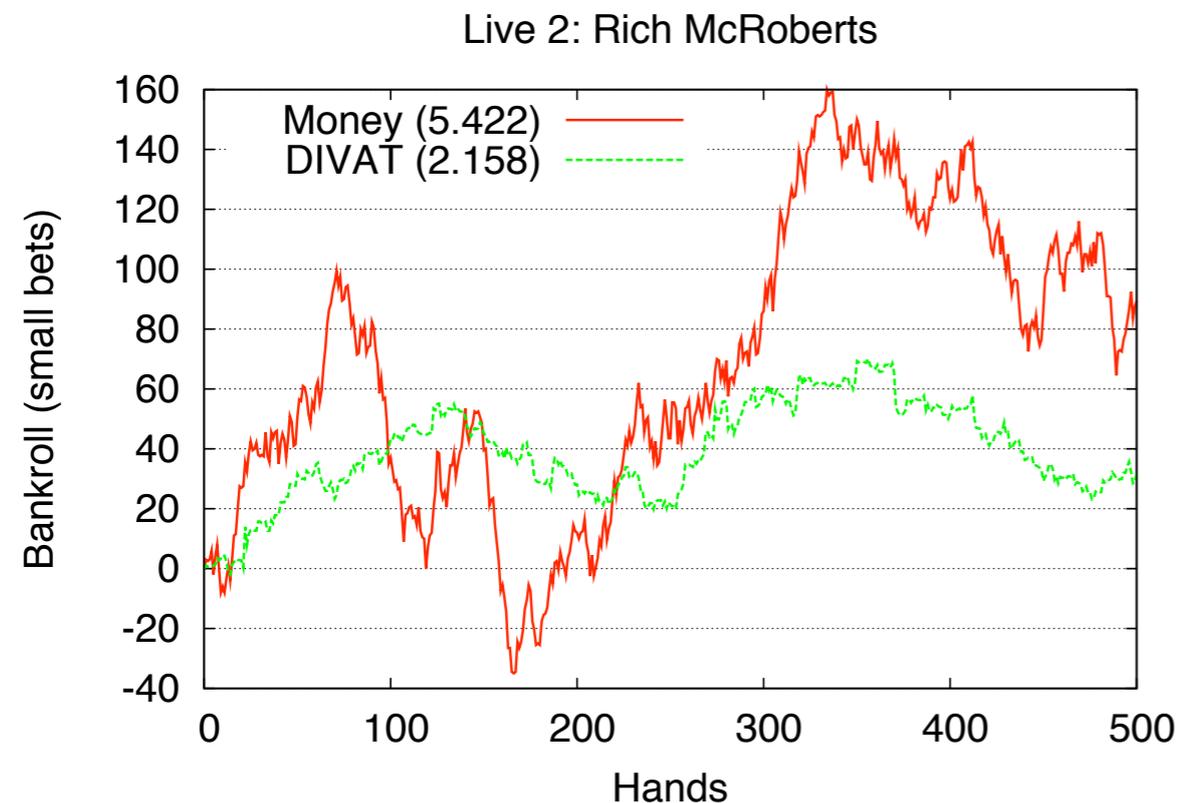
- Polaris up by 120 bets — Polaris wins!
  - Not statistically significant, p-value = 0.09
- DIVAT showed Polaris up by 92 bets
  - Not statistically significant, p-value = 0.07

# Live Match 1: Kyle-Nick



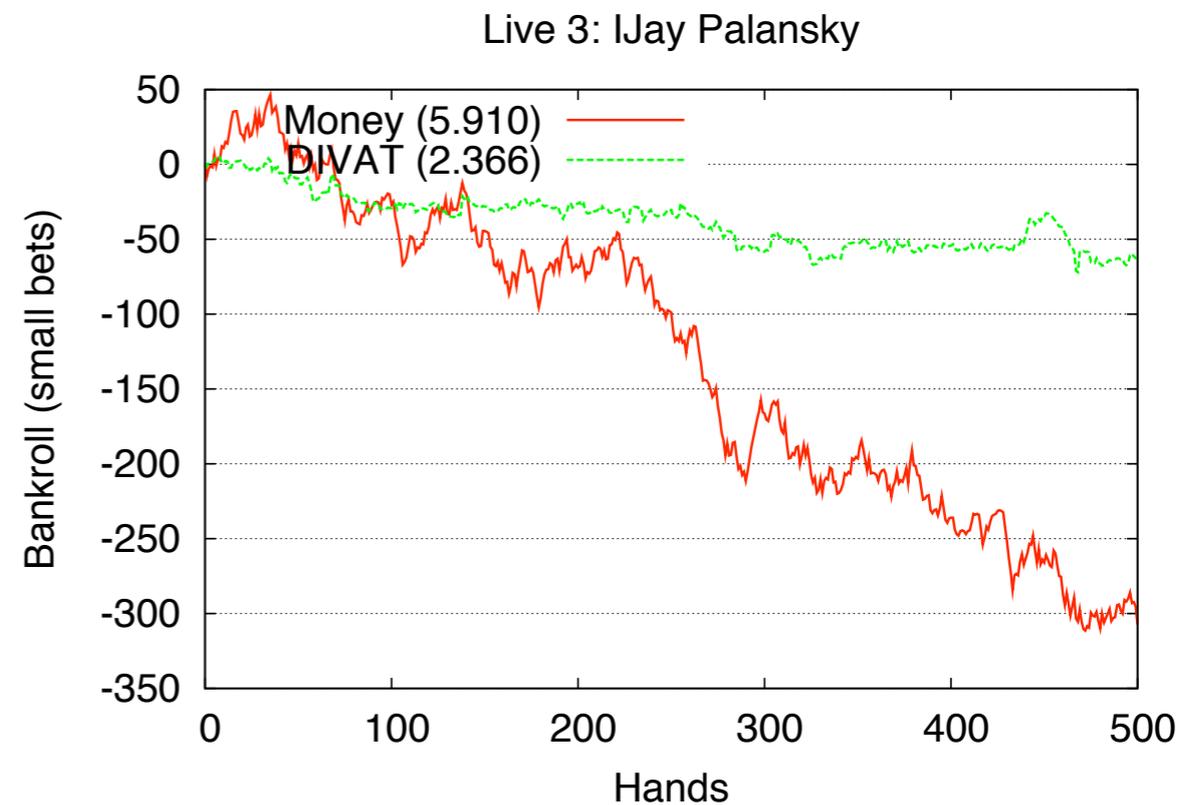
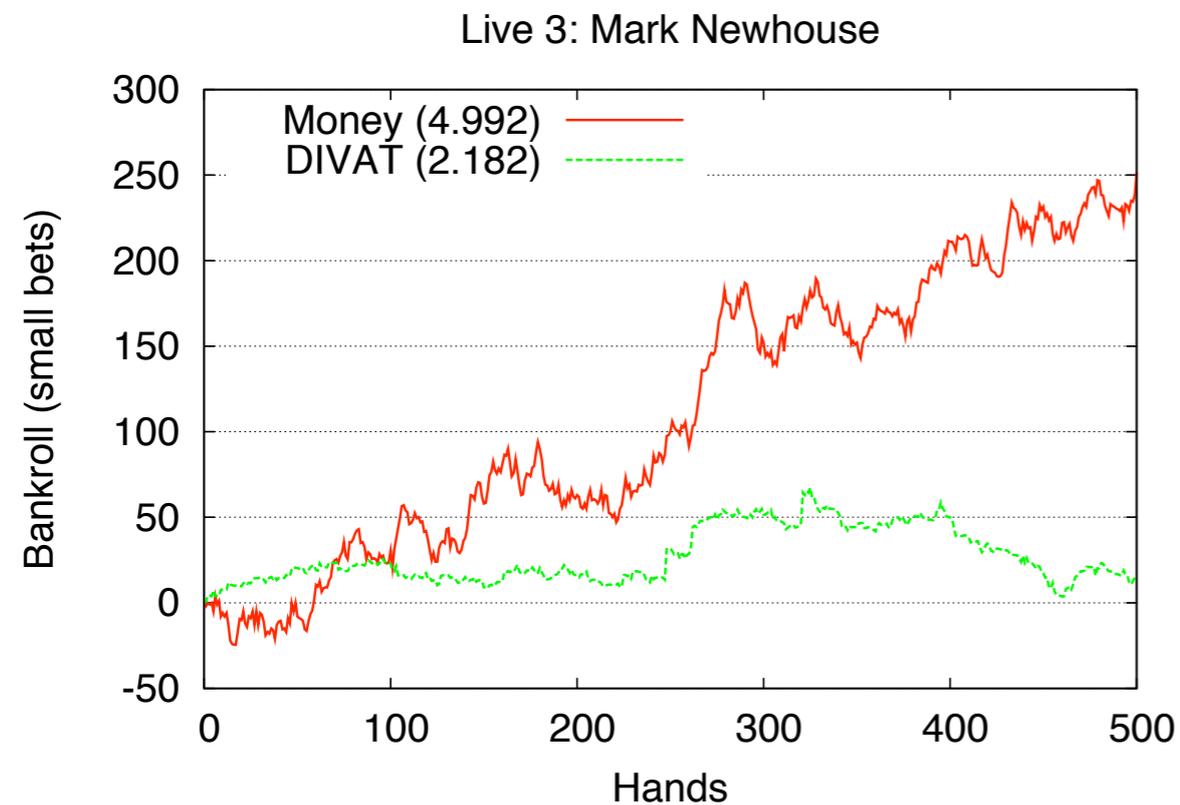
- Polaris up by 5 bets — Draw
  - Not statistically significant, p-value = 0.48
- DIVAT showed Humans up by 10 bets
  - Not statistically significant, p-value = 0.44

# Live Match 2: Rich-Victor



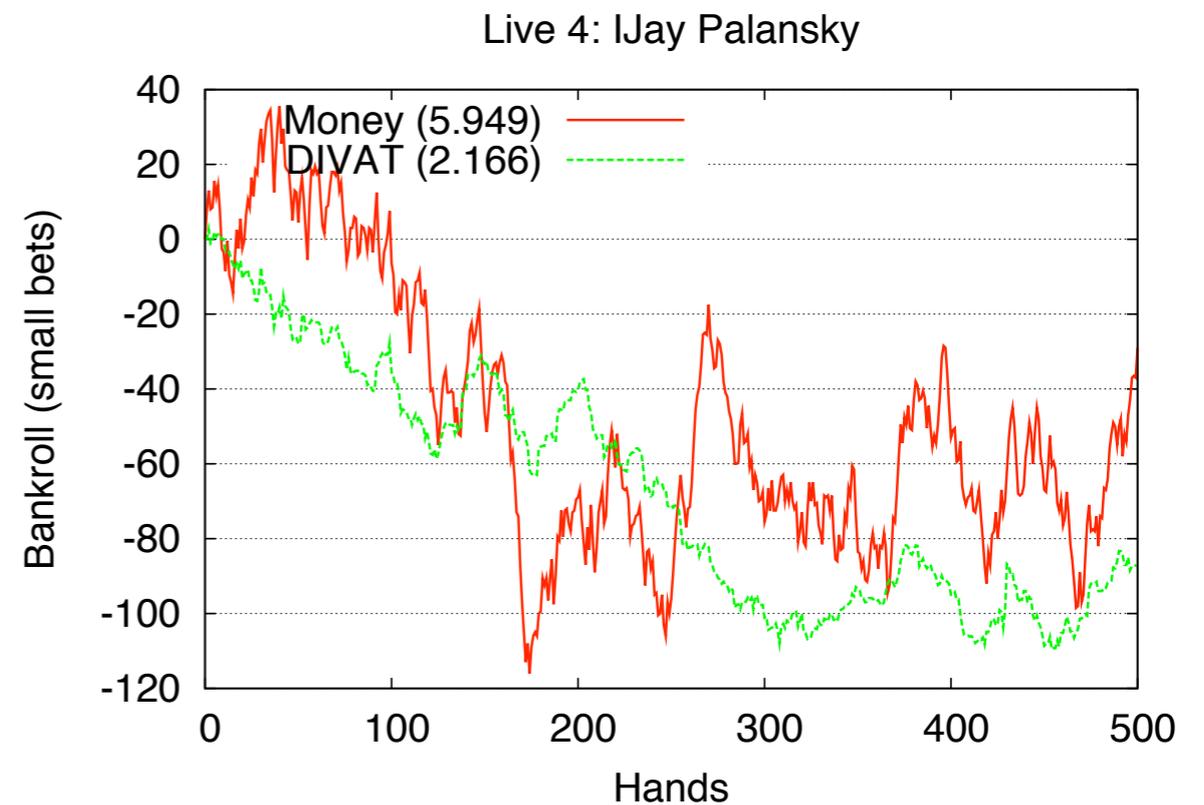
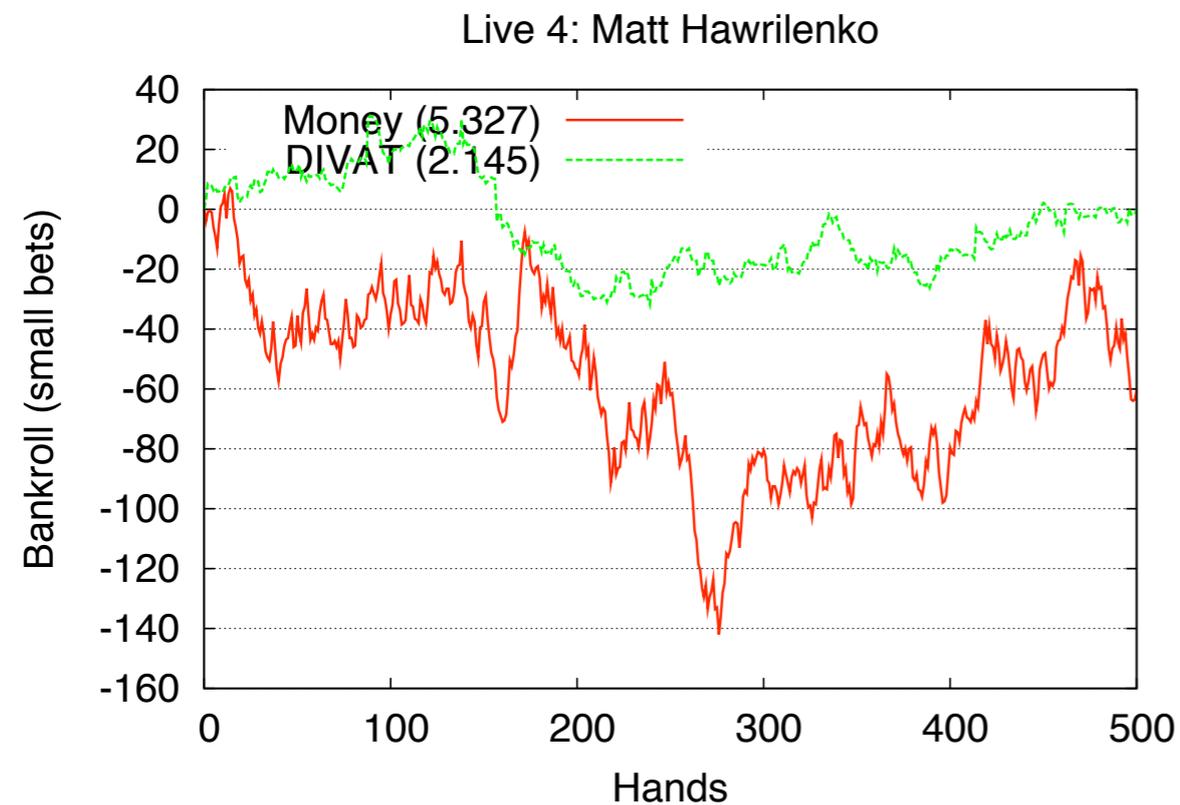
- Humans up by 50 bets — Polaris loses!
  - Not statistically significant, p-value = 0.27
- DIVAT showed Humans up by 16 bets
  - Not statistically significant, p-value = 0.40

# Live Match 3: Mark-IJay



- Polaris up by 56 bets — Polaris wins!
  - Not statistically significant, p-value = 0.26
- DIVAT showed Polaris up by 47 bets
  - Not statistically significant, p-value = 0.23

# Live Match 4: Hoss-IJay



- Polaris up by 89.5 bets — Polaris wins!
  - Not statistically significant, p-value = 0.17
- DIVAT showed Polaris up by 87 bets
  - Not statistically significant, p-value = 0.08

# The Results

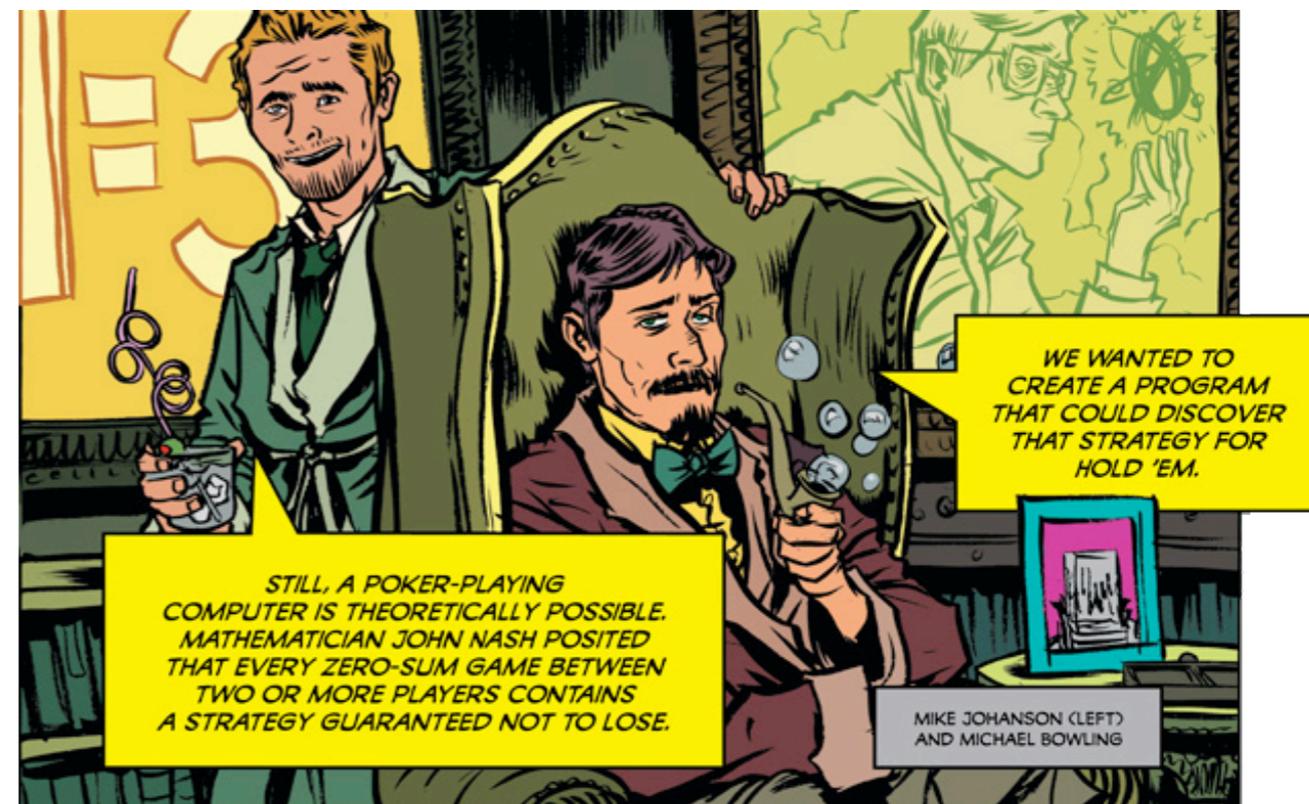
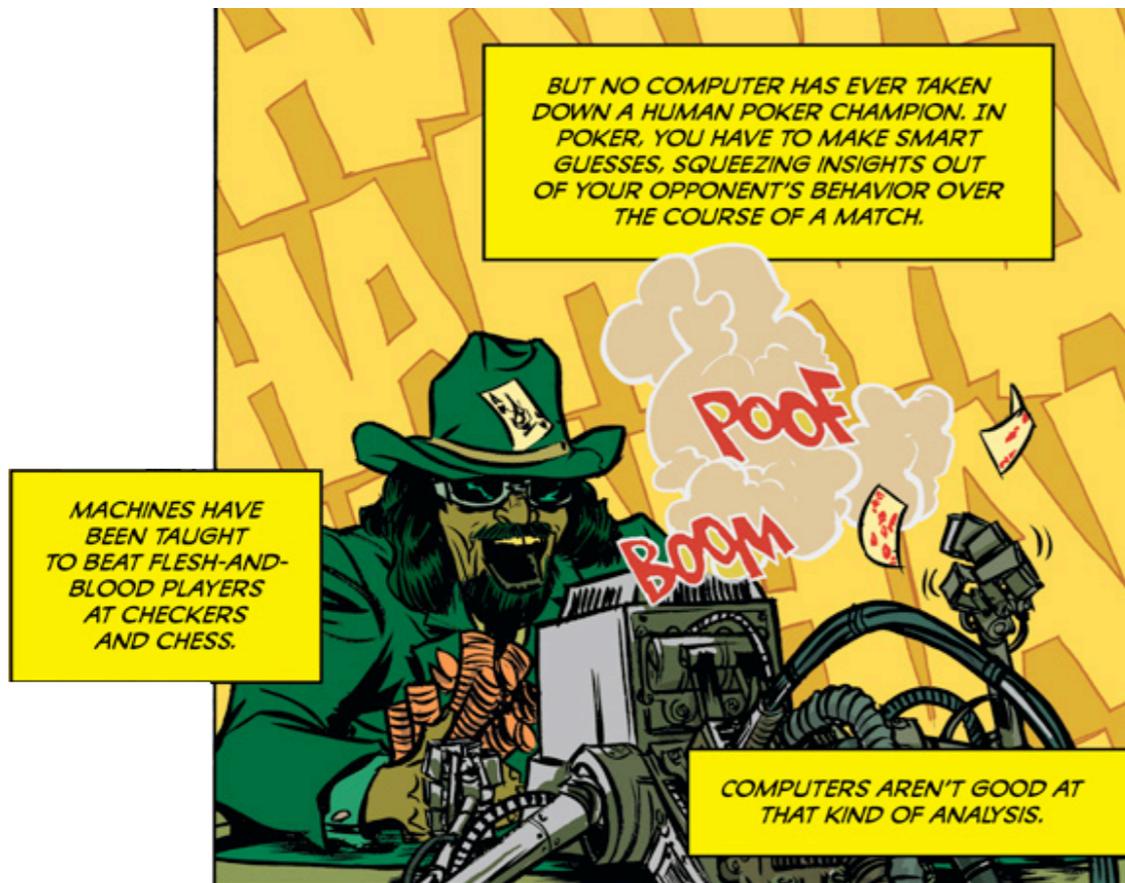
---

- Final Result: 3 Wins, 2 Losses, 1 Draw
- Not one match was statistically significant, even with DIVAT

# The Results

---

- Final Result: 3 Wins, 2 Losses, 1 Draw
- Not one match was statistically significant, even with DIVAT
- Is this the end of the story for the heroic scientists?



# Evaluating Poker Players

---

- Money is a terrible estimator!
  - ~20K hands needed for statistical significance
- Duplicate helps
  - ~2000 pairs of hands
- DIVAT+Duplicate is almost enough
  - ~1200 pairs of hands

# Removing More Luck: Imaginary Observations

---

# Removing More Luck: Imaginary Observations

---

- After every hand, imagine you saw many more hands

$$V = \sum_{h \in U} \text{MONEY}(h) \frac{P(h)}{\Pr(h \in U)}$$

$$= \sum_h I(h \in U) \text{MONEY}(h) \frac{P(h)}{\Pr(h \in U)}$$

$$E[V] = \sum_h \Pr(h \in U) \text{MONEY}(h) \frac{P(h)}{\Pr(h \in U)}$$

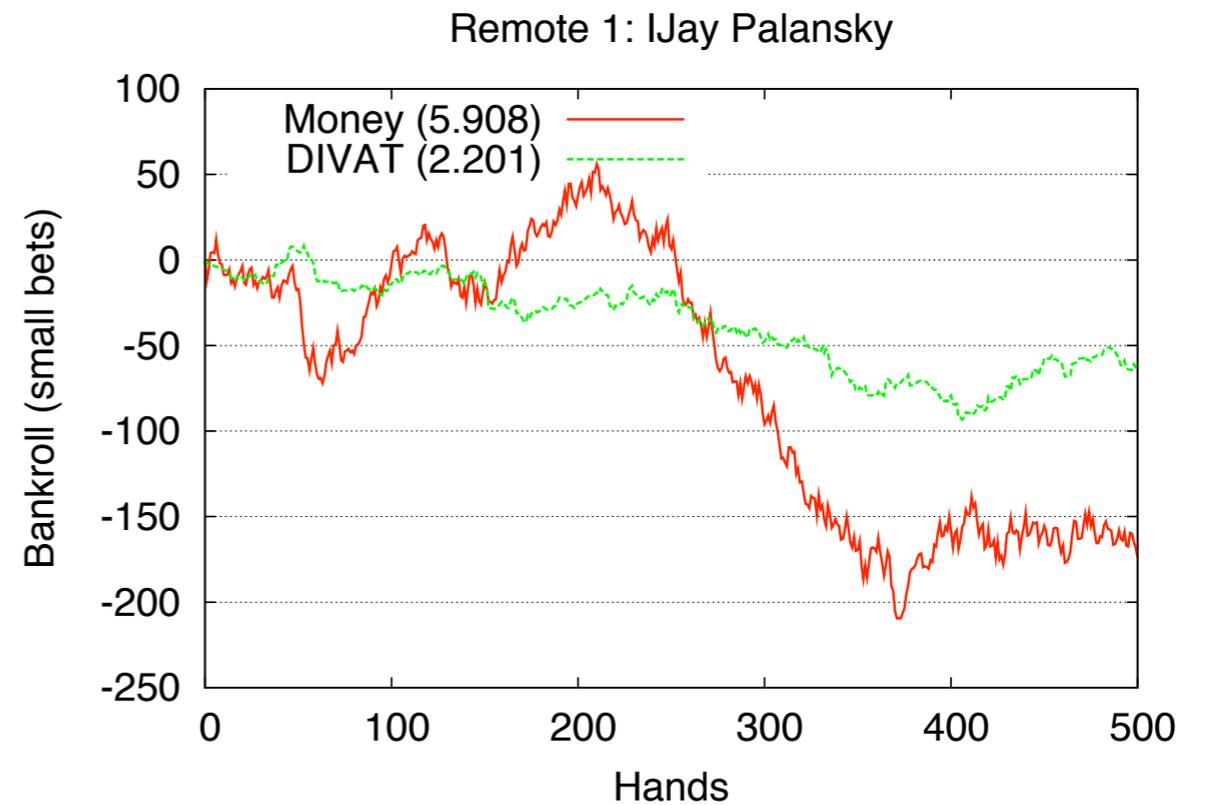
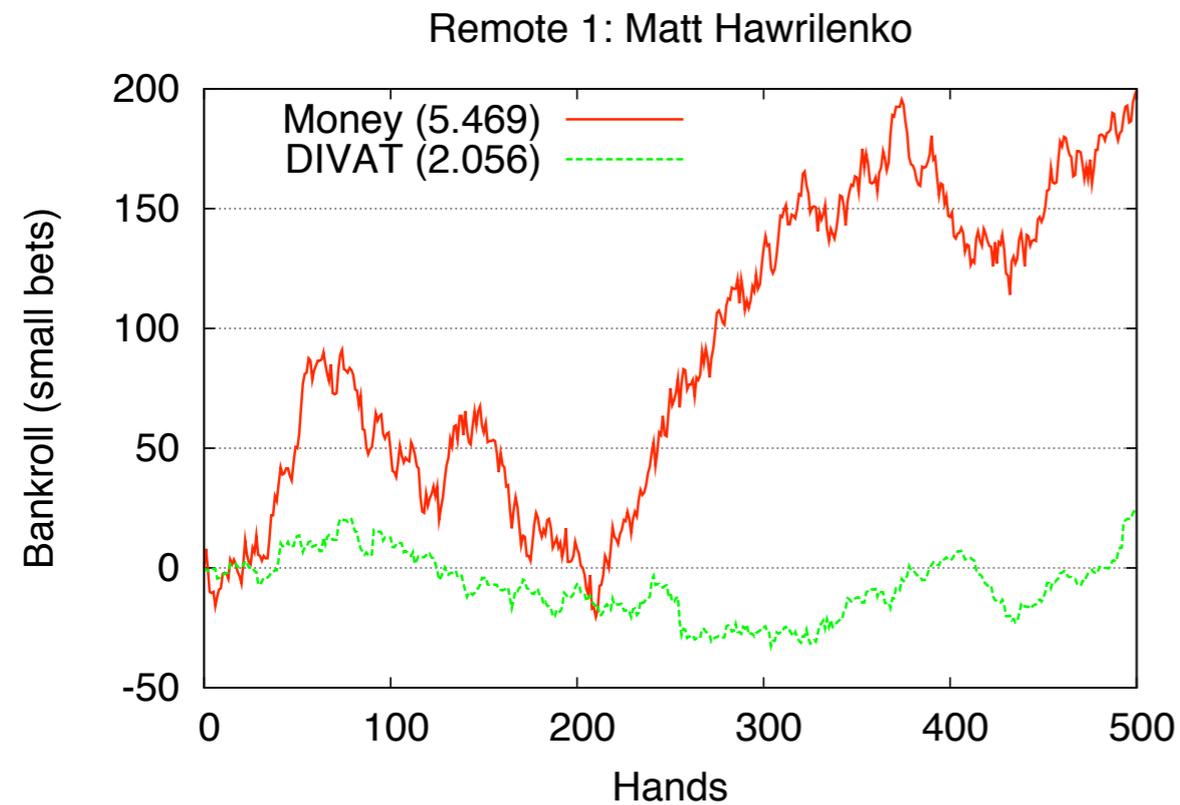
$$= \sum_h \text{MONEY}(h) P(h) = E[\text{MONEY}]$$

# Imaginary Observations

---

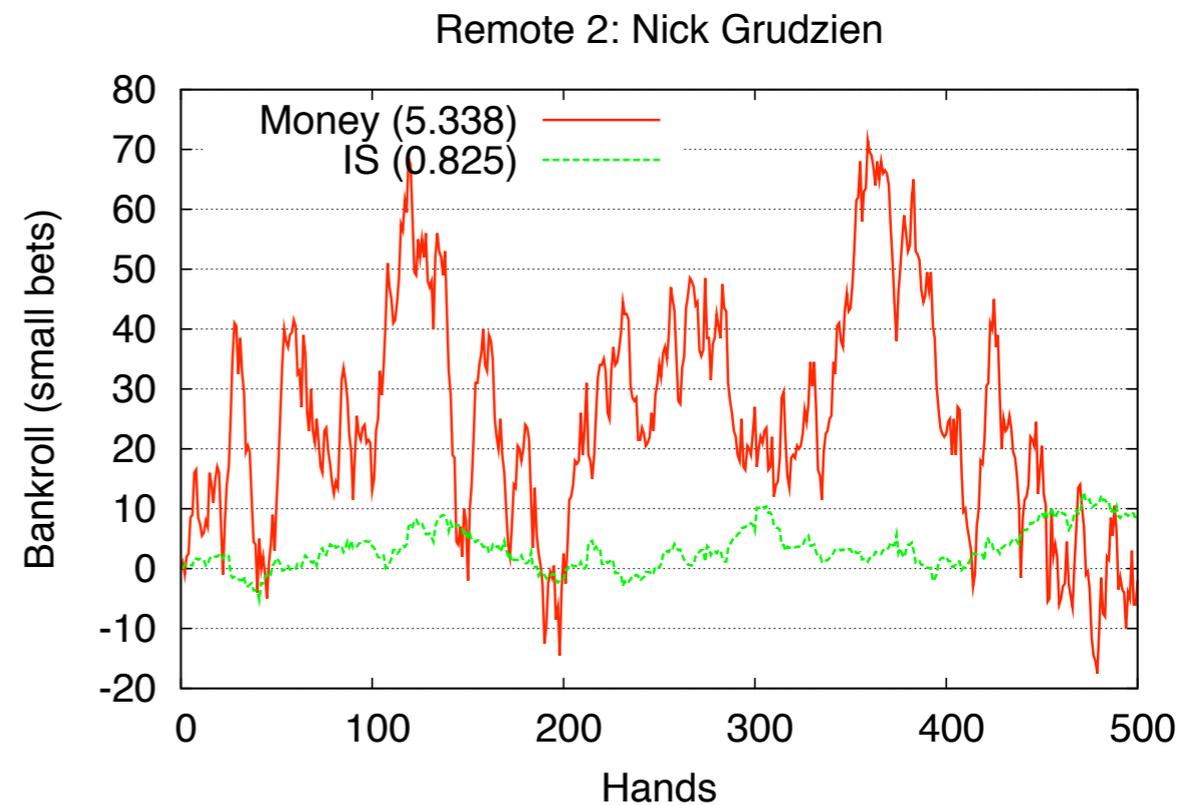
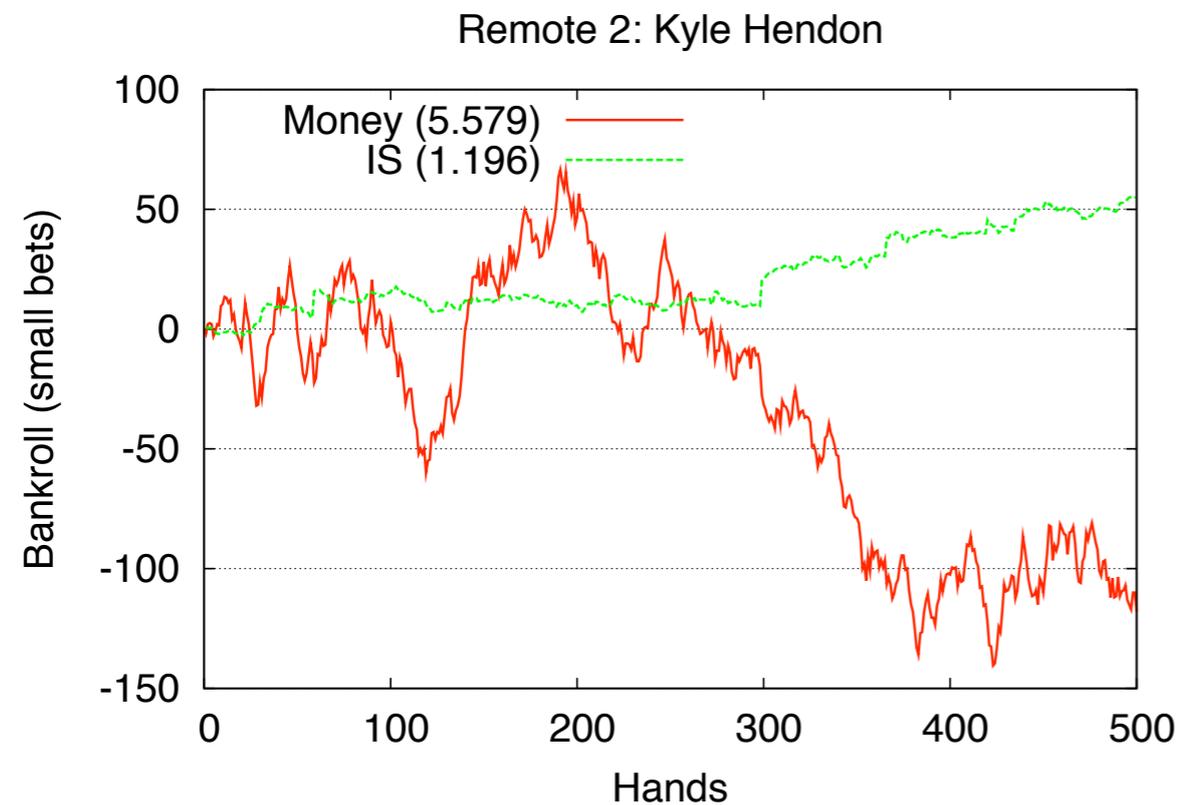
- Imaginary Observations
  - ~10K hands needed for statistical significance
- Imaginary Observations + DIVAT
  - ~1800 pairs of hands
- Imaginary Observations + DIVAT + Duplicate
  - ~500 pairs of hands

# Remote Match 1: Hoss-IJay



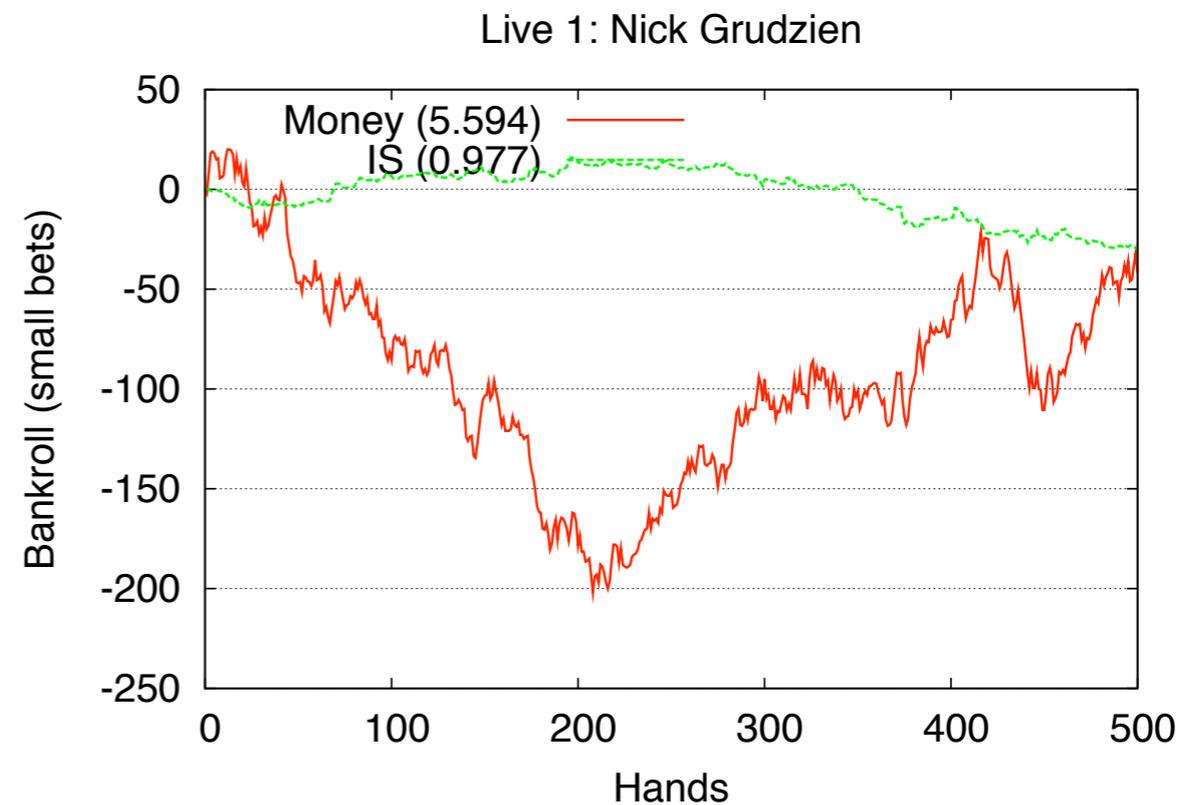
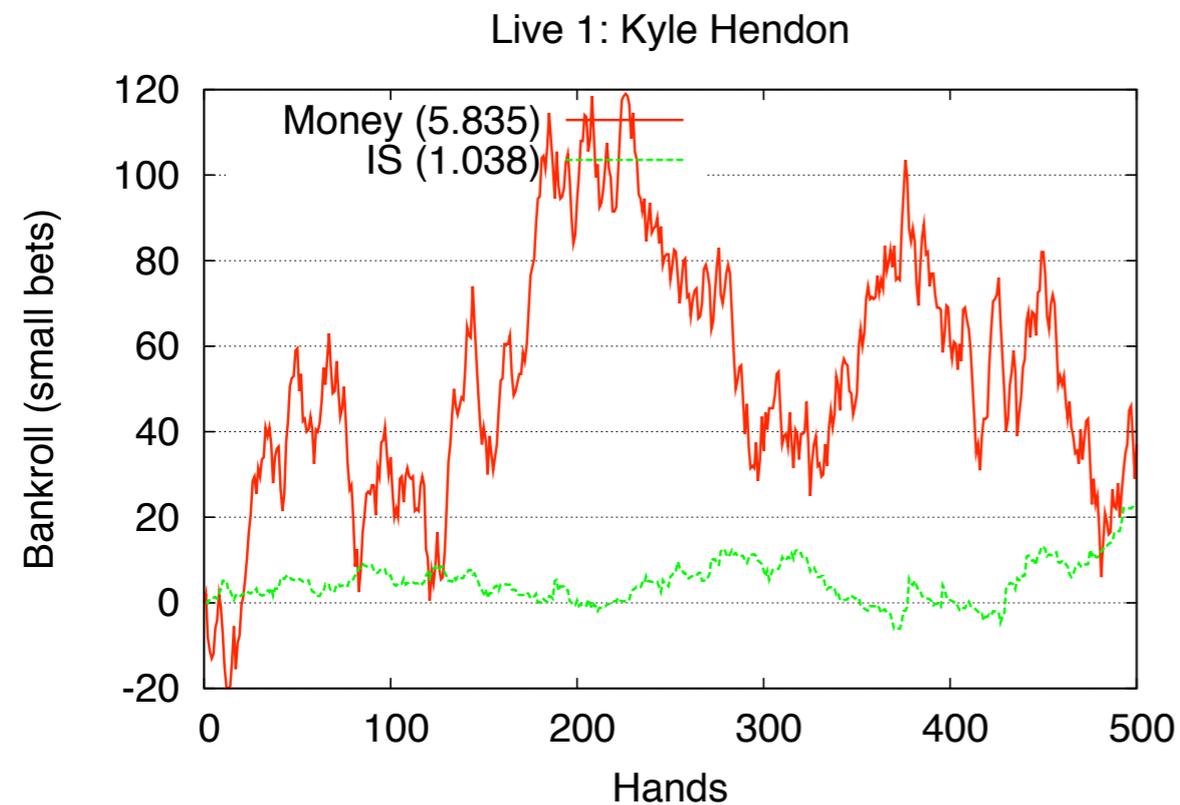
- IO-DIVAT showed Polaris up by 41 bets
  - Not statistically significant, p-value = 0.26

# Remote Match 2: Kyle-Nick



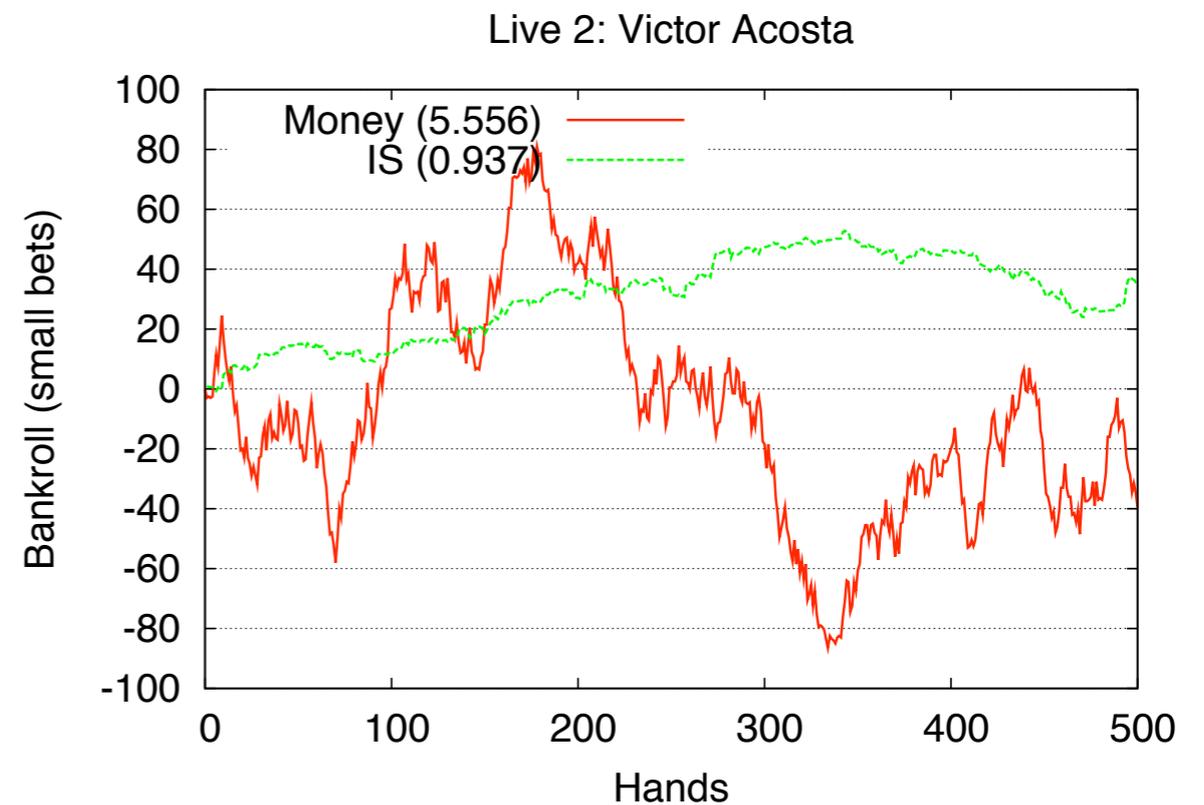
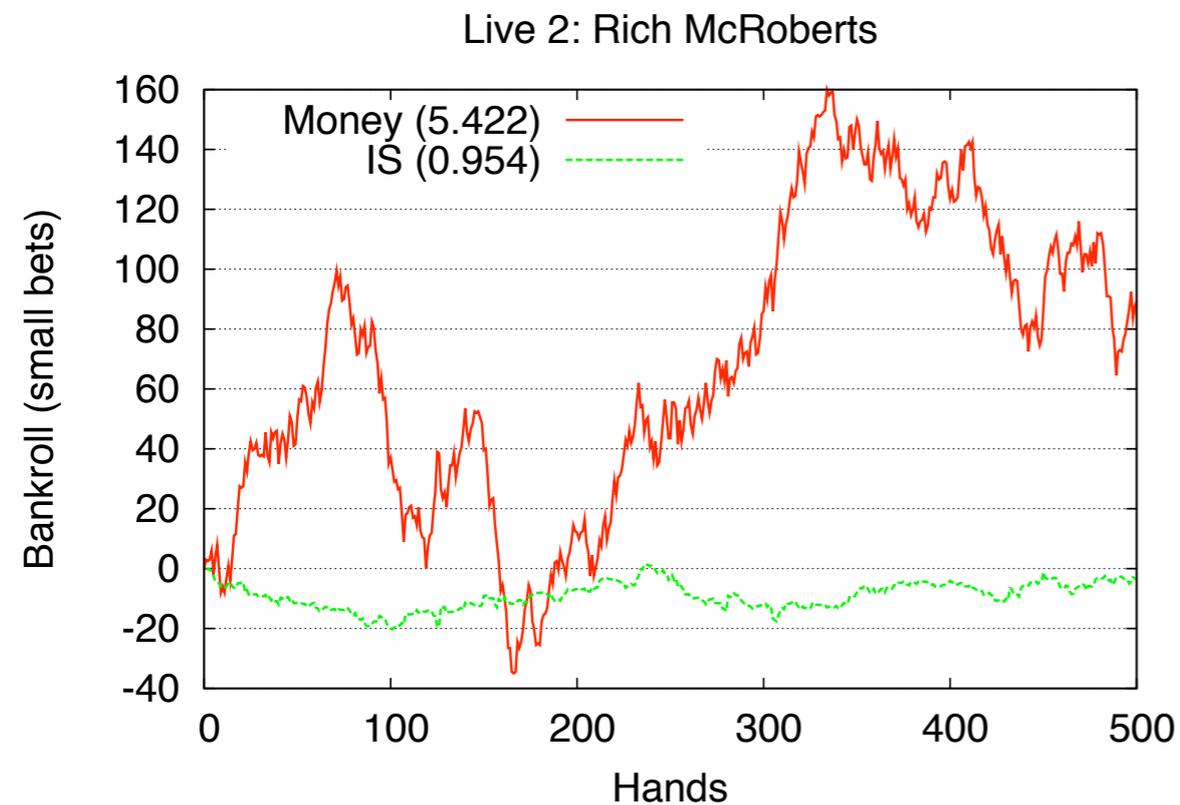
- IO-DIVAT showed Polaris up by 32 bets
  - Statistically significant, p-value = 0.019

# Live Match 1: Kyle-Nick



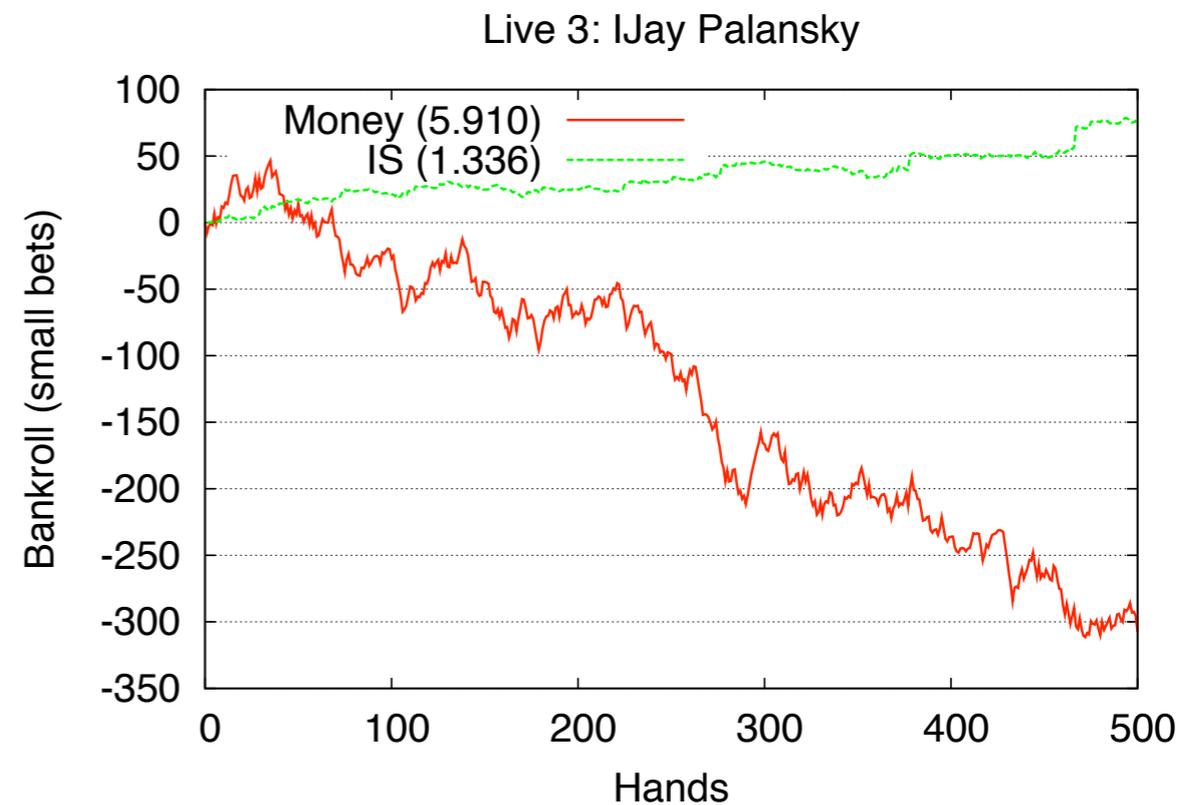
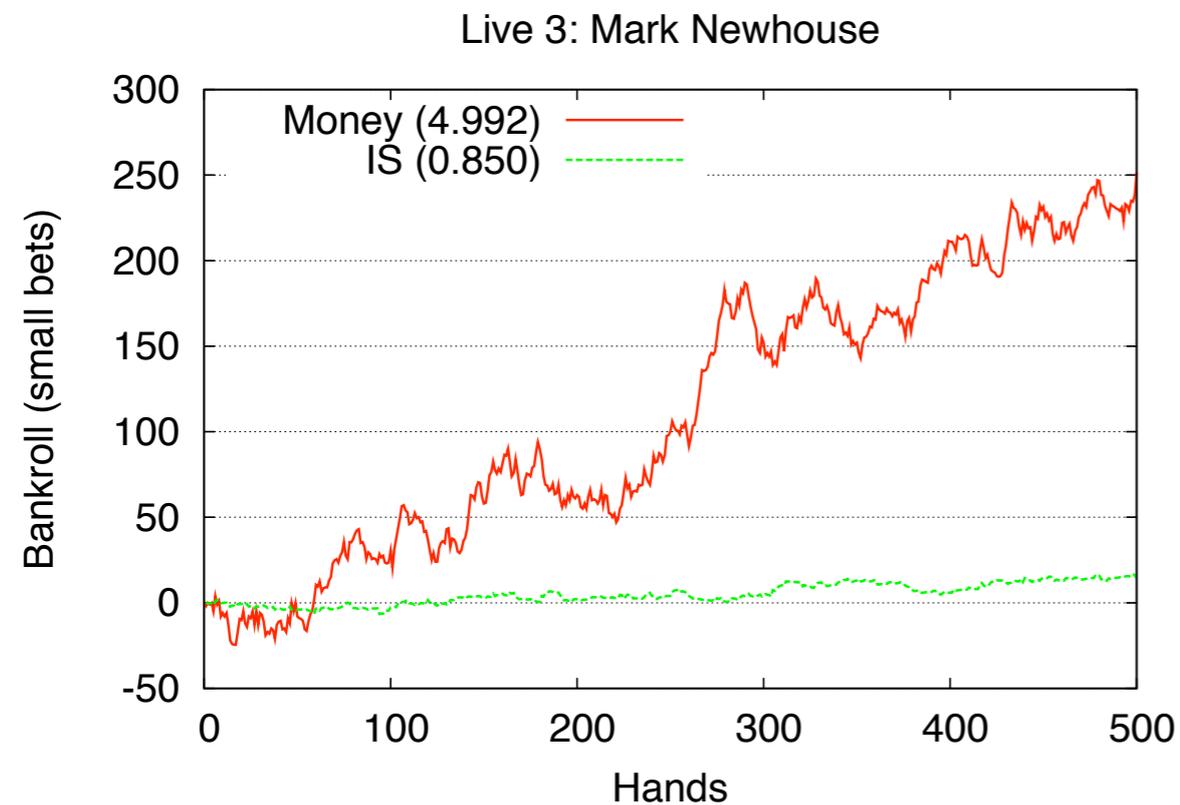
- IO-DIVAT showed Humans up by 3 bets
  - Not statistically significant, p-value = 0.42

# Live Match 2: Rich-Victor



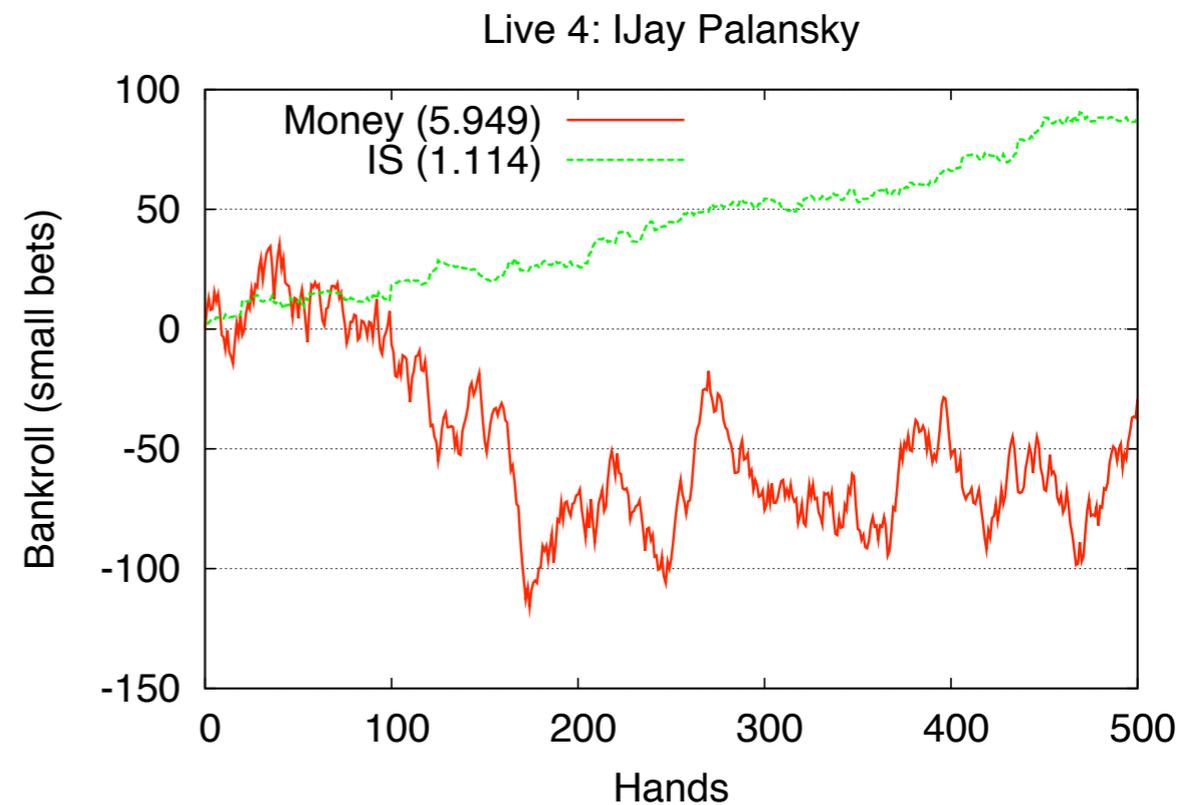
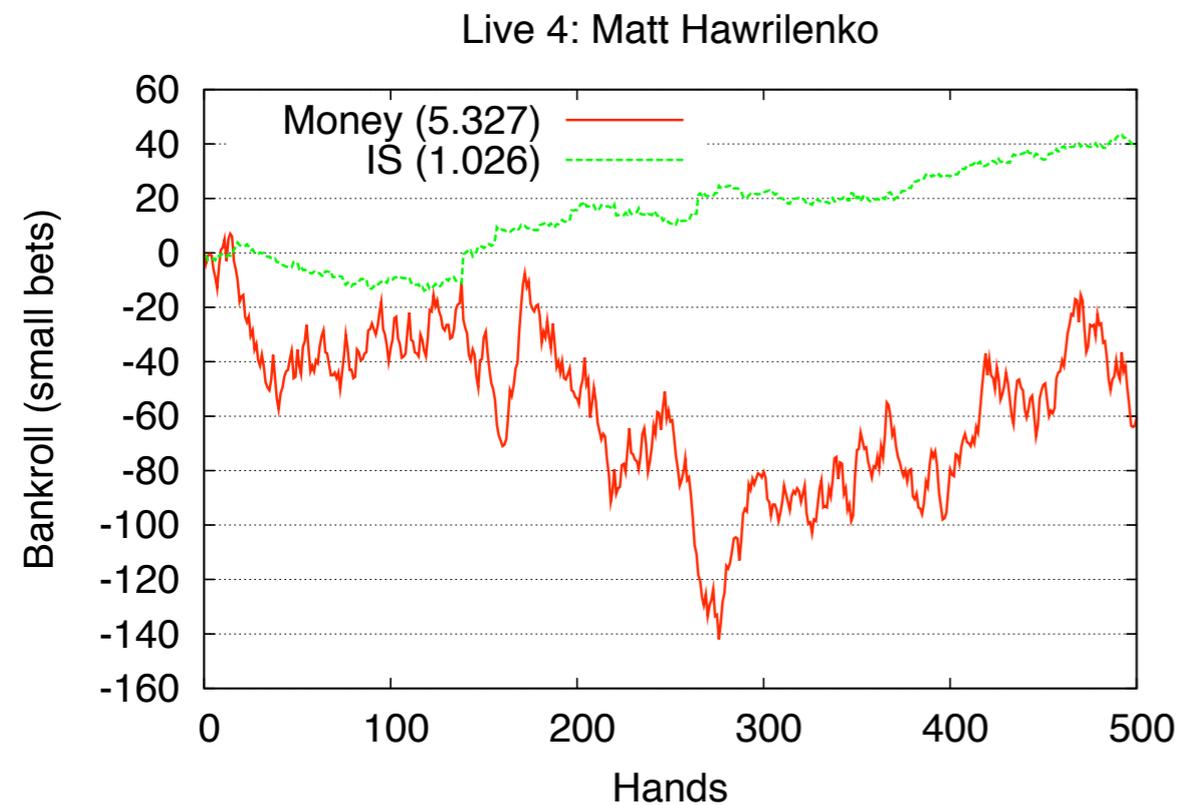
- IO-DIVAT showed Polaris up by 16.5 bets
  - Not statistically significant, p-value = 0.14

# Live Match 3: Mark-IJay



- IO-DIVAT showed Polaris up by 44.5 bets
  - Statistically significant, p-value = 0.002

# Live Match 4: Hoss-IJay



- IO-DIVAT showed Polaris up by 63.5 bets
  - Statistically significant, p-value = 0.00002

# The Results

---

# The Results

---

- Final Result: 3 Statistically Significant Wins, 0 Losses, 3 Draws

# The Results

---

- Final Result: 3 Statistically Significant Wins, 0 Losses, 3 Draws
- Winning rate of over 3 BB/100!

# The Results

---

- Final Result: 3 Statistically Significant Wins, 0 Losses, 3 Draws
- Winning rate of over 3 BB/100!
- Against the best in the world

# Sweet Victory

---

“I sat with the best in the world...  
and I won.”

Mike McDermott, *Rounders* (1998)

# Questions?

---