

Shifty Evidence and Shifty Books

Bob Beddor

Forthcoming in *Analysis*

Abstract

Are all epistemic notions—including evidence and rational credence—sensitive to practical considerations? A number of philosophers have argued that the answer must be ‘No’, since otherwise rational agents would be susceptible to diachronic dutch books (Greco 2013; Rubin 2015; Schroeder 2018). After unpacking this challenge, I show how it can be resisted by appealing to an analogy between shifting stakes and memory loss. The upshot: pervasive epistemic shiftiness may be tenable after all.

1 Preview

According to ‘shifty epistemology,’ the truth-values of knowledge ascriptions vary with practical factors. For example, subject sensitive invariantists hold that whether you know p depends on the practical consequences of relying on p . Contextualists advance similar theses, couched as claims about the extension of the word ‘knows’ rather than knowledge itself.¹

This paper investigates an under-explored challenge for shifty epistemology. The challenge is this: the main motivations for getting shifty about knowledge carry over to other epistemic notions, such as *evidence* and *rational credence*. But it has been argued that a shifty treatment of evidence and rational credence has disastrous implications for decision theory (Greco 2013; Rubin 2015; Schroeder 2018). In particular, a clever bookie can construct a series of bets with increasing stakes, thereby manipulating the bettor’s evidence in a way that ensures they come out poorer. Since dutch bookability is widely thought to be a sign of irrationality, this seems to be a *reductio* of any epistemology that posits pervasive shiftiness.

While a number of philosophers have developed versions of this challenge, shifty epistemologists have said little by way of response. This paper rises to their defense. Appealing to cases of memory loss, I argue that diachronic dutch bookability does not always entail irrationality. It only does so *when the agent’s evidence remains the same across the sequence of bets*. This gives shifty epistemologists a compelling response to the dutch book challenge.

¹I borrow the ‘shifty epistemology’ label from Fantl and McGrath (2012). For the sake of expository convenience, I will focus on subject sensitive invariantism. However, everything I say can be readily adapted to apply to contextualist theories. All that is needed is to resort to semantic ascent, reformulating first-order epistemological claims as claims about the extension of our epistemic terms in a context.

By their lights, changes in stakes induce changes in evidence. So while evidential shiftiness leads to a form of dutch bookability, it's a perfectly rational form.

2 From Shifty Knowledge to Shifty Evidence

Here's a common way of arguing for a shifty view of knowledge. Start with the idea that knowledge is abundant: we know lots of things about the external world. Combine this with the assumption that knowledge provides a secure basis for action: someone who knows p can rationally act as if p is true. Finally, observe that whether it is rational to act on the basis of some proposition depends on the stakes of the decision. From these premises, subject sensitive invariantists draw the conclusion that knowledge depends on stakes. In ordinary circumstances, knowledge is abundant. But once the stakes go up, it becomes scarce (Stanley 2005; Fantl and McGrath 2002, 2009; Ross and Schroeder 2014). Contextualists such as DeRose (1992) often argue along similar lines, though for them what shifts is the extension of the word, 'knows'.

While this argument is formulated in terms of *knowledge*, it can be extended to other epistemic notions, such as *evidence* and *rational credence*. To see how it would work in the case of evidence, start with the idea that evidence—like knowledge—is abundant: lots of propositions about the external world can form part of your evidence. To make things concrete, suppose you've just arrived at the office; you see your colleague Constance walk in the door ahead of you. Plausibly, you thereby acquire the proposition *Constance is in the office* ('OFFICE', for short) as part of your evidence.

Next, combine this with the idea that evidence—like knowledge—provides a secure basis for action. This idea is implicit in decision theory: decision theory says to maximize expected utility, and—on at least one natural interpretation—expected utilities are calculated using evidential probabilities. From this it follows that if p is part of your evidence, then it is rational to take p for granted in practical reasoning—i.e., to perform whatever actions maximize expected utility conditional on p .

Finally, observe that whether it is rational to take p for granted depends on the stakes. For example, suppose that having just seen Constance enter the office, a bookie materializes in front of you and offers a bet on whether OFFICE is true. If the bet is low-stakes, it seems rational to accept. But if the bookie asks you to bet your life on this proposition, the rational course is to decline.

A shifty view of evidence offers to reconcile these observations. According to subject sensitive invariantism about evidence, which propositions are part of your evidence is itself a function of the stakes. When the bookie presents a low-stakes bet on OFFICE, your evidence still includes OFFICE, hence it would be rational to accept. But once the stakes are raised, you lose this proposition as evidence, hence you shouldn't bet your life on it. Here too, this view can be given a contextualist twist. On this contextualist variant, shifts in stakes induce shifts in which propositions the phrase, 'my evidence' picks out in your context.

No advocates of shifty knowledge face pressure to go in for evidential shiftiness. No

surprise, then, that many shifty epistemologists have gone this route.² But it's a route that leads directly to the dutch book challenge...

3 The Dutch Book Challenge

There are a few ways of constructing a dutch book against an agent whose evidence shifts with stakes. Here's one:

Stakes Manipulation As before, you have just seen your colleague Constance enter the office. You think that, conditional on Constance being in the office, there's a 50% chance she parked on Main Street, but there's no chance she parked on Main Street if she's not in the office. As before, a bookie appears before you. This time, they start by offering you the following bet: you win \$1 if Constance parked on Main Street, and you lose \$1 otherwise. Since OFFICE is still part of your evidence, you regard this bet as fair, so you take it. Next, the bookie offers you an evidence-destroying bet: a penny if OFFICE is true, your life otherwise. Given a shifty view of evidence, you now lose OFFICE as evidence; consequently, you decline the bet. But now your first bet—concerning Constance's parking—starts to look unwise. After all, your new evidential state leaves open the possibility that Constance is not in the office. Noticing your remorse, the bookie offers to cancel your earlier bet, for a penny surcharge. You gratefully accept.³

Our bookie has constructed a series of offers, each of which you regard as fair, but which—if accepted—are guaranteed to make you poorer.

This seems like an uncomfortable result. It's often maintained that dutch bookability is a hallmark of irrationality. So if the shifty view of evidence—when combined with standard decision theory—says that you are behaving rationally in **Stakes Manipulation**, it would seem that this view of evidence must be wrong.

²Stanley is particularly explicit about this, maintaining that 'all epistemic notions are interest-relative' (2005: 182). Shifty views of evidence are also advocated—in somewhat different forms, on somewhat different grounds—in Greco (2017), Weatherson (2018), and Beddor (2020). This argument also extends naturally to motivate shifty views of rational credence, given the assumption that rational agents assign credence 1 to their evidence. For examples of this sort of view, see Clarke (2013) and Gao (2019). Of course, not all shifty epistemologists countenance evidential and credal shiftiness; see Fantl and McGrath (2009) and Schroeder (2018) for exceptions.

³This example is based on Greco (2013: §4.2). For alternative ways of developing a dutch book that exploits epistemic shiftiness, see Reed (2012); Rubin (2015); Schroeder (2018). I should note that some philosophers have questioned the use of high stakes betting scenarios in the pragmatic encroachment debate (e.g., de Prado Salas 2019; Fassio 2020). For example, some have suggested that we may have prudential reasons to avoid betting our lives on anything; there may also be reason to worry about the trustworthiness of a bookie who would offer such a bet. For the purposes of this paper, I will set aside such worries.

4 Resisting the Challenge

But closer examination shows that this conclusion is too hasty. Consider an analogous case:

Memory Manipulation As before, you've just seen Constance walk into the office. As before, our bookie starts by offering you a low-stakes bet, at even odds, that Constance parked on Main Street. You regard the bet as fair, so you take it. Immediately thereafter, the bookie gives you a sip of water. Bad news: the water was taken from the river Lethe, which induces forgetfulness. This causes you to forget whether you saw Constance enter the office. Post-potation, your earlier bet looks a lot riskier. Noticing your remorse, the bookie offers to cancel your earlier bet—again for a penny. You gratefully accept.

Once again, you've been dutch booked. But you acted rationally at every step.

Cases like **Memory Manipulation** carry an important lesson. Contrary to what is often claimed, is not necessarily irrational to accept a diachronic dutch book.⁴ But this does not mean that there is no connection between diachronic dutch bookability and irrationality. The following restricted connection remains extremely plausible:

RESTRICTED DUTCH BOOK PRINCIPLE It is irrational for *S* to accept a diachronic dutch book *provided S's evidence remains constant across the bets.*

This allows that in **Memory Manipulation**, your behavior is rational. When you imbibe the Lethean libation, you lose evidence that you possessed when taking the first bet.

But if we say this about **Memory Manipulation**, we can say the same thing about **Stakes Manipulation**. When the bookie raises the stakes on the second bet, you lose evidence that you possessed when taking the first bet. In both cases, your exploitability—while unfortunate—does not make you irrational.⁵

5 Reviving the Challenge?

Some readers might have the lingering sense that there is something especially troubling about the sort of exploitability exemplified by **Stakes Manipulation**. But is there any normatively relevant difference between **Stakes Manipulation** and **Memory Manipulation**? Let's consider some candidates for what the difference might be.

⁴Cf. Christensen (1991) for a similar conclusion, arrived at via independent considerations.

⁵Contextualists can avail themselves of a version of the Restricted Dutch Book Principle, provided that they recast it as a claim about the extension of the phrase 'S's evidence', rather than evidence itself. Here's one way this might go. Let $c(b)$ be the context *S* is in, when *S* is considering whether to accept bet b . And let $b_1 \dots b_n$ be a series of a bets. Contextualists can propose:

CONTEXTUALIZED RESTRICTED DUTCH BOOK PRINCIPLE If both (i) accepting $b_1 \dots b_n$ is guaranteed to result in a net loss, (ii) the extension of the phrase, 'S's evidence' remains constant across $c(b_1) \dots c(b_n)$, then it would be irrational for *S* to accept all of $b_1 \dots b_n$.

First, some might emphasize that the sort of exploitability in **Stakes Manipulation** is *foreseeable*. Perhaps this is what is really worrisome about the shifty view of evidence: it entails that you can know in advance that if the stakes are raised between bets, you will be led to accept offers that are guaranteed to result in a loss. But a moment's reflection reveals that this reply is inadequate, since the same could be said about **Memory Manipulation**. After all, you can know in advance that if you forget relevant information between bets, you will be condemned to accept a similarly impoverishing sequence of offers.

Second, some might maintain that dutch bookability is often just a symptom of a more basic rational failing. This idea is suggested by both Greco (2013: 100) and Rubin (2015: 27). They observe that if evidence shifts with stakes, agents will revise their evidential probabilities—and hence their credences—by procedures other than conditionalization, rendering them susceptible to dutch books. Perhaps it's the failure to conditionalize—rather than the dutch bookability *per se*—which makes these agents irrational. However, this reply is also inadequate, for much the same reason: in **Memory Manipulation**, you also revise your evidential probabilities by a procedure other than conditionalization—namely, *forgetting*. But this does not make you irrational.⁶

Finally, our objector might point out that in **Stakes Manipulation**, there is some sense in which you retain the information that Constance is in the office, even after receiving the evidence-destroying bet. After all, you certainly appear to remember seeing her walk in. If asked, you would be willing to say things like, 'Well, I'm fairly confident that I saw Constance walk in; I just wouldn't bet my life on it.' By contrast, in **Memory Manipulation**, you do not in any sense retain the relevant information.

However, this difference is inessential. Consider a variant of **Memory Manipulation** where your Lethean beverage does not completely erase your memory of Constance walking into the office; rather, it diminishes your confidence in this memory. In this 'memory-fading' variant, there's some sense in which you retain the information that Connie is in the office, even after imbibing. But since you are not certain of OFFICE, it is still rational for you to cancel your earlier bet (since the probability that she parked on Main Street, conditional on your current evidence, is less than .5). As before, you've been dutch booked, but your behavior was rationally impeccable. Once again we've failed to find any essential, normatively relevant difference between the two cases.

The upshot: it proves difficult to pinpoint anything especially concerning about the sort of exploitability that arises from shifty evidence. This should be welcome news for shifty epistemologists.⁷

⁶Various philosophers have pointed out that forgotten evidence poses a problem for conditionalization. And some of these philosophers have proposed variants of classical conditionalization designed to overcome this shortcoming (e.g., Skyrms 1983; Titelbaum 2014; Meacham 2016). For our purposes, we need not take a stand on which updating rule is correct. Whatever updating rule handles forgotten evidence should also extend to other forms of evidence loss, including evidence lost via raised stakes.

⁷Thanks to the editors of *Analysis* and two referees for very helpful comments.

References

- Bob Beddor. New Work for Certainty. *Philosophers' Imprint*, 20(8):1–25, 2020.
- David Christensen. Clever Bookies and Coherent Beliefs. *Philosophical Review*, 100(2):229–247, 1991.
- Roger Clarke. Belief is Credence One (In Context). *Philosophers' Imprint*, 13(11):1–18, 2013.
- Javier González de Prado Salas. Extreme Betting. *Ratio*, 32(1):32–41, 2019.
- Keith DeRose. Contextualism and Knowledge Attributions. *Philosophy and Phenomenological Research*, 52(4):913–929, 1992.
- Jeremy Fantl and Matthew McGrath. Evidence, Pragmatics, and Justification. *Philosophical Review*, 111(1):67–94, 2002.
- Jeremy Fantl and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, Oxford, 2009.
- Jeremy Fantl and Matthew McGrath. Arguing for Shifty Epistemology. In Brown and Gerken, editors, *Knowledge Ascriptions*. Oxford University Press, Oxford, 2012.
- Davide Fassio. Moderate Skeptical Invariantism. *Erkenntnis*, 85(4):841–870, 2020.
- Jie Gao. Credal Pragmatism. *Philosophical Studies*, 176(6):1595–1617, 2019.
- Daniel Greco. Prodigality and Probability. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, volume 4, pages 82–107. Oxford University Press, Oxford, 2013.
- Daniel Greco. Cognitive Mobile Homes. *Mind*, 126(501):93–121, 2017.
- Christopher J.G. Meacham. Ur-Priors, Conditionalization, and Ur-Prior Conditionalization. *Ergo*, 3(17), 2016.
- Baron Reed. Resisting Encroachment. *Philosophy and Phenomenological Research*, 85(2):465–472, 2012.
- Jacob Ross and Mark Schroeder. Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88(2):259–288, 2014.
- Katherine Rubin. Total Pragmatic Encroachment. *Pacific Philosophical Quarterly*, 96:12–38, 2015.
- Mark Schroeder. Rational Stability Under Pragmatic Encroachment. *Episteme*, 15(3):297–312, 2018.
- Brian Skyrms. Three Ways to Give a Probability Function a Memory. In Earman, editor, *Testing Scientific Theories, Minnesota Studies in the Philosophy of Science*, volume 10, pages 157–161. Minnesota University Press, 1983.
- Jason Stanley. *Knowledge and Practical Interests*. Oxford University Press, Oxford, 2005.
- Michael Titelbaum. *Quitting Certainties: A Bayesian Framework for Modeling Degrees of Belief*. Oxford University Press, Oxford, 2014.
- Brian Weatherson. Interests, Evidence, and Games. *Episteme*, 15:329–244, 2018.