THE BEST LAID SCHEMES OF MICE AND ME

RANDALL G. MCCUTCHEON

ABSTRACT. Adam Elga has argued that holders of imprecise credences fall prey to missed arbitrages, so that rational credences should be sharp. A decision rule proposed by Rohan Sud, *Forward Looking*, enables imprecise Bayesians to sidestep missed arbitrages and other "bad books" in isolated fixed-evidence binary betting sequences such as those of Elga. We show that *Forward Looking* imprecise Bayesians are committed to a bad book of bets when faced with a particular 3-bet variable-evidence binary betting sequence.

A. Elga (2010) maintains that it is irrational to hold imprecise credences. In particular, it is claimed that the imprecise rational agent cannot respond appropriately to a sequence of bets such as the following:

Bet A If H is true, you lose \$10. Otherwise you win \$15.
Bet B If H is true, you win \$15. Otherwise you lose \$10.

An agent with sufficiently imprecise credences, operating with "permissive" choice rules (rules allowing her to reject bets when acceptance would not decrease expected utility under at least one credence function in some non-singleton *representor*) might reject both bets (by employing a μ in said representor with $\mu(H) \geq .6$ when evaluating the expectation of **Bet A** and employing a ν for which $\nu(H) \leq .4$ when evaluating the expectation of **Bet B**). But rejecting both bets is irrational, since taken together they constitute an arbitrage opportunity.

Elga explores (in order to reject) several imprecisionist options for dealing with the above thought experiment. One of these is of interest to us here:

> *Plan:* Whenever a rational agent with unsharp credences performs an action, she simultaneously forms a plan governing her later actions. That plan requires her later actions to cohere with the action she just performed. If nothing unforeseen happens, she then follows through on her plan. In particular, in the Bet A/Bet B situation, whenever a rational agent rejects Bet A, she also simultaneously plans to accept Bet B. Later, when she is offered Bet B, she implements her plan.

Elga's argument against *Plan* involves an agent in a contrasting situation in which only Bet B is offered. It's not plausible, the argument goes, that the agent in the contrasting situation would be free to reject Bet B while the agent in the original situation is required (having rejected bet A) to accept it.

R. Sud (2014) grants that the imprecise Bayesian must not make decisions, over all problems, identical to a precise Bayesian, must not reject both Bet A and Bet B, and must not be influenced by past decisions. Still, he thinks that the imprecise Bayesian has resources to handle the current thought experiment. We needn't consider the decision rule (which he calls *Forward Looking*) in its complete generality. Rather we will content ourselves with an informal rendition. Sud has his imprecise, *Forward Looking* agent reason:

> Given that I can't control my future self's decision, the safest thing to do is to just accept (Bet A)–that guarantees me an acceptable sequence of actions.

Sud calls *Forward Looking* an "impressively robust decision rule", noting that:[1]

> Faced with simple 'good books' like Elga's, in which a miser is presented with a finite series of bets on some proposition, each of which the agent can accept or reject, *Forward Looking* will *never* allow the agent to pass up a good book of bets for a strictly worse book of bets.

Sud proves this claim for isolated sequences of bets made under the auspices of *Forward Looking*; in particular, for isolated sequences of binary (e.g., accept-or-reject) bets on a single proposition throughout which the agent's epistemic situation is held constant. On the other hand, Sud acknowledges the difficulty of generalizing the scheme to more complicated (e.g. non-binary) sequences. For suppose (to simplify an example from Sud 2014), one is offered:

Bet C If H is true, you win \$3. Otherwise, you lose \$3.

Bet D1 If H is true, you win \$2. Otherwise, you lose \$4. Or (but not both)
Bet D2 If H is true, you lose \$4. Otherwise, you win \$2.

Bets D1 and D2 are offered simultaneously; one may reject both or accept one. (So together these constitute a "trinary bet".) If the sufficiently imprecise Bayesian rejects Bet C, she risks accepting Bet D1 (dominated by Bet C); if she accepts Bet C, she risks accepting Bet D2 (dominated by rejecting all bets).

Sud reminds us that *Forward Looking* is silent about such cases: the rule is only written for binary decision sequences. Accordingly, they give no cause to think that *Forward Looking* is *wrong*; for all they bring into relief, an agent committed to it might–via further, careful augmentations ("backward" looking, perhaps) to her theory–still avoid "bad books". If such a task could be completed successfully, one could perhaps then view *Forward Looking* as an important first step on the road to vindication for the imprecise utility miser.

[1] Sud takes a book A of bets to be "strictly worse" than a book B when it affords lesser expected gains for every member of the agent's representor. He writes: "A sequence, s, (of decisions) is admissible iff according to at least one probability function c in your representor, the expected utility of enacting s is greater than or equal to all other sequences...".

Alas, this project must fail. For let H be a proposition for which a permissive and *Forward Looking* agent has sufficiently imprecise credences (supported on $[\frac{1}{3}, \frac{2}{3}]$, say), and suppose also that a fair coin is tossed independently of H. Let

$$Q = (H \wedge \ tails) \vee (\neg H \wedge heads).^2$$

The agent knows that she will be presented with the following sequence of bets:

Bet 1 If Q is true, you win \$11. Otherwise you lose \$9.
Bet 2 If Q is true, you lose \$12. Otherwise you win \$8.
Bet 3 If Q is true, you win \$13. Otherwise you lose \$7.

So far, none of these bets look to concern a proposition for which the agent has imprecise credences. Indeed, denote by μ_x any member of the agent's representor with $\mu_x(H) = x$, where $x \in [\frac{1}{3}, \frac{2}{3}]$. Since the coin is uncontroversially fair and independent of H, one will have $\mu_x(H \wedge \ tails) = \frac{x}{2}$ and $\mu_x(\neg H \wedge heads) = \frac{1-x}{2}$. Hence $\mu_x(Q) = \frac{x}{2} + \frac{1-x}{2} = \frac{1}{2}$. So the agent must accept Bet 1, as its expected value is $+1$ by every member of her representor. (Cf. footnote 1.)

We now introduce an addendum to the protocol.

Toss disclosure: The agent knows that between her decision on **Bet 1** and the offer of **Bet 2**, she will learn whether the coin landed *heads* or *tails*.

If the agent learns *tails*, then since $\mu_x(Q|tails) = x$ for each $x \in [\frac{1}{3}, \frac{2}{3}]$, the agent's credence in Q will "dilate", becoming imprecise with support $[\frac{1}{3}, \frac{2}{3}]$. If the agent learns *heads*, meanwhile, then since $\mu_x(Q|heads) = 1 - x$, the agent's credence in Q will similarly dilate.[3] The upshot is that in the agent's deliberations as to whether or not to accept **Bet 2**, *Forward Looking* kicks in and mandates that she must accept. For she may very well reject **Bet 3**, and to reject both **Bet 2** and **Bet 3** would be to pass on an arbitrage opportunity.

At this point, however, the agent has accepted both **Bet 1** and **Bet 2**. This pair constitutes a bad book, for it entails a sure loss of one where the agent could easily have guaranteed a payoff of zero (by refusing both bets). Worse, the agent knew in advance what she would do, yet did it anyway. Checkmate: to maintain imprecise Bayesian status, the agent *must* renounce *Forward Looking*.

---

[2] A proposition similar to Q was employed (to different effect) by Roger White (2010). Thanks to an anonymous referee for pointing this out.

[3] White (2010) argued that such dilation is in violation of something somewhat along the lines of the so-called Reflection Principle: if the agent knows that her credence in Q will be imprecise supported on $[\frac{1}{3}, \frac{2}{3}]$ after learning the result of the toss, it ought to be imprecise before learning the result of the toss. As is clear from the text, however, if it's rational to have imprecise credences in the first place then dilation does occur. (See also Joyce 2011.)

References

Elga, Adam. 2010. Subjective Probabilities should be Sharp. *Philosophers' Imprint* 10(5):1-11

Joyce, James. 2011. A Defense of Imprecise Credences in Inference and Decision Making. *Philosophical Perspectives* 24:281-323.

Sud, Rohan. 2014. A forward looking decision rule for imprecise credences. *Philosophical Studies* 167:119-139.

White, Roger. 2010. Evidential Symmetry and Mushy Credence. *Oxford Studies in Epistemology* 3:161-186.

Department of Mathematical Sciences, University of Memphis
E-mail address: rmcctchn@memphis.edu