Character Matching and the Envelope of Belief

§1. Probabilism and The Lockean Thesis

The Lockean thesis maintains that an individual fully believes a proposition ϕ just when she has a high level of confidence in ϕ. The trouble with the Lockean thesis, according to probabilists like Richard Jeffrey, is that it licenses throwing away perfectly good information. Any numerically determinate degree of belief that happens to fall above the Lockean's threshold for acceptance is abandoned in favor of a qualitative label, 'full belief'. What strict probabilists want to know is whether there is anything gained in the exchange.

To be sure, some probabilists take the case against full belief too far: the idea that full belief can be entirely replaced by credal probability, for example, is a non-starter. Probability presupposes possibility, and judgements of possibility--even if only to pick *this* Boolean algebra rather than *that* one for your probability structure--are categorical. Further still, there are alternatives to probabilism. Henry Kyburg's lottery paradox was a campaign button for his ε-acceptance theory of evidential probability, which is deeply at odds with probabilism. David Makinson's puzzle about prefaces preceded his important work on systems that facilitate genuine inductive expansion (Makinson 2005), yet inductive expansion is notoriously difficult to square with probabilism (Levi 1980, Paris and Simmonds 2009).

Nevertheless, Kyburg's lottery paradox and Makinson's paradox of the preface have developed a life of their own as Rorschach tests for the metaphysics of rational belief. In what sense, if any, is belief categorical? In what sense, if any, is it gradable? How weird does the logic have to be for us to have it both ways? The strict probabilist says *none*, *completely*, and *too weird*. The mainstream Lockean says that rational full belief summarizes information conveyed by a high degree of credal probability and the logic for full belief is classical except that it will not be closed under conjunction. Here then is a threshold acceptance view that is purported to be in harmony with the spirit of probabilism: 'Full belief' is simply code for 'has high credal probability'. But then the harmonious Lockean faces the question that started us off: why bother?

§2. Character Matching

In the course of defending a mission for harmonious Lockeanism, Scott Sturgeon introduces a normative principle called *character matching* which maintains that the character of a belief should match the character of the evidence upon which it is based.

> When evidence is essentially sharp, it warrants a sharp or exact

> attitude; when evidence is essentially fuzzy...it warrants at best a fuzzy attitude. In a phrase: evidential precision begets attitudinal precision; and evidential imprecision begets attitudinal imprecision (Sturgeon 2008: 159).

Sturgeon's idea is that everyday evidence is often imprecise in character; so, by the principle of character matching, everyday evidence will seldom rationalize numerically determinate credences but instead will typically rationalize only regions of credal space. Yet, when imprecise evidence for a claim occupies the range [θ, 1], for a suitably high threshold value θ between 1/2 and 1, Sturgeon claims that 'lending that confidence to [the] claim functions exactly like believing it in a threshold-based way' (160). Harmonious Lockeanism is thus conceived to be an extension of probabilism to cover the majority of cases in which everyday evidence is strong but imprecise. So, the harmonious Lockean is not throwing away information when he fully believes; he is making the best of the information he's got.

That commonplace evidence is more often imprecise than precise is surely right. I may observe that A is more frequent than B either by rough observation of the occurrence of A's versus the occurrence of B's or by more careful methods of statistical estimation. Neither case yields numerically determinate frequencies, however, unless of course my estimation is gotten directly from a complete description of the occurrence of A's and the occurrence of B's in the population. If I know that the proportion of red balls in the urn is exactly 0.85 because I loaded the balls into the urn myself, say, then I know that a sample set of 100 balls drawn from the urn should contain roughly 85 red balls. In this way I can calibrate my sampling mechanism because I have a complete description of the proportion of red balls in the population, i.e., the urn. If I do not know the proportion of red balls in the urn but I draw several 100-ball sample sets randomly, never observe less than 80 red balls in a sample and likewise never observe more than 90 red balls in a sample, then I may estimate (conservatively) that the proportion of red balls in the urn is between 0.8 and 0.9. In this way I learn about the frequency of red balls in the population by observing samples from the population. That's how we get much of our evidence about the world, and we may form rational beliefs on the basis of such evidence. Why? Because it is extremely unlikely to go through an exercise like this one, observe the same proportions of red balls from a set of 100-ball samples, yet have it be the case that I am drawing from an urn whose proportion of red balls is in fact outside of the interval [0.8, 0.9]. We may fully believe that the actual proportion of red balls in the urn is between .8 and .9 because, by manipulating the number of samples drawn from the population, we may control the risk faced from adopting this attitude in error. We may judge that the proportion of red balls in the urn is within [0.8,0.9], or that the proportion of red balls is roughly 0.85, or that it is roughly [0.8, 0.9]; analogously, we may judge that the length of a table leg is 71 +/-.03 cm, or that it is roughly 71 cm, or roughly between 70.97 and 71.03 units centimeter in length.

But while it is clear that much of our evidence is imprecise, it is far less clear what bearing this fact has on rational full belief. Sturgeon maintains that high imprecise credence is identical to full belief, and the linchpin to his argument is the normative principle that 'epistemic perfection demands character match between evidence and attitude' (160).

One might probe whether the character matching principle is true by considering cases where precise evidence appears to at best warrant only imprecise belief. Imagine that Fen the fence is trying to sell You[1] a rigged lottery ball machine. Fen tells You that the machine is calibrated to dispense a red ball 70% of the time and to dispense an even numbered ball 60% of the time, and that is all that You are told about the machine. A ball is dispensed. Arguably, before observing the ball Your credence that it is red is 0.7 and Your credence that the ball is even is 0.6. However, Your credence that the ball is both red *and* even is indeterminate, taking any value between 0.3 and 0.7.[2] Here the character of Your evidence is precise but incomplete: knowing the probability that a dispensed ball is red and the probability that a dispensed ball is even warrants at best an imprecise attitude about the ball being both red and even.

The example is not a problem for Sturgeon, however. The reason that Your disposition toward the proposition [*the ball is both red and even*] is a credal *state* bound by the closed interval [0.3, 0.7] rather than a numerically precise credal *probability* is due directly to Your evidence, which is that the probability of the ball being both red and even is between 0.3 and 0.7. Thus, rather than being an objection to the character matching principle, the example instead appears to be a shining example: Your imprecise belief that the ball is both even and red matches the imprecise character of Your available evidence.

Although the character of evidence and attitude appear to match in this example they do not match as a rule. To see this, imagine that it is the middle of the night and Claudius cannot sleep because of a recurring cough. Groggy, he makes his way in the dark to the medicine cabinet for cough syrup, finds a bottle, unscrews its top, and swallows a liquid that tastes like cough medicine but burns rather than soothes his throat. Panicked, he gropes for the light to see what he has ingested, knocking to the floor and shattering two bottles, A and B. Claudius determines that he drank from one of the two, neither of which contains cough syrup. Although unable to determine which bottle he drank from, he believes that D [*it is more likely that Claudius drank from A than from B*].

---

1 Following a long tradition that includes Walley, de Finetti, and Good, I sometimes use 'You' in examples to denote an intentional system and invite you, the reader, to play along in the role.

2 In general, for arbitrary propositions A and B, if p(A) and p(B) are defined with respect to a probability structure M, then with respect to M we have

1. $p(A \wedge B)$ lies in the interval $[\mathbf{max}(0, p(A) + p(B) - 1), \mathbf{min}(p(A), p(B))]$, and
2. $p(A \vee B)$ lies in the interval $[\mathbf{max}(p(A), p(B)), \mathbf{min}(p(A) + p(B), 1)]$.

For a proof, see (omitted).

At the hospital the attending physician, Phy, tells Claudius that A-poisoning is best treated by rest, whereas the recommended treatment for B-poisoning is to have ones stomach pumped. Pumping is inadvisable for A-poisoning, however, and while sleeping off B-poisoning won't kill him, it entails a much longer recovery period than stomach pumping. To make the example concrete, suppose Phy's utilities for these intervention options are given in the following table.

	A	B
Rest	1	0.1
Pump	0	0.9

Suppose that Claudius has told Phy the entire story so far except that he has withheld reporting that it is more likely that he ingested A than B. Based on the partial information that Phy has about the case, Phy's preference about which intervention is best for Claudius is highly indeterminate, ranging within [0.1,1] for rest and within [0,0.9] for the pump. This is so because at this stage Phy has no evidence whatsoever about whether it was A or B that Claudius ingested: for Phy, the probability that Claudius drank A is within [0,1] and the probability that Claudius drank B is within [0,1].

Now Claudius offers up to Phy the missing piece of information. With D, Phy's evidence is that the probability that Claudius has A-poisoning is within (.5,1] and the probability that Claudius has B-poisoning is within [0,.5]. This is very imprecise evidence for A-poisoning, and is well outside the 'bel-region' that Sturgeon imagines is the top 'five to fifteen percent of the scale' for credal probability that is supposed to sustain threshold-based belief (160-1). Even so, notice that this imprecise piece of evidence proves decisive in favor of rest.[3] That is, *any* classical probability function p that is consistent with the complete evidence will entail that rest is the best treatment. Thus, Phy's attitude toward the proposition R [*Rest is the correct treatment for Claudius*] changes from an uncommitted attitude to full rational belief solely on an item of evidence, D, that is highly imprecise.[4]

Curiously, the doctor's full belief that R is in a sense made stronger rather than weaker by the imprecise character of D. Were Phy a strict probabilist and insist upon extracting a numerically precise credence from Claudius about his level of confidence that he drank from bottle A, Phy's request would add nothing to support her full belief that rest is best.

3 To simplify matters, one may calculate expected utility (eu) on the closed interval [.5,1] instead of the clopen interval (.5,1] without loss of generality. (See note 4.) There are four expected utilities to calculate: (1) eu(rest) = 1, when A is maximal and B minimal; (2) eu(rest) = 0.55, when A is minimal and B is maximal; (3) eu(pump) = 0.45, when B is maximal; (4) eu(pump) = 0. Because (pump,A) = 0 in the utility table, we may ignore those probability values in the last pair of calculations for pump.

4 The example works for some values for p(A) < .5 as well; specifically, for all probability functions p such that p(A) is strictly greater than $.44\bar{4}$. Since the Lockeans typically restrict the range of admissible thresholds to reals greater than 0.5, the example cannot be explained away by arguing that the context licenses taking a minimal threshold value of 0.5.

Indeed, we might reasonably doubt Phy's good judgement if she bothered to ask.

In short, Phy's rational full belief that R rests on evidence about A-poisoning whose imprecision extends beyond the bel-region alleged to correspond to high-threshold full belief. What's more, we observed as an aside that there is decisive evidence in this case whose lower-bound for A-poisoning is strictly below 1/2. The upshot is that if Claudius either: (1) reported that he more likely drank from A than from B, (2) reported that it is slightly more likely that he ingested B than A, or (3) reported any numerically precise credence greater than 0.445 that he drank from A, then Phy would have formed the same categorical belief that R. None of the evidence in these three cases is in the bel-region, nor is the region of credal probability associated with R necessarily in the bel-region.

§3. The Envelope of Belief

Full belief and credence do not necessarily match in character nor should they. Credal probability (degrees of belief) encodes a disposition to make a collection of bets on the truth of ϕ. (That's Ramsey.) Credal states are a set of admissible credal probability functions, where the conditions for admissibility are a subtle affair. (That's Levi.) The disposition to fully believe ϕ, on the other hand, is the disposition to act as if ϕ were true relative to some specified range of actions. This is not the same thing as a disposition to accept bets on ϕ.[5] We imagined that Phy, once she has the last piece of evidence about Claudius's case, acts as if it is true that rest is the best treatment for him. It is in this respect that full belief is contextually dependent, since an agent's disposition to act as if ϕ is true may vary across contexts even when both the possible set of actions and the evidence for ϕ are held constant.

Some think that this observation about full belief means that the disposition to fully believe depends on the magnitude of the stake. 'The more that you care the less that you know,' as Jason Stanley has put it.[6] And there are plenty of examples that seem to support this view. A farmer might act as if a vaccine for flu is non-toxic to his pigs, but refuse to act as if the same vaccine is non-toxic to his children. In short, he may fully believe that the injection is safe for his pigs but not for his kids. This may be so even if the vaccine's fatality rate for children is less than the fatality rate for pigs. The farmer simply values his children much more than he values his pigs, viewing the risk of error acceptable for his pigs but not for his kids.

But focusing on the magnitude of the risk of being wrong is only half of the story; we must

5 The probability of a sequence of 50 flips of a fair coin landing all heads is 8.882×10^{-16} and ordinarily we act as if this outcome will not occur. But no one is willing to *offer* odds of $1 to $1.126 quadrillion against seeing 50 straight heads; Wall Street brokers are a cautionary exception. Furthermore, few would bother to take those odds even if offered.

6 But the idea goes back at least to R. B. Braithwaite (1949).

also focus on the potential reward from being right. To illustrate, suppose that an antibiotic for a relatively new strain of *Escherichia coli* is developed to treat pigs and children, and the antibiotic poses an identical risk to both, killing 1 out of every 15 in each population. However, whereas pigs almost always recover from this strain of *E. coli* on their own, children rarely survive it. Faced with this situation the farmer would give the antibiotic to his children but not to his pigs. That is, in this case the farmer would fully believe that the treatment is non-toxic to his children but toxic to his pigs. Although the farmer values his children much more than he values his pigs, he still values his pigs.[7]

What this last example illustrates is that although full belief is context dependent, it does not depend upon the total magnitude of the stake put at risk, as some would have it, but instead depends on the ratio of the amount put at risk (r) by acting as if ϕ is true when ϕ is false to the amount gained (w) by acting as if ϕ is true when ϕ is true.[8] Intuitively, we say that an agent fully (r:w) believes that ϕ in some context just in case all available actions the agent bases on ϕ in that context have a ratio of risk-to-reward less than r/w. This suggests the following definition.

Full (r:w) belief: An agent fully (r: w) believes ϕ in context C iff: for any act A in C, if
1. the agent judges A to cost r* if ϕ is false,
2. the agent judges A to payout w* if ϕ is true, and
3. r*/w* < r/w, then

the agent acts as if ϕ is true.

If an agent fully (r: w) believes that ϕ in C, then it follows that he fully (w: r) disbelieves ~ϕ in C. Since to risk r is to potentially gain –r, to gain w is to have risked –w, –r/–w = w/r, and ~ϕ is false if and only if ϕ is true, it follows that full (w: r) disbelief that ~ϕ in C holds just in case, for any act A in C, if (1) the agent judges to cost w* if ~ϕ is false, (2) to payout r* if ~ϕ is true, and (3) w*/r* > w/r (equivalently, –r*/–w* < r/w), then the agent acts as if ~ϕ is false. An agent (r: w) suspends judgement with respect to ϕ in C if the agent neither fully (r:w) believes ϕ in C, nor fully (w: r) disbelieves ϕ in C. From the conjugacy relation between belief and disbelief, i.e., full (r: w) belief that ϕ in C if and only if full (w: r) disbelief that ~ϕ in C, we say that the range [w/r, r/w] forms an *envelope of belief*.

To illustrate, the range [1/50, 50/1] might be a candidate for the envelope of belief in ordinary situations. This means that an agent who fully (1: 50) disbelieves the proposition T [*lottery ticket #591 wins*] judges his loss to be w* if T is false, judges the payout to be r* if T is true, and the risk-reward ratio w*:r* associated with his acting 'as if T is false' is strictly greater than 1/50. However, suppose that 1000 lottery tickets are sold for $1.00 each and the house takes a penny from each ticket sale, leaving $990 as the payoff to a single, fairly drawn ticket. Then it will cost the agent $1.00 if T is false, and his reward is $990 if T is true. Since 1/990 is less than 1/50, the agent *does not* disbelieve that ticket T losses.

7 A version of this example and of the risk-reward theory of full belief appears in (Kyburg 1990: 244-54).
8 We assume that an agent has a cardinal utility function over states of the world; thus utilities here are linear.

Likewise, the agent will not act as if T is true since 990/1 is greater than 50/1. He will, like most of us would under ordinary circumstances, suspend judgement about T until the outcome of the lottery is announced. Here the character of the evidence for ~T is in the bel-region, but we neither fully believe ~T nor fully disbelieve it: we have every ordinary reason to suspend judgement.

Given the same evidence about T an agent's attitude may change from suspended judgement to full belief either by increasing the envelope of full belief with respect to ~T, say to 991: 1, or by boosting the reward from believing ~T when ~T is true. It might be difficult to reconcile a [1/991, 1/991] envelope of belief with an ordinary lottery drawing, however, and adjusting the risk-reward ratio to drag ~T inside the envelope of ordinary belief may be viewed as changing the example altogether. Fair enough. But the reluctance to fully believe ~T provides no sound basis for generalization: there are plenty of 'lottery propositions' (Hawthorne 2004) that are consonant with an envelope of ordinary belief. Consider again the length of the table leg mentioned earlier. Unlike lottery tickets, suspending judgement until the announcement of the leg's true length is not an option, and there are appreciable rewards to fully believing L [*Table leg #591 is 71 +/- .03 cm*] when L is true: someone's acting as if L were true figured in the table's construction, for instance. Here too an agent's attitude may change from 'full belief' to 'suspended judgement' by altering parameters in the model to place the proposition L outside of the envelope of ordinary belief, which may easily occur without altering the character of the evidence.[9]

References

Braithwaite, R. B. 1946. Belief and action. *Proceedings of the Aristotelian Society* 20: 1-19.

Hawthorne, J. P. 2004. *Knowledge and Lotteries*. Oxford: Clarendon Press.

Kyburg, Jr., H. E. 1990. *Science and Reason*. New York: Oxford University Press.

Levi, I. 1980. *The Enterprise of Knowledge*. Cambridge, Mass.: MIT Press.

Makinson, D. 2005. *Bridges from Classical to Nonmonotonic Logic*. London: King's College Publications.

Paris, J. B. and Simmonds, R. 2009. O is not enough. *The Review of Symbolic Logic* 2(2): 298-309.

Ramsey, F. P. 1926. Truth and probability. In his*The Foundations of Mathematics and other Logical Essays.* Harcourt, Brace and Co. 1999 electronic edition. http://ideas.repec.org/h/hay/hetcha/ramsey1926.html.

Sturgeon, S. 2008. Reason and the grain of belief. *Noûs* 42(1): 139-65.

9 Thanks to ...