

Dutch Book Arguments and Imprecise Probabilities

Seamus Bradley

August 10, 2012

1 For and against imprecise probabilities

I have an urn that contains 100 marbles. 30 of those marbles are red. The remainder are yellow. What sort of bets would you be willing to make on the outcome of the next marble drawn from the urn? What odds would you accept on the event “the next marble will be yellow?”. A reasonable punter should be willing to accept any betting quotient up to 0.7. I define “betting quotient” as the ratio of the stake to the total winnings. That is the punter should accept a bet that, for an outlay of 70 cents, guarantees a return of £1 if the next marble is yellow. And the punter should obviously accept bets that cost less for the same return, but what we are really interested in is the *most* the punter would pay for a bet on an event.

I am making some standard simplifying assumptions here: agents are risk neutral and have utility linear with money; the world of propositions contemplated is finite. The first assumption means that expected monetary gain is a good proxy for expected utility gain and that maximising monetary gain is the agents’ sole purpose. The second assumption is made for mathematical convenience.

Now consider a similar case. This case is due originally to Daniel Ellsberg;¹ this is a slightly modified version of it.² My urn still contains 100 marbles, 30 of them red. But now the remainder are either yellow or blue, in some unknown proportion. Is it rational to accept bets on Yellow at 0.7? Presumably not, but what is the highest betting quotient the punter should find acceptable? Well, you might say, there are 70 marbles that could be yellow or blue; his evidence is symmetric so he should split the difference:³ a reasonable punter’s limiting betting quotient should be 0.35. Likewise for Blue. His limiting betting quotient for Red should be 0.3.

What this suggests is that this punter considers Yellow more likely than Red, since he’s willing to pay more for a bet on it. So, as a corollary, he should

¹Ellsberg 1961

²It appears in Halpern 2003

³I am studiously avoiding mentioning the “principle of indifference” since I use “indifference” to mean something else in the main text.

prefer a bet on Yellow to a bet on Red. And thus, if offered the chance to bet on Red or to bet on Yellow, for the same stakes, he should prefer the bet on Yellow.

Empirical studies show that many people prefer the bet on Red to the bet on Yellow, but are indeed indifferent between Yellow and Blue.⁴ This behaviour seems to contradict the good classical Bayesian story I have been telling above. And it seems that preferring to bet on Red has some intuitive appeal: you know more about the red marbles; you are more certain of their number.

In the first example, there was uncertainty⁵ about which marble would be drawn. In the second example, as well as that uncertainty, there was *ambiguity* about what the chance set-up was. This is uncertainty of a different kind. It is accommodating this second kind of uncertainty that motivates studies of “imprecise probabilities”. Instead of the standard Bayesian approach of representing uncertainty by a probability measure, the advocate of imprecise probabilities represents uncertainty by a *set* of probability measures. This sort of approach has been explored by, among others,⁶ Isaac Levi,⁷ Peter Walley,⁸ and Joseph Halpern.⁹

The precise probabilist has his belief represented by a probability function, for example $\text{pr}(R) = 0.3, \text{pr}(Y) = \text{pr}(B) = 0.35$ for the “split the difference” probabilist. The imprecise punter has a *set* of probability measures \mathcal{P} representing her belief. $\mathcal{P}(Y)$ is the set of values those probability measures give the event Yellow. For example, if the imprecise probabilist considers possible every possible combination of yellow and blue marbles, her credal state might be characterised as follows: $\mathcal{P}(R) = \{0.3\}, \mathcal{P}(Y) = \mathcal{P}(B) = \left\{ \frac{0}{100}, \frac{1}{100}, \dots, \frac{70}{100} \right\}$.

Some advocates of imprecise probabilities — Levi, for example — insist that the $\mathcal{P}(X)$ should be a convex set. That is, they would demand that $\mathcal{P}(Y) = [0, 0.7]$ the whole interval between the least and the most the probability might be. I don’t subscribe to this view. Consider representing my uncertainty in whether a strongly biased coin will land heads: if I don’t know which way the bias goes then any convex credal state will include a 0.5 chance. But this is exactly the sort of chance I know I can rule out, since I know the coin is biased.¹⁰

One might reason that each pr in \mathcal{P} is equally likely, so using a uniform

⁴Camerer and Weber 1992

⁵Economists are wont to distinguish “risk” and “uncertainty”; the former being where the probabilities are known, the latter where they are unknown. I prefer to use “uncertainty” as a catch-all term for ways one might fail to be certain, reserving “ambiguity” for cases of unknown, or incompletely known probabilities.

⁶It would be futile to try and list all those who have contributed to this area, so I list only those whose work informs the current paper

⁷Levi 1974, 1986

⁸Walley 1991, 2000

⁹Halpern 2003, 2006

¹⁰See §4 of Kyburg and Pittarelli 1992

“second-order probability” I can derive a single probability. This is a more formal version of the “split the difference” intuition. I think it makes an unwarranted assumption about the chance set up when it assumes that each pr is equally likely.

One criticism that has been levelled at this approach is that the imprecise probabilist is vulnerable to a Dutch book, and is therefore irrational. A Dutch book is a set of bets that always lose you money. A punter is vulnerable to the Dutch book if there is a set of bets that she considers acceptable — which she would take — that is a Dutch book. Accepting a set of bets that always lose you money is clearly an irrational thing to do, so avoiding Dutch books is an indicator of rationality.

The plan is to outline exactly how this Dutch book challenge is supposed to go, and show that it is flawed: that one of the premises it rests on is too strong. First I outline the Dutch book argument, making clear its premises and assumptions. Then I argue that one of the conditions on rational preference is too strong in the presence of ambiguity and that therefore the imprecise probabilist is not vulnerable to a Dutch book. This leads on to a discussion of decision-making with imprecise probabilities, and I defend the imprecise view against related criticisms.

2 The Dutch book argument

In this section, I set out a fairly detailed characterisation of the Dutch book theorem. Note that I am concerned only with a *synchronic* Dutch book in this paper. All bets are offered and accepted before any marbles are drawn from the urn. Once there is learning, things become much more tricky. Indeed, learning in the imprecise framework brings with it its own problems.¹¹

Before we can discuss the argument, we need some formal structure. We need a characterisation of formal theories of degree-of-belief, of betting and of preference among bets.

2.1 Formalising degrees of belief

We have an algebra of events: \mathbb{E} . I take \mathbb{E} to be a set of propositions¹² closed under negation, disjunction and conjunction (formalised \neg, \vee, \wedge respectively). One might also take the algebra of events to be a collection of sets of possible worlds closed under complementation, union and intersection.¹³ These are the events the punter is contemplating bets on. Red, Blue and Yellow are the events contemplated in the examples at the beginning. There are two important events: the necessary event and the impossible event. These are formalised as \top and \perp respectively.

¹¹See: Seidenfeld and Wasserman 1993; Wheeler m.s.

¹²Strictly speaking, we need the Lindenbaum algebra of the propositions: we take equivalence classes of logically equivalent propositions...

¹³The two views are more or less equivalent, see [Appendix A](#)

We are interested in functions that represent *degree of belief*. As a first approximation of this idea of modelling degree of belief, consider functions that map events to real numbers. The larger the number, the stronger the belief. Let \mathbf{B} be the set of all functions $\mathbf{b}: \mathbb{E} \rightarrow \mathbb{R}$. The question becomes which subsets of \mathbf{B} are of particular interest? It is typically claimed that the probability functions are the only rational ones.

One class of functions that will be of particular interest are the truth valuations. The function $\omega: \mathbb{E} \rightarrow \{0, 1\}$ is a truth valuation if, for all X, Y :

- $\omega(X \vee Y) = \max\{\omega(X), \omega(Y)\}$
- $\omega(X \wedge Y) = \min\{\omega(X), \omega(Y)\}$
- $\omega(\neg X) = 1 - \omega(X)$

Call the set of functions that satisfy these constraints \mathbf{V} . I will sometimes call ω a “world”, since specifying truth values of all propositions singles out a world. One particular world is actualised, and this determines which bets pay out. So if a red marble is drawn from the urn, the world that has $\omega(R) = 1, \omega(B) = \omega(Y) = 0$ is actualised. And so, for example, $\omega(R \vee B) = \max\{1, 0\} = 1$, as one would expect.

Another class of functions of particular importance are the probability functions. These also map events to real numbers and satisfy the following restrictions for all X, Y :

- $\mathbf{pr}(\top) > \mathbf{pr}(\perp)$
- $\mathbf{pr}(X)$ is in the closed interval bounded by $\mathbf{pr}(\top)$ and $\mathbf{pr}(\perp)$
- $\mathbf{pr}(X \vee Y) + \mathbf{pr}(X \wedge Y) = \mathbf{pr}(X) + \mathbf{pr}(Y)$

What is nice about this non-standard characterisation¹⁴ is that it makes clear that setting the probability of the necessary event to 1 is a matter of convention, not mathematical necessity. The important aspects of probability theory as a model of belief are that the functions are bounded and additive: setting $\mathbf{pr}(\top)$ to 1 gives us the standard probability axioms. Let \mathbf{PR} be the collection of all functions satisfying these constraints. $\mathbf{V} \subset \mathbf{PR} \subset \mathbf{B}$.

But *which* probability measure to take as one’s degree of belief in the Ellsberg case seems underdetermined. The “split the difference” reasoning used by the precise probabilist seems to go beyond his evidence of the situation. I claim that modelling belief by *sets* of probability functions is often better than using a single function. Instead of resorting to “split the difference” reasoning to home in on one probability function to represent my uncertainty, I think it better to represent that ambiguity by the set of probability functions consistent with the evidence.

But why ought the functions in that set be probability functions, rather than any old functions in \mathbf{B} ? Because probability functions are still a kind of

¹⁴Due to Joyce 2009

regulative ideal: the more evidence I accumulate the sharper my imprecise probabilities should become. That is, the more evidence I have, the narrower the range of values my set of probabilities should assign to an event. In the ideal limit, I should like to have a probability function; in the absence of ambiguity I should have a probability function.¹⁵

2.2 Bets and betting

Now we know how we are characterising degree of belief, let's turn to how to represent betting. This framework is due to Joseph Halpern¹⁶. A bet is, for our purposes, an ordered pair of an event in \mathbb{E} and a “betting quotient”. Bets will be ordered pairs of the form (X, α) where $X \in \mathbb{E}$ and $\alpha \in \mathbb{R}$. What is relevant about a bet is the betting quotient and the event in question. The higher the α the punter would accept the more likely she thinks the event in question is. The greater the proportion of the winnings a punter is willing to risk on a bet, the more likely she thinks the event is.

A bet (X, α) pays out £1 if X turns out true by the light of truth valuation ω and pays out nothing otherwise. Or more succinctly, (X, α) pays out $\omega(X)$. The bet costs α and you don't get your stake returned when you win. So the net gain of the bet (X, α) is $\omega(X) - \alpha$. The bet $(\neg X, 1 - \alpha)$ is called the complementary bet to (X, α) . Think of the complementary bet $(\neg X, 1 - \alpha)$ as “selling” the bet (X, α) . Whenever the punter takes a bet (X, α) , the bookie is effectively taking on the complementary bet $(\neg X, 1 - \alpha)$. Table 1 illustrates the “mirror image” quality that the payoffs of complementary bets have.

	$\omega(X) = 1$	$\omega(X) = 0$
(X, α)	$1 - \alpha$	$-\alpha$
$(\neg X, 1 - \alpha)$	$-(1 - \alpha)$	α

Table 1: Payoffs for a bet and its complement

A set of bets $B = \{(X_i, \alpha_i)\}$ costs $\sum \alpha_i$ and pays out $\sum \omega(X_i)$ in world ω . That is, you get 1 for every event that you bet on that ω makes true. So the value of a set of bets B at world ω is $\tau_\omega(B) = \sum (\omega(X_i) - \alpha_i)$. The value of the bet at a world is how much it pays out minus what the bet cost.

For set of bets B let its complement¹⁷ be $B^C = \{(\neg X_i, 1 - \alpha_i)\}$. It is easy to show that $\tau_\omega(B^C) = -\tau_\omega(B)$. The “mirror image” quality of Table 1 also holds for sets of bets.

A Dutch book in this context is a set of bets, B such that, for every $\omega \in \mathbf{V}$, we have $\tau_\omega(B) < 0$. That is, the pay out for the bet is negative *however the world turns out*.

¹⁵For another argument in favour of probabilities as epistemically privileged, see: Joyce 1998

¹⁶Döring 2000

¹⁷This is a lazy way of talking, B^C is not the complement of B in the sense of set-theoretic complement in the set of bets, but rather the set of bets complementary to those in B .

2.3 Constraints on rational betting preference

We are interested in preference among bets, so define a relation “ $A \succeq B$ ” which is interpreted as meaning “ A is at least as good as B ”, where A and B are bets. We will later put constraints on what preferences are reasonable. As I said above, what is of particular interest is the punter’s maximum willingness to bet. For an event X , define α_X by $\sup\{\alpha : (X, \alpha) \succeq (\neg X, 1 - \alpha)\}$. These maximum betting quotients are interpreted as characterising the punter’s belief state and it will often be useful to talk about the belief function corresponding to these α_X s. Define $\mathbf{q}(X) := \alpha_X$. You are vulnerable to a Dutch book unless your $\mathbf{q} \in \mathbf{PR}$. That is, unless your (limiting) betting quotients have the structure of a probability function, there is a set of bets — acceptable by the lights of your \mathbf{q} — that guarantees you a loss of money.

Halpern sets out four constraints on what sort of preferences it is rational to have among bets. These are sufficient to force any agent satisfying them to have betting quotients that have the structure of a probability measure. That is, failing to satisfy the axioms of probability makes your betting preferences incompatible with the constraints. In the strict Bayesian picture, \mathbf{q} and degree of belief \mathbf{pr} are used more or less interchangeably. It will be important in what follows that one’s willingness to bet and one’s degrees of belief are distinct, if related concepts.

Strictly speaking, the preference is among sets of bets, so when discussing single bets I should say “ $\{(X, \alpha)\} \succeq \{(Y, \beta)\}$ ”, but the preference relation induces an obvious relation among singletons, so I omit the braces. I don’t make much of a distinction in what follows between a bet and a set of bets.

The first of Halpern’s requirements says that if one bet B_1 always pays out more money than another B_2 , then you should prefer B_1 .

If, for all $\omega \in \mathbf{V}$ we have $\tau_\omega(B_1) \geq \tau_\omega(B_2)$ then: $B_1 \succeq B_2$ (DOMINANCE)

Note that this condition does not force the punter to prefer bets with higher *expected* value: only to prefer bets with a higher *guaranteed* value. Preferring bets guaranteed to give you more money seems eminently reasonable.

The second of Halpern’s conditions is simply that the preference relation be transitive.

If $B_1 \succeq B_2$ and $B_2 \succeq B_3$ then $B_1 \succeq B_3$ (TRANSITIVITY)

Again, this condition seems reasonable.

The third of Halpern’s conditions — the one I will take issue with in [section 3](#) — is **COMPLEMENTARITY**.

For all $X \in \mathbb{E}$ and $\alpha \in \mathbb{R}$ either,
 $(X, \alpha) \succeq (\neg X, 1 - \alpha)$ or $(\neg X, 1 - \alpha) \succeq (X, \alpha)$ (COMPLEMENTARITY)

Note that this is weaker than what is often assumed of rational preference: **COMPLEMENTARITY** does not require that the punter’s preference relation be

complete or total.¹⁸ It need only be complete with respect to complementary bets, but this is still too much for me. I will discuss why I find this too strong a condition in the next section. Note that this condition is specified in terms of single bets, but in the presence of **PACKAGE** below, it extends somewhat to sets.

The final condition, sometimes known as the “package principle”¹⁹ is, as Halpern²⁰ puts it, that “preferences are determined pointwise”.

$$\begin{aligned} \text{If } (X_i, \alpha_i) \succeq (Y_i, \beta_i) \text{ for each } 1 \leq i \leq n \text{ then:} & \\ \{(X_i, \alpha_i)\} \succeq \{(Y_i, \beta_i)\} & \quad (\text{PACKAGE}) \end{aligned}$$

Note this is quite a restricted principle. For example, it does not in general allow that if $A \succeq B$ and $C \succeq D$ then $A \cup C \succeq B \cup D$.

The Dutch book theorem says that if a punter’s preference among bets satisfies **DOMINANCE**, **TRANSITIVITY**, **COMPLEMENTARITY** and **PACKAGE**, then that punter’s betting quotients \mathbf{q} will have the structure of a probability function. Or to put it another way, if the punter’s betting quotients violate the axioms of probability, then this leads to a preference incompatible with the above conditions.

3 Ambiguity and complementarity

We have seen what is necessary in order to prove the Dutch book theorem (the proof itself is in **Appendix A**). In this section I argue that one particular premise of the theorem — **COMPLEMENTARITY** — is too strong. It is not warranted in the case of ambiguity.

I use preferring a bet to its complement as a proxy for acceptance of a bet. “This seems unintuitive,” one might say, “I prefer lukewarm coffee to cold coffee, but I wouldn’t accept either of them by choice.” But remember, we are dealing with preference for one bet *over its complement*. So the analogous example would be something like “preferring lukewarm coffee to no lukewarm coffee”, and here it seems that *that* preference is tantamount to accepting lukewarm coffee.

So why is **COMPLEMENTARITY** unreasonable? First let’s see why it does seem reasonable in the first example of **section 1**. Recall that there we had 100 marbles in an urn, 30 red and the remainder yellow. The punter’s maximum betting quotient on red, $\mathbf{q}(R) = 0.3$. That is, 0.3 is the largest value for which a bet on red is preferred to a bet against red. You have confidence that Red will come up about 30% of the time, and since if it’s not Red, it’s Yellow, $\mathbf{q}(Y) = 0.7$.

Compare this to the Ellsberg case, where instead of the remainder being yellow, the remainder are yellow and blue in some unknown proportion.

¹⁸A relation R is total or complete when xRy or yRx for all x, y

¹⁹See e.g. Schick 1986

²⁰Halpern 2003, p. 22

The probabilist splits the difference²¹ and sets his betting quotients as follows: $q(R) = 0.3, q(Y) = q(B) = 0.35$. The imprecise probabilist claims we can act differently. $q(R) = 0.3$ still seems acceptable, our evidence about Red hasn't changed. But it seems that "ambiguity aversion" about evidence for Yellow and Blue suggests that the punter's maximal betting quotients for each should be lower. Perhaps even $q(Y) = q(B) = 0$, since if you entertain the possibility that the chance of Yellow might be 0, then you should refuse to buy bets on Yellow. I focus on this extreme case. But first, two caveats. q is not a representation of belief. It is a representation of willingness to bet. Part of the conceptual baggage of precise probabilism that the imprecise probabilist needs to get away from is too close a connection between willingness to bet and belief. Obviously they are related, but not as tightly as they are in the precise framework. The belief is represented by the set of probabilities \mathcal{P} . Also, I am not endorsing as rational the extreme view ($q(Y) = q(B) = 0$), merely using it for illustrative purposes.

Say $q(Y) = q(B) = 0$. Then COMPLEMENTARITY demands that $q(\neg Y) = 1$. Or, in other words, if $\alpha_Y = 0$ COMPLEMENTARITY demands that $(\neg Y, 1-0.1)$ is preferred to its complement. Similarly for B . This bet should also be acceptable: $(\neg R, 1-0.4)$. Together, these bets form a Dutch book (see Table 2). Call this set D .

	R	B	Y
$(\neg R, 1-0.4)$	-0.6	1-0.6	1-0.6
$(\neg B, 1-0.1)$	1-0.9	-0.9	1-0.9
$(\neg Y, 1-0.1)$	1-0.9	1-0.9	-0.9
Total	-0.4	-0.4	-0.4

Table 2: Dutch booking an imprecise probabilist

But to demand that the imprecise probabilist conform to COMPLEMENTARITY (and therefore, accept these bets) is to misunderstand the nature of the uncertainty being encoded in $q(Y) = 0$. If there is ambiguity — uncertainty about the chance set-up itself — low confidence in an event does not translate into high confidence in its negation. There is an important distinction between the balance of evidence and the weight of evidence:²² how conclusively the evidence tells in favour of a proposition (balance) versus how much evidence there is for the conclusion (weight). COMPLEMENTARITY assumes that the unwillingness to bet on an event is due to the balance of evidence telling against it and if this is the case then it is a reasonable condition. If on the other hand the refusal to accept the bet is due to the lack of weight of evidence, then the condition is not reasonable. Because of the

²¹ Of course, a subjectivist could set any particular probabilistically coherent value to the events, but what is objectionable to the imprecise probabilist is the suggestion that $q(B) = 0.7 - q(Y)$

²²Joyce 2005

ambiguous nature of the chance set-up (the lack of weight of evidence), the punter is simply not willing to bet either way most of the time. So the imprecise probabilist will not want to conform to **COMPLEMENTARITY** and therefore, will not be subject to a Dutch book in this way. In **Appendix A**, I show that in the absence of **COMPLEMENTARITY** it is reasonable to have betting quotients satisfying restrictions weaker than those demanded of probabilities.

One might still argue that if the punter were forced to choose one side or the other of any given bet (X, α) — that is, if the punter were forced to obey **COMPLEMENTARITY** — then she would obey the probability axioms or be subject to a Dutch book. This is true, but I don't see how this procedure elicits a fair reflection of the punter's credal state. If I forced the punter to take bets where the α s were all 1s or 0s, the resulting betting quotients would be a valuation function or the punter would be open to a Dutch book: that does not mean that the punter's degrees of belief are truth valuations. The punter's actions only reflect the punter's belief when her actions are not too restricted. So the conditions on betting preference have to be independently reasonable for them to form the basis of a Dutch book argument. **COMPLEMENTARITY** is not independently reasonable unless the chance set-up is unambiguous.

4 Decision with imprecision

To further explore this issue, we need to say something about what decision theory looks like from the imprecise probabilities perspective. The precise probabilist acts in accordance with the rule “maximise expected value with respect to \mathbf{pr} ”; where \mathbf{pr} is the precise punter's probability. Let's say the expectation of the set of bets B is $E(B)$. Restricted to the case of choosing whether to buy a set of bets B or its complement B^C , this amounts to accepting B if $E(B) = \sum_i (\mathbf{pr}(X_i) - \alpha_i) > 0$ and taking B^C otherwise. It's easy to show that $E(B^C) = -E(B)$ so one and only one bet has positive value in any pair of complementary bets, unless both bets have 0 expected value. So the precise probabilist always prefers one bet to the other, unless he is indifferent between them. This just follows from his being opinionated.

What about the imprecise probabilist? How is she to decide? I don't intend to suggest a fully worked out decision theory for imprecise probabilities, but simply offer enough of a sketch to explain how the imprecise probabilist avoids Dutch books. So let's turn the discussion in the last section on its head and start from a punter's belief state and derive what decisions that punter would make.

Recall the imprecise punter's credal state \mathcal{P} is a set of probability measures, $\mathcal{P}(Y)$ is the set of values assigned to Y by members of \mathcal{P} . This $\mathcal{P}(Y)$ is already a “summary statistic” in a sense: it doesn't make clear that for any $\mathbf{pr} \in \mathcal{P}$ whenever $\mathbf{pr}(Y)$ is high, $\mathbf{pr}(B)$ is low and vice versa. So it is \mathcal{P} that represents the punter's belief state, and $\mathcal{P}(Y)$ and so on are only short-hands, summaries that miss out some information. This is an important

point, and one that is not often made.

We can define expectation of set of bets $B = \{(X_i, \alpha_i)\}$ for the imprecise probabilist as follows:

$$\mathcal{E}(B) = \left\{ \sum \mathbf{pr}(X_i) - \alpha_i : \mathbf{pr} \in \mathcal{P} \right\}$$

This is the set of expected values of the bet with respect to the set of probabilities in \mathcal{P} . So the idea of maximising expected value isn't well defined for the imprecise probabilist.

Another standard summary statistic for imprecise probabilities are the lower and upper envelopes:

- $\underline{\mathcal{P}}(X) = \inf\{\mathbf{pr}(X) : \mathbf{pr} \in \mathcal{P}\}$
- $\overline{\mathcal{P}}(X) = \sup\{\mathbf{pr}(X) : \mathbf{pr} \in \mathcal{P}\}$

Likewise we can define $\underline{\mathcal{E}}$ and $\overline{\mathcal{E}}$ to be the infimum and supremum of \mathcal{E} respectively.

These summary statistics have their own interesting formal properties.²³ For example, $\underline{\mathcal{P}}$ is superadditive. That is, for incompatible X, Y we have $\underline{\mathcal{P}}(X \vee Y) \geq \underline{\mathcal{P}}(X) + \underline{\mathcal{P}}(Y)$.

Again, these are summarising some aspect of the punter's credal state, but they are misrepresentative in other ways: considering the upper and lower envelopes to represent the agent's credal state is a mistake.²⁴

So, how should a punter bet? If $\underline{\mathcal{E}}(B) > 0$ then B looks like a good bet: every probability measure in \mathcal{P} thinks that this bet has positive expected value. Likewise, if $\overline{\mathcal{E}}(B) < 0$ then B^c looks promising. So any decision rule for imprecise probabilities should take into account these two intuitions. But this still leaves a lot of room for manoeuvre: what should the punter do when $\underline{\mathcal{E}}(B) < 0 < \overline{\mathcal{E}}(B)$?

The more general question of what decision rule the imprecise probabilist should use is left open. She could maximise $\underline{\mathcal{E}}$, she could use Levi's rule,²⁵ she could use the Hurwicz criterion²⁶... I leave this bigger problem unanswered for now. I don't need to offer a fully worked out decision rule to show how the imprecise probabilist can avoid the Dutch book.

As a first approximation, let's imagine an extreme case of the ambiguity averse imprecise probabilist. She refuses to take either side of any bet (X, α) if $\underline{\mathcal{E}}(X, \alpha) < 0 < \overline{\mathcal{E}}(X, \alpha)$. That is, she has no preference between (X, α) and its complement if α is in the interval $[\underline{\mathcal{P}}(X), \overline{\mathcal{P}}(X)]$. And she will obey the two concerns discussed above: take bets on X if α is low enough, bet against X — bet on $(\neg X, 1 - \alpha)$ — if α is big enough. This punter is disobeying **COMPLEMENTARITY**, but can act in accordance with the other three conditions. She is disobeying **COMPLEMENTARITY** because there are values of α for which she has no preference between the complementary bets: she would accept neither.

²³Halpern 2003; Paris 1994

²⁴Joyce offers nice examples of this in: Joyce 2011

²⁵Levi 1986

²⁶Hurwicz 1951

A punter who obeys the three other conditions but not **COMPLEMENTARITY** has her limiting betting quotients have the structure of a lower envelope like $\underline{\mathcal{P}}$. I prove this in [Appendix A](#). This is not to say that $\underline{\mathcal{P}}$ represents that punter’s beliefs, but rather that this particular elicitation procedure (betting) can only give us so much information about the punter’s credal state.

However, one might object that this maximally ambiguity-averse punter is still irrational in not accepting a collection of bets that guarantees her a positive value whatever happens: she will not accept what Alan Hájek calls a “Czech book”.²⁷ This is related to the idea in economics and finance of “arbitrage”. Consider the set of bets: $C = \{(Y, 0.1), (B, 0.1), (R, 0.3)\}$ in the Ellsberg example. Whatever colour marble is picked out of the urn, C makes you a profit of 0.5, so it would be crazy not to take it! However, our imprecise probabilist punter will not want to accept either of the first two bets, since $\underline{\mathcal{P}}(Y) < 0.1 < \overline{\mathcal{P}}(Y)$ and similarly for Blue. So does this mean she will refuse a set of bets guaranteed to make her money? Isn’t this just as irrational as accepting a collection of bets that always lose you money?

Let’s say this punter still conforms to **DOMINANCE** which guarantees that she will prefer bets that always pay more. This condition means that the imprecise punter will still accept the Czech book, even if every stake is between $\underline{\mathcal{P}}$ and $\overline{\mathcal{P}}$. This is because $\tau_\omega(C) = 0.5 > 0$ for any ω , so this set of bets is preferred to its complement by **DOMINANCE**. So if we take **DOMINANCE** as a necessary condition on imprecise decision rules, the imprecise punter can accept Czech books. Perhaps the other two conditions from [section 2.3](#) that I accept should also be taken into account when thinking about imprecise decision rules.

This brings home the point that the lower and upper envelopes are not all there is to an imprecise probabilist’s belief state: every $\mathbf{pr} \in \mathcal{P}$ assigns positive expected value to C (indeed, positive guaranteed value) so whatever the chance set up, this is a good bet. But if the punter were acting just in accordance with her upper and lower probabilities, this fact might get lost in the summarising. So again we see that just focusing on the spread of values misses out important information about the punter’s credal state.

The imprecise decision problem has been discussed extensively and many solutions have been proposed. Brian Weatherson²⁸ reviews some solutions, and argues that many of them fail. Indeed, he argues that no decision rule which relies solely on $\underline{\mathcal{E}}$ and $\overline{\mathcal{E}}$ can ever be plausible. The whole project of imprecise probabilities has been criticised on the grounds that it cannot offer *any* adequate decision rules.²⁹ There is obviously a lot more to be said on the subject of decision making in the imprecise framework.

Avoiding Dutch books is taken to be necessary for being rational. So being vulnerable to Dutch books is sufficient for being irrational. I have shown

²⁷Hájek 2008

²⁸Weatherson m.s.

²⁹Elga 2010

that the imprecise probabilist isn't as vulnerable to Dutch books as is sometimes suggested. So, I claim, this particular avenue for criticising imprecise models isn't available. I have also suggested that more work needs to be done in exploring how best to make imprecise decisions. I think the increased expressive power of imprecise frameworks, and the fact that we *do* have to deal with several kinds of uncertainty means that imprecise probabilities are a worthwhile area for research.

A Proof

Recall that α_X is defined as the most that a punter will buy a bet on X for. The Dutch book theorem claims that the function $\mathbf{q}(X) = \alpha_X$ is a probability function. What I prove here is that $\mathbf{q}(X) = \alpha_X$ is a lower probability³⁰ without using **COMPLEMENTARITY**. That is, I show that \mathbf{q} acts like one of the summary statistics of the “sets of probabilities” approach I favour. Define $\beta_X = \inf\{\beta : (\neg X, 1-\beta) \succeq (X, \beta)\}$. I also show that $\bar{\mathbf{q}}(X) = \beta_X$ is an upper probability.

If $X \models Y$ then $(Y, \alpha) \succeq (X, \alpha)$ and $(\neg X, 1-\alpha) \succeq (\neg Y, 1-\alpha)$. This follows from **DOMINANCE** and the fact that if $X \models Y$ then $\neg Y \models \neg X$. Remember, $(X, \alpha_X) \succeq (\neg X, 1-\alpha_X)$ by definition, so by the above result and **TRANSITIVITY**, if $\gamma \leq \alpha_X$ then $(X, \gamma) \succeq (\neg X, 1-\gamma)$. This is in fact an “if and only if” result, since α_X is *defined* as the largest value for which this preference holds. A similar result holds for β_X . The above results allow us to show that if $X \models Y$ then $\alpha_X \leq \alpha_Y$ and $\beta_X \leq \beta_Y$.

α_X and β_X are related to each other. $(\neg\neg X, 1-\beta_{\neg X}) \succeq (\neg X, \beta_{\neg X})$ by definition. Since $\neg\neg X = X$ it follows that $1-\beta_{\neg X} \leq \alpha_X$. We also have $1-(1-\alpha_X) = \alpha_X$, so by definition $(\neg\neg X, 1-(1-\alpha_X)) \succeq (\neg X, 1-\alpha_X)$ so $1-\alpha_X \geq \beta_{\neg X}$. Thus $1-\beta_X \geq \alpha_X$. These two inequalities together imply that $\alpha_X = 1-\beta_{\neg X}$.

The set of bets $\{(X, \alpha_X), (\neg X, 1-\beta_X)\}$ is preferred to its complement by **PACKAGE**. This set always pays out 1, and costs $\alpha_X + 1-\beta_X$. So the net gain of this bet is always $\beta_X - \alpha_X$. If $\alpha_X > \beta_X$, the net gain would be negative, so the net gain of the complementary bets would be positive. So the preference for this set over its complement would contradict **DOMINANCE**. So for any X we know that $\alpha_X \leq \beta_X$.

For logically incompatible propositions, X, Y consider the bet $B = (X \vee Y, \alpha_X + \alpha_Y)$. Now compare this with $C = \{(X, \alpha_X), (Y, \alpha_Y)\}$. These bets always have the same payout. So by **DOMINANCE**, we know that $B \succeq C$. We also know that $C^c \succeq B^c$ for the same reason. Now, $C \succeq C^c$ by **PACKAGE**. $B \succeq C \succeq C^c \succeq B^c$ so, by **TRANSITIVITY** we know that $B \succeq B^c$. By definition, $\alpha_{X \vee Y}$ is the maximum value for which $B \succeq B^c$. So $\alpha_{X \vee Y} \geq \alpha_X + \alpha_Y$. A similar chain of reasoning leads to the conclusion that $\beta_{X \vee Y} \leq \beta_X + \beta_Y$.

This demonstrates that $\mathbf{q}(X) = \alpha_X$ is superadditive and $\bar{\mathbf{q}}(X) = \beta_X$ is subadditive, as is characteristic of lower and upper probabilities. This proof makes no use of the **COMPLEMENTARITY** condition. However, to make the con-

³⁰In the terminology of Cozman n.d.

nection with § 7.1 of Cozman’s characterisation of lower probabilities, we also need the assumption that $\alpha_{\perp} = 0$ and $\alpha_{\top} = 1$. This isn’t needed once **COMPLEMENTARITY** is in place. Using this condition it is easy to show that $\alpha_X = \beta_X$ for all X and thus that $\underline{\mathbf{q}}(X) = \bar{\mathbf{q}}(X) = \mathbf{q}(X)$ and that this function is additive, non-negative and normalised: a probability measure.

One can, in fact, construct a mass function out of the α_X s and show that \mathbf{q} is a Dempster-Shafer belief function (this again without using **COMPLEMENTARITY**).

The last part is easier to do in terms of sets of worlds, rather than in terms of propositions, so I sketch the translation between the two paradigms here. Define $[X]$ as the set of valuations that make X true: $[X] = \{\omega \in \mathbf{V} : \omega(X) = 1\}$. $[X] \subseteq [Y]$ if and only if $X \models Y$, so we have an structure preserving bijection between propositions and sets of “worlds”. I drop the square brackets in what follows.

Define a mass function as follows:

$$m(X) = \alpha_X - \sum_{Y \subsetneq X} m(Y)$$

That is, $m(X)$ picks up all the mass not assigned to subsets of X . This should, strictly speaking be an inductive definition on the size of X , but I take it that it is obvious what is meant here. If we now consider the following equation:

$$\underline{\mathbf{q}}(X) = \sum_{Y \subseteq X} m(Y)$$

This is equivalent to the above characterisation of $\underline{\mathbf{q}}(X) = \alpha_X$. That \mathbf{q} has this mass function associated with it means that it is a Dempster-Shafer belief function. This is a particular kind of lower probability: it is an infinite-monotone lower probability.

Acknowledgements

Thanks to Jonny Blamey, Chris Clarke, Erik Curiel and Katie Steele for comments on a previous version of this paper and to Roman Frigg and Richard Bradley for comments on this version. Thanks also to participants at the following conferences where versions of this talk were presented: Philosophy of Probability II conference at LSE 8–9 June 2009; DGL Paris 9–11 June 2010; “Pluralism in Foundations of Statistics” ESF workshop Canterbury 9–10 September 2010.

References

Camerer, C. and M. Weber (1992). “Recent Developments in Modeling Preferences: Uncertainty and Ambiguity”. In: *Journal of Risk and Uncertainty* 5, pp. 325–370.

- Cozman, F. (n.d.). *A Brief Introduction to the Theory of Sets of Probability Measures*. <http://www.poli.usp.br/p/fabio.cozman/Research/CredalSetsTutorial/quasi-bayesian.html>.
- Döring, F. (2000). “Conditional Probability and Dutch Books”. In: *Philosophy of Science* 67, pp. 391–409.
- Elga, A. (2010). “Subjective Probabilities should be Sharp”. In: *Philosophers’ Imprint* 10.
- Ellsberg, D. (1961). “Risk, ambiguity and the Savage axioms”. In: *Quarterly Journal of Economics* 75, pp. 643–696.
- Hájek, A. (2008). “Arguments for—or against—probabilism?” In: *British Journal for the Philosophy of Science* 59, pp. 793–819.
- Halpern, J. Y. (2003). *Reasoning about uncertainty*. MIT press.
- (2006). *Using sets of probability measures to represent uncertainty*. [arXiv: cs/0608028v1](https://arxiv.org/abs/cs/0608028v1).
- Hurwicz, L. (1951). *Optimality Criteria for Decision Making under Ignorance*. Cowles Commission Discussion Paper Statistics 370.
- Joyce, J. M. (1998). “A Nonpragmatic Vindication of Probabilism”. In: *Philosophy of Science* 65, pp. 575–603.
- (2005). “How Probabilities Reflect Evidence”. In: *Philosophical Perspectives* 19, pp. 153–178.
- (2009). “Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Belief”. In: *Degrees of Belief*. Ed. by F. Huber and C. Schmidt-Petri. Springer, pp. 263–297.
- (2011). “A Defense of Imprecise Credence”. In: *Oxford Studies in Epistemology* 4. forthcoming.
- Kyburg, H. and M. Pittarelli (1992). *Set-based Bayesianism*. Tech. rep. UR CSD;TR407. <http://hdl.handle.net/1802/765>. University of Rochester, Computer Science Department.
- Levi, I. (1974). “On Indeterminate probabilities”. In: *Journal of Philosophy* 71, pp. 391–418.
- (1986). *Hard choices: decision making under unresolved conflict*. Cambridge University Press.
- Paris, J. (1994). *The uncertain reasoner’s companion*. Cambridge University Press.
- Schick, F. (1986). “Dutch Bookies and Money Pumps”. In: *Journal of Philosophy* 83, pp. 112–119.
- Seidenfeld, T. and L. Wasserman (1993). “Dilation for sets of probabilities”. In: *Annals of Statistics* 21, pp. 1139–1154.
- Walley, P. (1991). *Statistical Reasoning with Imprecise Probabilities*. Vol. 42. Monographs on Statistics and Applied Probability. Chapman and Hall.
- (2000). “Towards a unified theory of imprecise probabilities”. In: *International Journal of Approximate Reasoning* 24, pp. 125–148.
- Weatherson, B. (m.s.). “Decision Making with Imprecise Probabilities”. Available <http://brian.weatherson.org/vdt.pdf>.
- Wheeler, G. (m.s.). *Demystifying Dilation*.