

Pragmatic Considerations on Comparative Probability¹

Thomas F. Icard, III

Imagine an explorer searching for a buried treasure among a trio of remote islands A , B , and C . The explorer reasons as follows. The treasure is more likely to be on island B than on C , because B is much larger than C , so there is more land area where it could be buried. For the same reason it is more likely to be on A than on B . At the same time, on the explorer's map there are markings of what appear to be Xs on both B and C , and the marking on C is especially pronounced. Since there does not seem to be an X on A , the explorer judges that this makes C more likely than A , making up for the difference in size, though this is not enough to overturn the respective judgments about A and B , or about B and C . The explorer's judgments about relative likelihoods in this scenario are therefore intransitive: $A \succ B \succ C \succ A$.

What, if anything, is wrong with such a set of judgments? The explorer is able to provide some explanation for why each one individually seems reasonable, and it is perhaps not patently obvious that any kind of inconsistency or incoherence arises from holding them together. Yet transitivity, or at least non-intransitivity, is nearly always taken as axiomatic in the study of uncertain reasoning. What exactly is the rationale for this assumption? In what sense is intransitivity defective, in that any reasonable person would want to avoid it, all else being equal?

One rather obvious remark is that, given intransitive judgments, there can be no numerical probability measure P that agrees with these judgments, in the sense that $P(E) > P(F)$ whenever E is taken to be more likely than F , and $P(E) = P(F)$ whenever E and F are taken to be equally likely. In particular, these judgments could not be derived from any probability measure in this way. Intransitivity is just one example of a property of orderings that prevents probabilistic agreement. Call an ordering, or set of comparative judgments, that admits of an agreeing numerical probability measure *representable*. We might therefore ask more generally, what, if anything, is wrong with a set of comparative probability judgments that is not representable? Is there any desirable feature—e.g., some notion of coherence or reasonableness—of representable orderings that is lacking in non-representable orderings? Put otherwise, is there

¹Draft of September 9, 2014. Acknowledgements to be added.

any kind of wide scope¹ rational requirement that one’s comparative probability judgments be probabilistically representable?

The aim of this paper is to demonstrate that a person will forestall a blatant kind of pragmatic defect just in case that person’s comparative probability judgments are, in totality, representable by a standard, finitely additive probability measure. In the background is a theorem, the details of which appear in the technical Appendix, showing that in a class of “canonical” situations an agent avoids strict dominance, in the decision theoretic sense, if and only if that agent’s judgments are probabilistically representable. The theorem is in fact a systematization and generalization of a variation on the well known Money Pump argument, which can be thought of as a sort of dramatization of the main dominance result for the special case of intransitivity. The majority of the paper is dedicated to clarifying this main result, justifying the assumptions behind it, and above all, explaining why this makes for a compelling case against intransitivity, and indeed any feature of orderings that prevents probabilistic agreement.

A number of important figures in the history of probability—including Keynes, de Finetti, Koopman, and others²—have taken comparative probability judgments to be more basic and intuitive than numerical probability judgments, partly motivating a series of attempts to characterize representability axiomatically in terms of properties that can simply be read off an ordering. This work resulted in a number of representation theorems using relatively technical axioms,³ whose intuitive justification many have found questionable.⁴ While we will be appealing directly only to the general semantic notion of representability, our work can also be seen as offering justification for the axioms proposed in this literature, in so far as they are shown to be extensionally equivalent to representability. Indeed, while probabilistic representability has typically been taken as desirable in its own right—and in turn, there are well known arguments for why explicitly numerical judgments ought to satisfy the standard probability axioms⁵—the work presented here appears to be the first direct argument for why a person violating any one of these qualitative axioms is rationally criticizable.⁶

After discussing several concrete examples of non-representable sets of judgments, as well as some of the historical precedents in the literature anticipating

¹See, e.g., Broome (2007). A requirement of this sort is wide scope in the sense that it concerns an entire set of judgments taken in totality, and conformity with the requirement does not require a person to make or to refrain from making any particular comparative judgment.

²Keynes (1921), de Finetti (1937), and Koopman (1940), *inter alia*.

³E.g., Kraft et al. (1959) and Scott (1964). For overviews of this literature, see Krantz et al. (1971, §5,§9), and Fishburn (1986).

⁴For instance Fine (1973, p. 23), explicitly denies that such axioms have a rational basis.

⁵Viz. Dutch book arguments (Ramsey, 1926; de Finetti, 1937), purely epistemic arguments (Maher, 1993; Joyce, 1998), axiomatic derivations (Jaynes, 2003), and others.

⁶However, see Fitelson and McCarthy (2014) for an insightful discussion of comparative probability from the perspective of alethic and evidential requirements, in the tradition of Joyce (1998) on “accuracy based” norms. See also Fishburn (1986, pp. 337-338), who makes an off-hand comment pointing in the direction of our result. Cf. Fn 17 below.

the current work, we will have set the stage for outlining the background assumptions and general setup of the theorem, followed by responses to some potential criticisms and objections. The proof of the main result is given in the Appendix.

THE MONEY PUMP ARGUMENT

The argument in this paper draws inspiration from a well known argument for the irrationality of intransitivity—the Money Pump argument, made famous by Davidson, McKinsey, and Suppes (1955)—which was first given in the context of value-based preferences between outcomes, or states of affairs, rather than comparative probability judgments. The framework can be easily adapted to our setting, however. To take the original example from Davidson et al., consider the following three possible outcomes for a hypothetical assistant professor named S :

- (a) A full professorship for S and a salary of \$5,000.
- (b) An associate professorship and a salary of \$5,500.
- (c) An assistant professorship and a salary of \$6,000.

S is assumed to maintain the following intransitive preferences:

$$(a) \succ (b) \succ (c) \succ (a).$$

The story then begins with a clever department head H who reasons that if S prefers outcome (b) to (c) , then that preference must be worth some monetary amount to S . With this in mind H agrees to give S the associate professorship and \$500 salary reduction for some appropriate monetary compensation. Upon realizing that S also prefers (a) to (b) , H offers to allow S the full professorship and a further salary reduction, again for some small compensation. Recognizing in turn S 's preference for (c) over (a) , H agrees to offer (c) instead for some fee, which puts S back in the same initial position, only poorer. Continuing on in this way, H has turned S into a “money pump” by capitalizing on S 's intransitive, indeed cyclic, preferences.

While Davidson et al. intended the Money Pump scenario mainly as a dramatization of the observation that anyone with intransitive preferences will necessarily be incapable of choosing an outcome to which none other is preferred,⁷ the particular scenario has generated a sizable literature in philosophy and economics. Much of the philosophical discussion has focused on the allegedly diachronic nature of the scenario. For instance, both Schick and McClennen have argued that, whereas someone with a preference for (c) over (a) would perhaps be rationally compelled to pay something to receive (c) in place of (a) , all else being equal, this does not necessarily extend to a situation in which other exchanges have already taken place.⁸ There may be good practical reason for S to stop accepting the

⁷See Davidson et al. (1955): p. 145. This way of putting it is from Gustafsson (2013).

⁸Schick (1986) and McClennen (1990, §5.4). The differences between their arguments and criticisms will not be important here. Also note that we do not intend to respond in detail to all of the challenges identified by these authors. Cf. Gustafsson (2013), *inter alia*.

exchanges at some point, even if S 's pairwise preferences do not change.

Before addressing this general concern, we might first ask whether an argument analogous to the Money Pump can be made for intransitivity of comparative probability judgments, our main topic. Consider again our explorer searching islands A , B , and C , and imagine that the three islands are situated in the shape of a triangle. Our explorer is approaching island C with B slightly in the distance, and reasons that, because B is more likely to harbor the treasure than C , it is perhaps worth traveling the extra distance to B and exploring there first. However, upon approaching B , the explorer catches a glimpse of A in the distance. Recalling that A is much larger than B , and therefore more likely to harbor the treasure, the explorer decides to take the extended trek to island A . Upon reaching A , with C again visible on the horizon, the explorer ends up back at C , having merely wasted several hours, and no closer to finding the treasure. Thus, the style of argument does seem to carry over.

It will be useful to isolate several assumptions underlying this argument pattern, which will form the core of what is to follow. The first assumption provides a fundamental link between comparative probability judgments and choice. Suppose we have two outcomes o_{good} and o_{bad} , such that o_{good} is strictly preferred to o_{bad} . Suppose moreover we have two “gambles” that depend on events A and B :

GAMBLE 1 — Outcome o_{good} if A obtains, and o_{bad} if A fails to obtain.

GAMBLE 2 — Outcome o_{good} if B obtains, and o_{bad} if B fails to obtain.

For instance, in our explorer scenario, o_{good} would correspond to finding the treasure, o_{bad} to not finding the treasure; A is the event in which the treasure is on island A , while B is the event in which the treasure is on island B . In adapting the Money Pump to the comparative probability setting, we essentially assumed:

ASSUMPTION 1: Whenever event A is judged strictly more likely than event B , GAMBLE 1 will be strictly preferred to GAMBLE 2.

That is, one will prefer a favorable gamble on A to one on B whenever A is judged more probable. This assumption obviously applies to literal gambling situations, but as the explorer scenario makes clear, it also applies to arbitrary choices under uncertainty, in which a favorable outcome is directly linked to which of two propositions about the world obtains. When deciding whether to search island A or B , the explorer is essentially choosing between GAMBLES 1 and 2. Thus the word ‘gamble’ should not be understood too literally. The content of ASSUMPTION 1 is that this choice ought to reflect the comparative probability judgments in the obvious way. This assumption is arguably bedrock.⁹

⁹ASSUMPTION 1 could perhaps even be offered as an operational definition of comparative probability judgments. For instance, in the other direction—from preferences between gambles or acts to comparative probability—Krantz et al. (1971, §5.2.4) extract principles of comparative probability from Savage’s (1954) axioms on gambles, treating ASSUMPTION 1 as giving a definition of the agent’s comparative probability judgments.

A less obvious but equally important assumption behind Money Pump style arguments is that the preference in question ought to be worth something. In our setting, this means a person ought to be willing to give something up in order to receive GAMBLE 1 over GAMBLE 2. This might involve monetary payment as in the original story from Davidson et al., but it might also involve other valuable resources such as time, effort, energy, etc., as in the explorer scenario. The cost of exploring island A instead of B consists in the additional time and effort required to travel there, taking away from the time spent searching for the treasure.

ASSUMPTION 2: One ought to be willing to give something up, however small, to guarantee GAMBLE 1 in place of GAMBLE 2.

ASSUMPTION 2 is perhaps not unassailable, and we will discuss it further below. For now, let us assume that this is also simply implicit in what it means to have a judgment that A is more likely than B . It is important that this cost can be measured along the same scale as the utility of outcomes o_{good} and o_{bad} .¹⁰

A third, apparently additional, assumption—that the gambles specified in ASSUMPTIONS 1 and 2 ought to be accepted in succession—is the one that Schick, McClennen, and others have found objectionable. The objection is that it would simply be unreasonable for S to continue circling the islands, and this is so no matter whether S actually abandons any of the three comparative judgments individually. There are independent reasons why the behavior exhibited in the Money Pump scenario is irrational, so it is wrong to say that S *ought to* accept the successive gambles. For instance, as McClennen might say, S ought to adopt a rational plan for how to explore the islands under various contingencies, and no such plan would include circling the trio to arrive back at the origin. Thus, so the objection goes, the adapted Money Pump does not actually tell us anything about the irrationality of these judgments *per se*.

As will become evident in the next section, there is nothing essential about the temporal aspect of the Money Pump argument; perhaps the diachronic version even obfuscates the more general phenomenon. The gambles need not be taken in succession, for example. However, in so far as representability is a wide scope notion—it applies to sets of judgments, rather than individual comparative judgments taken alone—we will need some assumption to the effect that these gambles should be accepted all together. This will be made more precise in the next section, after first introducing another example of a non-representable ordering. This example is intended in part to demonstrate why the simultaneity—or more generally, *agglomerativity*—assumption that we will make is a natural one.

RAIFFA ON ELLSBERG

Imagine an urn with 90 balls, 30 of which are red, and 60 of which are black or yellow in unknown proportion. Suppose a ball has been drawn at random, and consider the following four propositions:

¹⁰Thus, we are assuming basic utility theory, along the lines of the *locus classicus*, von Neumann and Morgenstern (1953, §1.3, §12.66).

- P1 — The ball is red.
- P2 — The ball is black.
- P3 — The ball is either red or yellow.
- P4 — The ball is either black or yellow.

Suppose, for whatever reason, a person S makes the following judgments:

$$P1 \succ P2 \quad \text{and} \quad P4 \succ P3 .$$

Evidently, these judgments are not probabilistically representable.¹¹ Our task is thus to understand what is defective about them.

In keeping with ASSUMPTION 1, S will correspondingly prefer GAMBLE 1 over GAMBLE 2, and likewise GAMBLE 4 over GAMBLE 3, in the following decision problem, where receiving 100 utiles (or perhaps some more concrete desirable commodity, such as dollars) is preferable to receiving 0.

	Red	Black	Yellow
GAMBLE 1	100	0	0
GAMBLE 2	0	100	0
GAMBLE 3	100	0	100
GAMBLE 4	0	100	100

Note that these are exactly the preferences exhibited in Ellsberg’s (1961) much discussed commentary on Savage (1954).¹² Of course, S may come to prefer GAMBLE 1 over 2, and 4 over 3, for any number of reasons; and S may hold these preferences without explicitly assenting to the judgments $P1 \succ P2$ and $P4 \succ P3$.¹³ At any rate, because we have assumed that these preferences between gambles are implicated by S ’s comparative probability judgments, we would like to know in what sense these preferences are flawed, if indeed they are.

In a comment on Ellsberg (1961), Raiffa (1961) suggested the following argument, slightly altered for current purposes. Consider a variation on the decision problem above, whereby a fair coin is flipped to determine whether the choice will be between GAMBLES 1 and 2, or between GAMBLES 3 and 4. On Option 1, GAMBLE 1 would be played in case of Heads, while GAMBLE 4 would be played in case of Tails; GAMBLES 2 and 3 would be played on Option 2 in these cases.

	Heads	Tails
Option 1	GAMBLE 1	GAMBLE 4
Option 2	GAMBLE 2	GAMBLE 3

¹¹Specifically, they violate a basic axiom of comparative probability proposed already in the 1940s by de Finetti, known as quasi-additivity: if $(X \cup Y) \cap Z = \emptyset$, then $X \succ Y$ if and only if $X \cup Z \succ Y \cup Z$. See, e.g., Krantz et al. (1971, §5).

¹²The example is offered as a counterexample to Savage’s “Sure-Thing Principle,” introduced and discussed in Savage (1954, Chapter 6).

¹³For example, Ellsberg himself suggested that the apparently problematic preference between gambles depended on a natural—indeed, not at all irrational—reaction to ambiguity, or “Knightian uncertainty.”

Given that GAMBLE 1 is strictly preferred to GAMBLE 2, and GAMBLE 4 is strictly preferred to GAMBLE 3, Option 1 strictly dominates Option 2. Thus, anyone who has such preferences will evidently prefer Option 1 to Option 2. However, it is easily seen that for both options, the objective probability of receiving the payoff of 100 is a half. Thus, Raiffa concludes, a strict preference in this case is unreasonable.

The problem becomes especially poignant when we consider the consequence of adopting ASSUMPTION 2 in this scenario. Assume that S is willing to pay something to ensure that GAMBLE 1 is played over GAMBLE 2, and likewise for GAMBLE 4 over GAMBLE 3. When presented with the choice between Option 1 and Option 2, it thus seems clear that S would pay to receive Option 1 over Option 2. But that is tantamount to paying so as to switch from a gamble on one random process to another, with no apparent difference between the two.

Why does it follow from S 's willingness to pay for each of GAMBLES 1 and 4 individually that S should be willing to pay for Option 1? Imagine the situation from S 's point of view. Given a choice between GAMBLES 1 and 2, S would opt for 1, and indeed would pay something for this preference, say c_1 . Similarly, given a choice between GAMBLES 3 and 4, S would pay something to ensure GAMBLE 4, say c_2 . Now S learns that only one of these choices will come into play. A random process will determine whether S 's choice between 1 and 2, or that between 3 and 4, will become a live gamble, and this process is known to be completely independent of S 's choice. What should S agree to pay to ensure 1 is played in case of Heads, and 4 is played in case of Tails? This should at least be worth the minimum of c_1 and c_2 . Such reasoning is completely general, and the fact that the resulting preference in the larger game in this case is patently irrational—in effect, S prefers a strictly dominated act—just shows why this set of judgments is deeply defective. The upshot of the theorem to be presented is that this kind of patent irrationality may result from any non-representable set of judgments (and conversely, will not result from any representable judgments).

What is the nature of the agglomerativity assumption we are now making? It is evidently quite different from the diachronic assumption present in the original the Money Pump argument. The decision problem in Raiffa's scenario, for instance, is a one-shot problem. In the general case, we will assume that, for each comparative judgment made, there is some *weight* associated with that pair of events, and the payoffs in the respective gambles are multiplied by these weights. The pairwise preferences between individual gambles should extend to weighted combinations of these gambles considered together; and following ASSUMPTION 2, we require that one be willing to give something up to ensure this weighted combination over any other combination with the same weights.

ASSUMPTION 3: For all comparisons, one ought to prefer the gambles as specified in ASSUMPTION 1 simultaneously, with payoffs weighted appropriately. Moreover, one ought to be willing to give something up to ensure this combination of gambles over any other combination.

In many natural cases, the weights can be viewed as giving probabilities that different choices of gambles will become effective, as in the above example from Raiffa. But we leave the formulation more general, so that they could also stand in for, e.g., relative importance of different gambles when all taken simultaneously. In another special case the weights are all unity, and the payoffs for the combination are given by the sums of those for the individual gambles. We will require only that the weights are all given positive real values.

How would this style of argument, now invoking our agglomerativity ASSUMPTION 3, apply to the case of intransitivity? Take again our initial explorer example, supposing now that the explorer will first arrive at island A , B , or C , say, each with equal probability. Again, upon reaching island C , B will be visible in the distance, so the question will arise of whether to stay at C or travel the extra distance to B , and the same for B and A , and for A and C , *mutatis mutandis*. In line with ASSUMPTIONS 1 and 2, we can assume that in each case there is some distance that the explorer would tolerate in order to reach the preferred island. Thus, for instance, upon reaching C , let us suppose, B is sufficiently more likely to warrant the additional trek. In line with ASSUMPTION 3, the explorer accepts each of these simultaneously. What this means is that, no matter which island the explorer meets first, time that could be spent searching that island will invariably be spent moving on to another island. Whatever chance the explorer had of finding the treasure originally, it can now only be worse.

If these look like mistakes that no reasonable person would make, that is exactly the point. ASSUMPTION 3 allows us to agglomerate individual preferences in such a way that what might locally look like reasonable judgments give perverse results when taken in concert. Finding an agent with such a pattern of judgments, we could either suggest that the agent ought to revise in some way, or perhaps we could deny ASSUMPTION 3. The fact that these problems resulting from agglomeration affect all and only those orderings that are not representable (Theorem 1 below) speaks in favor of revision. Moreover, it is unclear on what ground ASSUMPTION 3 could be questioned, at least if one is convinced by ASSUMPTIONS 1 and 2. That judgments are, or should be, agglomerative is a commonly accepted norm in both the theoretical and practical domains.¹⁴ Moreover, we have removed any potentially problematic diachronic aspect from the framework. The intuition, for instance, in case the weights correspond to probabilities, is that, while the subject knows that some decision problem or other will present itself—and the right action will depend on the subject’s comparative judgments—there is uncertainty about what this problem will be, and thus which judgments will matter. As in the explorer case, it is clear that one of the three islands will be met first, and thus the question will arise of whether to move on from *there*. It is uncertain which island that will be.

¹⁴For an argument that beliefs should be agglomerative, see, e.g., Velleman (2000); for an argument that intentions should be agglomerative (in a way that does not depend on agglomerativity of belief), see Bratman (2009). We will not argue for the assumption any further in this paper, as it strikes us as especially evident in the current setting.

KRAFT, PRATT, AND SEIDENBERG'S EXAMPLE

Lest one suspect that non-representable orderings only arise in relatively trivial cases, viz. intransitivity, Ellsberg-like scenarios, etc., consider the following set of judgments by a person S about the winner of an upcoming World Cup:

Denmark is more likely to win than either of Argentina or China.

One of Argentina or England is more likely to win than China or Denmark.

One of Brazil or China is more likely than one of Argentina or Denmark.

One of Argentina, China, or Denmark is more likely than Brazil or England.

Is this a reasonable set of judgments for S to maintain? Or are these judgments somehow incoherent? As it happens, this ordering is not probabilistically representable.¹⁵ We must therefore show why the ordering leads to a pragmatic defect, indeed to preference of a strictly dominated act, just as in the previous cases. In this case we illustrate the argument in a particularly perspicuous—if also somewhat more contrived—way: we simply consider a simultaneous system of bets.¹⁶ Consider the following eight gambles:

G1: Payoff of 1 on Denmark, 0 if they lose.

G2: Payoff of 1 on Argentina or Chile.

G3: Payoff of 1 on Argentina or England.

G4: Payoff of 1 on China or Denmark.

G5: Payoff of 1 on Brazil or China.

G6: Payoff of 1 on Argentina or Denmark.

G7: Payoff of 1 on Argentina, China, or Denmark.

G8: Payoff of 1 on Brazil or England.

Evidently, the judgments above would lead S to the following pairwise preferences: G1 over G2, G3 over G4, G5 over G6, and G7 over G8. In particular, following ASSUMPTION 3, if S were offered a free set of gambles consisting of G2, G4, G6, and G8 (all weighted equally), S would agree to pay some positive amount to trade these in for G1, G3, G5, and G7 instead. But this is absurd, since no matter which of the five teams wins, S 's net payoffs under the two sets of gambles are guaranteed to be identical.

Observe that, unlike other types of pragmatic arguments such as Dutch book theorems, the structure of the argument here does not depend on one's attitude toward risk or betting. S may, for example, face a scenario in which payoff

¹⁵Essentially this example was used by Kraft et al. (1959) to refute a conjecture of de Finetti, who proposed that quasi-additivity (cf. Footnote 11), together with some other obvious axioms, would be sufficient to ensure representability. They showed that this is the simplest example of an ordering that satisfies quasi-additivity while not admitting any probabilistic representation.

¹⁶Actually, this is perhaps the most natural illustration for beliefs about World Cup matches.

N corresponds to receiving N tickets in a lottery, the winner of which keeps their job. This is a scenario S would undoubtedly not enter willingly, but it goes without saying that S ought gladly trade in for gambles that are perceived as more favorable. The explorer scenario also makes this point clear, since the explorer need not be risk seeking in any sense. Perhaps finding the treasure is one’s only possibility for survival, to take an extreme case.¹⁷

THE GENERAL CASE

In the previous two sections we have tried to motivate the intuitive assumptions underlying the result to be offered now more formally. The examples have ranged from literal betting scenarios to more ordinary practical decision problems, encompassing cases where all the gambles are taken together with summed payoffs, as well as cases where the payoffs from gambles are weighted, e.g., with probabilities of being “actualized.” In the background is a theorem, showing that the types of pragmatic defects exposed in these scenarios arise systematically for agents with non-representable comparative probability judgments, whereas agents with representable judgments are in a sense immune from these pathologies.

Let Ω be any finite set of atomic states, and let $\mathcal{E} \subseteq \wp(\Omega) \times \wp(\Omega)$ be any set of pairs of events, to be compared by our subject. The subset \mathcal{E} of events could be the set of all pairs, but we allow it to be a proper subset.¹⁸ In the former case we can consider only a single weak relation \succsim , from which \succ and \sim can be defined. Otherwise, we must keep track of \succ and \sim individually. Thus, we are assuming that a subject makes a number of judgments of the form $E \succ F$ —read as “Event E is more likely than event F ”—and of the form $E \sim F$ —read as “Events E and F are equally likely.” We therefore identify subjects with relation pairs $\{\succ, \sim\}$, whereby for each event pair $(E, F) \in \mathcal{E}$, either $E \succ F$ or $E \sim F$ holds.¹⁹

Enumerating the pairs of events from \mathcal{E} as $(E_1, F_1), \dots, (E_k, F_k)$, such that either $E_i \succ F_i$ or $E_i \sim F_i$, for $i \leq k$, let X index the pairs such that $E_x \succ F_x$, and let Y index the pairs $E_y \sim F_y$. We assume $X \neq \emptyset$, so that at least one pair is strictly compared by our subject.

Let Σ be the set of all functions σ sending each $i \leq k$ to either E_i or F_i . Our *canonical decision problem* consists of a set $\Omega = \{\omega_1, \dots, \omega_m\}$ of states and an action space Σ . A pair $\{\succ, \sim\}$ can be associated with a specific mixed act over Σ , namely that which randomizes over all pure acts σ for which $\sigma(x) = E_x$ for all $x \in X$.²⁰ Intuitively, this corresponds to the idea that such a subject prefers those acts that select E_x whenever $E_x \succ F_x$, but is indifferent between E_y and

¹⁷This is one of the main differences between the proposal made here and the off-hand remark by Fishburn (1986), who also observes that combinations of gambles like those above are guaranteed a net gain of exactly 0. We view these cases as arising in the course of ordinary decision making under uncertainty, not just in artificial gambling scenarios.

¹⁸Our result is therefore compatible with, but independent of, the claim that one must be able to compare any two options in terms of relative probability.

¹⁹We do not assume anything about \succ and \sim . Thus, e.g., \sim need not even be reflexive. All the usual requirements will come out of our result, Theorem 1 below.

²⁰Recall a *mixed* act is given by a probability distribution over *pure* or “deterministic” acts.

F_y whenever $E_y \sim F_y$. Let us write Σ^\succ for the set of such acts, and Q^\succ for the particular mixed act that selects each $\sigma \in \Sigma^\succ$ with equal probability $1/|\Sigma^\succ|$.

It remains only to specify the payoff function. In line with ASSUMPTIONS 1, 2, and 3, we presume first that there are two outcomes o_{good} and o_{bad} for every pair $i \leq k$, and that these outcomes can be assigned real number utilities u_{i+} and u_{i-} , respectively, with $u_{i+} > u_{i-}$. We also presume that there are real number weights $w_i > 0$ for each pair, and a cost $c > 0$ that our subject is willing to pay to ensure preferred gambles. The payoff function is derived directly from these assumptions. For an action $\sigma \in \Sigma$ and a state $\omega \in \Omega$, we define the *pre-utility* \mathcal{U}° (not yet incorporating costs) as:

$$\mathcal{U}^\circ[\sigma, \omega] = \sum_{i \leq k} w_i \cdot \begin{cases} u_{i+} & \text{if } \omega \in \sigma(i) \\ u_{i-} & \text{if } \omega \notin \sigma(i) \end{cases} .$$

That is, the payoff for an act σ is given by the (weighted) sum of utilities u_{i+} for those $\sigma(i)$ that come out true, and u_{i-} for those $\sigma(i)$ that come out false. For an action $\sigma \in \Sigma \setminus \Sigma^\succ$ incompatible with the orderings $\{\succ, \sim\}$, the utility coincides with the pre-utility: $\mathcal{U}[\sigma, \omega] = \mathcal{U}^\circ[\sigma, \omega]$. However, for $\sigma \in \Sigma^\succ$, the utility is c less than the pre-utility: $\mathcal{U}[\sigma, \omega] = \mathcal{U}^\circ[\sigma, \omega] - c$. This codifies the presumption that a subject whose judgments are expressed by $\{\succ, \sim\}$ would give up c to ensure playing any of the acts in Σ^\succ , including any mixed act over Σ^\succ , and in particular the uniform mixed act Q^\succ . We incorporate this assumption right into the decision problem. Thus, our canonical decision problem might look as follows:

	ω_1	ω_2	\dots	ω_m
σ_1	$\mathcal{U}_{1,1} - c$	$\mathcal{U}_{1,2} - c$	\dots	$\mathcal{U}_{1,m} - c$
σ_2	$\mathcal{U}_{2,1}$	$\mathcal{U}_{2,2}$	\dots	$\mathcal{U}_{2,m}$
\vdots	\vdots	\vdots	\ddots	\vdots
σ_{2^k}	$\mathcal{U}_{2^k,1} - c$	$\mathcal{U}_{2^k,2} - c$	\dots	$\mathcal{U}_{2^k,m} - c$

where $\mathcal{U}_{i,j} = \mathcal{U}^\circ[\sigma_i, \omega_j]$, and for instance, $\sigma_1, \sigma_{2^k} \in \Sigma^\succ$, while $\sigma_2 \in \Sigma \setminus \Sigma^\succ$.

We can now state our first and main lemma:

Lemma 1. Given a probability distribution P on Ω , the uniform mixed act Q^\succ maximizes P -expected utility for some $c > 0$, if and only if P agrees with $\{\succ, \sim\}$.

The proof of Lemma 1 is given in the Appendix. As a principle of rationality, expected utility maximization is not universally accepted.²¹ But in the case of an agent whose comparative probability judgments are not representable, the agent's preferred strategy does not maximize expected utility under *any* probability distribution. This allows us to borrow a technique from game theory, to show that any such act is in fact *dominated* by some (mixed) act. The second lemma, folklore in game theory, is stated in terms of two-player games:

²¹The references are far too many to list. Of course, Ellsberg (1961) provides an early example. See also Buchak (2013) for a recent argument that risk attitudes may justify behavior that does not maximize expected utility. Cf. also the discussion below.

Lemma 2. In a two-player game, a strategy σ for player 1 is a never-best response, if and only if σ is strictly dominated by a mixed strategy for player 1.

A proof of this result, using Nash’s Theorem on the existence of equilibria, is also given in the Appendix. Combining Lemmas 1 and 2, we can view the canonical decision problem as a game with our agent playing against nature, and nature’s utility simply the inverse of the agent’s utility. A pure strategy for nature is just a probability distribution over states, and a never-best response is an act that maximizes expected utility for no such distribution. We thus obtain our result.

Theorem 1. Act Q^c is *strictly dominated* in the canonical decision problem (for all $c > 0$), if and only if $\{\succ, \sim\}$ is not probabilistically representable.

The arguments given above in Sections 1, 2, and 3 can all be seen as special cases of this theorem, which generalizes these arguments to show that it is exactly the non-representable sets of judgments that are susceptible to this kind of defect.

Here is a summary of our proposal up to this point. A person with comparative probability judgments is assumed to reflect these judgments in decision making, by being indifferent between situations that depend on pairs of events judged equally likely, by preferring situations in which positive outcomes are associated with events judged more likely, and by agreeing to sacrifice something to be in a situation judged preferable. Provided these assumptions are met, Theorem 1 tells us that such a person will make choices that are worse in all possible ways things may turn out, that is, choices that are strictly dominated, just in case these judgments are incompatible with standard finitely additive probability.

For the rest of the paper, we offer several clarifications and responses to some potential objections.

COMPARISON WITH DUTCH BOOK ARGUMENTS

The argument for representability in this paper is similar in spirit to Dutch book style arguments for probabilism. Ramsey, de Finetti, and others held subjective probability to be of interest mainly to the extent that it makes a difference in action and behavior.²² Thus they were concerned to show that someone whose numerical judgments were not in accord with the standard probability axioms would be subject to some kind of pragmatic defect. This style of argument has come under criticism for being insufficiently “epistemic” in nature.²³ In particular such arguments do not deal directly with the truth, accuracy, or other distinctively epistemic features of the judgments themselves.

This fact should not detract from the interest of pragmatic arguments in general, however, even as applied to attitudes of purportedly mind-to-world direction of fit. One may of course be interested in the purely epistemic status of a set of

²²Ramsey (1926) famously claimed that “the differentia of belief lies in its causal efficacy”; de Finetti went so far as to say, “Probability [. . .] exists in that it serves to express, in a precise fashion, for each individual, his choice in his given state of ignorance” (de Finetti, 1974, p. 84).

²³See especially Joyce (1998) and many papers following.

judgments as such. But one may also be interested in how these judgments play out under reasonable assumptions about what people do with them. If our goal were, for instance, to convince someone of why their beliefs should satisfy a proposed postulate of rationality, it may be enough to point out that failure to do so would result in a systematic bias away from the truth. However, virtually anyone will be moved by the observation that failure to satisfy the postulate would result in systematically making choices that would be plainly foolish.²⁴

That said, there are more serious objections to Dutch book arguments in particular, and to the nature of the idealizations that seem to be required in applying them.²⁵ To what extent do these objections apply to the present proposal?

Using notation from the previous section, a typical Dutch book argument starts from the assumption that a subject assigns *previsions* $\{p_i\}_{i \leq n}$ to the elements of \mathcal{E} , with p_i assumed to capture simultaneously the subject's highest rate at which to buy, and lowest rate at which to sell, a bet on a proposition $E_i \in \mathcal{E}$. Assuming an agglomerativity principle similar to our ASSUMPTION 3, it is possible to show that such a subject averts a *sure loss* across all combinations of bets judged fair or favorable, if and only if the subject's previsions obey the standard probability axioms.

The most obvious difference between Dutch book arguments and the argument in this paper is that we have not assumed the subject assigns real valued previsions to the events in \mathcal{E} . The subject is only assumed to make qualitative judgments of relative probability. This is an importance difference in itself, but it also has significant repercussions for the link between judgments and action. Even if one is comfortable with the idea of making arbitrarily precise real valued judgments, that these judgments should match a subject's betting prices strikes many authors as implausible. Too many seemingly irrelevant prudential factors come into play: possibility of risk aversion, moral aversion to betting, and so on.

By contrast, in our setting the link between judgments and action is patent. We merely assume that a subject who judges E more likely than F would prefer to be in a situation where a positive outcome is attached to E being true, rather than to F being true; and moreover the subject would be willing to give something up, however small, to ensure being such a situation. This has the flavor of a gamble, but such scenarios abound in commonplace decision making. For instance, if S is at the airport to pick up a visitor, upon seeing that domestic airlines are one way and international the other, S 's choice of which way to go will depend wholly S 's judgments about which kind of airline the visitor is more likely to have taken. Importantly, as already mentioned earlier, such decision problems simply arise. There is no assumption that the subject must actively go in for betting, for example.

²⁴To quote de Finetti again: "Let this be said in order to make clear that such conditions, although *normative*, are not unjustified impositions of a criterion which their promoters consider 'reasonable': they merely assert that 'you must avoid this if you do not want ...' (and there follows the specification of something which is obviously undesirable)" (1974, p. 85).

²⁵See, e.g., Hájek (2008) for a relatively recent summary of many of them.

This difference is also manifested in the nature of the associated pragmatic defect. While Dutch book arguments concern sure loss—an incoherent subject is shown to go in for bets that guarantee a loss—our result concerns strict dominance. Some authors have claimed that, whereas the normative status of sure loss arguments is questionable, avoidance of strict dominance is incontrovertible as a normative standard.²⁶ Thus, the argument in this paper does have much in common with Dutch book arguments, but we also avoid many of the pitfalls and questionable assumptions that have been raised in the literature, which need not be seen as characteristic of pragmatic arguments in general.

STRICT DOMINANCE WITH MIXED ACTS

One of the purported advantages of our setting, in which subjects need only make qualitative comparative judgments, is that the result applies to a wider class of agents, including those who may not be capable of making sensible numerical judgments, for example. Yet, our results make reference to several probabilistic quantities, including, critically, the mixtures between acts invoked in the proof of Theorem 1. Indeed, the dominating act in the canonical decision problem will generally be a mixed act. Does this somehow diminish the interest of the result?

One can imagine at least two worries about this. First, a potential objection is that there may be agents who simply cannot understand claims about quantitative probability; for such agents our result will not be compelling. Second, and perhaps more seriously, it may seem as though we have smuggled in an assumption of expected utility maximization through the back door, so to speak, hiding it in the mixed act.²⁷ Let us take these two concerns in turn.

For the first, there are ways of making the result vivid without talking about numbers at all. Going back to our initial example of the explorer, suppose once again that we have no idea which island the explorer will reach first. The choice is between staying at that island and moving on to the next island. Following our assumptions, we claimed that the explorer with intransitive judgments would invariably move on to the next island, no matter which of the islands is met first. It does not take any sophisticated understanding of numbers, or of quantitative probability, to see: whatever the chances are that the explorer would find the treasure without moving, they can only be worse if the explorer takes the extra time to move to the next island in all possible eventualities.

Or, to return to Raiffa's comment on the Ellsberg scenario, suppose now that the payoffs represent not monetary units, but instead cubic centimeters of wax, from a large block of wax containing a small but very valuable diamond hidden somewhere in it.²⁸ As we saw, the two options promise the same amount of wax; yet, a person with the non-representable judgments in this case would evidently

²⁶For instance, Maher (1993, pp. 110-111) makes this argument. Though there are subtleties involved; see the discussion in the following section.

²⁷As Joyce (1998) points out, de Finetti's Dutch book argument does depend on the assumption that subjects maximize expected utility. It is important that our result does not.

²⁸This variation is inspired by Smith's (1961) discussion of "probability currency."

agree to give up some amount of wax to ensure one option over the other. In this case, we can literally *see* why such a choice is plainly irrational: the person has opted to receive less wax, and thus a worse chance of winning the diamond.

The second concern seems more serious. To see the force of the objection, consider the abstract form of the canonical decision problem in a representative instance of intransitive judgments. The relevant cells are pictured below:

	S1	S2	S3
A1	$1 - c$	$1 - c$	$1 - c$
A2	0	1	2
A3	1	2	0
A4	2	0	1

Our subject is assumed to prefer **A1** in this instance, even though this act is strictly dominated by the mixed act Q that selects **A2**, **A3**, and **A4**, each with equal probability. In expectation Q guarantees payoff 1 in all states, which is of course better than $1 - c$. But is it really always irrational to opt for **A1** in this kind of scenario? What if an agent simply wants to ensure some reasonably high utility ($1 - c$) instead of risking the (maybe disastrous?) outcome of utility 0? Indeed, exactly these types of decision problems have motivated some authors to call expected utility maximization into question as a rational *requirement*.²⁹

Let us grant, at least for the sake of argument, that in some such cases it may in fact be rationally permissible for a subject to choose an act dominated only by a mixed act. Of course, our argument does not depend on strict dominance *always* being objectionable. We already know that a number of assumptions need to be in place in order for dominance arguments to have force. For instance, if the state of the world is somehow dependent on the agent's act, then even dominance by a pure act may not be irrational.³⁰ Perhaps risk-related issues single out another class of cases where dominance, specifically by a mixed act, is unproblematic.

To be sure, one could not hope to show that a person who violates some maxim will always make foolish decisions. The most we can hope to show is that such foibles occur systematically for agents who violate the maxim, and are guaranteed not to occur for agents that satisfy it. We claim that, for agents whose judgments are not representable, such unambiguously unjustifiable choices do indeed occur systematically, and that for any non-representable set of judgments, there will be scenarios in which such foibles are clearly exhibited. Although this

²⁹See Buchak (2013) for a recent and comprehensive defense of *risk-weighted* utility maximization in particular, as a generalization of expected utility maximization.

³⁰One of the simplest illustrations of this is the following:

	S1	S2
A1	3	1
A2	4	2

If action A1 reliably leads to state S1, and A2 leads to S2, then even though A2 dominates A1, clearly A1 may nonetheless be the rational choice.

is a not a formal statement that can be proven, the examples given above—the explorer scenario, Raiffa’s example, and the World Cup gamble—lend support to the claim. In these examples, one cannot justify the dominated act by appeal to risk attitudes, or indeed any other considerations. This extends at least to all cases where the utilities in question measure some variant of “probability currency.” In any event, the normative force of our argument is proportional to the extent that these obviously objectionable consequences are in fact systematic.

PAYING THE COST

A final issue to address before closing is the role of cost in our argument. In presenting Theorem 1 we assumed that if a person judges E more likely than F , then that person will readily pay something to ensure a gamble on E over a gamble on F . Indeed, ASSUMPTION 2 is crucial for our result.³¹ Is this a reasonable assumption? Could one hold that judgment without also agreeing to give up some utility, however small, to guarantee the gamble judged favorable?

It is noteworthy that some quantitative representations of uncertainty seem to be incompatible with ASSUMPTION 2. For instance, if we take uncertainty to be represented by *real intervals* instead of real values,³² then if E is assigned the interval $(0.5, 1]$, one may reasonably judge E to be more likely than not, without agreeing to pay any real value to ensure a gamble on E over one on the negation of E . To take another example,³³ suppose we know that a certain real valued quantity is normally distributed around mean 0. Evidently, any reasonable person would judge value 0.01 to be more likely than 0.1, for instance. Yet, because both have zero probability, it would seem unreasonable to pay for the gamble on 0.01 over that on 0.1. Neither gamble is worth much of anything.

In as far as these cases arise, ASSUMPTION 2 may seem dubious. Not everyone has been convinced by such examples, however. In discussion of analogous arguments from Koopman and others, Savage (1954) writes:

I do not think it appropriate to suppose that the person would distinctly rather stake a gain on the line than on the null set [or on the value being 0.01 rather than 0.1]. The issue is not really either an empirical or a normative one, because the point and line in question are mathematical idealizations. (39)

The thought is that, whereas one can devise such examples in mathematically idealized scenarios, they do not arise in the course of everyday decision making, where the theory is supposed to apply most naturally. Implicit in Savage’s remark

³¹Incidentally, if one does not make this assumption, it is possible to alter the proof of Theorem 1 to show that an analogous result holds for the notion of *almost representability*. P almost represents an ordering pair $\{\succ, \sim\}$ if $E \succ F$ implies $P(E) \geq P(F)$, but not necessarily the stronger $P(E) > P(F)$. This is a rather weak notion; it does not even entail transitivity or irreflexivity of \succ , for instance (see Fishburn 1986 for discussion and axiomatization).

³²See, e.g., Smith (1961) for discussion and a variant on Dutch book arguments in this setting.

³³This example is from Koopman (1940): Fn 3.

is the assertion that one cannot really be said to find E more likely than F unless one is willing to pay. That is, one might argue that ASSUMPTION 2, like ASSUMPTION 1, is *constitutive* of the judgment that E is more likely than F .

We will not make this argument here. Instead, we note that considerations such as these are unlikely to call any of our results into question anyway. If a subject makes (a finite number of) relative probability judgments on the basis of, say, interval representations, or non-standard real number representations, then it is easily seen that those judgments are *ipso facto* probabilistically representable.³⁴ Hence, such a person would not succumb to dominance even with assumed cost. For a counterexample we would need a convincing case where a sensible person would not agree to any cost, and only thereby avoids dominance. Of course, any purported counterexample based solely on intuition will likely lead to a stalemate between those who, like Savage, insist that the person must not really assent to the comparative judgments, and those who insist otherwise. We leave open whether there are any genuine such counterexamples.

CONCLUSION

The aim of this paper has been to demonstrate the sense in which a person whose comparative probability judgments are incompatible with numerical probability suffers from a substantive pragmatic defect. Drawing on assumptions about the fundamental link between comparative judgments and choice, which we have argued are reasonable, our main result states that a person will be prone to choose strictly dominated acts across a wide array of choice problems, just in case the person's judgments are not probabilistically representable. The argument is similar in spirit to both Money Pump arguments and Dutch book style arguments, while at the same time avoiding many of the objections to each.

Of course, comparative probability judgments do not exhaust the range of attitudes under uncertainty that people may have. Sometimes more fine-grained distinctions, e.g., of degree, are apposite. In deciding whether to cross the street, for instance, it will not be enough to judge that it is more likely than not to be safe; one would like to know that it is *much* more likely. Our result implies that a person's judgments should be probabilistically representable, at least when restricting to their purely qualitative judgments of comparative probability. But the result applies to explicit numerical probability judgments only insofar as those judgments may engender problematic comparative judgments. It is possible to violate the standard quantitative axioms of probability even when the associated comparative judgments are representable by some other, coherent distribution.

In as far as the argument presented here avoids some of the drawbacks of other pragmatic arguments for quantitative principles, one may hope to adapt some of these ideas to the quantitative setting. Even if this proves difficult for numerical probability judgments in general, it may be possible for richer systems

³⁴Recall we were careful to state representability so as to allow for incomparabilities; thus representability by a single measure and by a set of measures coincides.

of representation in between, e.g., including judgments like “Event E is twice as likely as event F .” We leave such possibilities for future work.

At any rate, it bears mention once more that, while we as theorists have relied on quantitative methods to prove the main result in this paper, the result itself can be seen to appeal only to qualitative, comparative notions: a person with non-representable comparative judgments would often be *more likely* to achieve some desirable result if those judgments were instead representable.

REFERENCES

- Bratman, M. E. (2009). “Intention, Belief, Practical, Theoretical.” *Spheres of Reason: New Essays in the Philosophy of Normativity*, S. Robertson, ed., Oxford University Press.
- Broome, J. (2007). “Wide or narrow scope?.” *Mind*, 116, 359–370.
- Buchak, L. (2013). *Risk and Rationality*. Oxford University Press.
- Davidson, D., McKinsey, J. C. C., and Suppes, P. (1955). “Outline of a formal theory of value, I.” *Philosophy of Science*, 22, 140–160.
- Ellsberg, D. (1961). “Risk, ambiguity, and the Savage axioms.” *Quarterly Journal of Economics*, 75(4), 643–669.
- Fine, T. L. (1973). *Theories of Probability*. Academic Press.
- de Finetti, B. (1937). “La prévision: ses lois logiques, ses sources subjectives.” *Annales de l’Institut Henri Poincaré*, 7, 1–68.
- de Finetti, B. (1974). *Theory of Probability*, Vol. 1. Wiley, New York.
- Fishburn, P. C. (1986). “The axioms of subjective probability.” *Statistical Science*, 1(3), 335–358.
- Fitelson, B. and McCarthy, D. (2014). “Toward an epistemic foundation for comparative confidence”. Draft of 1/19/14.
- Gustafsson, J. E. (2013). “The irrelevance of the diachronic money-pump argument for acyclicity.” *Journal of Philosophy*, 110(8), 460–464.
- Hájek, A. (2008). “Dutch book arguments.” *The Oxford Handbook of Rational and Social Choice*, P. Anand, P. Pattanaik, and C. Puppe, eds., Oxford University Press.
- Jaynes, E. T. (2003). *Probability Theory: The Logic of Science*. Cambridge University Press.
- Joyce, J. M. (1998). “A nonpragmatic vindication of probabilism.” *Philosophy of Science*, 65, 575–603.

- Keynes, J. M. (1921). *A Treatise on Probability*. Macmillan.
- Koopman, B. O. (1940). “The axioms and algebra of intuitive probability.” *Annals of Mathematics*, 41(2), 269–292.
- Kraft, C. H., Pratt, J. W., and Seidenberg, A. (1959). “Intuitive probability on finite sets.” *The Annals of Mathematical Statistics*, 30(2), 408–419.
- Krantz, D. H., Luce, R. D., Suppes, P., and Tversky, A. (1971). *Foundations of Measurement*, Vol. 1. Academic Press, New York.
- Maher, P. (1993). *Betting on Theories*. Cambridge University Press.
- McClellenn, E. F. (1990). *Rationality and Dynamic Choice: Foundational Explorations*. Cambridge University Press.
- von Neumann, J. and Morgenstern, O. (1953). *Theory of Games and Economic Behavior*. Princeton University Press.
- Osborne, M. J. and Rubinstein, A. (1994). *A Course in Game Theory*. MIT Press.
- Raiffa, H. (1961). “Risk, ambiguity, and the Savage axioms: Comment.” *Quarterly Journal of Economics*, 75(4), 690–694.
- Ramsey, F. P. (1926). “Truth and probability.” *Foundations of Mathematics and Other Essays*, R. B. Braithwaite, ed., Routledge.
- Savage, L. J. (1954). *Foundations of Statistics*. Dover, 1972, 2nd revised edition.
- Schick, F. (1986). “Dutch bookies and money pumps.” *Journal of Philosophy*, 83, 112–119.
- Scott, D. (1964). “Measurement structures and linear inequalities.” *Journal of Mathematical Psychology*, 1, 233–247.
- Smith, C. A. B. (1961). “Consistency in statistical inference and decision.” *Journal of the Royal Statistical Society, Series B*, 23, 1–37.
- Velleman, J. D. (2000). *The Possibility of Practical Reason*. Oxford University Press.

APPENDIX

As a result of the definitions given in the main text, relative to a specific distribution P we can express the expected utility $\mathbb{E}\mathcal{U}(\sigma)$ of an act $\sigma \in \Sigma$ in a perspicuous way. For each $i \leq k$, let us write β_{E_i} for the following sum:

$$\beta_{E_i} = (P(E_i) \cdot u_{i+}) + (P(F_i) \cdot u_{i-}),$$

and likewise for β_{F_i} . Then for an act $\sigma \in \Sigma \setminus \Sigma^\succ$,

$$\mathbb{E}\mathcal{U}(\sigma) = \sum_{i \leq k} w_i \beta_{\sigma(i)}.$$

Likewise, the expected utility of an act $\sigma \in \Sigma^\succ$ can be written:

$$\mathbb{E}\mathcal{U}(\sigma) = -c + \sum_{i \leq k} w_i \beta_{\sigma(i)}.$$

The expected utility for a mixed act Q is simply $\mathbb{E}\mathcal{U}(Q) = \sum_{\sigma} Q(\sigma) \cdot \mathbb{E}\mathcal{U}(\sigma)$, where $Q(\sigma)$ is the probability of opting for σ under Q . In the particular case that P agrees with $\{\succ, \sim\}$, it turns out that all mixtures Q over Σ^\succ (including pure acts) have the same expected utility. This can be seen from the fact that, whenever $E_i \sim F_i$, we have

$$\beta_{E_i} = (1/2)(u_{i+}) + (1/2)(u_{i-}) = \beta_{F_i};$$

hence all weighted combinations are equivalent as well.

We are now in a position to prove the main lemma:

Lemma 3. Given a probability distribution P on Ω , the uniform mixed act Q^\succ maximizes P -expected utility for some $c > 0$, if and only if P agrees with $\{\succ, \sim\}$.

Proof. For the right-to-left direction, suppose that P agrees with $\{\succ, \sim\}$. To show that Q^\succ maximizes expected utility it suffices to show $\mathbb{E}\mathcal{U}(Q^\succ) \geq \mathbb{E}\mathcal{U}(\sigma)$ for all *pure* acts σ . If $\sigma \in \Sigma^\succ$, then as we noted above, $\mathbb{E}\mathcal{U}(Q^\succ) = \mathbb{E}\mathcal{U}(\sigma)$.

If $\sigma \in \Sigma \setminus \Sigma^\succ$, then for some $h \leq k$ we have $\sigma(h) = F_h$, while by assumption $\hat{\sigma}(h) = E_h$ for all $\hat{\sigma} \in \Sigma^\succ$. Since P agrees with $\{\succ, \sim\}$, we have $P(E_h) > P(F_h)$, and hence

$$\begin{aligned} \beta_{E_h} - \beta_{F_h} &= \left((P(E_h) - P(F_h)) \cdot u_{h+} \right) + \left((P(F_h) - P(E_h)) \cdot u_{h-} \right) \\ &= (P(E_h) - P(F_h)) \cdot (u_{h+} - u_{h-}) \\ &> 0. \end{aligned}$$

Choosing c to be any value less than $w_h(\beta_{E_h} - \beta_{F_h})$, we therefore have

$$\mathbb{E}\mathcal{U}(Q^\succ) - \mathbb{E}\mathcal{U}(\sigma) > -c + w_h(\beta_{E_h} - \beta_{F_h}),$$

and thus $\mathbb{E}\mathcal{U}(Q^\succ) > \mathbb{E}\mathcal{U}(\sigma)$.

For the left-to-right direction, suppose P does not agree with $\{\succ, \sim\}$. Again, we have two cases: this could either be because (i) $E_h \succ F_h$ while $P(F_h) \geq P(E_h)$, or (ii) $E_h \sim F_h$ while $P(E_h) > P(F_h)$.³⁵ In case (i), define a new pair $\{\succ^*, \sim^*\}$

³⁵The other subcase of (ii), where $P(F_h) > P(E_h)$, is of course symmetric.

to be exactly like $\{\succ, \sim\}$, except that $E_h \sim^* F_h$. It follows that, with probability $1/2$, act Q^{\succ^*} will opt for F_h instead of E_h , and not pay the cost c . Hence,

$$\mathbb{E}\mathcal{U}(Q^{\succ^*}) - \mathbb{E}\mathcal{U}(Q^{\succ}) = \frac{1}{2} \left(c + w_h (\beta_{F_h} - \beta_{E_h}) \right).$$

Since $c, w_h > 0$, and $\beta_{F_h} - \beta_{E_h} \geq 0$, we have $\mathbb{E}\mathcal{U}(Q^{\succ^*}) > \mathbb{E}\mathcal{U}(Q^{\succ})$.

In case (ii), define a new ordering $\{\succ^*, \sim^*\}$ to be exactly like $\{\succ, \sim\}$, except that, instead of $E_h \sim F_h$, we have $E_h \succ^* F_h$. Note that the support of Q^{\succ^*} , i.e., Σ^{\succ^*} , is a strict subset of Σ^{\succ} . In particular, both suffer cost c . On the other hand, since $P(E_h) > P(F_h)$, someone who plays Q^{\succ^*} can be expected to fare better:

$$\mathbb{E}\mathcal{U}(Q^{\succ^*}) - \mathbb{E}\mathcal{U}(Q^{\succ}) = w_h (\beta_{F_h} - \beta_{E_h}).$$

Because $\beta_{F_h} - \beta_{E_h} > 0$, we have once again that $\mathbb{E}\mathcal{U}(Q^{\succ^*}) > \mathbb{E}\mathcal{U}(Q^{\succ})$. \dashv

The following lemma is folklore in game theory (see, e.g. Osborne and Rubinstein (1994): Lemma 60.1), but we include a proof here for completeness.

Lemma 4. Suppose $(\mathcal{A}^1, \mathcal{A}^2, \mathcal{U}^1, \mathcal{U}^2)$ is a two-player game with pure strategy sets \mathcal{A}^1 and \mathcal{A}^2 , and utilities \mathcal{U}^1 and \mathcal{U}^2 , and associated mixed strategy sets \mathcal{M}^1 and \mathcal{M}^2 , for players 1 and 2, respectively. Then an act $a \in \mathcal{A}^1$ is a never-best response, if and only if a is strictly dominated by $m \in \mathcal{M}^1$.

Proof. (\Leftarrow) If a is dominated by x , then for any $m \in \mathcal{M}^2$, $\mathcal{U}^1(x, m) > \mathcal{U}^1(a, m)$, so a is never a best response.

(\Rightarrow) If a is a never-best response, then there is a function $\beta : \mathcal{M}^2 \rightarrow \mathcal{A}^1 \setminus \{a\}$ such that $\mathcal{U}^1(\beta(m), m) > \mathcal{U}^1(a, m)$ for all $m \in \mathcal{M}^2$. Define a zero-sum game $(\mathcal{A}^1, \mathcal{A}^2, \bar{\mathcal{U}}^1, \bar{\mathcal{U}}^2)$, such that $\bar{\mathcal{U}}^1(x, y) = \mathcal{U}^1(x, y) - \mathcal{U}^1(a, y)$ and $\bar{\mathcal{U}}^2(x, y) = -\bar{\mathcal{U}}^1(x, y)$. Note that

$$\bar{\mathcal{U}}^1(z, y) > \bar{\mathcal{U}}^1(x, y) \quad \Leftrightarrow \quad \mathcal{U}^1(z, y) > \mathcal{U}^1(x, y).$$

Let (x^*, y^*) be a Nash equilibrium of this game. We claim x^* dominates a :

$$\begin{aligned} \bar{\mathcal{U}}^1(x^*, m) &\geq \bar{\mathcal{U}}^1(x^*, y^*) \\ &\geq \bar{\mathcal{U}}^1(\beta(y^*), y^*) \\ &> \bar{\mathcal{U}}^1(a, y^*) \\ &= 0 \\ &= \bar{\mathcal{U}}^1(a, m). \end{aligned}$$

for any $m \in \mathcal{M}^2$. Thus, for any $m \in \mathcal{M}^2$, $\mathcal{U}^1(x^*, m) > \mathcal{U}^1(a, m)$. \dashv