Ramsey without Ethical Neutrality: A New Representation Theorem

Edward Elliott
Australian National University
Edward.elliott@anu.edu.au

Frank Ramsey's 'Truth and Probability' sketches a proposal for the empirical measurement of credences, along with a corresponding set of axioms for a (somewhat incomplete) representation theorem intended to characterize the preference conditions under which this measurement process is applicable. There are several features of Ramsey's formal system which make it attractive and worth developing. However, in specifying his measurement process and his axioms, Ramsey introduces the notion of an ethically neutral proposition, the assumed existence of which plays a key role throughout Ramsey's system. A number of later representation theorems have also appealed to ethically neutral propositions. The notion of ethical neutrality has often been called into question — in fact, there seem to be good reasons to suppose that no ethically neutral propositions exist. In this paper, I present several new, Ramsey-inspired representation theorems that avoid any appeal to ethical neutrality. These theorems preserve the benefits of Ramsey's system, without paying the cost of ethical neutrality.

1. Introduction

In his posthumously published 'Truth and Probability', Frank Ramsey (1931b) sketches a proposal for the empirical measurement of credences, along with a corresponding set of axioms for a (somewhat incomplete) representation theorem intended to characterize the preference conditions under which this measurement process is applicable. Ramsey's formal approach is distinctive, first deriving a utility function to represent an agent's utilities, and then using this to construct their credence function. In specifying his measurement process and his axioms, Ramsey introduces the notion of an *ethically neutral proposition*, the assumed existence of which plays a key role throughout his system.

The existence of such propositions has often been called into question. Ramsey's own definition of ethical neutrality presupposes the philosophically suspect theory of logical atomism. On other common ways of defining the notion, it is frequently noted that we lack good

Mind, Vol. 126 . 501 . January 2017
doi:10.1093/mind/fzv180 Advance Access publication 14 October 2016

reasons for supposing that ethically neutral propositions exist, and in some cases we find that there are very good reasons for supposing that they cannot exist. Any system which relies on the existence of such propositions ought to be rejected.

However, we can have a representation theorem which is essentially Ramseyian in character that does not succumb to these objections. In §4, I will develop several new representation theorems inspired by Ramsey's own which avoid appeal to ethically neutral propositions in any problematic sense, substantially improving upon Ramsey's proposal. In §5, I will note several features that the new theorems share with Ramsey's system which make them independently attractive, especially in comparison to later works such as Savage (1954) and Jeffrey (1983). First, however, I will discuss Ramsey's original proposal in some detail (§2), which will help us see why Ramsey thought he needed to introduce the notion of ethical neutrality, why it is problematic, and how those problems can be circumvented (§3).

2. Ramsey's proposal

2.1 Background

One of Ramsey's main goals in 'Truth and Probability' was to argue that the laws of probability supply for us the 'logic of partial belief' (1931b, p. 166). His argument proceeds by first attempting to say what credences *are*, and on the basis of that understanding, showing that they conform to the laws of probability.

Regarding the first step, Ramsey clearly had operationalist sympathies, asserting that the notion 'has no precise meaning unless we specify more exactly how it is to be measured' (1931b, p. 167). To be measured as having probabilistically coherent credences *is* (more or less) on this picture to have probabilistically coherent credences, and anyone who can be measured through Ramsey's procedure at all will have credences conforming to the laws of probability.[1] Note that the

[1] Setting operationalism aside, it is easy to see in 'Truth and Probability' an early statement of something like functionalism: degrees of belief are to be understood through their causal role with respect to behaviour when considered in conjunction with a total desire state. Ramsey writes that 'the degree of a belief is a causal property of it, which we can express vaguely as the extent to which we are prepared to act on it' (1931b, p. 169). Ramsey argues against characterizing degrees of belief in terms of some introspectively accessible feeling had by a subject upon considering the relevant proposition. These arguments go well beyond operationalism, though I will not recapitulate them here. He concludes that 'intensities of belief-feelings … are no doubt interesting, but … their practical interest is entirely due to their

procedure was intended to be applicable to ordinary agents — Ramsey was not trying to define degrees of belief for some ideally rational being, but for the everyday person on the street (albeit not without some unavoidable idealization).

With this goal in mind, Ramsey proposes to take as the theoretical basis of his measurement system a particular theory of decision making — that is, the theory that 'we act in the way that we think most likely to realize the objects of our desires, so that a person's actions are completely determined by his desires and opinions' (1931b, p. 173). His idea is to assume the basic truth of something like subjective expected utility theory, and on that assumption, use empirical information about an agent's preferences to work out what her credences and utilities must be. Ramsey was entirely aware of the empirical difficulties that this theory faces, writing that:

> [it] is now universally discarded, but nevertheless comes, I think, fairly close to the truth in the sort of cases with which we are most concerned … This theory cannot be made adequate to all the facts, but it seems to me a useful approximation to the truth particularly in the case of our self-conscious or professional life, and it is presupposed in a great deal of our thought. (1931b, p. 173)

We will return shortly to exactly what Ramsey meant by 'the sort of cases with which we are most concerned', and exactly what he needed to assume to get his measurement process off the ground.

At several points Ramsey notes that a similar kind of reasoning is used frequently by physicists and other hard scientists in the development of systems for the measurement of non-psychological quantities. Measurement systems are never developed in a theoretical vacuum; the actual data we receive are always interpreted through the lens of some presupposed theory or another — and quite frequently, as Ramsey notes, such a theory might 'like Newtonian mechanics … still be profitably used even though it is known to be false' (1931b, p. 173), so long as it is accurate for the cases at hand. Indeed, at several points Ramsey notes that measurement 'cannot be accomplished without introducing a certain amount of hypothesis or fiction … if it is allowable in physics it is allowable in psychology also' (1931b, p. 168; cf. Krantz *et al* 1971, pp. 26–31 on the role of idealizing assumptions in the construction of measurement systems). For Ramsey, this hypothesis or

position as the hypothetical causes of beliefs *qua* bases of action' (1931b, p. 172). On this more charitable interpretation, Ramsey's representation theorem can be seen as spelling out precisely the relevant causal roles associated with credence states.

fiction is that ordinary folk are more or less expected utility maximizers, at least in the right circumstances.

With this in mind, we can summarize Ramsey's measurement procedure as follows:

(a) determine an agent's preferences over outcomes and gambles;

(b) define a relation of equal difference in utilities;

(c) locate ethically neutral propositions of probability ½;

(d) construct an interval scale representation of the agent's preferences;

(e) use that representation to define a probability function, which is taken to represent the agent's credences.

I will discuss each step in turn. For the sake of simplicity, I have neglected to discuss one important aspect of Ramsey's procedure — *viz.*, the use of preferences over complex gambles to define conditional probabilities, which are used to show that the measured credences constitute a probability function. (This part of Ramsey's procedure is outlined in Bradley 2001.) It is also worth emphasizing that much of what follows is a rational reconstruction — Ramsey's own remarks are sketchy at best, and he rarely explains the motivations for any of the steps he makes.

2.2 Determining a preference ordering

The first stage of Ramsey's procedure is to determine the agent's preferences. We begin with the agent's preferences over what I will call *outcomes*. This is, according to Ramsey, relatively straightforward:

> If … we had the power of the Almighty, and could persuade our subject of our power, we could, by offering him options, discover how he placed in order of merit all possible courses of the world. In this way all possible worlds would be put in an order of value … (1931b, p. 176)

Ramsey here makes it clear that he takes outcomes to be 'different possible totalities of events … the ultimate organic unities' (1931b, pp. 177–8); that is, possible worlds. It is unclear here exactly what Ramsey intended by 'possible worlds'. For instance, it is unclear whether we are considering metaphysically or conceptually possible worlds, and whether we are only interested in worlds consistent with history up to the event of the choice.

In any case, we will designate the set of outcomes with $\boldsymbol{W} = \{w_1, w_2, \ldots\}$. Importantly, within only a few paragraphs of referring to the outcomes as worlds, Ramsey goes on to note that with respect to at least one proposition P, and some outcomes w_1 and w_2, '[w_1] and [w_2] must be supposed so far undefined as to be compatible with both P and $\neg P$' (1931b, p. 178, Fn. 1). The most natural interpretation of this seems to be that w_1 and w_2 ought to be considered *not quite* as worlds, but rather as propositions maximally specific with respect to everything except P. (See §3.1 for further discussion.) Generalizing, we can take Ramsey's outcome space $\boldsymbol{W}$ to be a set of very highly specific propositions, some—but not all—of which may be maximally specific. We will say that $w_1 \succ w_2$ just in case the agent prefers w_1 to w_2; $w_1 \sim w_2$ just in case the agent is indifferent between the two options. The first task in Ramsey's measurement procedure is, then, to determine a preference ordering over $\boldsymbol{W}$.

We are required also to empirically determine how the agent ranks *gambles*. Once again, Ramsey asks us to imagine that we have convinced our subject of our power, but this time we make offers of the following kind: 'Would you rather have world [w_3] in any event, or world [w_1] if P is true, and world [w_2] if P is false?' (1931b, p. 177). Let us represent the latter option, the gamble w_1 *if P is true, and* w_2 *if P is false*, as simply $(w_1, P; w_2)$.[2] With his background assumption in mind, Ramsey then notes that:

> If … [the agent] were certain that P was true, he would simply compare [w_1] and [w_3] and choose between them as if no conditions were attached; but if he were doubtful his choice would not be decided so simply. (1931b, p. 177)

Here, Ramsey looks to compare an outcome with a gamble, so we are to assume that gambles and outcomes are comparable. It is also evident from the axioms he later provides that we need to consider agents' preferences between gambles. In sum, if $\boldsymbol{G}$ is the set of all gambles, then Ramsey requires us to empirically determine a $\succcurlyeq$-ordering on $\boldsymbol{W} \cup \boldsymbol{G}$.

[2] It is unclear exactly how we ought to understand Ramsey's gambles. For reasons outlined in Joyce (1999, pp. 62–3), 'w_1 if P is true, and world w_2 if P is false' should not be understood using material conditionals. It is not obvious whether either subjunctive or indicative conditionals would fare any better; though Sobel (1998, p. 239) suggests that $(w_1, P; w_2)$ is just a conjunction of subjunctives, $(P \square\!\!\rightarrow w_1)$ & $(\neg P \square\!\!\rightarrow w_2)$. Part of the exegetical difficulty here is due to Ramsey's lack of specificity regarding the nature of the outcome set, and to the interpretive difficulties present in Ramsey's proposed system discussed shortly below.

There are a number of interpretive difficulties with Ramsey's proposed system that might be raised at this point. For one thing, it is unclear how outcomes as highly specific as Ramsey suggests can be 'offered' to any ordinary human subject; the power to conceptualize even one possible world in all its detail seems beyond the average person. Perhaps one might chalk this up as a harmless idealization, of the kind frequently appealed to in many measurement systems. Somewhat more worrying, however, is that convincing a subject that 'we had the power of the Almighty' would surely drastically alter her doxastic state prior to measuring it, as Jeffrey (1983, pp. 158–60) has noted. Likewise, when a subject is offered the choice of either w_3 or (w_1, P; w_2), we must not suppose that her credence in P is in any way changed by the offer or this would ruin the measurement.

Interestingly, Ramsey himself objects to another proposed bet-based measurement system on the grounds that 'the proposal of the bet may inevitably alter [the subject's] state of opinion' (1931b, p. 172). Either Ramsey did not recognize that the same objection applies with greater force to his own account, or he believed that the worry could be addressed. Bradley (2001, pp. 285–8) suggests one way in which it might be addressed: rather than making the subject believe in our godlike powers, we simply ask her to choose amongst options *as if* they were genuinely available to her (with perhaps the added proviso that she ought not to change her credences in any relevant propositions). To the extent that such a request can be satisfied, this reconstrual of Ramsey's methodology may help to minimize any changes to subjects' credences prior to measurement.

In any case, we can now say precisely what Ramsey meant when he referred to the accuracy of expected utility theory in 'the sort of cases with which we are most concerned'. We are to limit our attention to conscious, deliberate and presumably reflective judgements of preference between worlds and worlds, gambles and gambles, and worlds and gambles. Plausibly, Ramsey would have also held that we are not to consider cases where the subject is intoxicated, or under any kind of substantial physical or emotional duress. Ramsey does not need to assume anything as strong as the truth of subjective expected utility theory *tout court*, nor even its approximate truth across a wide range of cases — he only needs to assume that it is accurate in this particular kind of case. It would not be misleading, therefore, to interpret Ramsey's use of $\succ$ as, roughly: $\alpha \succ \beta$ relative to an agent S just in case S prefers α to β after consciously deliberating on the matter, while neither under physical or emotional distress, nor under the influence of any intoxicating substances. On this interpretation, Ramsey's assertion that expected utility theory is broadly accurate in 'the sort of cases with which we are most concerned' is essentially the claim that an ordinary agent's reflective preference ordering over worlds and gambles is what we would expect of an expected utility maximizer. Put like this, Ramsey's assumption appears quite plausible.

2.3 Defining an equal difference relation

Ramsey's first step has us empirically determine how the agent ranks outcomes and gambles. However, a simple preference ordering on outcomes and gambles only suffices for an ordinal scale representation of an agent's utilities. For Ramsey, this is unsatisfactory: 'There would be no meaning in the assertion that the difference in value between [w_1] and [w_2] was equal to that between [w_3] and [w_4]' (1931b, p. 176). We usually suppose that it is meaningful to say such things as 'I desire α *much more* than β; more so than α over γ' and 'I prefer α over β to the same extent that I prefer γ over δ'. A merely ordinal representation is unable to capture this information, representing as it does only the relative positions of outcomes in the preference order. Thus Ramsey sets himself the task of characterizing an *equal difference* (in utilities) relation between pairs of outcomes wholly in terms of preferences over gambles. If he can do this, then on the basis of well-known results from the mathematical theory of measurement, he can construct a richer representation of our utilities.

Let us say that $(w_1, w_2) =^{d} (w_3, w_4)$ holds just in case the difference in value for the agent between w_1 and w_2 is equal to the difference in value between w_3 and w_4. Ramsey's goal of defining $=^{d}$ in terms of preferences over gambles then sets up a certain difficulty to be overcome. According to the assumed background theory, an agent's preferences over gambles are determined by two factors: their *utilities* and their *credences*. Whether an agent prefers $(w_1, P; w_2)$ to $(w_3, Q; w_4)$, for example, depends partly on the utilities that she attaches to the outcomes w_1, w_2, w_3, w_4, and partly on the credences she has with respect to P and Q. However, whether $(w_1, w_2) =^{d} (w_3, w_4)$ holds for that agent should depend *solely* on the utilities she attaches to w_1, w_2, w_3, w_4. In order to define $=^{d}$ in terms of preferences over gambles, then, Ramsey needs some way of factoring out any confounding influences, so that whether the agent prefers one of the relevant gambles to another depends *only* on the utilities attached to the outcomes involved.

Ramsey's central innovation here is to define, in terms of preference, what it is for a proposition to have probability ½, and then to use this to define $=^{d}$. Let us suppose for simplicity that whether an agent prefers $(w_1, P; w_2)$ to $(w_3, Q; w_4)$ depends *only* on the utilities the agent has for w_1, w_2, w_3, w_4, and the credences she has for P and Q. More specifically, assume the following:

Naïve Expected Utility Theory
If *des* is a real-valued interval scale that measures the agent's utilities for outcomes (i.e., $des(w_1) \geq des(w_2)$ if and only if $w_1 \succcurlyeq w_2$), and *bel* is a probability function that represents the agent's credences, then $(w_1, P; w_2) \succcurlyeq (w_3, Q; w_4)$ if and only if $des(w_1).bel(P) + des(w_2).(1 - bel(P)) \geq des(w_3).bel(Q) + des(w_4).(1 - bel(Q))$.

We will note shortly that Ramsey did *not* assume Naïve Expected Utility Theory; but for now it suffices to explain the reasoning behind his definitions.

Suppose that the agent is indifferent between the following two gambles: $(w_1, P; w_2)$ and $(w_2, P; w_1)$. According to Naïve Expected Utility Theory, there are only two (not mutually exclusive) ways in which this might come about: either w_1 and w_2 have exactly the same utility for the agent, or the agent's credence in P is exactly 0.5. To rule out the former possibility, we consider a pair of gambles $(w_3, P; w_4)$ and $(w_4, P; w_3)$, where we know that the agent is not indifferent between w_3 and w_4. If we find that the agent is indifferent between $(w_3, P; w_4)$ and $(w_4, P; w_3)$, we will have established that $bel(P) = 0.5$. If the probability of P were any other way, then the agent would not have been indifferent between $(w_3, P; w_4)$ and $(w_4, P; w_3)$. For instance, if $w_3 \succ w_4$ and $bel(P) > 0.5$, then any minimally rational agent would prefer $(w_3, P; w_4)$ to $(w_4, P; w_3)$.

With this in place, we are then able to say that $(w_1, w_2) =^{d} (w_3, w_4)$ holds just in case $(w_1, P; w_4) \sim (w_2, P; w_3)$, where P has probability ½. The reasoning behind this is not immediately obvious, but goes as follows: from the assumption of Naïve Expected Utility Theory, we have that $(w_1, P; w_4) \sim (w_2, P; w_3)$ holds if and only if:

$$des(w_1).bel(P) + des(w_4).(1 - bel(P)) = des(w_2).bel(P) + des(w_3).(1 - bel(P))$$

We have also already established that $bel(P) = 0.5 = 1 - bel(P)$, so we can drop the constant factor leaving us with:

$$des(w_1) + des(w_4) = des(w_2) + des(w_3)$$

which holds just in case:

$$des(w_1) - des(w_2) = des(w_3) - des(w_4)$$

This just states that the difference between w_1 and w_2 is equal to the difference between w_3 and w_4; so if $bel(P) = 0.5$, then $(w_1, P; w_4) \sim (w_2, P; w_3)$ if and only if $(w_1, w_2) =^d (w_3, w_4)$. Given Naïve Expected Utility Theory, by appealing to probability ½ propositions, we are able to use preferences to define an equal difference relation that depends solely on the differences in utility of the outcomes.

2.4 Locating ethically neutral propositions of probability ½

Before moving on to measuring utilities, however, Ramsey makes the following note:

> There is first a difficulty which must be dealt with; the propositions like P… which are used as conditions in the [gambles] offered may be such that their truth or falsity is an object of desire to the subject. This will be found to complicate the problem, and we have to assume that there are propositions for which this is not the case, which we shall call ethically neutral. (1931b, p. 177)

This is the entirety of what Ramsey writes regarding his motivation for introducing ethically neutral propositions.

The idea is clear enough: Naïve Expected Utility Theory is mistaken, as it fails to take into account the utility that may attach to the gamble's condition, and how the condition might influence the valuation of outcomes. Assuming that w_1 is consistent with both P and $\neg P$, it is possible that an agent might attach a different value to $(w_1 \& P)$ than to $(w_1 \& \neg P)$. These are potentially quite different states of affairs with potentially different utilities, and the truth or falsity of P might make a great deal of difference to how the outcome w_1 is valued. As a rough example, suppose that in w_1 the agent has a puppy as a pet, while in w_2 she instead keeps a kitten, and let P be *puppies spread disease but kittens do not*; plausibly, $(w_1 \& P)$ will be valued quite differently than $(w_1 \& \neg P)$, and likewise for $(w_2 \& P)$ and $(w_2 \& \neg P)$.

Instead of Naïve Expected Utility Theory, and supposing w_1, w_2, w_3, and w_4 are each compatible with the relevant propositions, instead of Naïve Expected Utility Theory we should instead have a principle that implies that:

$$(w_1, P; w_2) \geq (w_3, Q; w_4)$$

if and only if:

$$des(w_1\&P).bel(P) + des(w_2\&\neg P).(1 - bel(P)) \geq des(w_3\&Q).bel(P) + des(w_4\&\neg Q).(1 - bel(Q))$$

It is easy to see that this fact invalidates the reasoning behind both the definition of a ½ probability proposition and the definition of $=^d$, for now we can no longer say that the agent's preferences between (w_1, P; w_2) and (w_3, Q; w_4) depend on their credences in P and Q and the utilities the agent has for w_1, w_2, w_3, and w_4. Rather, they actually depend on the agent's credences in P and Q and their utilities for (w_1 & P), (w_2 & $\neg P$), (w_3 & Q), and (w_4 & $\neg Q$).

Ramsey's solution to this difficulty is the ethically neutral proposition — a kind of proposition the truth or falsity of which is of absolutely no concern to the agent. Ramsey provides us with a problematic definition of the notion, which I will discuss further in §3.2. The apparent purpose of the notion, however, is that if P is ethically neutral, then the conjunction of P with an outcome has the same utility as the outcome itself, and similarly for the conjunction of $\neg P$ and the outcome. Setting aside Ramsey's own definition, we can say that P is *ethically neutral* whenever $w \sim (w \,\&\, P) \sim (w \,\&\, \neg P)$, for any outcome $w \in \boldsymbol{W}$ that is compatible with both P and $\neg P$.[3]

It is clear what the upshot of introducing ethically neutral propositions is supposed to be: so long as we are considering gambles conditional on ethically neutral propositions, we can without risk of error apply Naïve Expected Utility Theory. Thus Ramsey (1931b, pp. 177–8) happens upon the following two definitions:

Definition 1: Ethically neutral proposition of probability ½
P is an ethically neutral proposition of probability ½ if and only if P is ethically neutral, and for some $w_1, w_2 \in \boldsymbol{W}$, $\neg(w_1 \sim w_2)$, and $(w_1, P; w_2) \sim (w_2, P; w_1)$[4]

[3] If there are no outcomes compatible with both P and $\neg P$, then P is trivially ethically neutral by this definition. Ramsey likely would have required non-trivially ethically neutral propositions, though the issues here are complex. See §3.1 and §4.3 below for discussion relating to this issue.

[4] Ramsey states that 'We assume by an axiom that if this is true of any one pair [w_1], [w_2], it is true of all such pairs' (1931b, pp. 177–8). This axiom is never actually listed, and it does not seem to follow from any of (RAM 1)–(8) either separately or in conjunction.

And:

Definition 2: Equal difference relation
$(w_1, w_2) =^{d} (w_3, w_4)$ if and only if $(w_1, P; w_4) \sim (w_2, P; w_3)$, where P is an ethically neutral proposition of probability ½

2.5 Measuring utilities
At this point (1931b, pp. 178–9), Ramsey lists eight axioms, and states (but does not prove) that their satisfaction enables an appropriately rich representation of the agent's preferences. That is:

Ramsey's representation conjecture
If (RAM 1)–(8) hold, then there exists a real-valued function *des* on $\boldsymbol{W}$ such that for all $w_1, w_2, w_3, w_4 \in \boldsymbol{W}$, $des(w_1) - des(w_2) = des(w_3) - des(w_4)$ if and only if $(w_1, w_2) =^{d} (w_3, w_4)$; furthermore, *des* is unique up to positive linear transformation

We will not consider whether Ramsey's axioms successfully ensure the desired representation result, or how they might be fleshed out to do so if not—though see Bradley (2001) for relevant work in this regard. It is at least clear that something in the vicinity of Ramsey's axioms should suffice, however.

I here reproduce Ramsey's axioms, albeit with slightly improved notation. Let $\boldsymbol{P}$ be a set of propositions, $\boldsymbol{W}$ the set of outcomes, and $\boldsymbol{G}$ the set of gambles. There are multiple options for the formalization of $\boldsymbol{G}$. One might treat $\boldsymbol{G}$ as a set of functions from pairs of mutually exclusive and exhaustive propositions to $\boldsymbol{W}$ (as in Bradley 2001, p. 273), or simply take $\boldsymbol{G}$ to be a subset of $\boldsymbol{W} \times \boldsymbol{P} \times \boldsymbol{W}$ (as I do below, §4); the differences between these options need not concern us here. $\succ$ and $\sim$ are defined on $\boldsymbol{W} \cup \boldsymbol{G}$. The very first axiom is the most distinctive aspect of Ramsey's theorem:

(RAM 1) There is at least one ethically neutral proposition of probability ½

The importance of (RAM 1) for the rest of Ramsey's formal system should not be understated. Most of the axioms to follow are stated in terms of $=^{d}$, which is defined in terms of ethically neutral propositions. If (RAM 1) is false, those axioms will be in some cases false, in others trivial; in either case, the system as a whole collapses without this foundational assumption.

The next three axioms are each obviously necessary for Ramsey's desired representation result: for all $P, Q \in \boldsymbol{P}$, $w_1, w_2, w_3, w_4, w_5, w_6 \in \boldsymbol{W}$, $(w_1, P; w_2), (w_3, P; w_4) \in \boldsymbol{G}$, and $\alpha, \beta, \gamma \in \boldsymbol{W} \cup \boldsymbol{G}$, (RAM 2) (i) If P, Q, are both ethically neutral propositions with probability ½, and $(w_1, P; w_2) \sim (w_3, P; w_4)$, then $(w_1, Q; w_2) \sim (w_3, Q; w_4)$; (ii) if $(w_1, w_2) =^{d} (w_3, w_4)$, then $w_1 > w_2$ iff $w_3 > w_4$, and $w_1 \sim w_2$ iff $w_3 \sim w_4$

(RAM 3) $\alpha \sim \beta$ and $\beta \sim \gamma$ only if $\alpha \sim \gamma$

(RAM 4) If $(w_1, w_2) =^{d} (w_3, w_4)$ and $(w_3, w_4) =^{d} (w_5, w_6)$, then $(w_1, w_2) =^{d} (w_5, w_6)$

The role of (RAM 2) is to ensure that the definition of $=^{d}$ is coherent. (RAM 3) simply states that $\sim$ is transitive, while (RAM 4) states that $=^{d}$ is transitive — an obvious requirement given the intended interpretation of $=^{d}$. Given the definition of the equal difference relation above, (RAM 4) is equivalent to the following:

If $(w_1, P; w_4) \sim (w_2, P; w_3)$ and $(w_3, P'; w_6) \sim (w_4, P'; w_5)$, then $(w_1, P''; w_6) \sim (w_2, P''; w_5)$, where P, P', and P'' are ethically neutral propositions of probability ½

Together, (RAM 2)–(RAM 4) help to ensure that $=^{d}$, which holds between pairs of outcomes, mirrors the behaviour of the equals relation between the differences of pairs of real numbers.

The following two existential axioms are stated in terms of what Ramsey calls *values* — essentially, $\sim$-equivalence classes of outcomes. Formally:

Definition 3*: The value of w*
For every outcome $w \in \boldsymbol{W}$, let $\underline{w} = \{w' \in \boldsymbol{W}: w' \sim w\}$

The value of an outcome w, denoted $\underline{w}$, is thus the set of all outcomes in $\boldsymbol{W}$ with the same desirability as w. Ramsey's next two axioms are then:

(RAM 5) For all $\underline{w_1}$, $\underline{w_2}$, $\underline{w_3}$, there is exactly one $\underline{w_4}$ such that $(w_1, w_4) =^{d} (w_2, w_3)$

(RAM 6) For all $\underline{w_1}$, $\underline{w_2}$, there is exactly one $\underline{w_3}$ such that $(w_1, w_3) =^{d} (w_3, w_2)$

(RAM 5) implies that there is always at least one outcome w_4 such that the difference between w_1 and w_4 is equal to the difference between w_2 and w_3, for *any* choice of outcomes w_1, w_2, and w_3. In a manner of speaking, (RAM 6) says that for any pair of worlds w_1 and w_2, there is a third world w_3 with a utility which is exactly half-way between the utilities of w_1 and w_2. Given (RAM 1) (which implies the non-triviality of $\succ$ on $\boldsymbol{W}$), this entails a certain denseness to the agent's preference structure, and correspondingly, that $\boldsymbol{W}$ is infinite.

Finally, Ramsey lists two other axioms, which are not spelled out in any detail:

(RAM 7) 'Axiom of continuity: — Any progression has a limit (ordinal)' (1931b, p. 179)

(RAM 8) Archimedean axiom

What Ramsey intended for (RAM 7) is something of a mystery. One guess (cf. Sobel 1998 and Bradley 2001) would be that for every gamble $(w_1, P; w_2)$, there is an outcome w_3 such that $w_3 \sim (w_1, P; w_2)$. An axiom to this effect seems to be required to ensure that every real number can be mapped to at least one world's value.

Ramsey does not specify the character of (RAM 8), though it is easy to guess its role — like other so-called Archimedean axioms in various representation theorems, it is supposed to rule out any one outcome or gamble being incomparably better or worse than another. More specifically, (RAM 8) is intended to ensure that the numerical representation satisfies the Archimedean property of real numbers: for any positive number x, and any number y, there is an integer n such that $n + x \geq y$. This ensures that we do not require infinite desirability values to measure the agent's utilities. Were one to spell out (RAM 8), it is likely that it would need to look much like (ADS 5) of *Definition 13* below.

2.6 Measuring credences

Suppose that we have our function *des*. Ramsey then argues that:

> Having thus defined a way of measuring value we can now derive a way of measuring belief in general. If the option of $[w_2]$ for certain is indifferent with that of $[(w_1 \text{ if } P; w_3)]$, we can define the subject's degree of belief in P as the ratio of the difference between $[w_2]$ and $[w_3]$ to that between $[w_1]$ and $[w_3]$ … This amounts roughly to defining the degree of belief in P by the odds at which the subject would bet on P, the bet being conducted in terms of differences of value as defined. (1931b, pp. 179–80)

In a footnote, Ramsey adds that '$[w_1]$ must include the truth of P, $[w_3]$ its falsity; P need no longer be ethically neutral' (1931b, p. 179, Fn. 1). We are led to the following definition:

Definition 4: Ramsey's credence function
For all contingent propositions P and outcomes w_1, w_2, w_3 such that w_1 implies P, w_3 implies $\neg P$, $\neg(w_1 \sim w_3)$, and $w_2 \sim (w_1, P; w_3)$, $bel(P) = (des(w_2) - des(w_3)) \,/\, (des(w_1) - des(w_3))$

This definition does not require P to be ethically neutral. Ramsey mistakenly states that it 'only applies to partial belief and does not include certain beliefs' (1931b, p. 180), though plausibly he meant that the definition does not apply if P is non-contingent. In this case, we simply stipulate that $bel(P) = 1$ if P is necessary, 0 if P is impossible. Note that, because ratios of differences are preserved across positive linear transformations of the *des* function, $bel(P)$ so-defined is unique.

The reasoning behind this final step is again left to the reader, though again it follows from Ramsey's background assumption of the descriptive adequacy of subjective expected utility theory. Note first of all that if w_1 entails P, then the conjunction of P and w_1 is just equivalent to w_1, so $des(w_1) = des(w_1 \,\&\, P)$. Ramsey clearly presupposes, then, that the desirability of a prospect $(w_1, P; w_2) \sim w_3$, where w_1 entails P and w_2 entails $\neg P$, is given by the following equality:

$$des((w_1, P; w_2)) = des(w_1).bel(P) + des(w_2).(1 - bel(P)) = des(w_3)$$

This is then rearranged to give us the definition of $bel(P)$ as above.

Ramsey does note that for his definition to work we require two more assumptions. The first of these is that the value of $bel(P)$ does not depend on the choice of worlds and gambles satisfying the stated conditions. In effect, this is to place restrictions directly upon *bel* after it has been defined in terms of preferences. (See *Condition 1*, §4.4 below.) The second assumption is that for any gamble $(w_1, P; w_2)$ we will always be able to find some world w_3 such that $w_3 \sim (w_1, P; w_2)$. Both of these assumptions are clearly required if *bel* is to be coherently defined for each contingent proposition.

Ramsey (1931b, pp. 180ff) goes on to define conditional probabilities using preferences over more complicated gambles, and he argues that *bel* satisfies the laws of probability, though I will not recapitulate that argument here: it is enough that Ramsey provides a credence function, bel: $\boldsymbol{P} \mapsto [0,1]$, that supposedly represents the agent's degrees of belief — after all, it combines with the agent's utilities for outcomes to determine their preference ordering for two-outcome gambles in more or less the manner we pre-theoretically expect degrees of belief to. For our purposes, it is incidental whether *bel* satisfies the axioms of the probability calculus.

In summary, Ramsey's proposal is as follows. We begin with the background assumption of the empirical adequacy of subjective expected utility theory, and an empirically determined preference ordering on the set of outcomes and two-outcome gambles. Then, referring specifically to gambles conditional on an ethically neutral proposition of probability ½ that is assumed to exist (and determined through observation of the agent's preferences), we construct a function *des* that measures the agent's utilities. Having determined *des*, and given some further assumptions, we use *des* to determine the function *bel*. It is then argued that *bel* has certain attractive properties that lend support to its interpretation as a measure of the agent's degrees of belief.

3. Ethical neutrality

3.1 Why Ramsey introduced ethical neutrality

Ramsey was right to reject Naïve Expected Utility Theory. If the outcomes w_1 and w_2 are each compatible with both P and $\neg P$, then it is entirely possible that the agent values (w_1 & P) more (or less) than (w_1 & $\neg P$), and similarly for (w_2 & P) and (w_2 & $\neg P$). Any rational agent ought to take this into account when deliberating between gambles conditional on P with w_1 and w_2 as outcomes. For example, contrary to Naïve Expected Utility Theory, it is possible that the agent could be indifferent between $w_1 \sim w_2$ without thereby being indifferent between (w_1, P; w_2) and (w_2, P; w_1), if the truth or falsity of P makes a difference to how the agent values w_1 or w_2.

However, this point is conditional on w_1 and w_2 being each compatible with both P and $\neg P$. If instead we suppose that w_1 implies P, then (w_1 & P) is logically equivalent to w_1—and for Ramsey, if w_1 implies P, then the desirability of w_1 is just the desirability of (w_1 & P).[5] Ramsey's characterisation of the *bel* function relies on this very fact.[6]

[5] If w_1 implies $\neg P$, then (w_1 & P) = $\varnothing$. It is unclear whether Ramsey intended his utility function to assign any value to $\varnothing$, as his definition of the credence function only requires the utility function to be defined on $\boldsymbol{W} \cup \boldsymbol{G}$, which does not contain $\varnothing$.

[6] While descriptively questionable, I will not discuss this idealizing assumption in any further detail. Such assumptions are standardly made in many descriptive decision theories and measurement systems—if it is a problem for Ramsey, then it is a problem for a great deal of the work in this area.

So, inasmuch as w_1 implies P and w_2 implies $\neg P$,

$$\begin{aligned} des((w_1, P; w_2)) &= des(w_1 \& P).bel(P) + des(w_2 \& \neg P).(1 - bel(P)) \\ &= des(w_1).bel(P) + des(w_2).(1 - bel(P)) \end{aligned}$$

Note that this holds regardless of whether P is ethically neutral. In other words, if w_1 implies P and w_2 implies $\neg P$, then we can apply Naïve Expected Utility Theory to the gamble $(w_1, P; w_2)$.

Interestingly, Ramsey originally describes his outcome set $\boldsymbol{W}$ as a set of possible worlds, and it is part of Ramsey's background theory that every world individually determines the truth or falsity of any proposition. In particular, Ramsey assumed a broadly Wittgensteinian logical atomism (though he believed it possible to reformulate his theorem without these commitments; see his 1931b, p. 177, fn. 1). Ramsey's discussion does not commit him to the entirety of Wittgenstein's theory. Instead, Ramsey assumes what Sobel (1998) calls a *thin* logical atomism, in the following sense: we are to suppose that there exists a class of atomic propositions such that no two worlds are exactly identical with respect to the truth of these propositions, every atomic proposition can be true or false entirely independently of any others, and for every world w and atomic proposition P, there is another world w^* that differs only with respect to the truth of P. Every possible world on this picture is determined by the set of atomic propositions true at that world. Even setting aside the assumption of logical atomism, on an orthodox conception of propositions as sets of worlds, for any given (determinate) proposition, a given world either is or is not a member of that proposition. Every world therefore determines either the truth or falsity of any proposition, and no possible world is compatible with both P and $\neg P$.

This leaves us with something of a puzzle: why did Ramsey alter his characterisation of the outcome set (as noted in §2.2)? It seems that if he instead limited his attention to gambles like $(w_1, P; w_2)$, where w_1 implies P and w_2 implies $\neg P$, then he could have used preferences over *these* to define $=^{\mathrm{d}}$ without needing to introduce the notion of ethical neutrality. Let us call any gamble $(w_1, P; w_2)$ *impossible* if either w_1 implies $\neg P$ or w_2 implies P. A gamble is *possible* if and only if it is not impossible. If outcomes are possible worlds, then every possible gamble $(w_1, P; w_2)$ must be such that w_1 implies P and w_2 implies $\neg P$, and so $des(w_1) = des(w_1 \ \& \ P)$ and $des(w_2) = des(w_2 \ \& \ \neg P)$. We can therefore always apply Naïve Expected Utility Theory to possible gambles if the outcomes are worlds. So why did Ramsey not stick to his original characterisation of outcomes as worlds, and simply use preferences over possible gambles to define $=^{\mathrm{d}}$?

The answer to this question can be discovered by considering again how Ramsey defines ½ probability propositions. In particular, to determine whether P has a probability ½, we need to consider preferences over two gambles of the form $(w_1, P; w_2)$ and $(w_2, P; w_1)$. The definition Ramsey gives us *only* makes sense if the outcomes w_1 and w_2 are *not* possible worlds. If w_1 and w_2 are possible worlds, then at least one of the two gambles is impossible, and if either gamble is impossible, then the reasoning behind the assignment of probability ½ to the proposition P is no longer valid. Indeed, Ramsey recognized the difficulty here, and for this reason wrote that, at least for those outcomes w_1 and w_2 required for his definition of ½ probability propositions, '$[w_1]$ and $[w_2]$ must be supposed so far undefined as to be compatible with both P and $\neg P$' (1931b, p. 178, Fn. 1). Supposing for simplicity that P is atomic, we are presumably to take the outcomes w_1 and w_2 as *near-worlds*, which we can understand as propositions that are just shy of being maximally specific. Given thin logical atomism, for every world w and every atomic proposition P, there is a proposition that *nearly* uniquely identifies w but does not specify whether P is true. In Ramsey's framework, a near-world with respect to an atomic proposition P is a disjunction of two worlds w^{P} and $w^{\neg\mathrm{P}}$ that are identical with respect to all of their atomic propositions except for P.

The answer to our puzzle, then, is that Ramsey's set of outcomes cannot quite be the set of possible worlds given his strategy for defining $=^{\mathrm{d}}$. For the pair of possible gambles $(w_1, P; w_2)$ and $(w_2, P; w_1)$ referred to in *Definition 1*, neither w_1 nor w_2 can imply either P or $\neg P$. It follows for the reasons given, then, that we cannot in general apply Naïve Expected Utility Theory to such gambles, unless we appeal to ethically neutral propositions.[7]

[7] An anonymous referee suggests an alternative interpretation of what it is for w_1 and w_2 to be 'so far undefined as to be compatible with both P and $\neg P$': outcomes are in all cases possible worlds, while $(w_1, P; w_2)$ should be understood as (i) having outcome w_1 if P and w_1 implies P, or w_1' if P and w_1 implies $\neg P$, and (ii) likewise, w_2 if $\neg P$ and w_2 implies $\neg P$, or w_2' if $\neg P$ and w_2' implies P, where w_1' and w_2' are like w_1 and w_2 respectively in all atomic details except P. On this interpretation, Ramsey kept his original characterisation of outcomes as worlds, introducing considerable complexity into his understanding of a gamble instead.

It is clear why the notion of ethical neutrality is needed on this picture: if P is ethically neutral then $w_1 \sim w_1'$ and $w_2 \sim w_2'$, which would justify applying Naïve Expected Utility Theory to $(w_1, P; w_2)$ and $(w_2, P; w_1)$. On either interpretation, then, the main upshot of the present discussion is retained: Ramsey needed to introduce ethical neutrality largely as a result of his strategy for defining $=^{\mathrm{d}}$.

3.2 Problems with ethical neutrality

In looking at whether the notion of ethical neutrality is viable, we ought first to start with Ramsey's own definition:

> *Definition 5: Ethical neutrality (Ramsey's original)*
> P is ethically neutral if and only if (i) if P is atomic, then $w^{P} \sim w^{\neg P}$, for all pairs of worlds w^{P}, $w^{\neg P}$ identical with respect to all their atomic propositions except for P, (ii) if P is non-atomic, then all of Ps atomic truth arguments are ethically neutral

So, an atomic proposition P is *ethically neutral* for an agent just in case any two possible worlds differing in their atomic propositions only in the truth of P are always equally valued by that agent, and ethical neutrality for non-atomic propositions is understood in terms of atomic propositions. Ramsey here demonstrates commitment to another aspect of Wittgensteinian atomism: every non-atomic proposition can be constructed from atomic propositions using truth-functional connectives.[8] We are able to locate such a proposition, if it exists, by considering the agent's preferences over worlds. As just noted, for some gambles $(w_1, P; w_2)$ and $(w_2, P; w_1)$, Ramsey requires that w_1 and w_2 are compatible with both P and $\neg P$. If we suppose for simplicity that P is atomic, then w_1 and w_2 are near-worlds with respect to P. It follows from Ramsey's definition then that $(w_1 \;\&\; P) \sim (w_1 \;\&\; \neg P)$ and $(w_2 \;\&\; P) \sim (w_2 \;\&\; \neg P)$. It does *not* yet follow that $(w_1 \;\&\; P) \sim (w_1) \sim (w_1 \;\&\; \neg P)$, which Ramsey also required. However, we can take this as an unstated background assumption: if $(w_1 \;\&\; P) \sim (w_1 \;\&\; \neg P)$, then $(w_1 \;\&\; P) \sim (w_1) \sim (w_1 \;\&\; \neg P)$.

Sobel (1998, p. 241) has argued that there are few or no ethically neutral propositions in this sense. Consider the proposition *there are an even number of hairs on Dan Quayle's head*. Sobel argues that this can be ethically neutral for 'almost no one':

> Though it is true that I do not care about Quayle's hair, there are worlds that differ regarding the truth of that proposition that, just because of that difference, differ in their values for me. I am thinking of worlds in which I have bet money on this proposition! The argument … can be readdressed to atomic propositions, if such there be, to the conclusion that *no* atomic proposition is

[8] Interestingly, Ramsey's definition suggests that ethically neutral propositions can be logically equivalent to non-ethically neutral propositions. Suppose that P is ethically neutral, but Q is not. Then P is logically equivalent to $(P \;\&\; (Q \vee \neg Q))$, but $(P \;\&\; (Q \vee \neg Q))$ is not ethically neutral. Thanks to Rachael Briggs for noting this.

> Ramsey-ethically-neutral for any of us. Ramsey's existence axiom for ethically neutral atomic propositions, even if coherent, severely curtails the applicability of his theory. (Sobel 1998, p. 248, emphasis in original)

There seem to be two concerns here. The first appears to be something like the following: for any proposition whatsoever, we should be able to find a set of otherwise similar possible worlds where we have entered into a bet conditional on that proposition with desirable outcomes if things turn out one way, and undesirable outcomes if things turn out another way. Since we care about the outcomes of the bet, we will value the relevant worlds differently. However, this objection seems to have no hold given Ramsey's view: the relevant worlds are supposed to differ at the atomic level *only* with respect to the proposition in question. In all other respects — including, importantly, the payouts for any bets we may enter into — the worlds are supposed to be identical.

The second and more obvious worry is that Ramsey's conception of ethical neutrality requires the assumption of logical atomism for its cogency. Ramsey built his theory upon the assumption of logical atomism so that he could make sense of the idea of two worlds differing *only* with respect to a particular proposition. The notion is of little use to contemporary philosophers who by and large reject that aspect of Wittgenstein's view.

In his atomism-free reconstruction of Ramsey's theorem, Bradley (2001) supplies the following definition, intended to achieve the same purpose:

> *Definition 6: Ethical neutrality (atom-free)*
> P is ethically neutral if and only if for all propositions Q that are compatible with both P and $\neg P$, $(P \,\&\, Q) \sim Q \sim (\neg P \,\&\, Q)$

Tautological and impossible propositions will be trivially ethically neutral according to this definition. Clearly, however, we are interested only in non-trivially ethically neutral propositions. A common suggestion is that propositions such as *the tossed coin will land heads* constitute ethically neutral propositions of probability ½. Part of the reason why we use coin tosses occasionally when making decisions is because we have no intrinsic interest in whether the coin lands heads or tails. If Q is something like *there are dogs*, and P is *the tossed coin will land heads*, then it seems plausible that $(P \,\&\, Q) \sim Q \sim (\neg P \,\&\, Q)$.

However, there are strong reasons to think that *no* contingent propositions will be ethically neutral in the sense of *Definition 6*, for any minimally rational subject. Let P be *the tossed coin will land heads*, and take Q to be the proposition:

(*the tossed coin will land heads* & *I receive $100,000*) *or* (*the tossed coin will not land heads* & *I get kicked in the shins*)

Q is obviously compatible with both P and $\neg P$. However, (P & Q) is equivalent to *the tossed coin will land heads* & *I receive $100,000* while ($\neg P$ & Q) is equivalent to *the tossed coin will not land heads* & *I get kicked in the shins*. But for some very strange preference orderings, it is certainly not the case that (P & Q) $\sim Q \sim$ ($\neg P$ & Q). The point here generalizes easily; there are no non-trivially ethically neutral propositions in this sense. Note that the issue here is not that no contingent proposition satisfies the definition *exactly*, while there may nevertheless be some propositions which *approximate* ethical neutrality. Rather, the upshot is that no proposition even comes *close* to satisfying the requirements of ethical neutrality. We will always be able to find countlessly many propositions Q that falsify the indifference requirements.[9]

A refinement of *Definition 6* might be useful here. Instead of requiring (P & Q) $\sim Q \sim$ ($\neg P$ & Q) for all *propositions* Q compatible with both P and $\neg P$, Ramsey only requires the following:

Definition 7: Ethical neutrality (atom-free, refined)
P is ethically neutral if and only if $w \sim (w \,\&\, P) \sim (w \,\&\, \neg P)$, for any outcome $w \in \boldsymbol{W}$ that is compatible with both P and $\neg P$

If there are no outcomes compatible with both P and $\neg P$, then P is trivially ethically neutral by this definition. Again, we can set such propositions aside; we are interested in non-trivially ethically neutral propositions. (See §4.3 for more discussion.) If Q is not in the outcome set $\boldsymbol{W}$, then there are no relevant gambles with Q as an outcome and we do not need to concern ourselves over whether (P & Q) $\sim Q \sim$ ($\neg P$ & Q). If we assume that there are far fewer propositions in $\boldsymbol{W}$ than in $\boldsymbol{P}$, then the foregoing objection is blocked. This will certainly be true if the outcomes in $\boldsymbol{W}$ are highly specific, as is the case in Ramsey's system.

With that said, it is still not obvious that any non-trivially ethically neutral propositions exist even in this weaker sense. Why should we

[9] Thanks to Rachael Briggs and Alan Hájek for these points and related discussion.

suppose that there are *any* propositions P such that (non-trivially), $w \sim (w \,\&\, P) \sim (w \,\&\, \neg P)$ for all $w \in \boldsymbol{W}$ compatible with P and $\neg P$? Indeed, without knowing the exact nature of the outcome space $\boldsymbol{W}$, we cannot even know whether there *are* any outcomes compatible with both P and $\neg P$, for an arbitrarily chosen proposition P. Ramsey explicitly stipulates that there must be at least one pair of outcomes compatible with *some* ½ probability ethically neutral proposition and its negation — but this stipulation is meaningless inasmuch as we do not already know what proposition that may be. Unfortunately, Ramsey's discussion leaves the nature of $\boldsymbol{W}$ quite vague, making the matter impossible to judge.

We can circumvent this concern by stipulating that $\boldsymbol{W}$ contains, for each of a very wide range of propositions in $\boldsymbol{P}$, highly specific outcomes that are undefined with respect to that proposition. But even then, Ramsey gives us little reason to suppose that ethically neutral propositions exist relative to a given agent's preference ordering — still less to suppose that there are any such propositions that satisfy *Definition 1*. (RAM 1) clearly cannot be defended as a condition of rationality, and it does not follow from Ramsey's background assumption of the descriptive adequacy of expected utility theory. Ramsey's aim in the first instance was to develop a procedure for the measurement of degrees of belief, so unlike other intended uses for decision theoretic representation theorems he did not require his axioms to be constraints of practical rationality; nevertheless, if the process is to be viable then it ought at least to be *applicable*. It may not be impossible for a rational agent to satisfy the axiom, but we still require good reasons to believe that most do — yet no reasons are forthcoming.

A related issue regards Ramsey's proto-functionalist attempt to define degrees of belief in terms of his measurement procedure: a definition of credences which relies on a dubitable and unjustified existential assumption is, at best, of very limited interest. Are we to suppose that agents who falsify (RAM 1) do not have degrees of belief? Ultimately, given his reliance upon ethically neutral propositions, Ramsey's system was not sufficient to establish the main upshot of 'Truth and Probability', namely, that the laws of probability provide for us the logic of partial belief. Even if it is understood in terms of *Definition 7*, (RAM 1) is a very shaky foundation for a measurement procedure, and still worse for a definition of credences.

Many expected utility representation theorems developed since Ramsey's original work have also made use of ethically neutral propositions, either explicitly or implicitly. Davidson and Suppes (1956) develop a representation theorem similar to Ramsey's wherein they explicitly characterize and axiomatize the existence of ethically neutral propositions. Others make implicit appeal to ethically neutral propositions, in the sense that they figure in the intended interpretation of the formal system, rather than being formalized directly. In this capacity we find ethical neutrality in the theorems of Davidson, Siegel, and Suppes (1957), Debreu (1959), Suppes (1956), and Fishburne (1967). Each of these works appears to require an understanding of ethical neutrality in something like the senses of *Definition 6* or *7* (each for essentially the same reason that Ramsey required the notion), and thus they inherit the problems associated with those definitions. We need a firmer foundation for the measurement of credences.

4. Refining Ramsey's system

4.1 Preliminaries

Ramsey's motivation for introducing the idea of ethical neutrality arises ultimately from his strategy for defining ½ probability propositions and $=^{d}$. But we are not forced to use Ramsey's definitions. It is possible to avoid introducing ethical neutrality in any of the senses defined so far, if we can develop alternative definitions.

One obvious possibility would be to introduce a second qualitative relation into our formal system, à la Joyce (1999). Let $=^{b}$ be a *qualitative probability relation* defined on the set of propositions: $P =^{b} Q$ if and only if the agent judges P to be exactly as likely as Q. Now we might say that a proposition P has probability ½ just in case $P =^{b} \neg P$. Developing this line of thought would lead us to a two-primitive (*preference* and *qualitative probability*) representation theorem, though it is also likely to render much of Ramsey's approach redundant. For instance, if the agent has a sufficiently rational qualitative probability ranking $\geq^{b}$ on the set of propositions, then we can apply something like de Finetti's (1931) representation theorem to arrive at a cardinal representation of her credences without needing to consider her preferences over gambles at all.

Let us suppose, then, that we have no straightforward empirical access to agents' qualitative probability rankings. We are then in Ramsey's original position, aiming to develop a system for the measurement of credences where we only have empirical access to agents' preferences over outcomes and gambles.

Another possibility would be to construct a *relative* notion of ethical neutrality. Ramsey's definition of $=^{\mathrm{d}}$ in terms of gambles does not need ethical neutrality *tout court*; instead, the conditioning proposition need only be ethically neutral with respect to the outcomes involved in the relevant gambles. That is, for Ramsey, $(w_1, w_2) =^{\mathrm{d}} (w_3, w_4)$ if and only if $(w_1, P; w_4) \sim (w_2, P; w_3)$, where P is an ethically neutral proposition of probability ½, but here it is enough if P is *neutral with respect to* w_1, w_2, w_3, and w_4. Sobel (1998, pp. 268–9) suggests a modification to Ramsey's system along these lines, which forces the replacement of (RAM 1) with:

(RAM 1′) For every quartet of worlds there is a proposition believed to degree ½ that is ethically neutral with respect to each world in the quartet

The difficulty here then becomes spelling out the notion of relative ethical neutrality in such a manner as to render (RAM 1′) plausible. Sobel's own suggested definition also presupposes logical atomism, as it was intended to fit with the rest of Ramsey's system. A more straightforward strategy, however, is to bypass the issue of ethical neutrality altogether; we can have an essentially Ramseyian representation theorem without needing to first construct a notion of ethical neutrality at all, rendering attempts to define the notion pointless.

Before we move on, I will note two assumptions that underlie the rest of the discussion. First, I will follow Ramsey's initial suggestion and take the space of outcomes, $\boldsymbol{W}$, to be the set of all possible worlds.[10] This assumption is for simplicity only; nothing of importance to the formal result is altered if we take outcomes to be propositions. (See the discussion in §4.3.) Secondly, and far more importantly, I will limit my attention to *possible* gambles only. I am doubtful that it makes much sense to assert that an agent can have interesting preferences with respect to impossible gambles, and in any case their inclusion comes with a great deal of added complexity. Preferences over impossible gambles play no role in deriving the

[10] I leave open the kind of modality being employed here. One might take $\boldsymbol{W}$ to be the set of metaphysically possible worlds, though this may lead to some counterintuitive results. My preference is to construe $\boldsymbol{W}$ as the set of epistemically possible worlds; in the sense of Chalmers (2011).

forthcoming representation result. I axiomatize this limitation in (GRS 1) below.

The key idea of the theorem below is that, while Ramsey used the same *outcomes* in two different gambles to define what it is for a proposition to have probability ½, this is unnecessary — it is enough if we instead use outcomes with exactly the same *desirability*. Suppose that $w_1 \sim w'_1$, $w_2 \sim w'_2$, and $\neg(w_1 \sim w_2)$, but:

$$(w_1, P; w_2) \sim (w'_2, P; w'_1)$$

We assume that each of these two gambles is possible; that is, w_1 and w'_2 each imply P, and w_2 and w'_1 each imply $\neg P$. Given the Ramseyian background assumption, this is possible only if:

$$des(w_1).bel(P) + des(w_2).(1 - bel(P)) = des(w'_2).bel(P) + des(w'_1).(1 - bel(P))$$

Since $w_1 \sim w'_1$ and $w_2 \sim w'_2$, we know that $des(w_1) = des(w'_1) = x$ and $des(w_2) = des(w'_2) = y$; and because $\neg(w_1 \sim w_2)$, we know that $x \neq y$. Let $bel(P) = z$. We are left with:

$$xz + y(1 - z) = yz + x(1 - z)$$

Regardless of the specific values of x and y, this is possible only if $z = (1 - z)$; thus, $bel(P) = 0.5$. This is essentially the same reasoning that Ramsey used in developing his definition of ½ probability propositions, with the only difference being that we have appealed to equally valued outcomes rather than using the same outcomes for the two gambles. Since $(w_1, P; w_2)$ and $(w'_2, P; w'_1)$ are both possible gambles with worlds as outcomes, we can apply Naïve Expected Utility Theory. There is no reason to require that P is ethically neutral. A set Π of ½ probability propositions is characterized precisely in *Definition 8* below.

Making this modification, in light of the two assumptions noted above, forces a number of changes to Ramsey's axiomatisation. There are two particularly important changes that I will note here, before laying out the main theorem in full. First, we can no longer employ Ramsey's definition of $=^{\mathrm{d}}$. (Instead of defining $=^{\mathrm{d}}$, I will instead define $\geq^{\mathrm{d}}$.) However, here we can employ the same trick: there is no reason why $(w_1, w_2) \geq^{\mathrm{d}} (w_3, w_4)$ must be defined using the worlds w_1, w_2, w_3, and w_4. It is enough if we use worlds with exactly the same desirability. Furthermore, there is no reason why we need to use the same ½ probability proposition in both gambles. Instead, we can say that $(w_1, w_2) \geq^d (w_3, w_4)$ holds just in case, if $w_1 \sim w_1'$, $w_2 \sim w_2'$, $w_3 \sim w_3'$, $w_4 \sim w_4'$, and both $(w_1', P; w_4')$ and $(w_2', P'; w_3')$ are possible gambles where P and P' both have a probability of ½, then:

$$(w'_1, P; w'_4) \geq (w'_2, P'; w'_3)$$

The reasoning behind this is essentially identical to that outlined in §2.3. Note that w_1 need not be distinct from w_1'; the only thing we require is that w_1' implies P (and similarly for the other outcomes, *mutatis mutandis*). $\geq^d$ is characterized precisely in *Definition 9* below.

Secondly, we need to ensure that there are enough worlds for this definition of $\geq^d$ to work. That is, we need to assume that we will always be able to find the required gambles $(w_1', P; w_4')$ and $(w_2', P';$ $w_3')$. This is not obviously going to be the case, if we are limiting our attention to possible gambles only. In effect, we need to assume that for every pair of worlds w_1 and w_2, there will always exist at least one proposition P of probability ½ such that for some worlds $w_1' \sim w_1$ and $w_2' \sim w_2$, w_1' implies P and w_2' implies $\neg P$. This is a reasonably strong assumption. It requires that every value contains multiple members, and that at least two of these members will disagree with respect to P for some P of probability ½. In effect, this assumption replaces Ramsey's axiom (RAM 1). It is formalized as (GRS 2) below.

4.2 Main representation theorem

In this section I reproduce the main formal result of the paper: a representation theorem for the construction of an interval scale *des* on the set of gambles $\boldsymbol{G}$ and outcomes $\boldsymbol{W}$ such that for all α, $\beta \in \boldsymbol{W} \cup \boldsymbol{G}$ and all $w_1, w_2, w_3, w_4 \in \boldsymbol{W}$,

$\alpha \succcurlyeq \beta$ if and only if $des(\alpha) \geq des(\beta)$
$(w_1, w_2) \geq^d (w_3, w_4)$ if and only if $des(w_1) - des(w_2) \geq des(w_3) - des(w_4)$

In §4.4 I will introduce some further conditions that suffice for the construction of a credence function *bel*. Assume, in all that follows:

$\boldsymbol{W}$ is a non-empty set of possible worlds
$\boldsymbol{P}$ is an algebra of sets on $\boldsymbol{W}$ (i.e., a set of propositions)
$\boldsymbol{G} \subseteq \boldsymbol{W} \times \boldsymbol{P} \times \boldsymbol{W}$
$\succcurlyeq$ is a binary relation defined on $\boldsymbol{W} \cup \boldsymbol{G}$; $\succ$ and $\sim$ are defined in the usual way

Note that $\boldsymbol{P}$ need not contain all sets of worlds. For instance, the system to be developed is consistent with supposing that $\boldsymbol{P} = \{\varnothing, \boldsymbol{W}, P, \neg P\}$, where P (or $\neg P$) has probability ½. Given (GRS 1), $\boldsymbol{G}$ will be restricted to all and only possible gambles.[11]

In the sequel, I have adopted the notational convention that sameness of subscript for outcomes implies sameness of desirability (but the reverse need not hold). For instance, it should be assumed in all that follows that w_1' and w_1'' each refer to outcomes with the same desirability as w_1 (i.e., $w_1 \sim w_1'$ and $w_1' \sim w_1''$). It should not be assumed, however, that either w_1' or w_1'' is necessarily distinct from w_1. Likewise, $(w_1, P; w_2)$ should be understood as a variable for gambles with outcome w_1 if P, w_2 otherwise; and $(w_1', P; w_2')$ for gambles conditional on P with outcomes *equal in value* to w_1 and w_2. Again, the pair $(w_1, P; w_2)$ and $(w_1', P'; w_2')$ need not be distinct.

We first define the set of ½ probability propositions:

Definition 8*: The set of ½ probability propositions,* Π
$\Pi = \{P \in \boldsymbol{P}$: there are $w_1, w_2 \in \boldsymbol{W}$ such that $(w_1, P; w_2), (w_2', P; w_1') \in \boldsymbol{G}$, $\neg(w_1 \sim w_2)$, and $(w_1, P; w_2) \sim (w_2', P; w_1')\}$

Henceforth, I will use π, π', and so on, to designate propositions within Π. It should not be assumed that $\pi \neq \pi'$. Given this, I will use $(w_1, \pi; w_2)$ specifically for gambles conditional on *some* π in Π (with outcomes w_1 and w_2).

We can now define $\geq^{d}$:

Definition 9*: The difference between* w_1*,* w_2 *is at least as much as between* w_3*,* w_4
$(w_1, w_2) \geq^{d} (w_3, w_4)$ if and only if $(w_1', \pi; w_4') \succcurlyeq (w_2', \pi'; w_3')$ for all $(w_1', \pi; w_4'), (w_2', \pi'; w_3') \in \boldsymbol{G}$

For the purposes of characterizing the Archimedean axiom, we will also need to define a strictly bounded standard sequence. We can break this notion down into two concepts:

Definition 10*: Standard sequence*
$w_1, w_2, \ldots, w_i, \ldots$ is a standard sequence if and only if (i) for all $(w_2', \pi; w_1'), (w_1'', \pi'; w_1''') \in \boldsymbol{G}$, $\neg((w_2', \pi; w_1') \sim (w_1'', \pi'; w_1'''))$, and (ii) for every w_i, w_{i+1} in the sequence, $(w_{i+1}', \pi; w_2') \sim (w_1', \pi'; w_i')$ for all $(w_{i+1}', \pi; w_2'), (w_1', \pi'; w_i') \in \boldsymbol{G}$

[11] Note that $(w_1, P; w_2)$ and $(w_2, \neg P; w_1)$ are distinct elements of $\boldsymbol{W} \times \boldsymbol{P} \times \boldsymbol{W}$; here I leave it undetermined whether they represent distinct objects of preference. See §4.4, *Theorem 4*.

In light of the axioms to be characterized shortly, it will turn out that w_1, $w_2, \ldots, w_i, \ldots$ is a standard sequence just in case (i) $(w_2, w_1) \neq^d (w_1, w_1)$ and (ii) $(w_{i+1}, w_i) =^d (w_2, w_1)$ for all w_i, w_{i+1} in the sequence. So, for instance, the sequence w_1, w_2, w_3, w_4 is a standard sequence just in case:

$$(w_2, w_1) \neq^d (w_1, w_1) \text{ and } \left(w_4, w_3\right) =^d \left(w_3, w_2\right) =^d (w_2, w_1)$$

The idea, of course, is that the (nonzero) difference in desirability between any two adjacent members in the sequence is always equal to the difference in desirability between any other two adjacent members.

Definition 11*: Strictly bounded standard sequence*
$w_1, w_2, \ldots, w_i, \ldots$ is a strictly bounded standard sequence if and only if $w_1, w_2, \ldots, w_i, \ldots$ is a standard sequence and there exists w_a, $w_b \in \boldsymbol{W}$ such that for all w_i in the sequence, $(w'_a, \pi; w'_i) \succ (w'_1, \pi'; w'_b)$ and $(w''_i, \pi''; w''_b) \succ (w''_a, \pi'''; w''_1)$, for all $(w'_a, \pi; w'_i)$, $(w'_1, \pi'; w'_b)$, $(w''_i, \pi''; w''_b)$, $(w''_a, \pi'''; w''_1) \in \boldsymbol{G}$

In other words, any standard sequence $w_1, w_2, \ldots, w_i, \ldots$ is strictly bounded if there are $w_a, w_b \in \boldsymbol{W}$ such that for any w_i in the sequence, $(w_a, w_b) >^d (w_i, w_1) >^d (w_b, w_a)$. Essentially, regardless of the size of the interval between w_i and w_1, we can find outcomes in $\boldsymbol{W}$ that are spaced even further apart.

The coherence of the foregoing definitions will be ensured by the axioms (GRS 1)–(9), which we can now specify:

Definition 12*: Generalized Ramsey structure*
$<\boldsymbol{W}, \boldsymbol{P}, \boldsymbol{G}, \succcurlyeq>$ is a *generalized Ramsey structure* if and only if $\boldsymbol{W}$ is non-empty, $\boldsymbol{P}$ is an algebra of sets on $\boldsymbol{W}$, $\boldsymbol{G} \subseteq \boldsymbol{W} \times \boldsymbol{P} \times \boldsymbol{W}$, $\succcurlyeq$ is a binary relation on $\boldsymbol{W} \cup \boldsymbol{G}$, and for all $w_1, w_2 \in \boldsymbol{W}$, all sequences $w_1, w_2, \ldots, w_i, \ldots \in \boldsymbol{W}$, all $P \in \boldsymbol{P}$, and all $(w_1, P; w_2)$, $(w_1, \pi; w_2)$, $(w'_2, \pi'; w'_1)$, $(w_1, \pi; w_4)$, $(w_2, \pi'; w_3)$, $(w_3, \pi''; w_6)$, $(w_4, \pi'''; w_5) \in \boldsymbol{G}$, the following nine conditions hold:

(GRS 1) $(w_1, P; w_2) \in \boldsymbol{G}$ iff $w_1, w_2 \in \boldsymbol{W}$, $P \in \boldsymbol{P}$, and $w_1 \in P$, $w_2 \in \neg P$

(GRS 2) For every pair of worlds $w_1, w_2 \in \boldsymbol{W}$, there exists a $\pi \in \Pi$ such that for some $w'_1, w'_2 \in \boldsymbol{W}$, $w'_1 \in \pi$ and $w'_2 \in \neg\pi$

(GRS 3) $< \boldsymbol{W} \cup \boldsymbol{G}, \succcurlyeq>$ is a weak order

(GRS 4) If $(w_1, \pi; w_2)$, $(w'_2, \pi'; w'_1) \in \boldsymbol{G}$, then $(w_1, \pi; w_2) \sim (w'_2, \pi'; w'_1)$

(GRS 5) If $(w_1, \pi; w_4) \succcurlyeq (w_2, \pi'; w_3)$ and $(w_3, \pi''; w_6) \succcurlyeq (w_4, \pi'''; w_5)$, then, for all $(w_1', \pi^*; w_6'), (w_2', \pi^+; w_5') \in \boldsymbol{G}$, $(w_1', \pi^*; w_6') \succcurlyeq (w_2', \pi^+; w_5')$

(GRS 6) For every triple $w_1, w_2, w_3 \in \boldsymbol{W}$, there is a $w_4 \in \boldsymbol{W}$ such that for some $(w_1', \pi; w_3'), (w_4, \pi'; w_2') \in \boldsymbol{G}$, $(w_1', \pi; w_3') \sim (w_4, \pi'; w_2')$

(GRS 7) If $w_1, w_2, \ldots, w_i, \ldots$ is a strictly bounded standard sequence, then it is finite

(GRS 8) $w_1 \succcurlyeq w_2$ iff for all $(w_1', P; w_2') \in \boldsymbol{G}$, $w_1 \succcurlyeq (w_1', P; w_2') \succcurlyeq w_2$

(GRS 9) For each $(w_1, P; w_2) \in \boldsymbol{G}$, there is a $w_3 \in \boldsymbol{W}$ such that $(w_1, P; w_2) \sim w_3$

We can now state the main representation theorem:

Theorem 1*: Generalized Ramseyian utility measurement*
If $<\boldsymbol{W}, \boldsymbol{P}, \boldsymbol{G}, \succcurlyeq>$ is a generalized Ramsey structure then there is a function *des*: $\boldsymbol{W} \cup \boldsymbol{G} \mapsto \mathbb{R}$ such that for all α, β in $\boldsymbol{W} \cup \boldsymbol{G}$ and all $w_1, w_2, w_3, w_4 \in \boldsymbol{W}$,

(a) $\alpha \succcurlyeq \beta$ if and only if $des(\alpha) \geq des(\beta)$

(b) $(w_1, w_2) \geq^{d} (w_3, w_4)$ if and only if $des(w_1) - des(w_2) \geq des(w_3) - des(w_4)$

Furthermore, *des* is unique up to positive linear transformation

A proof is provided in the appendix. We now turn to a discussion of the axioms (GRS 1)–(9) before turning to how to get from generalized Ramsey structures to credence functions.

4.3 Generalized Ramsey structures

The strategy for proving *Theorem 1* is closely connected to Ramsey's process; given the agent's preferences over outcomes and two-outcome gambles, we first determine the relation $\geq^{d}$ between pairs of outcomes and on that basis construct an interval scale measurement of the agent's preferences. The most important step here is establishing that if $<\boldsymbol{W}, \boldsymbol{P}, \boldsymbol{G}, \succcurlyeq>$ is a generalized Ramsey structure, then $<\boldsymbol{W} \times \boldsymbol{W}, \geq^{d}>$ is an *algebraic difference structure*:

Definition 13*: Algebraic difference structure*
$<\boldsymbol{A} \times \boldsymbol{A}, \succcurlyeq^*>$ is an *algebraic difference structure* if and only if $\boldsymbol{A}$ is non-empty, $\succcurlyeq^*$ is a binary relation on $\boldsymbol{A} \times \boldsymbol{A}$, and for all $a_1, a_2, a_3,$

a_4, a'_1, a'_2, $a'_3 \in \boldsymbol{A}$, and all sequences $a_1, a_2, \ldots, a_i, \ldots \in \boldsymbol{A}$, the following five conditions hold:

(ADS 1) $<\boldsymbol{A} \times \boldsymbol{A}, \succcurlyeq^* >$ is a weak order

(ADS 2) If $(a_1, a_2) \succcurlyeq^* (a_3, a_4)$, then $(a_4, a_3) \succcurlyeq^* (a_2, a_1)$

(ADS 3) If $(a_1, a_2) \succcurlyeq^* (a_4, a_5)$ and $(a_2, a_3) \succcurlyeq^* (a_5, a_6)$, then $(a_1, a_3) \succcurlyeq^* (a_4, a_6)$

(ADS 4) If $(a_1, a_2) \succcurlyeq^* (a_3, a_4) \succcurlyeq^* (a_1, a_1)$, then there exist a_5, $a_6 \in \boldsymbol{A}$ such that $(a_1, a_5) \sim^* (a_3, a_4) \sim^* (a_6, a_2)$

(ADS 5) If $a_1, a_2, \ldots, a_i, \ldots$ is such that $(a_{i+1}, a_i) \sim^* (a_2, a_1)$ for every a_i, a_{i+1} in the sequence, $\neg((a_2, a_1) \sim^* (a_1, a_1))$, and there exist $a', a'' \in \boldsymbol{A}$ such that $(a', a'') \succ^* (a_i, a_1) \succ^* (a'', a')$ for all a_i in the sequence, then it is finite

This allows us to invoke the following theorem:

Theorem 2: Algebraic difference measurement
If $<\boldsymbol{A} \times \boldsymbol{A}, \succcurlyeq^*>$ is an algebraic difference structure, then there exists a real-valued function f on $\boldsymbol{A}$ such that, for all a_1, a_2, a_3, $a_4 \in \boldsymbol{A}$, $(a_1, a_2) \succcurlyeq^* (a_3, a_4)$ if and only if $f(a_1) - f(a_2) \geq f(a_3) - f(a_4)$; furthermore, f is unique up to positive linear transformation

For a proof of *Theorem 2*, see Krantz *et al* (1971, Ch. 4).

Though none of the axioms (GRS 1)–(9) is identical to any of Ramsey's listed axioms, the majority of them bear a close resemblance to the axioms and assumptions mentioned by Ramsey in his paper. The differences are largely due to the assumption that the set of gambles includes only possible gambles, and the altered definition of ½ probability propositions. In this section I will discuss the axioms individually. It is worth noting first that none of the axioms is intended to be independently plausible *qua* norms of practical rationality, though at least a few may seem to have this status.[12] As with Ramsey's axiom system, the goal here is to establish conditions for the possibility of measurement under the assumption of the broad

[12] Plausibly, (GRS 3), (4), and (9) are constraints of practical rationality. I am inclined to take (GRS 5) as a rationality constraint, though this is difficult to justify without presupposing the representation result. The status of the Archimedean axiom (GRS 7) is unclear, though representation theorems that forego an Archimedean axiom can be developed, e.g., Bartha (2007). The existential axioms (GRS 2), (6) and (8) are not plausibly rationality constraints. (GRS 1) is a purely structural axiom, and places no constraints on any agent whether ideal or not.

descriptive adequacy of something like expected utility theory — we are not directly interested in establishing foundations for a prescriptive decision theory.

The purely structural axiom (GRS 1) does not correspond to any of Ramsey's axioms or any of the further assumptions he mentions, and his discussion is too sparse to know with any certainty whether he implicitly assumed anything like it. The move from worlds to near-worlds as outcomes suggests that he desired to avoid impossible gambles. However, given how Ramsey proposed to define $=^{\mathrm{d}}$, without changes elsewhere in his system he also required either that we have preferences over impossible gambles, or that every outcome in $\boldsymbol{W}$ was compatible with both the truth and the falsity of some ethically neutral proposition. The argument for this proceeds by first noting that if $des(w_1) - des(w_2) = des(w_3) - des(w_4)$, then it should be the case that $(w_1, w_2) =^{\mathrm{d}} (w_3, w_4)$. Suppose that $w_1 \sim w_1'$, so $des(w_1) - des(w_1') = des(w_1') - des(w_1)$. From *Definition 2*, we know that $(w_1, w_1') =^{\mathrm{d}} (w_1', w_1)$ is only defined if the agent has preferences over some pair of gambles of the form $(w_1, P; w_1)$ and $(w_1', P; w_1')$ for some ethically neutral P of probability ½. It follows that either w_1 is compatible with P and $\neg P$, and similarly for w_1', or at least one of these two gambles is impossible.

We might suppose that Ramsey was happy to deal with preferences over impossible gambles. This would have forced him to assume that there is an interesting difference between two impossible propositions $(w_1 \,\&\, P)$ and $(w_2 \,\&\, P)$, where both w_1 and w_2 entail $\neg P$ but $\neg(w_1 \sim w_2)$. For suppose that Ramsey had only one impossible proposition, $\varnothing$. Then $des(w_1 \,\&\, P) = des(w_2 \,\&\, P) = des(\varnothing)$, but $des(w_1) \neq des(w_2)$. For whatever value we take $des(\varnothing)$ to have, it is clear that this will lead to problems. For instance, suppose that $des(\varnothing) \neq des(w_1)$; w_1 and w_2 each imply P; w_3 implies $\neg P$; and $des(w_1) = x$, $des(w_2) = des(w_3) = y$. We require that $(w_1, w_2) =^{\mathrm{d}} (w_1, w_3)$, for obviously $x - y = x - y$. However, the justification for the definition of $=^{\mathrm{d}}$ in terms of preferences fails under these conditions:

$$(w_1, w_2) =^{\mathrm{d}} (w_1, w_3) \text{ if and only if } (w_1, P; w_3) \sim (w_2, P; w_1)$$

But this holds just in case:

$$des(w_1 \& P).bel(P) + des(w_3 \& \neg P).(1 - bel(P)) = des(w_2 \& P).bel(P) + des(w_1 \& \neg P).(1 - bel(P))$$

Supposing P is ethically neutral and has probability ½, this equals:

$$0.5x + 0.5y = 0.5y + 0.5des(\varnothing)$$

It follows that $des(\varnothing) = x = des(w_1)$, which contradicts our initial assumption.

The only consistent way that Ramsey could have involved impossible gambles in his system would have been to treat different impossible propositions as different objects of desire. Perhaps an appeal to impossible worlds would suffice for this purpose: the impossible prospect of being a married bachelor might be desired to a greater degree than the prospect of being a square circle.

I do not want to judge whether Ramsey intended to countenance impossible gambles; for my own part, restricting our attention to possible gambles seems the better option. It is not obvious how we ought to treat preferences with respect to impossible propositions, if indeed there is more than one such proposition. For instance, it is implicit in Ramsey's system that if $w_1 \succcurlyeq w_2$, then $w_1 \succcurlyeq (w_1, P; w_2) \succcurlyeq w_2$. Without some such assumption he is unable to show that the function *bel* is a credence function. (See the proof in the appendix.) However, if we know that (w_1 & P) is impossible, it is not obvious why this should be the case: it seems at least as plausible that $w_2 \succcurlyeq (w_1, P; w_2)$ in this case, as we know we are not going to receive w_1 in the event that P and there is only a potentially very small probability $bel(\neg P) \leq 1$ of receiving w_2. Because it explicitly avoids any impossible gambles and the issues that come with them, *Theorem 1* constitutes an advance over Ramsey's system.

As noted earlier, (GRS 2) plays a very similar foundational role to (RAM 1); indeed, it is involved in every major step of the theorem's proof. It is worth noting again that, on the assumption that Ramsey wanted to avoid impossible gambles, his system requires something stronger than just (RAM 1). In fact, given his definition of $=^{d}$, Ramsey's system needs the stronger claim that every *outcome* is compatible with the truth or falsity of at least one ethically neutral proposition of probability ½. This can be phrased as an existential condition that replaces (RAM 1):

(RAM 1*) For every $w \in \boldsymbol{W}$, there is at least one ethically neutral proposition of probability ½, P, such that w is compatible with P and $\neg P$

(GRS 2) posits a substantially more plausible requirement than this. It postulates the existence of a set of propositions, Π, such that all have probability ½, but *none* has to be ethically neutral in any of the senses defined above.[13] It is not required that every outcome has to be compatible with both the truth and the falsity of at least one proposition in Π; in fact *none* will be, if $\boldsymbol{W}$ is the set of possible worlds. Instead, it implies that for every w in $\boldsymbol{W}$, there is a proposition π in Π such that for some pair of outcomes w' and w'' with the same desirability as w, w' implies π and w'' implies $\neg\pi$. Given this, independent of its intrinsic plausibility, the use of (GRS 2) as the basis for a representation theorem constitutes a substantial advance over Ramsey's system.

(GRS 2) is nevertheless likely to be a somewhat contentious axiom. It implies, for instance, that *every* value $\underline{w}$ contains two worlds w and w' that differ with respect to some ½ probability proposition. It is plausible that for *many* values — perhaps even most — we will be able to find such a proposition. Consider, for instance, the following situation. World w is a world much like our own, where at some point in history far away a fair coin was tossed and it landed heads. No bets were ever made on the outcome of the toss, indeed nobody within the world paid attention to the outcome, and the toss had no significant impact on history in any way. World w' is essentially just like w, but the coin lands tails in it. Our subject has no intrinsic interest in the outcomes of coin tosses. Let P be the proposition *the tossed coin lands heads*. Plausibly, for our subject, $w \sim w'$, while $w \in P$ and $w' \in \neg P$. Importantly, it does not matter here that at some worlds (or even *most* worlds) the truth or falsity of P might make a difference to how our subject values that world. The propositions required by (GRS 2) might be highly relevant to the agent's subjective valuation of a world.

The case just given suggests that for most pairs of worlds w_1 and w_2, we should be able to find a ½ probability proposition which satisfies the conditions of (GRS 2). The axiom seems to be at least approximately satisfied in this sense — for any world in which there are coin tosses, we should be able to find a world which is equivalent in all respects that the agent cares about but where the outcome of some fair

[13] Suppose we define a new notion, *ethical substitutability: P* is ethically substitutable with respect to a value $\underline{w}$ iff $\underline{w}$ is compatible with P and $\neg P$, and $(P \,\&\, \underline{w}) \sim \underline{w} \sim (\neg P \,\&\, \underline{w})$. Given the (very plausible) assumption that for all values $\underline{w}$, $w \sim \underline{w}$ and if Q is a non-empty subset of $\underline{w}$, then $Q \sim \underline{w}$, (GRS 2) implies the existence of at least one ethically substitutable proposition. This is hardly problematic, however: all that is required for P to be ethically substitutable relative to some $\underline{w}$ is the existence of two equally valued worlds that differ with respect to P. Supposing that ethically substitutable propositions exist would not leave us with the problems discussed in §3.2. Thanks to Rachael Briggs and David Chalmers here.

coin toss is altered. If we come to a world where coin tosses matter (or where there are no coins), we should usually be able to find some other even-probability event of no intrinsic interest that we can appeal to — or some other pair of worlds with the same value which do differ with respect to some ½ probability proposition.

But it is still not obviously the case that this holds for *every* value $\underline{w}$. Perhaps there are some worlds which are unique in their desirability ranking, being equal in value to no other worlds; or perhaps there are some values which contain multiple worlds, but none of which disagree with respect to any ½ probability proposition. However, this circumstance would seem to be rare if it occurs at all, and if so it would not be a devastating problem: it would primarily mean that sometimes, $\geq^{\mathrm{d}}$ on $\boldsymbol{W} \times \boldsymbol{W}$ is undefined. Some pairs of worlds might be left out of the $\geq^{\mathrm{d}}$ comparison, but the relation would nevertheless still be a well-defined order on the others. It would likely be possible (though not without substantial added complexity) to prove a weaker representation result, which leaves certain utility values for worlds (and correspondingly, probability values for propositions) unspecified or within certain constrained intervals.

(GRS 3) corresponds closely to (RAM 3), and it is a standard necessary condition in decision theoretic representation theorems. The role of (GRS 4) is complex, and while no axiom like it shows up in Ramsey's system, amongst other things it plays many of the same roles as (RAM 2). Like (GRS 2), this axiom shows up in every major step of the proof that $<\boldsymbol{W} \times \boldsymbol{W}, \geq^{\mathrm{d}}>$ is an algebraic difference structure. In a manner of speaking, this axiom says that the rational agent treats the same way all prospects with similarly valued outcomes conditional on any ½ probability proposition. It tells us that we can substitute one world w_1 for another w_1' within a gamble, or one ½ probability proposition for another, so long as the outcomes have the same desirabilities and the substitution results in a possible gamble. So, for example, if w_1 and w_1' have the same desirability and both are compatible with the ½ probability propositions π and π', then $(w_1, \pi; w_2) \sim (w_1', \pi'; w_2)$. It also allows that we can change the order of outcomes, in the sense that if $(w_1, \pi; w_2)$ and $(w_2', \pi; w_1')$ are both possible gambles, then $(w_1, \pi; w_2) \sim (w_2', \pi; w_1')$. (GRS 4) helps to ensure the coherence of the definitions of $\Pi, \geq^{\mathrm{d}}$, and of the *bel* function.

(GRS 5) is designed to play the same role as (RAM 4): in light of the other axioms, it effectively asserts that $\geq^{\mathrm{d}}$ is transitive, which is crucial for establishing that $<\boldsymbol{W} \times \boldsymbol{W}, \geq^{\mathrm{d}}>$ satisfies (ADS 1) and (ADS 3) of *Definition 13*. (GRS 6) is essentially a restatement of (RAM 5). Its role is limited to establishing that $<\boldsymbol{W} \times \boldsymbol{W}, \geq^{\mathrm{d}}>$ satisfies (ADS 4), a non-necessary structural condition. (GRS 7) is the Archimedean axiom; appropriately translated, it basically asserts that $<\boldsymbol{W} \times \boldsymbol{W}, \geq^{\mathrm{d}}>$ satisfies (ADS 5).

(GRS 1)–(7) are sufficient to establish that $<\boldsymbol{W} \times \boldsymbol{W}, \geq^{\mathrm{d}}>$ is an algebraic difference structure, which entails the existence of a real-valued function *des* on $\boldsymbol{W}$ with the property that $(w_1, w_2) \geq^{\mathrm{d}} (w_3, w_4)$ if and only if $des(w_1) - des(w_2) \geq des(w_3) - des(w_4)$. (GRS 8)–(9) are used to ensure that *des* represents $\succcurlyeq$ on $\boldsymbol{W} \cup \boldsymbol{G}$ in the sense that for all $\alpha, \beta \in \boldsymbol{W} \cup \boldsymbol{G}$, $des(\alpha) \geq des(\beta)$ if and only if $\alpha \succcurlyeq \beta$. These final two axioms also play central roles in the construction of a credence function *bel* (§4.4).

(GRS 8) does not correspond directly to any of Ramsey's stated axioms or any of the assumptions he otherwise mentions, though he clearly presupposed something like it in any case. It essentially states that the desirability of a gamble $(w_1, P; w_2)$ is somewhere weakly between the desirability of w_1 and the desirability of w_2. This ensures that *des* is such that for all $w_1, w_2 \in \boldsymbol{W}$, $des(w_1) \geq des(w_2)$ if and only if $w_1 \succcurlyeq w_2$. It also helps to ensure that *bel* does not result in probabilities less than 0 or greater than 1.

The sole formal role of (GRS 9) is to ensure that we can extend *des* on $\boldsymbol{W}$ to $\boldsymbol{W} \cup \boldsymbol{G}$; it is perhaps identical to what Ramsey intended for his (RAM 7). It necessitates the existence, for each gamble, of an outcome that is directly comparable with that gamble. This ensures that for all $\alpha, \beta \in \boldsymbol{W} \cup \boldsymbol{G}$, $des(\alpha) \geq des(\beta)$ if and only if $\alpha \succcurlyeq \beta$. Given the non-triviality of $\succ$ on $\boldsymbol{W} \cup \boldsymbol{G}$ (ensured by (GRS 2)) and given that, if $w_1 \succ w_2$, then $w_1 \succ (w_1, \pi; w_2) \succ w_2$, (GRS 9) forces the set of outcomes to be infinite. In this respect, it is similar to Ramsey's (RAM 6), though it plays a quite different role from that which Ramsey had intended for his axiom. This is also likely to be a contentious axiom; though here it is noteworthy that the assumption is not necessary for the main representation result. The failure of (GRS 9) would open up the possibility of *bel* being undefined for some $P \in \boldsymbol{P}$. Other means of extending *des* to $\boldsymbol{W} \cup \boldsymbol{G}$ are also likely possible in lieu of (GRS 9).

Although *Theorem 1* takes as outcomes a set of worlds, an essentially identical theorem exists where $\boldsymbol{W}$ is a set of propositions: $\boldsymbol{W} \subseteq \boldsymbol{P}$. On this interpretation, $\boldsymbol{W}$ might refer to sets of $\sim$-equivalence classes of worlds. Alternatively, $\boldsymbol{W}$ might be taken to be an arbitrary set of propositions, perhaps more or less equivalent to $\boldsymbol{P}$ itself. Formally, the difference between the theorems essentially just involves substituting every instance of $\in$ between worlds and propositions to a $\subseteq$ relation (between two propositions). For instance, (GRS 2) becomes:

(GRS 2*) For every pair of propositions $w_1, w_2 \in \boldsymbol{W}$, there exists a $\pi \in \Pi$ such that for some $w_1' \sim w_1$ and $w_2' \sim w_2$, $w_1' \subseteq \pi$ and $w_2' \subseteq \neg\pi$

The representation result can then be established along precisely the same lines as in *Theorem 1, mutatis mutandis*. Of course, this would alter the interpretation of the axioms. (GRS 2*), for instance, implies that for every proposition w in $\boldsymbol{W}$, there exist at least two propositions w' and w'' and a ½ probability proposition π such that $w \sim w' \sim w''$ and $w' \subseteq \pi$ and $w'' \subseteq \neg\pi$. The modified version of the first axiom, (GRS 1), would imply that $\boldsymbol{G}$ is the set of gambles with outcomes that entail their conditions rather than merely being compatible with them.

4.4 Constructing bel

Let us suppose that $<\boldsymbol{W}, \boldsymbol{P}, \boldsymbol{G}, \succcurlyeq>$ is a generalized Ramsey structure; our goal then is to construct a representation of the agent's credences. Closely following Ramsey's suggestion, we can define the function *bel* as follows:

Definition 14*: bel*
For all $P \in \boldsymbol{P}$, if $P = \varnothing$, then $bel(P) = 0$; and if $P = \boldsymbol{W}$, then $bel(P) = 1$; otherwise, if $w_1, w_2 \in \boldsymbol{W}$ are such that $\neg(w_1 \sim w_2)$ and $(w_1, P; w_2) \in \boldsymbol{G}$, then $bel(P) = (des((w_1, P; w_2)) - des(w_2)) \,/\, (des(w_1) - des(w_2))$

As with Ramsey's definition, *bel* so defined is unique, due to the fact that ratios of differences are preserved across admissible transformations of the *des* function.

We will also need the following condition to ensure the coherence of the definition:

Condition 1*: bel coherence*
For all $(w_1, P; w_2), (w_3, P; w_4) \in \boldsymbol{G}$ where $\neg(w_1 \sim w_2)$ and $\neg(w_3 \sim w_4)$, $(des((w_1, P; w_2)) - des(w_2)) \,/\, (des(w_1) - des(w_2)) = (des((w_3, P; w_4)) - des(w_4)) \,/\, (des(w_3) - des(w_4))$

Condition 1 is a formal restatement of one of the conditions that Ramsey briefly mentions are required to ensure the coherence of the *bel* function. We can understand what it says as follows. *Definition 14* tells us that $bel(P)$ is, say, 0.75, if it is the case that $w_1 \succ w_2$ and the value of the gamble $(w_1, P; w_2)$ sits exactly three quarters of the way from the values of w_2 to w_1. *Condition 1* then tells us that for all w_3, w_4 such that $w_3 \succ w_4$, if the gamble $(w_3, P; w_4)$ exists then it also sits three quarters of the way between w_4 and w_3 in the agent's desirability scale (and if $w_4 \succ w_3$, then $(w_3, P; w_4)$ is one quarter of the distance between w_4 and w_3). This directly implies that the value $bel(P)$ does not depend on which outcomes and gambles we choose to consider, which is an obvious coherence requirement.[14] The other condition mentioned by Ramsey — that we will always be able to find outcomes satisfying the conditions of *Definition 14* — follows already from (GRS 1)–(3).

Theorem 3: Generalized Ramseyian credence and utility measurement If $<\boldsymbol{W}, \boldsymbol{P}, \boldsymbol{G}, \succcurlyeq>$ is a generalized Ramsey structure and *Condition 1* holds, then there is a function des: $\boldsymbol{W} \cup \boldsymbol{G} \mapsto \mathbb{R}$, and a function bel: $\boldsymbol{P} \mapsto [0,1]$, such that for all α, β in $\boldsymbol{W} \cup \boldsymbol{G}$, all $w_1, w_2, w_3, w_4 \in \boldsymbol{W}$, and all $(w_1, P; w_2) \in \boldsymbol{G}$,

(a) $\alpha \succcurlyeq \beta$ if and only if $des(\alpha) \geq des(\beta)$

(b) $(w_1, w_2) \geq^{d} (w_3, w_4)$ if and only if $des(w_1) - des(w_2) \geq des(w_3) - des(w_4)$

(c) $bel(\boldsymbol{W}) = 1$ and $bel(\varnothing) = 0$

(d) $des((w_1, P; w_2)) = des(w_1).bel(P) + des(w_2).(1 - bel(P))$

Furthermore, bel is unique and des is unique up to positive linear transformation

There are some important points to note about this representation result, the proof of which is in the appendix. The first is that the function bel need not be a probability function, though it could be. It *is* a credence function in the sense that bel maps propositions to some value within [0,1]; thus it is plausible to take $bel(P)$ as a representation of the degree of belief the agent associates with the proposition P.

[14] I have chosen to state *Condition 1* in terms of *des* as there is no apparent straightforward means of stating it purely in terms of preferences. Since *des* is constructed entirely from preferences, *Condition 1* is equivalent to *some* (perhaps infinitary) condition on preferences. Importantly, *Condition 1*'s content is more transparent when expressed in terms of *des*, which requires that *des* has already been characterized. Davidson and Suppes's (1956) axiom *A10* achieves a similar purpose as my *Condition 1* without referring to the intended representation, but only through a complicated series of definitions that serve to obscure its content.

The reason for *bel*'s permissiveness is that (GRS 1)–(9) and *Condition 1* jointly place very few restrictions on preferences for gambles conditional on propositions outside Π. For instance, suppose that neither P nor Q are in Π, $P \subset Q$, $\neg(w_1 \sim w_2)$, and the agent is to rank the two gambles $(w_1, P; w_2)$ and $(w_1, Q; w_2)$. Only (GRS 3), (GRS 8)–(9), and *Condition 1* can have any impact on how these gambles are ranked, as the other conditions are either purely existential or refer to gambles conditional on ½ probability propositions. (GRS 9) only asserts the existence of some w_3 and w_4 such that $w_3 \sim (w_1, P; w_2)$ and $w_4 \sim (w_1, Q; w_2)$, while (GRS 8) only asserts that $(w_1, P; w_2)$ and $(w_1, Q; w_2)$ must be valued somewhere between w_1 and w_2. Each of these, along with (GRS 3), can clearly be satisfied even if $(w_1, P; w_2) \succ (w_1, Q; w_2)$. Finally, *Condition 1* only restricts the relative rankings of gambles conditional on the *same* proposition, so it is also consistent with $(w_1, P; w_2) \succ (w_1, Q; w_2)$. Assuming all the other conditions to be satisfied, it follows immediately that if $(w_1, P; w_2) \succ (w_1, Q; w_2)$, then $bel(P) > bel(Q)$. Hence, *bel* in this instance is not a probability function.

It is of course possible to state further conditions to ensure that *bel* satisfies certain desirable structural properties. For instance, an extremely plausible condition is that the order in which outcomes are presented in a gamble makes no difference to their value:

Condition 2: Order indifference
For all $(w_1, P; w_2), (w_2, \neg P; w_1) \in \boldsymbol{G}$, $(w_1, P; w_2) \sim (w_2, \neg P; w_1)$

Another way to motivate *Condition 2* would be to say that $(w_1, P; w_2)$ and $(w_2, \neg P; w_1)$ are merely notational variants each representing the same object of preference; in this case, the condition will be satisfied by all preference orderings necessarily. In light of the other conditions, it implies that *bel* is such that $bel(P) = 1 - bel(\neg P)$. From (GRS 1), we know that $(w_1, P; w_2) \in \boldsymbol{G}$ if and only if $(w_2, \neg P; w_1) \in \boldsymbol{G}$; however, the axioms (GRS 1)–(9) (plus *Condition 1*) are compatible with the agent's valuing them differently (for essentially the reasons just outlined). This leads us to the following result:

Theorem 4: Generalized Ramseyian measurement with order indifference
If $<\boldsymbol{W}, \boldsymbol{P}, \boldsymbol{G}, \succcurlyeq>$ is a generalized Ramsey structure and *Condition 1* and *2* hold, then there is a function des: $\boldsymbol{W} \cup \boldsymbol{G} \mapsto \mathbb{R}$, and a function bel: $\boldsymbol{P} \mapsto [0,1]$, such that for all α, β in $\boldsymbol{W} \cup \boldsymbol{G}$, all $w_1, w_2, w_3, w_4 \in \boldsymbol{W}$, all $(w_1, P; w_2) \in \boldsymbol{G}$, and all $P \in \boldsymbol{P}$,

(a) $\alpha \succcurlyeq \beta$ if and only if $des(\alpha) \geq des(\beta)$

(b) $(w_1, w_2) \geq^{d} (w_3, w_4)$ if and only if $des(w_1) - des(w_2) \geq des(w_3) - des(w_4)$

(c) $bel(\boldsymbol{W}) = 1$ and $bel(\varnothing) = 0$

(d) $des((w_1, P; w_2)) = des(w_1).bel(P) + des(w_2).(1 - bel(P))$

(e) $bel(P) = 1 - bel(\neg P)$

Furthermore, *bel* is unique and *des* is unique up to positive linear transformation

Note that the fact that $(w_1, P; w_2)$ and $(w_2, \neg P; w_1)$ are distinct entities in $\boldsymbol{G}$ is an artefact of how we have formalized gambles, that is, as elements of $\boldsymbol{W} \times \boldsymbol{P} \times \boldsymbol{W}$. Without changing the axioms and with almost no changes to the proof, it is entirely possible to characterize $\boldsymbol{G}$ as a set of one- or two-valued functions from $\boldsymbol{W}$ to $\boldsymbol{W}$. On this formalization, $(w_1, P; w_2)$ and $(w_2, \neg P; w_1)$ will be the same object—that is, each will be a function f: $\boldsymbol{W} \mapsto \boldsymbol{W}$ such that $f(w)$ equals w_1 if $w \in P$, and w_2 if $w \in \neg P$. From this, *Condition 2* follows immediately from (GRS 3). I have chosen to use the more general formalization, $\boldsymbol{G} \subseteq \boldsymbol{W} \times \boldsymbol{P} \times \boldsymbol{W}$, as it allows for the more general representation result of *Theorem 3*. If we desire that *bel* satisfies condition (e), *Condition 2* is only a very weak additional constraint on preferences.

It would also be possible to ensure that *bel* is monotonic, in the sense that if $P \subseteq Q$, then $bel(P) \leq bel(Q)$, by adding the following condition for all $P, Q \in \boldsymbol{P}$:

Condition 3: Monotonicity
If $P \subseteq Q$, then for all $w_1, w_2 \in \boldsymbol{W}$, (i) if $w_1 \succ w_2$ and some $(w_1, Q; w_2)$, $(w_1', P; w_2')$ exists in $\boldsymbol{G}$, then $(w_1, Q; w_2) \succcurlyeq (w_1', P; w_2')$; (ii) if $w_2 \succ w_1$ and some $(w_1, Q; w_2)$, $(w_1', P; w_2')$ exists in $\boldsymbol{G}$, then $(w_1', P; w_2') \succcurlyeq (w_1, Q; w_2)$

If *Condition 3* holds and $\boldsymbol{P}$ is the power set of $\boldsymbol{W}$, then *bel* is a Choquet capacity; that is, a function f: $2^{\boldsymbol{W}} \mapsto [0,1]$ such that $f(\varnothing) = 0$, $f(\boldsymbol{W}) = 1$, and $f(Q) \geq f(P)$ whenever $P \subset Q$. (GRS 1)–(9) plus *Condition 1* and *3* thus form the basis for a version of Choquet expected utility theory (cf. Schmeidler 1989):

Theorem 5: Ramseyian Choquet expected utility theory
If $<\boldsymbol{W}, \boldsymbol{P}, \boldsymbol{G}, \succcurlyeq>$ is a generalized Ramsey structure where $\boldsymbol{P}$ is the power set of $\boldsymbol{W}$ and *Condition 1* and *3* hold, then there is a function *des*: $\boldsymbol{W} \cup \boldsymbol{G} \mapsto \mathbb{R}$, and a function *bel*: $\boldsymbol{P} \mapsto [0,1]$, such that such that

for all α, β in $\boldsymbol{W} \cup \boldsymbol{G}$, all $w_1, w_2, w_3, w_4 \in \boldsymbol{W}$, all $(w_1, P; w_2) \in \boldsymbol{G}$, and all $P, Q \in \boldsymbol{P}$,

(a) $\alpha \succcurlyeq \beta$ if and only if $des(\alpha) \geq des(\beta)$

(b) $(w_1, w_2) \geq^{d} (w_3, w_4)$ if and only if $des(w_1) - des(w_2) \geq des(w_3) - des(w_4)$

(c) $bel(\boldsymbol{W}) = 1$ and $bel(\varnothing) = 0$

(d) $des((w_1, P; w_2)) = des(w_1).bel(P) + des(w_2).(1 - bel(P))$

(f) If $P \subseteq Q$, $bel(P) \leq bel(Q)$

Furthermore, *bel* is unique and *des* is unique up to positive linear transformation

Condition 2 and *Condition 3* are independent; if both hold (and $\boldsymbol{P}$ is the power set of $\boldsymbol{W}$), then *bel* will satisfy both conditions (e) and (f).

Substantially more is required to ensure that *bel* is additive, and it is difficult to come up with conditions which are not simply direct assertions of the additivity requirement itself, or which do not involve gambles more complex than the simple two-outcome gambles we have considered so far. I am inclined to take *bel*'s potential lack of structure as a feature, not a bug. Plausibly, ordinary agents do not have probabilistically coherent degrees of belief, so any representation of credences which requires such coherence is flawed. The representation result of *Theorem 3* is compatible with an extremely wide range of credence functions — although it does require that if $P \in \Pi$, then $bel(P) = 0.5$, and there must be at least one $P \in \Pi$.[15] The result does imply that $bel(\varnothing) = 0$ and $bel(\boldsymbol{W}) = 1$, but it is noteworthy that these values for $\varnothing$ and $\boldsymbol{W}$ are entirely stipulative. Without this stipulation, *bel*'s uniqueness condition does not hold; however, *bel* on $(\boldsymbol{P} - \boldsymbol{W}) - \varnothing$ remains unique. It would be possible to set aside that stipulation and leave $bel(\varnothing)$ and $bel(\boldsymbol{W})$ undefined.

It is also important to note that the foregoing three representation results do not imply that the agent who satisfies the preference conditions can be represented as an expected utility maximizer generally. The simple reason for this is that satisfying (GRS 1)–(9) and *Condition 1*, *Condition 2*, or *Condition 3* is compatible with all kinds of preference patterns over more complex gambles. For instance, suppose that

[15] The reasoning of §2.3 essentially counts as a proof that if P is in Π, then $bel(P) = 0.5$. That there is at least one proposition in Π follows immediately from (GRS 2) and *Definition 8*.

$\{P_1, P_2, P_3\}$ is a partition of $\boldsymbol{W}$. Nothing about any of the above theorems implies that the desirability of the more complex gamble, $(w_1, P_1; w_2, P_2; w_3, P_3)$, must be a function of the *des* values of the outcomes and the *bel* values of the propositions. In particular, the theorem does not imply that:

$$des\big((w_1, P_1; w_2, P_2; w_3, P_3)\big) = des(w_1).bel(P_1) + des(w_2).bel(P_2) + des\big(w_3\big).bel\big(P_3\big)$$

In line with Ramsey's suggestion that we only presuppose the adequacy of the theory of subjective expected utility maximisation for the kinds of cases that we are concerned with, the representation result is compatible with high degrees of irrationality for choices more complex than those between simple gambles of the form $(w_1, P; w_2)$. Again, I take this to be a feature rather than a bug. Ordinary agents are not expected utility maximizers across the board, so it should not turn out that they can only be represented as such.

5. Conclusions

Despite its very early inception, there are several features that make Ramsey's system attractive, especially in comparison to later works. The theory of cardinal utility developed by John von Neumann and Oskar Morgenstern (1944) was in some respects a rediscovery of ideas already present in 'Truth and Probability', but its appeal to extrinsic probabilities limits its applicability, whereas Ramsey's system makes no such appeal and instead serves to characterize the agent's subjective probabilities.

Savage's (1954) theorem was also founded on Ramseyian ideas, and like Ramsey, Savage attempted to derive both a utility function and subjective probabilities purely from an agent's preferences. But Savage's system suffers from defects not present in Ramsey's system. Notably, Ramsey's system does not seem to require the assumption of *state neutrality*, which is a crucial element of Savage's system. Savage assumes that the desirabilities associated with outcomes (which are arbitrary objects or states of affairs in Savage's system) are entirely independent of the wider circumstances in which those outcomes are instantiated. If an outcome were, say, *going for a swim in the sea*, then state neutrality requires that we would value this state of affairs to a particular degree regardless of whether it is cold out, whether we have been swimming recently, or whether there have been recent reports of vicious shark attacks nearby. The analogue of assuming state neutrality in a Ramseyian system would be the assumption of Naïve Expected Utility Theory, and we have seen that Ramsey does not make that mistake. Indeed, it is precisely because Ramsey recognized the falsity of this assumption that he introduced the notion of ethical neutrality.

Furthermore, given the plausible assumption that Ramsey wanted to avoid impossible gambles (§4.3), the outcomes of a gamble are always consistent with the conditions of the gamble; consequently, Ramsey's system seems to avoid the *constant acts* problem that plagues Savage's formalization. Savage posits a set of functions from an infinite set of states into a distinct set of outcomes (intended to represent possible actions an agent might take), but he crucially assumes that so-called *constant acts* exist — that there are actions we might perform which result in a particular outcome regardless of the state of the world we are in, even if that state of the world is *incompatible* with the outcome's obtaining.[16]

Of course, this is not to say that Ramsey's gambles have an unequivocally clear interpretation. (See §2.2, above.) If we construe Ramsey's measurement system as involving merely *hypothetical* judgements, as Bradley (2001) suggests, then perhaps we might read $(w_1, P; w_2)$ as just a conjunction of subjunctive conditionals:

If P were the case, then w_1, *and if* $\neg P$ *were the case, then* w_2

At least one problem arises from this interpretation, however. It seems plausible, as Jeffrey (1983) suggests, that the utility of a proposition should be a (credence) weighted average of the utilities of the worlds within it. Yet there is nothing in Ramsey's eight axioms to guarantee that agents' utilities for worlds (or near worlds) will align nicely with their utilities for the propositions that they are a part of. Of course, the same concern seems to apply to *Theorem 1*, perhaps suggesting the need for a further *utility consistency* condition if this interpretation is adopted.

Finally, on the assumption that his representation conjecture holds, another attractive feature of Ramsey's proposal is that it provides us with strong uniqueness results: if an agent satisfies Ramsey's axioms and the further stated conditions, then there is a unique credence function *bel*, and a utility function *des* unique up to positive linear transformation, which jointly represents her preferences for two-outcome

[16] See Joyce (1999, Ch. 3) for a damning discussion of the problem of constant acts in Savage's work.

gambles in expectational form. We might contrast this with Jeffrey's (1983) expected utility theorem, where his pair of functions, <*P*, *U*> (analogous to our *bel* and *des*, respectively), is only unique up to a fractional linear transformation.

All of this is achieved, however, on the basis of the highly problematic (RAM 1) (or (RAM 1*)). For this reason, *Theorem 1* constitutes a substantial advance over Ramsey's proposed axiom system for the measurement of utilities, while *Theorem 3*, *Theorem 4*, and *Theorem 5* together show that (as Ramsey suggested but never proved) we can use such utility scales to measure degrees of belief with only minimal additional commitments. These theorems have all the attractive features just noted of Ramsey's system, without requiring the existence of ethically neutral propositions in any problematic sense. Certainly, there is no presupposition of the existence of ethically neutral propositions as Ramsey defined them (*Definition 5*), for we have made no use of logical atomism in setting out the theorems. And no proposition need non-trivially satisfy the highly questionable notion of ethical neutrality given in *Definition 6*; that is, there need not exist any contingent propositions P such that for all propositions Q that are compatible with both P and $\neg P$, $(P \,\&\, Q) \sim Q \sim (\neg P \,\&\, Q)$. Indeed, there need not even be any propositions that non-trivially satisfy the weaker *Definition 7*.

There are, naturally, further problems to deal with in the construction of Ramsey-style measurement systems. We have just seen that there are some concerns regarding the best interpretation of the elements of $\boldsymbol{G}$. It is also clear that there is still plenty of de-idealization required: the application of the theorem requires knowledge of preferences between far too many (indeed, *infinitely* many) pairs of gambles and outcomes for us ever to gain empirical access to all of them — and the longer we take to empirically discover an agent's preferences, the more likely her beliefs and desires are to change.[17] Furthermore, the outcomes that make up $\boldsymbol{W}$ are far too fine-grained to ever be fully conceptualized or described; we could never genuinely ask any person to consider whether they would prefer $(w_1, P; w_2)$ over $(w_3, P; w_4)$, if w_1, w_2, w_3, and w_4 are possible worlds. Relatedly, we have to contend with the idealizing assumption that if w_1 implies P, then $w_1 \sim (w_1 \,\&\, P)$. This requires a kind of logical omniscience on the part of our agent: a capacity to always recognize the logical equivalence of

[17] The Ramsey-style theorem of Davidson and Suppes (1956) was developed in part to deal with this problem.

w_1 and (w_1 & P). Finally, we will need to say more about those circumstances where the existential commitments (GRS 2) and (GRS 9) fail, and what kind of representations can be had in such cases.

It is difficult to see how we might complicate our axioms to encompass more realistic scenarios. Importantly, though, these are not issues specific to Ramsey's proposal — in fact they arise quite generally for any attempt to construct a measurement procedure on the basis of an expected utility representation theorem. We do know, however, that ethical neutrality is not a problem for this programme — we simply do not need a notion of ethical neutrality to measure degrees of belief and desire in much the way that Ramsey suggested.[18]

6. Appendix

The proof of *Theorem 1* proceeds as follows. First, we show that (GRS 1)–(7) jointly entail that $< \boldsymbol{W} \times \boldsymbol{W}, \geq^{d} >$ is an algebraic difference structure, allowing us to invoke *Theorem 2* giving us *des* on $\boldsymbol{W}$. (GRS 8) and (GRS 9) are then used to extend *des* to $\boldsymbol{W} \cup \boldsymbol{G}$, and it is shown that this provides us with an interval scale representation of $\succcurlyeq$ on $\boldsymbol{W} \cup \boldsymbol{G}$.[19]

It will be helpful to establish some lemmas first:

(Lemma A):
For every pair $w_1, w_2 \in \boldsymbol{W}$, there is a $(w_1', \pi; w_2') \in \boldsymbol{G}$

(1) Follows immediately from (GRS 1) and (GRS 2). ■

We thus know that universally quantified statements about possible gambles conditional on some ½ probability proposition are never trivially satisfied; so, for instance, where a step says 'for all $(w_1', \pi; w_4')$, $(w_2', \pi'; w_3') \in \boldsymbol{G}$, $(w_1', \pi; w_4') \succcurlyeq (w_2', \pi'; w_3')$', (Lemma A) ensures that at least one such pair of gambles exists in $\boldsymbol{G}$. I will generally omit this step in what follows. Set memberships have been suppressed where obvious: henceforth we are only concerned with gambles in $\boldsymbol{G}$.

[18] Thanks to Rachael Briggs, David Chalmers, Alan Hájek, Hanti Lin, Wolfgang Schwarz, a PhilSoc audience at the ANU, a 2014 AAP conference audience in Canberra, and two anonymous referees at *Mind* for helpful feedback and discussion.

[19] Several of the steps in what follows owe much to Bradley (2001), especially (Lemma C) and the steps involving it.

(Lemma B):
If $(w'_1, \pi; w'_4) \succcurlyeq (w'_2, \pi'; w'_3)$ for some pair $(w'_1, \pi; w'_4)$, $(w'_2, \pi'; w'_3)$, then $(w_1, w_2) \geq^{d} (w_3, w_4)$

(1) Suppose that $(w'_1, \pi; w'_4) \succcurlyeq (w'_2, \pi'; w'_3)$ for some such pair.

(2) By (Lemma A), some $(w''_4, \pi''; w''_1)$ exists, and by successive iterations of (GRS 4), $(w'_1, \pi; w'_4) \sim (w''_4, \pi''; w''_1)$ and $(w''_4, \pi''; w''_1) \sim (w'''_1, \pi'''; w'''_4)$ for all such pairs. Because $\sim$ is an equivalence relation (GRS 3), $(w_1, \pi; w_4) \sim (w'''_1, \pi'''; w'''_4)$ for all such pairs.

(3) By the same steps, we know that $(w'_2, \pi'; w'_3) \sim (w''_2, \pi^*; w''_3)$ for all such pairs.

(4) So given our starting supposition, $(w'''_1, \pi'''; w'''_4) \succcurlyeq (w''_2, \pi^*; w''_3)$ for all such pairs, which is just the right hand side of *Definition 9*. ■

(Lemma C):
If $(w_1, w_2) \geq^{d} (w_3, w_4)$, then $(w_4, w_3) \geq^{d} (w_2, w_1)$, and $(w_1, w_3) \geq^{d} (w_2, w_4)$

(1) Suppose $(w_1, w_2) \geq^{d} (w_3, w_4)$, so $(w'_1, \pi; w'_4) \succcurlyeq (w'_2, \pi'; w'_3)$ for all such gambles.

(2) (Lemma A) ensures some $(w''_4, \pi^*; w''_1)$, $(w''_3, \pi^+; w''_2)$ exist, and by (GRS 4), $(w''_4, \pi^*; w''_1) \sim (w'_1, \pi; w'_4)$ and $(w''_3, \pi^+; w''_2) \sim (w'_2, \pi'; w'_3)$. Substituting for equally valued gambles, we get $(w''_4, \pi^*; w''_1) \succcurlyeq (w''_3, \pi^+; w''_2)$, which given (Lemma B) implies $(w_4, w_3) \geq^{d} (w_2, w_1)$.

(3) Likewise, $(w'_1, \pi; w'_4) \succcurlyeq (w''_3, \pi^+; w''_2)$, so $(w_1, w_3) \geq^{d} (w_2, w_4)$. ■

We can now show that (ADS 1)–(5) follow from (GRS 1)–(7). (ADS 2) is simply the first part of (Lemma C). Next we will prove that $\geq^{d}$ on $\boldsymbol{W} \times \boldsymbol{W}$ is complete:

(1) From (Lemma A), for any two (w_1, w_4), (w_2, w_3), there exist $(w'_1, \pi; w'_4)$, $(w'_2, \pi'; w'_3)$. From (GRS 3), either $(w'_1, \pi; w'_4) \succcurlyeq (w'_2, \pi'; w'_3)$ or $(w'_2, \pi'; w'_3) \succcurlyeq (w'_1, \pi; w'_4)$.

(2) Given (Lemma B), if the former then $(w_1, w_2) \geq^{d} (w_3, w_4)$, and if the latter then $(w_3, w_4) \geq^{d} (w_1, w_2)$. So either $(w_1, w_2) \geq^{d} (w_3, w_4)$ or $(w_3, w_4) \geq^{d} (w_1, w_2)$. ■

We also prove that $\geq^{\mathrm{d}}$ on $\boldsymbol{W} \times \boldsymbol{W}$ is transitive:

(1) Suppose that $(w_1, w_2) \geq^{\mathrm{d}} (w_3, w_4)$ and $(w_3, w_4) \geq^{\mathrm{d}} (w_5, w_6)$.

(2) From *Definition* 9, for all the relevant gambles, this implies that $(w_1', \pi; w_4') \succcurlyeq (w_2', \pi'; w_3')$ and $(w_3', \pi^*; w_6') \succcurlyeq (w_4', \pi^+; w_5')$.

(3) For any pair of gambles $(w_1'', \pi''; w_6'')$, $(w_2'', \pi'''; w_5'')$, (GRS 5) then requires that $(w_1'', \pi''; w_6'') \succcurlyeq (w_2'', \pi'''; w_5'')$, and $(w_1, w_2) \geq^{\mathrm{d}} (w_5, w_6)$ follows from (Lemma B). ■

So $<\boldsymbol{W} \times \boldsymbol{W}, \geq^{\mathrm{d}}>$ is a weak order and (ADS 1) is satisfied. Next we show that (ADS 3) is satisfied:

(1) Suppose $(w_1, w_2) \geq^{\mathrm{d}} (w_4, w_5)$ and $(w_2, w_3) \geq^{\mathrm{d}} (w_5, w_6)$.

(2) The second part of (Lemma C) applied to each conjunct entails $(w_1, w_4) \geq^{\mathrm{d}} (w_2, w_5)$ and $(w_2, w_5) \geq^{\mathrm{d}} (w_3, w_6)$. Because $\geq^{\mathrm{d}}$ is transitive, $(w_1, w_4) \geq^{\mathrm{d}} (w_3, w_6)$. So from (Lemma C) again, $(w_1, w_3) \geq^{\mathrm{d}} (w_4, w_6)$. ■

(ADS 4) is satisfied:

(1) Suppose $(w_1, w_2) \geq^{\mathrm{d}} (w_3, w_4) \geq^{\mathrm{d}} (w_1, w_1)$.

(2) From (GRS 6), for every triple w_1, w_3, w_4, there is a w_5 such that for some $(w_1', \pi; w_4')$, $(w_5, \pi'; w_3')$ (ensured by (Lemma A)), $(w_1', \pi; w_4') \sim (w_5, \pi'; w_3')$. Applying (Lemma B), we see that there must be a w_5 such that $(w_1, w_5) =^{\mathrm{d}} (w_3, w_4)$.

(3) Likewise, for every triple w_3, w_4, w_2, there is a w_6 such that $(w_3', \pi; w_2') \sim (w_6, \pi'; w_4')$ for some such pair; so there is a w_6 such that $(w_3, w_4) =^{\mathrm{d}} (w_6, w_2)$. ■

And (ADS 5) is also satisfied. The proof of this is trivial given (GRS 7), and *Definition* 9, and the definition of a strictly bounded standard sequence; it has therefore been left unstated. (GRS 1)–(7) therefore imply that $<\boldsymbol{W} \times \boldsymbol{W}, \geq^{\mathrm{d}}>$ is an algebraic difference structure, which ensures the existence of *des* on $\boldsymbol{W}$ (unique up to positive linear transformation), such that:

$$(w_1, w_2) \geq^{d} (w_3, w_4) \text{ if and only if } des(w_1) - des(w_2) \geq des(w_3) - des(w_4)$$

We appeal primarily to (GRS 8) to show that $des(w_1) \geq des(w_2)$ if and only if $w_1 \succcurlyeq w_2$:

(1) From (GRS 8), $w_1' \sim w_1$ if and only if, for all $(w_1'', P; w_1''')$, $w_1 \sim (w_1'', P; w_1''')$; and similarly, $w_2' \sim w_2$ if and only if, for all $(w_2'', P; w_2''')$, $w_2 \sim (w_2'', P; w_2''')$.

(2) Given (GRS 3) then, $w_1 \succcurlyeq w_2$ if and only if $(w_1'', \pi; w_1''') \succcurlyeq (w_2'', \pi'; w_2''')$ for all such gambles, which holds if and only if $(w_1, w_2) \geq^{d} (w_2, w_1)$.

(3) From *Theorem 2*, $(w_1, w_2) \geq^{d} (w_2, w_1)$ if and only if $des(w_1) - des(w_2) \geq des(w_2) - des(w_1)$, which can only be if $des(w_1) \geq des(w_2)$. So $w_1 \succcurlyeq w_2$ if and only if $des(w_1) \geq des(w_2)$. ■

We further require that *des* is defined on $\boldsymbol{W} \cup \boldsymbol{G}$. From (GRS 9), we know that for every $(w_1, P; w_2)$ there is a w_3 such that $(w_1, P; w_2) \sim w_3$. We can achieve the desired extension by making the following stipulation:

For all $w_3, (w_1, P; w_2) \in \boldsymbol{W} \cup \boldsymbol{G}$, $des((w_1, P; w_2)) = des(w_3)$ if and only if $(w_1, P; w_2) \sim w_3$

The proof that condition (a) of *Theorem 1* then holds is trivial and left unstated. The uniqueness properties of *des* on $\boldsymbol{W}$ will also clearly hold for *des* on $\boldsymbol{W} \cup \boldsymbol{G}$. The foregoing thus establishes *Theorem 1*.

To prove *Theorem 3*, we need to show that *Definition 14* provides us with a unique function *bel* that satisfies the stated properties. There are three parts to proving this. First, we need to show that we will always be able to find worlds and gambles satisfying the definition's conditions. Second, we need to show that the *bel*(*P*) is independent of the choice of worlds and gambles satisfying the definition; along with the first step this will ensure that *bel* is a function with domain $\boldsymbol{P}$. Finally, we show that the range of the function is [0,1].

(1) By stipulation, if $P = \varnothing$ then $bel(P) = 0$; and if $P = \boldsymbol{W}$, then $bel(\boldsymbol{W}) = 1$. Thus we are only concerned with *bel*(*P*) for contingent propositions, which must contain at least one member.

(2) Given (GRS 2) and *Definition 8*, $\succ$ on $\boldsymbol{W}$ is non-trivial. Let w_1 and w_2 be two outcomes such that $\neg(w_1 \sim w_2)$. There are then three possibilities: for any contingent proposition *P*, either *P*

contains both w_1 and w_2, or $\neg P$ does, or P and $\neg P$ contain one of w_1 and w_2 each.

(3) If either P contains both outcomes, given (GRS 3), its negation must contain at least one world w_3 such that either $\neg(w_1 \sim w_3)$ or $\neg(w_2 \sim w_3)$. Likewise if $\neg P$ contains both outcomes. If P and $\neg P$ contain one of w_1 and w_2 each, we already know that there are worlds in P and in $\neg P$ such that the agent is not indifferent between them. In any case, then, for any contingent proposition P, given (GRS 1), there will be at least one gamble of the form $(w_1, P; w_2)$ where $\neg(w_1 \sim w_2)$.

(4) That $bel(P)$ is independent of the choice of worlds and gambles satisfying the antecedent conditions follows immediately from *Condition 1*.

(5) The range of *bel* is [0,1]: from (GRS 8) and (GRS 3), for all $(w_1, P; w_2)$, either $w_1 \succcurlyeq w_2$ and $w_1 \succcurlyeq (w_1, P; w_2) \succcurlyeq w_2$, or $w_2 \succcurlyeq w_1$ and $w_2 \succcurlyeq (w_1, P; w_2) \succcurlyeq w_1$. Given the established properties of *des*, we know $des(w_1, P; w_2)$ always sits somewhere weakly between $des(w_1)$ and $des(w_2)$. It follows that the ratio of the difference between $des((w_1, P; w_2))$ and $des(w_2)$ and the difference between $des(w_1)$ and $des(w_2)$ will always be within [0,1]. ■

We now prove that *bel* is such that $des((w_1, P; w_2)) = des(w_1).bel(P) + des(w_2).(1 - bel(P))$:

(1) Suppose first that $w_1 \sim w_2$; then, by reasoning noted above, $des(w_1) = des(w_2) = des((w_1, P; w_2))$. Let $des(w_1) = x$. The required equality then holds just in case $x = x.bel(P) + x.(1 - bel(P))$; we have already noted that $bel(P) \in [0,1]$, so this is true regardless of the value of $bel(P)$.

(2) Suppose next that $\neg(w_1 \sim w_2)$. From *Definition 14*, $bel(P) = (des((w_1, P; w_2)) - des(w_2)) / (des(w_1) - des(w_2))$, which holds if and only if

$(des(w_1) - des(w_2)).bel(P) = des((w_1, P; w_2)) - des(w_2)$ if and only if

$des((w_1, P; w_2)) = des(w_1).bel(P) - des(w_2).bel(P) + des(w_2)$ if and only if

$$des((w_1,\ P;\ w_2)) = des(w_1).bel(P) + des(w_2).(1 - bel(P)).\ ■$$

Finally, *bel* is unique:

(1) From the earlier proofs and the fact that ratios of differences are preserved across admissible transformations of *des*, we know that there is only one function *bel* on $\boldsymbol{P}$ such that $bel(\boldsymbol{W}) = 1$, $bel(\varnothing) = 0$, and for any contingent P, if w_1, w_2 are such that $\neg(w_1 \sim w_2)$ and $(w_1,\ P;\ w_2)$, then $bel(P) = des((w_1,\ P;\ w_2)) - des(w_2)\ /\ des(w_1) - des(w_2)$.

(2) We have also already established that the previous equality holds if and only if $des((w_1,\ P;\ w_2)) = des(w_1).bel(P) - des(w_2).bel(P) + des(w_2)$. Since there is only one function satisfying the left-hand side, only one satisfies the right-hand side. ■

The foregoing establishes *Theorem 3*. Next we prove *Theorem 4*. Clearly, assuming that $<\boldsymbol{W}, \boldsymbol{P}, \boldsymbol{G}, \succcurlyeq>$ is a generalized Ramsey structure and *Condition 1* holds, the only thing we need to prove here is that if *Condition 2* holds then $bel(P) = 1 - bel(\neg P)$.[20]

(1) $bel(\boldsymbol{W}) = 1$ and $bel(\varnothing) = 0$ by stipulation, so we are only interested in contingent P.

(2) As already shown, for all contingent P, there is some $(w_1,\ P;\ w_2)$ such that the ratio $(des((w_1,\ P;\ w_2)) - des(w_2))\ /\ (des(w_1) - des(w_2))$ is defined (i.e., such that $\neg(w_1 \sim w_2)$). From *Theorem 3*, this is the value of $bel(P)$.

(3) From (GRS 1) and since $\boldsymbol{P}$ is closed under negation, if $(w_1,\ P;\ w_2)$ is in $\boldsymbol{G}$ then $(w_2,\ \neg P;\ w_1)$ is in $\boldsymbol{G}$; thus $bel(\neg P) = (des((w_2,\ \neg P;\ w_1)) - des(w_1))\ /\ (des(w_2) - des(w_1))$. Multiplying the denominator and the numerator by -1 gets us $bel(\neg P) = (des(w_1) - des((w_2,\ \neg P;\ w_1)))\ /\ (des(w_1) - des(w_2))$.

(4) *Condition 2* ensures $(w_1,\ P;\ w_2) \sim (w_2,\ \neg P;\ w_1)$, so $des((w_1,\ P;\ w_2)) = des((w_2,\ \neg P;\ w_1))$.

(5) Let $des((w_1,\ P;\ w_2)) = x$. Given the foregoing, $bel(P) + bel(\neg P)$ is equal to:

$$((x - des(w_2))\ /\ (des(w_1) - des(w_2))) + ((des(w_1) - x)\ /\ (des(w_1) - des(w_2)))$$

[20] Thanks to Rachael Briggs for the main outline of the following proof.

$= (x - w_2 + w_1 - x) \ / \ (w_1 - w_2) = (x - x + w_1 - w_2) \ / \ (w_1 - w_2)$

$= (w_1 - w_2) \ / \ (w_1 - w_2) = 1$. Condition (e) of *Theorem 4* follows immediately. ■

Finally, we can prove *Theorem 5*. As before, given *Theorem 3*, the only thing we need to prove here is that if *Condition 3* holds along with the other conditions then *bel* is monotonic.

(1) If $P = \varnothing$, then by stipulation, $bel(P) = 0$ and for all Q, $bel(P) \leq bel(Q)$. If $Q = \mathbf{W}$, then $bel(Q) = 1$ and for all P, $bel(P) \leq bel(Q)$. Thus suppose that P and Q are contingent and $P \subseteq Q$.

(2) From *Condition 3*, for all w_1, w_2, if some $(w_1, Q; w_2)$, $(w_1', P; w_2')$ exist, then: if $w_1 \succ w_2$, $(w_1, Q; w_2) \succcurlyeq (w_1', P; w_2')$, and if $w_2 \succ w_1$, $(w_1', P; w_2') \succcurlyeq (w_1, Q; w_2)$.

(3) We know from steps proven above that for all contingent Q, we will find at least one pair w_1, w_2 such that either (a) $w_1 \succ w_2$ and $(w_1, Q; w_2)$ exists, or (b) $w_2 \succ w_1$ and $(w_1, Q; w_2)$ exists. Given either (a) or (b), there are only two possibilities: $(w_1', P; w_2')$ exists in $\mathbf{G}$, or it does not.

(4) Suppose (a) is true, and that $(w_1', P; w_2')$ exists; thus $(w_1, Q; w_2) \succcurlyeq (w_1', P; w_2')$ and therefore $des((w_1, Q; w_2)) \geq des((w_1', P; w_2'))$ and $des(w_1) > des(w_2)$. With the established properties of *bel*, the former (weak) inequality implies $des(w_1).bel(Q) + des(w_2).(1 - bel(Q)) \geq des(w_1).bel(P) + des(w_2).(1 - bel(P))$. This can only hold given the latter (strict) inequality if $bel(Q) \geq bel(P)$.

(5) Suppose now that (b) is true, and that $(w_1', P; w_2')$ exists. By the same steps just given, $des(w_1).bel(P) + des(w_2).(1 - bel(P)) \geq des(w_1).bel(Q) + des(w_2).(1 - bel(Q))$ and $des(w_1) > des(w_2)$, which can only hold if $bel(Q) \geq bel(P)$.

(6) So given either (a) or (b) and the assumption that $(w_1', P; w_2')$ exists, $bel(Q) \geq bel(P)$.

(7) Assume now instead either (a) or (b), and that $(w_1', P; w_2')$ is not in $\mathbf{G}$. Because $(w_1, Q; w_2)$ is in $\mathbf{G}$ and all members of $\neg Q$ are members of $\neg P$, we know that there is some $w_2' \in \neg P$. Thus $(w_1', P; w_2') \notin \mathbf{G}$ only if there is no $w_1' \in P$. Since all

members of P are members of Q, Q must contain members of multiple values.

(8) All $w \in P$ must be such that for all $w' \in \neg Q$, $w \sim w'$; if instead $\neg(w \sim w')$, then since any member of P is also in Q and any member of $\neg Q$ is also in $\neg P$, we would have $w > w'$ and $(w, P; w')$, $(w, Q; w')$ in $\boldsymbol{G}$ or $w' > w$ and $(w, P; w')$, $(w, Q; w')$ in $\boldsymbol{G}$, contradicting the earlier supposition. Given this, Q must contain all $w \in \boldsymbol{W}$ such that $\neg(w \sim w')$, for any $w' \in P$.

(9) Given (GRS 2), all $\sim$-equivalence classes of worlds contain at least two members; thus, for any $w_1 \in Q$ and $w_2 \in P$ such that $w_1 \neq w_2$ and furthermore $\neg(w_1 \sim w_2)$, there is a $w_1' \in Q$ which is not in P. Let R be the proposition which contains all members of P, and w_1' but not w_1. Thus, $P \subset R \subset Q$. As $\boldsymbol{P}$ is the power set of $\boldsymbol{W}$, $R \in \boldsymbol{P}$.

(10) We now know that for some w_1, w_2', $w_1 > w_2'$ and $(w_1, Q; w_2')$, $(w_1', R; w_2')$ exist, or $w_2' > w_1$ and $(w_1, Q; w_2')$, $(w_1', R; w_2')$ exist. By the steps given earlier, $bel(Q) \geq bel(R)$. We also know that for some w_1, w_2, $w_1 > w_2$ and $(w_2, R; w_1)$, $(w_2, P; w_1)$ exist, or $w_2 > w_1$ and $(w_2, R; w_1)$, $(w_2, P; w_1)$ exist. By the same steps, then, $bel(R) \geq bel(P)$. By transitivity, $bel(Q) \geq bel(P)$. ■

References

Bartha, Paul 2007: 'Taking Stock of Infinite Value: Pascal's Wager and Relative Utilities'. *Synthese*, 54, pp. 5–52.

Bradley, Richard 2001: 'Ramsey and the Measurement of Belief'. In Corfield and Williamson 2001, pp. 263–90.

Chalmers, David 2011: 'Frege's Puzzle and the Objects of Credence'. *Mind*, 120, pp. 587–635.

Corfield, David and Jon Williamson (eds) 2001: *Foundations of Bayesianism*. Dordrecht: Kluwer.

Davidson, Donald, Sidney Siegel, and Patrick Suppes 1957: *Decision Making: An Experimental Approach*. Stanford, CA: Stanford University Press.

Davidson, Donald and Patrick Suppes 1956: 'A Finitistic Axiomatization of Subjective Probability and Utility'. *Econometrica*, 24, pp. 264–75.

Debreu, Gerard 1959: 'Cardinal Utility for Even-chance Mixtures of Pairs of Sure Prospects'. *Review of Economic Studies*, 26, pp. 174–7.

de Finetti, Bruno 1931: 'Sul Significato Soggettivo Della Probabilita'. *Fundamenta Mathematicae*, 17, pp. 298–329.

Fishburne, Peter C. 1967: 'Preference-Based Definitions of Subjective Probability'. *Annals of Mathematical Statistics*, 38, pp. 1,605–17.

Jeffrey, Richard 1983: *The Logic of Decision*, second edition. Chicago: University of Chicago Press, rpnt 1990.

Joyce, James 1999: *The Foundations of Causal Decision Theory*. New York: Cambridge University Press.

Krantz, David H., R. Duncan Luce, Patrick Suppes, and Amos Tversky 1971: *Foundations of Measurement Volume 1*. New York: Academic Press.

Ramsey, Frank P. 1931a: *The Foundations of Mathematics and other Logical Essays*. New York: Harcourt, Brace and Company.

Ramsey, Frank P. 1931b: 'Truth and Probability'. In his 1931a, pp. 156–98.

Savage, Leonard J. 1954: *The Foundations of Statistics*. New York: Dover.

Schmeidler, David 1989: 'Subjective Probability and Expected Utility without Additivity'. *Econometrica*, 57, pp. 571–87.

Sobel, Jordan Howard 1998: 'Ramsey's Foundations Extended to Desirabilities'. *Theory and Decision*, 44, pp. 231–78.

Suppes, Patrick 1956: 'The Role of Subjective Probability and Utility in Decision Making'. *Proceedings of the Third Berkeley Symposium on Mathematical Statistics and Probability*, 5, pp. 61–73.

von Neumann, John and Oskar Morgenstern 1944: *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press.

