Ramsey and the Measurement of Belief

Richard Bradley
Department of Philosophy, Logic and Scientific Method
London School of Economics
Houghton Street
London WC2A 2AE

November 23, 2007

Abstract

The problems of justifying the claims made by Bayesian decision theory and the problem of how to measure its main variables - degrees of belief and desire - are often confused. This paper argues that the main constituents of a solution to the latter (but not the former) are to be found in Ramsey's essay 'Truth and Probability'. Despite its influence, Ramsey's proposal for a method for measuring belief has received little detailed critical attention. This paper reconstructs the details of his argument in order to clarify its presuppositions. We argue that despite some technical problems, Ramsey's method remains unsurpassed and, in particular, that it may be interpreted in such a way as to avoid a well known objection to it due to Richard Jeffrey.

1 The Problem of Measurement

Bayesian decision theories are formal theories of rational agency that tell us both what the properties of a rational state of mind are and what action it is rational for an agent to perform, given her state of mind. What makes a theory of this kind *Bayesian* is commitment to the following claims:

1. The (only) factors relevant to rational decision making are the options available, the relative subjective desirability of the possible outcomes of choosing one or another option and the subjective likelihood that each outcome will be achieved, given the exercise of a particular option.

2. To act rationally is to choose the option (or one of the options) which has the best expected consequences, given one's partial beliefs and desires.

3. Rational degrees of belief are probabilities.

4. Rational degrees of desire are utilities.

Different versions of Bayesian decision theory, such as those of Savage and Jeffrey, differ with respect to the kinds of options they postulate, with respect to the precise interpretation of the second claim (the expected utility hypothesis) and with respect to further constraints they place on rational desire. But such differences will not affect our discussion here. (Nor will the question of whether Bayesians are committed to conditionalisation as means of revising partial beliefs).

This paper will discuss a particular problem for Bayesian decision theories that derives from the fact that its main variables - an agent's degrees of partial belief and desire - are not (directly) observable, but have to be inferred from what the agent says and does. It is my contention that there is much to be learnt in this regard from a proposal made by Frank Ramsey in his paper 'Truth and Probability'[10]. His proposal will be accordingly be examined in considerable detail and then, having established it viability, defended against some influential objections.

1.1 Two Kinds of Problems

Let us start by distinguishing two kinds of problems concerning the status of scientific theories which can be called the problems of justification and of measurement. The first is the problem of validating the claims made by a theory about the objects and relations of which it speaks. The second is the problem of determining, in a particular context, the values of the theory's variables. Of course, it is often the case that the same or similar observations will be used by scientists both to confirm a theory's claims and to determine the values of its variables. But the two problems are clearly distinct: the task of justifying Newton's first law of mechanics is different to that of measuring the forces, masses and acceleration of physical objects. Indeed the truth of theories may be presupposed in the interpretation of measurement observations, such as when Newton's law is used to calculate forces from masses and accelerations.

For a decision theory the problem of justification concerns both the claims it makes about the properties of consistent partial belief, desire and preference and the claims it makes about the rational way to act in the light of them. The measurement problem on the other hand boils down to the problem of determining an agent's partial beliefs and desires on the assumption that they have the properties postulated by the theory. Banal though the distinction between these problems may be, often enough it is neglected in debates within Bayesian decision theory. The reason is, I suspect, that Bayesian decision theories can be interpreted both as normative and as descriptive theories of agency and that its representation theorems can be construed as providing a solution to either the problem of justification or that of measurement, depending on the interpretation of the theory.

Let me explain in a bit more detail. Decision-theoretic representation theorems show that some set of conditions on an agent's preferences determines the existence of a quantitative representation of her partial beliefs and desires, consistent with her preferences. If the conditions on preference are naturally construed as conditions of rationality, then these theorems can serve to address the problem of justification in the sense that they show why the theory is normatively compelling. For they imply that by accepting that rational preference should satisfy the conditions in question, one is committed to accepting the theory's claims about the properties of rational partial belief and desire. The problem of justification then reduces to a defence of these qualitative rationality conditions on preference - supposedly an easier task.

I simplify considerably, of course. Some of the conditions that such theorems postulate are not rationality conditions but idealisations, designed to make the issue mathematically tractable. The ubiquitous completeness and continuity conditions are cases in point. But the thought is that these idealisations do not distort the main results, so that the relation between, for instance, incomplete preferences and incomplete probabilities and utilities should approximate the relation between complete ones. It would be good to see this rigorously demonstrated, but the claim has considerable plausibility. For why should making up my mind about matters concerning which I previously had no attitude affect, say, whether my beliefs about those that I do are probabilities or not?

Representation theorems also function as demonstrations of the measurability of the main decision theoretic variables - degrees of belief and desire. Or rather they show that if the idealised conditions postulated by the axioms of preference are realised then knowledge of agents' preferences suffice for knowledge of their degrees of belief and desire. In this context the way in which we evaluate the axioms is different from before. When justification was at issue we asked of the axioms of preference whether they were really rationality conditions. When the task is measurement, and our interest descriptive rather than normative, we ask whether the conditions are actually satisfied by agent's preferences or whether they can be made to satisfy them. These questions may be connected in practice. There may be evolutionary grounds, for instance, for supposing that the preferences of actual agents will by and large satisfy rationality conditions. But in principle, one is not required to solve the problem of justification in order to solve the measurement problem. In particular it is not necessary to establish the normative validity of Bayesian decision theory to demonstrate the measurability of the variables it postulates.

1.2 The Evidence Base

We now focus on the concern of this paper: the problem of measurement. To solve it we need to do two things. First we need to identify what sort of evidence can be used to determine agents' states of partial belief and desire. And second, we need show how and to what extent the evidence in question determines a measure of these states. We shall attend to the first question in this section and then devote most of the rest of the paper to the second, taking for granted a particular answer to the first.

1.2.1 Behaviourism

Historically discussion of this issue has been dominated by what I will call Epistemological Behaviourism: a rather puritanical doctrine of empiricist descent according to which the only acceptable evidence for an agent's mental states is intersubjectively observable behaviour. Evidence obtained by introspection, in particular, should not be countenanced as introspected states are not intersubjectively observable (even if 'observable' by the person whose states they are). As a consequence reports by someone on their mental states are to count as evidence only insofar as any perceptible effect of a mental cause can count as evidence. It does not follow from someone reporting that they believe x, that they do so, except if it has already been established through behavioural evidence that such belief reports are reliable.

Epistemological Behaviourism is to be distinguished from the associated metaphysical and semantic doctrines according to which mental states are nothing more than dispositions to behaviour or, more radically, than sets of observable behaviours (e.g. to desire that x is to act in ways which tend to bring it about that x, to prefer x to y is to choose x over y whenever both are available). Both forms of Behaviourism have been influential in the development of decision theory. On the metaphysical side, for instance, De Finetti and to some extent Savage saw themselves as giving behavioural meaning to the concept of probability, while Samuelson saw his axioms of weak preference as giving behavioural meaning to the concept of preference. Ramsey too gestures in this direction.

It is with the epistemological form of Behaviourism that I am concerned, however, because of its influence on Bayesian decision theorists' understanding of the problem of the measurement. In behaviourist spirit, Bayesian decision theorists typically take a satisfactory solution of this problem to be a demonstration that decision theoretic representations of the states of minds of rational agents are characterisable in terms of sets of intersubjectively observable behaviours, that there exist (so to speak) defining sets of observable behaviours for each possible state of mind. This attitude is exemplified Savage's claim regarding strengths of belief:

> "If the state of mind in question is not capable of manifesting itself in some sort of extraverbal behaviour, it is extraneous to our main interest. If , on the other hand, it does manifest itself through more material behaviour, that should, at least in principle, imply the possibility of testing whether a person holds one event to be more probable than another, by some behaviour expressing, or giving meaning to, his judgement."[1]

If such a reduction of belief and desire to preference were possible it would, of course, represent a solution of a kind, for it would imply that a set of observations of someone's behaviour with the right kinds of properties would suffice

[1] Savage [12, p. 27-28]. This passage is a little misleading: Savage's views seem on closer examination to be more like those that I will attribute to Ramsey. But it accurately represents the behaviourist tradition in decision theory.

to determine the values of the main decision theoretic variables. But contrary to what seems to be prevailing belief, such a reduction has never been successfully achieved and I doubt very much will ever be so. The argument for this is quite simple. An agent's state of indifference between two options is something that decision theories can (and should) represent. But no behaviour can ever attest in a conclusive manner to someone's indifference between two or more possibilities. So decision-theoretic models are necessarily underdetermined by behavioural evidence (even all possible behavioural evidence, elicited under ideal conditions).

Let me elaborate a little by consideration of the example of Savage's representation theorem, since this is probably the best known.[2] Formally what Savage shows is that a binary relation, $\geq$, on a set of actions (actions on his account being functions from events to outcomes) that satisfies Savage's axioms will determine the existence of a utility function, u, on outcomes (unique up to choice of scale) and a unique probability function, Pr, on events such that for all actions, a_1 and a_2, $a_1 \geq a_2 \iff EU(a_1) \geq EU(a_2)$, where $EU(a_1)$ is the expectation of utility given a_1, relative to u and Pr. If $\geq$ is interpreted as the 'at least as preferred as' relation then it is plausible to construe u and Pr as respectively measures of the agent's degrees of desire and belief. We are now very near to what we want. Having reduced quantitative mental states - degrees of belief and desire - to a qualitative one - preference - all we require now is a reduction of the 'at least as preferred as' mental relation to some behavioural correlate. And this seems like it should not be too difficult as preferences between actions will be directly manifested in the agent's choice of action under suitable circumstances.

So near and so yet so far. Even if we set aside obvious practical difficulties and grant that failure to perform an action indicates the presence of a preferred action rather than agent's ignorance of its availability, Savage's theory cannot be construed in a behaviouristically acceptable manner. For an agent's choices reveal, if anything, her strict preferences between options and not the 'at least as preferred as' relations that are the subject of Savage's rationality conditions. Moreover, strict preferences need not satisfy these conditions. In particular it would be wrong to require that strict preferences be complete, not just because this requires agents to have too much knowledge, but because agents can legitimately be indifferent between options.

Could indifference not be revealed by a failure to choose when confronted by a non-empty set of alternatives? But such behaviour has a number of possible interpretations. It could show that the agent prefers doing nothing to all of these options. Or that two or more of the options are equally preferred. Or even that some of the options are not comparable. Each possibility implies something different for the measurement process: the first that the 'do nothing' option has greater subjective expected utility than the others, the second that option set needs to be refined so as to determine which options it is that have equal subjective expected utility; and the third that the measurement cannot

[2] See Savage [12].

be completed since one of its preconditions (the comparability of the options) is not satisfied. Behavioural evidence may allow elimination of the first possibility, but in principle it could not discriminate between the second and third. One might of course eliminate the problem by forcing a choice in every circumstance, but the evidence so obtained could not then be used for the intended purpose without producing measurement artifacts.

1.2.2 The Use of Verbal Evidence

To my mind the search for a *reduction* of belief and desire to choice is misguided. We should accept that if Savage's theorem (or one like it) does what it says then it successfully reduces quantitative belief and desire to preference. We can also accept that observed choice constrains (but does not determine) attributions of preference to agents. But to attribute determinate preferences to an agent we need to accept evidence of a wider source and in particular the verbal reports of the agent concerning their preferences and (crucially) instances of their indifference between possibilities.

This is not to grant agents any kind of epistemic privilege with regard to their mental states. People certainly can be wrong about such things. But Behaviourism draws too sharp a contrast between the status of behavioural and verbal evidence, falsely identifying the former with what is observable and indubitable and the latter with what is not. In fact both in theory and in practice the distinction is not nearly so neat. What people report clearly *can* be evidence for what they really think and feel, certainly not indubitable and perhaps different in kind from the evidence gleaned by observing their behaviour, but not therefore better or worse. Indeed it is hardly conceivable in practice that we could do without verbal reports. For language has the great advantage of allowing precise formulation of the alternatives with respect to which we wish to determine agents' attitudes. Granted, this can raise the question as to whether the observer and the subject have a shared understanding of what is conveyed by particular linguistic expressions. But these difficulties are hardly more severe than those affecting the interpretation of their behaviour - it is just as easy to misinterpret the meaning of a wave of a hand as the meaning of an utterance.

The objection to introspection based evidence is not much more sustainable in theory than in practice. The supposed indubitability of behavioural evidence rests on the idea that because, perceptual illusions aside, we can *see* what someone is doing, there can be no doubting evidence statements of the 'The agent did such and such' kind. But in reality observations of behaviour (or at very least their linguistic representation) are always 'polluted' by interpretations: we don't see someone making circular motions with a cloth on the table, we see them wiping it clean; we don't see limbs describing particular trajectories, we see someone going for a walk; and so on. These interpretations are contestable and observers with different background knowledge will often interpret what someone is doing in different ways. Intersubjective observability does not mean either certainty or consensus.

Once one gives up on the goal of deriving all knowledge claims from indubitable foundations, we need not take an all or nothing attitude to the results of introspection. Under certain conditions introspection may reliably produce evidence of a particular kind, under others it may be less satisfactory. A decision about what sort of evidence to admit in the determination of agents' mental states must be made and motivated. Inevitably there will be trade-offs between the reliability of the evidence admitted and its richness. And depending on the attitude taken to the admissibility of different types of evidence, there will be different problems of measurement and different methods for solving it that depart from different trade-offs. For this reason, in addition to the usual ones, the solution defended here cannot be regarded as the only one deserving consideration.

1.2.3 Ramsey's Problem

Having granted that there is more than one way in which the problem of measurement can be formulated, let us now concentrate on the version of the problem found in Ramsey's essay "Truth and Probability"[10]. Ramsey's main concern in this paper is to explicate the concept of probability, but as he held that notions like this had no precise meaning unless a measurement procedure is specified, much of the paper is devoted to addressing our problem. We need not of course endorse Ramsey's operationalism in order to learn from his measurement methods.

Ramsey's thinking on the question of the measurement of belief seems at first to be very much in the behaviourist mould. He argues, for instance, that the idea that believing something more or less strongly was connected to a perceptible feeling of belief of a certain intensity is:

> "... observably false, for the beliefs which we hold most strongly are often accompanied by practically no feeling at all; no one feels strongly about things he takes for granted."[3]

And that:

> "when we seek to know what is the difference between believing more firmly and believing less firmly, we can no longer regard it as consisting in having more or less observable feelings; at least I personally cannot recognise any such feelings. The difference seems to me to lie in how far we should act on these beliefs ..."[4]

But Ramsey is not in fact rejecting introspection wholesale, only the particular use of introspection associated with the idea of measuring strength of belief in terms of the sensations or feelings that accompany it. Indeed in the argument just quoted he makes use of introspective evidence: his own failure to perceive a feeling corresponding to his belief. And further on, when he argues that although we may feel that:

[3] Ramsey [10, p. 65]

[4] ibid, p.66

> "we know how strongly we believe things and that we can only know this if we can measure our belief by introspection ... our judgement about the strength of our belief is really about how we should act in hypothetical circumstances." [5]

the judgement that he refers to - as to how we would act under hypothetical circumstances - is presumably an introspective one. In fact, as I shall argue in greater detail later on, not much sense can be made of Ramsey's measurement procedure unless introspective evidence of this kind is admitted.

Ramsey takes his arguments to show that we can although we might be able to introspect whether we do or do not believe something, there is no reliable way of introspecting the degree to which we do. It would appear that this suspicion of introspection, if not his arguments against it, extends to the possibility of qualitative judgements as to whether one believes one thing more strongly than another, despite the fact that they seem not dissimilar in nature to judgements as to how we would choose or act in particular circumstances. In any case, the upshot is that he admits only evidence as to the choices that an agent does or would make between specified alternatives and not their direct reports on their partial attitudes. In this respect Ramsey's approach is very much in line with the norm in decision theory. While this is not the only reasonable formulation of the problem, it has the advantage of expressing it as it is most commonly understood by Bayesian decision theorists.

1.3 Ramsey's Theory of Measurement

1.3.1 Conditional Prospects

The problem, as I have presented it, is to explain how the evidence relating to an agent's choices or actions determines a decision-theoretic representation of her partial belief and desires. Any approach that admits only evidence of this kind faces what is frequently termed the problem of the simultaneous determination of belief and desire: essentially that of untangling the respective contributions made by an agent's beliefs and desires to her choices. To solve it Ramsey makes use of two important ideas: that of a conditional prospect and that of an ethically neutral proposition.

Conditional prospects are possibilities such as that if its hot today, then it will rain tomorrow and, if not, it will snow. Such prospects have come to be termed 'gambles' or 'actions' in the literature on Ramsey, although he does not use either term. Both are misleading in some important ways, and I will largely avoid them.

Ramsey takes it for granted that the desirability of choosing a particular conditional prospect is related in a precise way to the desirability of the possible states of the world consistent with it. Namely, that the former is a weighted average of the latter, with the weights coming from the agent's degrees of belief. Formally, let $\Pr(P)$ be a measure of the degree to which the agent believes that

[5] Ibid, p.67

P and $U(\alpha)$ and $U(\beta)$ be measures of the degrees to which she desires that respectively α and β be the actual state of the world. Then the desirability, $U(\Gamma)$, of the conditional prospect Γ, that α be the case if P is and that β be the case if not, is its expectation of utility:

Proposition 1 $U(\Gamma) = U(\alpha).\Pr(P) + U(\beta).(1 - \Pr(P))$

Ramsey doesn't attempt to justify this assumption, arguing that although the theory upon which it rests:

> "cannot be made adequate to all the facts, ... [it is] a useful approximation to the truth, particularly in our self-conscious or professional life, and it is presupposed in a great deal of our thought."[6]

It might well be objected that Ramsey's assumption about the way in which expectations determine desires is a good deal more specific than the sort of 'folk-psychology' that arguably much of our thought presupposes. This objection would not, I think, worry Ramsey much as he held that concepts like partial belief or utility are at least partially defined by the procedures for measuring them.[7] And Ramsey quite explicitly ties his method to the measurement of partial attitudes qua bases or causal components of actions, the concept of partial attitude that has come to predominate in the social sciences. Ramsey seems on safe ground here in that the quantities implicitly defined by his measurement scheme turn out to have many of the properties commonly attributed to them by social scientists. This fact alone may suffice to motivate the assumptions that Ramsey makes about the nature of partial belief and desire, but it should be remembered that the motivation is only as strong as the consensus amongst social scientists from which it stems.

But we are straying into the issue of justification, which we have already undertaken to set aside. The important fact is that Proposition 1 expresses (albeit formally) no more than what I previously stipulated as one of the defining contentions of Bayesian decision theory: namely that they desirability of an option is given by the expected benefit of choosing it, given one's beliefs. In attempting to solve a measurement problem it is perfectly appropriate to assume the truth of the theory whose variables require measurement. So there is no requirement here to justify Proposition 1 any further.

1.3.2 Ethically neutral propositions

The second critical innovation of Ramsey's is the postulation of what he calls ethically neutral propositions. An ethically neutral proposition is simply one whose truth or falsity is a matter of indifference to the agent and does not affect their attitude to any other prospects e.g. the prospect of a dust storm on Mars does not influence any of my earthly concerns.

[6] ibid, p.69

[7] "the degree of belief is like a time interval; it has no precise meaning unless we specify how it is to be measured" ibid, p.63

Crucially the probabilities of some ethically neutral propositions can be inferred from an agent's preferences. Suppose, for instance, that an agent is not indifferent between the prospect of sun and that of snow, but indifferent between the prospect that if P is true, then it will be sunny, but if P is not, then it will snow, and the prospect that if P is true, then it will snow, but if P is not, then it will be sunny. Then we can infer that they regard P as likely to be true as not. For were it not, they should prefer one of the conditional prospects over the other (this follows directly from Proposition 1). The proposition that a toss of the coin in my hands will land heads up may be an example of such a proposition.

1.3.3 Ramsey's Method (Informally)

Suppose we want to determine Mary's attitudes to the various kinds of weather that the next day might bring: sun, snow, rain, and so on. Then if

> "we had the power of the Almighty, and could persuade our subject of our power" [8]

we could offer each kind of weather as an option to be exercised if she so chooses. By getting her to choose between any pair of them we obtain a ranking of all of her prospects and can assign some arbitrary number - say 1 and 0 for simplicity - to the top and bottom ranked ones.

Let us suppose that in Mary's case the results of coin tosses are indeed ethically neutral and of probability one-half and that the top ranked prospect is a sunny day tomorrow and the bottom ranked one is snow tomorrow. Now, Mary's attitude to the prospect that there will be sun tomorrow if the coin lands heads and snow if it lands tails will depend on the degree to which she respectively desires sunny and snowy days and the degree to which she believes that the coin will land heads or tails. But she considers the latter to be equiprobable, so she will regard the conditional prospect of sunny and snowy days in the event of heads or tails, as being midway in desirability between the prospects of a sunny day and that of a snowy one. (Again, this follows directly from Proposition 1). Relative then to our arbitrary choice of values for the top and bottom of the ranking, this means it has a utility of 0.5. And so too will any prospect ranked with this 'gamble'.

Just as we identified the value 0.5 with a particular conditional prospect we can, by compounding conditional prospects, identify any point in the interval from 0 to 1 with some 'gamble' on the top and bottom ranked prospects. For instance, 0.25 might be identified with the prospect that, in the event of the first toss of the coin coming up heads, there will be sun tomorrow if a second toss of the coin lands heads as well and snow if it lands tails, and that, in the event of it coming up tails on the first toss, there will be snow tomorrow. This gives us a scale with which to measure the utility of all prospects, whatever their form.

[8] ibid, p. 72. It is not in fact really necessary to have any powers other than those of persuasion.

We can also use Mary's choices amongst her options to determine her degrees of belief. Suppose that she is indifferent between the prospect of it being cloudy tomorrow and that of it being sunny tomorrow if it rains tonight and of it being snowy if it doesn't. Then the utility of the prospect of cloudy weather tomorrow, relative to that of sunny and snowy weather, indicates the degree to which Mary believes that it will rain tonight: the closer the utility of cloudy weather is to that of sunny weather, the stronger must be her belief that it will rain, the closer it is to the utility of snowy weather the weaker must be her belief. In general, if Mary is indifferent between the prospect that α and the prospect that β if P is the case and that γ if P is not, then by rearrangement of Proposition 1:

$$\Pr(P) = \frac{U(\alpha) - U(\gamma)}{U(\beta) - U(\gamma)}$$

So, on the assumption that the right conditional prospect can always be found, we have a means to determine Mary's degrees of belief in every proposition, whether ethically neutral or not.

2 Reconstructing Ramsey's Account

2.1 The Basic Framework

Our informal presentation suggests that Ramsey has an simple, elegant and effective method for measuring belief and desire to offer us. Unfortunately, Ramsey does not work all the details of his theory, claiming at one point:

> "this would, I think, be rather like working out to seven places of decimals a result only valid to two"[9].

But it is clearly important to determine whether his demonstration can work in principle. The literature on Ramsey is sadly lacking in this respect. Many decision theorists have drawn inspiration from Ramsey, without engaging with the details of his work or by reconstructing it in terms of ideas drawn from Savage and others.[10] What little expository literature there is tends to remain too distant from the details or too uncritical of Ramsey's claims.[11] This has meant that crucial assumptions have gone unnoticed. The only exception that I am aware of is Sobel [14], which carefully examines Ramsey's argument, filling in some of the details and making suggestions for amendments. Sobel does not however provide all the proofs that are missing (e.g. of the existence of a utility representation of preference) nor does he examine all the ones that Ramsey does give (e.g. of the additivity of degrees of belief).

[9] ibid, p. 76
[10] For instance, Davidson and Suppes [4], Jeffrey [6], Bolker [3]
[11] For instance, Sahlin, N.-E. (1990) [11]

2.1.1 Worlds and propositions

Ramsey makes a distinction between the objects of belief - propositions - and the objects of desire - prospects: possible courses of the world (worlds for short) and conditional prospects or 'gambles'. We denote propositions by italic upper case letters, worlds by lower case Greek letters and arbitrary prospects by upper case Greek letters. In Ramsey's framework conditional prospects are essentially functions from partitions of propositions to worlds, with a function from the partition $\{X_1, X_2, .., X_n\}$ to worlds α_1, α_2, ..., α_n being written as (α_1 if X_1)(α_2 if X_2)...(α_n if X_n). We take for granted that the ordering does not matter: (α_1 if X_1)(α_2 if X_2) is, for instance, the same conditional prospect as (α_2 if X_2)(α_1 if X_1).

As Ramsey speaks of worlds as being compatible with propositions or of including their truth, worlds must be the sorts of entities that can stand in logical relationships to propositions. He defines them as:

> "the different possible totalities of events between which our subject chooses - the ultimate organic unities"[12]

From this, and on the basis of what he does with worlds, it might seem that we should think of them as something like maximally specific propositions, so that for all worlds α and propositions X, α implies that X or α implies that not X. (I shall represent the fact of α implying that X in the manner of possible worlds semantics, by writing $\alpha \in X$). In one respect, however, this cannot be exactly right. In his definition of an ethically neutral proposition of probability one half, and elsewhere, Ramsey implicitly assumes the existence of worlds that are compatible with both the truth and falsity of ethically neutral propositions. So either we have to say that worlds are only maximally specific about things that matter to the agent (and, hence, qualify the claims of the previous paragraph) or reformulate the relevant definitions. One way of doing so is to introduce world-like entities with just the right lack of specificity concerning the truth values of ethically neutral propositions. The matter has been thoroughly explored in Sobel [14] and so I permit myself to ignore here the complications that it gives rise to.

2.1.2 Ethical Neutrality.

The concept of an ethically neutral proposition is intuitively clear: it is just a proposition whose truth or falsity does not affect an agent's attitude to any of her prospects. But to express this more formally Ramsey needs to be able to say what it is for two prospects to differ from another only with respect to the truth of the proposition in question. Worlds, however, cannot differ with respect to the truth of a single proposition only. To get around this problem Ramsey assumes the existence of a class of atomic propositions; propositions which are true or false independently of one another and such that no two worlds are exactly

[12] Ramsey [10, p. 72-72]

alike with regard to the truth of all of them. An atomic proposition P is then defined to be ethically neutral for some agent iff she is indifferent between any two worlds which differ only in respect to the truth of P. Finally, non-atomic propositions are said to be ethically neutral iff all their atomic truth-arguments are. This cost of this formulation is that it ties his account to Wittgenstein's logical atomism to a far greater extent than I imagine he would care.[13]

Are there any ethically neutral propositions? The standard candidates are propositions such as that the next card drawn will be an Ace or the coin will land heads. But the truth of such propositions do affect agents' attitudes to some prospects. Take any proposition X and suppose that the agent is not indifferent to prospect that A. Then the truth of X will be a matter of consequence to her attitude to the prospect that A if X, because in the event that X is true, the prospect that A if X amounts to that of A. So X is not ethically neutral. It follows that there are no propositions that are neutral with respect to all prospects.

On the other hand, Ramsey only assumes that ethically neutral propositions leave the agents' attitudes to *worlds* unaffected. So, as long as conditional prospects do not serve to discriminate worlds, the supposition of ethically neutral propositions will not seem unreasonable.[14] And Ramsey's definition of an ethically neutral proposition of probability one-half, one of the linchpins of his method, will be well-founded.

Definition 2 *Let P be any ethically neutral proposition and α and β be any worlds consistent with both the truth and the falsity of P and such that the agent prefers α to β. Then P is of* ***probability one-half*** *iff the agent is indifferent between the prospect of $(\alpha$ if $P)(\beta$ if $\neg P)$ and that of $(\beta$ if $P)(\alpha$ if $\neg P)$.*

2.2 The Existence of a Utility Measure

2.2.1 Axiomatising Preference

Under what conditions will the kind of measurement process that we previously described yield a measure of the strength of an agent's desires and to what extent will such a measure be unique? Ramsey answers this question by stating a representation theorem for preference orderings of prospects that establishes the existence of quantitative (utility) representations of the agent's degree of preference (or desire). He does not, however, give a full proof of this theorem. As far as I know nobody else has supplied one either. Indeed, despite the fact that Ramsey gives strong indications of how the proof should go, there seems to

[13] Sobel [14, p. 237] says that Ramsey is not committed to the claim that every proposition is a truth functional compound of atomic ones. But without this claim his definition of ethical neutrality is not completely general. See Sobel for a discussion of what he calls Ramsey's thin logical atomism.

[14] The question left open is: if worlds are not distinguished by conditional prospects, what is the relation between the two? This question, indeed the more general one of the logical status of conditional prospects is never addressed by Ramsey. We return to the problem right at the end of the paper.

be very little recognition of the fact that the basic strategy of the proof differs to a considerable extent from others to be found in the decision theoretic literature. Perhaps the problem is a natural tendency to read his work through the lenses of either von Neumann and Morgenstern or Savage.

I will begin with a statement of Ramsey's representation theorem, before returning to the question of his strategy for proving it. Though I will remain fairly close to the letter of Ramsey's account, I will deviate from it at points in the interests of clarity and ease of understanding. Let a preference relation, $>$, and an indifference relation, $\approx$, be defined over the set of all prospects, Γ: a set of worlds plus all conditional prospects defined with respect to them (relative to a given set of propositions). By $\Phi \geq \Psi$ is meant that $\Phi > \Psi$ or $\Phi \approx \Psi$. Although Ramsey does not explicitly postulate it, he takes it for granted in subsequent discussion that agents' preferences are complete i.e. that $\Phi > \Psi$ or $\Psi > \Phi$ or $\Phi \approx \Psi$.

Let us call the set of prospects equally preferred to Φ, its **value** and denote it by $\bar{\Phi}$, avoiding Ramsey's economical, but often confusing, practice of denoting values as well as worlds by Greek letters. We write $\bar{\Phi} \succeq \bar{\Psi}$ iff $\Phi \geq \Psi$. While Ramsey directly axiomatises the relation on values induced by the preference relation on prospects, I will state the axioms in terms of the latter. His way of doing it obscures some issues of importance to our discussion and is easily recovered from ours. Let Δ be a non-empty set of ethically neutral propositions of probability one-half and suppose that P belongs to Δ. Then Ramsey postulates:

R1 If $Q \in \Delta$ and $(\alpha \text{ if } P)(\beta \text{ if } \neg P) \geq (\gamma \text{ if } P)(\delta \text{ if } \neg P)$, then:
$(\alpha \text{ if } Q)(\beta \text{ if } \neg Q) \geq (\gamma \text{ if } Q)(\delta \text{ if } \neg Q)$.

R2 If $(\alpha \text{ if } P)(\delta \text{ if } \neg P) \approx (\beta \text{ if } P)(\gamma \text{ if } \neg P)$ then:
(i) $\alpha > \beta \Longleftrightarrow \gamma > \delta$
(ii) $\alpha \approx \beta \Longleftrightarrow \gamma \approx \delta$.

R3 If $\Phi \geq \Psi$ and $\Psi \geq \Theta$, then $\Phi \geq \Theta$.

R4 If $(\alpha \text{ if } P)(\delta \text{ if } \neg P) \geq (\beta \text{ if } P)(\gamma \text{ if } \neg P)$ and $(\gamma \text{ if } P)(\zeta \text{ if } \neg P) \geq (\delta \text{ if } P)(\eta \text{ if } \neg P)$, then $(\alpha \text{ if } P)(\zeta \text{ if } \neg P) \geq (\beta \text{ if } P)(\eta \text{ if } \neg P)$

R5 $\forall(\alpha, \beta, \gamma)[\exists(\delta)$: $(\alpha \text{ if } P)(\gamma \text{ if } \neg P) \approx (\delta \text{ if } P)(\beta \text{ if } \neg P)]$

R6 $\forall(\alpha, \beta)[\exists(\delta)$: $(\alpha \text{ if } P)(\beta \text{ if } \neg P) \approx (\delta \text{ if } P)(\delta \text{ if } \neg P)]$

R7 Axiom of Continuity

R8 Archimedean Axiom.

I have slightly strengthened Ramsey's first, third and fourth axiom by stating them in terms of the weak order '$\geq$' rather than the indifference relation '$\approx$'. In the presence of R2, my R5 and R6 are jointly equivalent to his fifth and sixth axioms. Ramsey doesn't say what he intends by R7 and R8, though R7 is presumably about the continuity of preferences and R8 about the values of worlds. In particular, what is required of R8, whatever its precise formulation, is that it allows the derivation of the Archimedean condition referred to in Definition 7 below.

Presumably R7 is meant to ensure that for every conditional prospect (α if X)(β if $\neg X$) there exists a world γ such that (α if X)(β if $\neg X$) $\approx \gamma$, so that every value contains a world. Apart from simplifying movement between values and worlds, this assumption plays a crucial role in his subsequent derivation of degrees of beliefs. Furthermore, as Ramsey doesn't assume (as we did in our informal presentation of his method) the existence of compound conditional prospects - prospects of the form (Φ if X)(Ψ if $\neg X$), where Φ and Ψ are conditional prospects - he does need something to ensure, for any two prospects, the existence of a conditional prospect whose value lies midway between theirs. This in turn is required for Ramsey to conclude that "These axioms enable the values to be correlated one-one with real numbers ...".[15]

2.2.2 Intervals of Values

Ramsey's next move is to "explain what is meant by the difference in value of α and β being equal to that between γ and δ" by defining it to mean that if P is an ethically neutral proposition of probability one-half then the agent is indifferent between (α if P)(δ if $\neg P$) and (β if P)(γ if $\neg P$). Although Ramsey seems to be speaking here of a relation (of difference in value) between *worlds*, what he really needs for his representation theorem is a definition of a difference relation between *values* of worlds.[16] So let us denote the difference between the values $\bar{\alpha}$ and $\bar{\beta}$ by $\bar{\alpha} - \bar{\beta}$ and define an ordering, $\succeq$, of **differences in values** as follows (the coherence of the definition is guaranteed by R1):

Definition 3 *If $P \in \Delta$ then $\bar{\alpha} - \bar{\beta} \succeq \bar{\gamma} - \bar{\delta}$ iff $\forall(\alpha \in \bar{\alpha}, \beta \in \bar{\beta}, \gamma \in \bar{\gamma}, \delta \in \bar{\delta})$, ($\alpha$ if P)(δ if $\neg P$) $\geq$ (β if P)(γ if $\neg P$).*

With the concept of difference of value in hand, it becomes much easier to understand the role played by Ramsey's preference axioms. Essentially they are there to ensure that it be possible to give a numerical representation of not only such facts as the agent preferring one thing to another, but also of the extent to which their preference or desire for one thing exceeds their desire for another. To this end axiom R1 ensures the coherence of the definition of differences in values, axioms R5, R6, R7 and R8 a correspondence between values and real numbers, and R2, R3, and R4 that the difference operation on values functions like the subtraction operation on real numbers.

With regard to the latter, note that Axiom R4, opaque in its original formulation, translates as an axiom of transitivity for differences in values:

R4* If $\bar{\alpha} - \bar{\beta} \succeq \bar{\gamma} - \bar{\delta}$ and $\bar{\gamma} - \bar{\delta} \succeq \bar{\eta} - \bar{\zeta}$, then $\bar{\alpha} - \bar{\beta} \succeq \bar{\eta} - \bar{\zeta}$

[15] Ramsey [10, p. 75]

[16] This ambiguity is reproduced without comment in most expositions of Ramsey.

while, as we now prove, it follows from the definitions of ethical neutrality and the difference operation that if $\bar{\alpha}-\bar{\beta} \succeq \bar{\gamma}-\bar{\delta}$ then $\bar{\beta}-\bar{\alpha} \succeq \bar{\delta}-\bar{\gamma}$ and $\bar{\alpha}-\bar{\gamma} \succeq \bar{\beta}-\bar{\delta}$.

Lemma 4 *If $\bar{\alpha}-\bar{\beta} \succeq \bar{\gamma}-\bar{\delta}$ then:*
(i) $\bar{\delta}-\bar{\gamma} \succeq \bar{\beta}-\bar{\alpha}$
(ii) $\bar{\alpha}-\bar{\gamma} \succeq \bar{\beta}-\bar{\delta}$

Proof. Suppose that $P \in \Delta$. Omitting explicit quantification over worlds where the meaning is obvious, we note that $\bar{\alpha}-\bar{\beta} \succeq \bar{\gamma}-\bar{\delta} \Leftrightarrow (\alpha \text{ if } P)(\delta \text{ if } \neg P) \geq (\beta \text{ if } P)(\gamma \text{ if } \neg P)$. But by Definition 2, (i) $(\alpha \text{ if } P)(\delta \text{ if } \neg P) \geq (\beta \text{ if } P)(\gamma \text{ if } \neg P) \Leftrightarrow (\delta \text{ if } P)(\alpha \text{ if } \neg P) \geq (\gamma \text{ if } P)(\beta \text{ if } \neg P) \Leftrightarrow \bar{\delta}-\bar{\gamma} \succeq \bar{\beta}-\bar{\alpha}$, and (ii) $(\alpha \text{ if } P)(\delta \text{ if } \neg P) \geq (\beta \text{ if } P)(\gamma \text{ if } \neg P) \Leftrightarrow (\alpha \text{ if } P)(\delta \text{ if } \neg P) \geq (\gamma \text{ if } P)(\beta \text{ if } \neg P) \Leftrightarrow \bar{\alpha}-\bar{\gamma} \succeq \bar{\beta}-\bar{\delta}$. ■

2.2.3 Proving the Representation Theorem

We are now in a position to state Ramsey's theorem establishing the existence of utility measures of agents' desires. Ramsey does not give a uniqueness theorem for such utility measures, but his subsequent discussion of the measurement of probabilities assumes that they are unique up to affine linear transformation (or choice of scale) i.e. that preferences are interval-scale measurable. We state below the theorem he requires.

Theorem 5 (Existence) *There exists a utility function, U, on the set of all values, $\bar{\Gamma}$, such that $\forall(\bar{\alpha}, \bar{\beta}, \bar{\gamma}, \bar{\delta} \in \bar{\Gamma})$, $\bar{\alpha}-\bar{\beta} \succeq \bar{\gamma}-\bar{\delta} \Leftrightarrow U(\bar{\alpha})-U(\bar{\beta}) \geq U(\bar{\gamma})-U(\bar{\delta})$*

Theorem 6 (Uniqueness) *If U' is another such a utility function, then there exists real numbers a and b, such that $a>0$ and $U' = a.U+b$.*

The key to understanding Ramsey's representation theorem is to recognise that it implicitly draws on the theory of measurement deriving from the work of the German mathematician Hölder (with which Ramsey would have been familiar). We begin with a statement of the relevant results in this area, drawing from their presentation in Krantz et al [7, chapter 4].[17]

Definition 7 *Let A be a non-empty set and $\succeq$ a binary relation on $A \times A$. Then $< A \times A, \succeq >$ is an* ***algebraic difference structure*** *iff $\forall(a, b, c, d, a', b', c' \in A)$:*
1. $\succeq$ is a complete and transitive
2. If $ab \succeq cd$, then $dc \succeq ba$
3. If $ab \succeq a'b'$ and $bc \succeq b'c'$ then $ac \succeq a'c'$
4. If $ab \succeq cd \succeq ba$ then there exists $x, x' \in A$, such that $ax \approx cd \approx x'b$
5. Archimedean condition

[17] The authors point out that Hölder's results can be applied to the problem of measurement of degrees of preference, but (oddly) make no attempt do so directly. Nor is there explicit recognition of the use that Ramsey makes of them.

Theorem 8 *If $< A \times A, \succeq >$ is an algebraic difference structure, then there exists a real-valued function, ϕ, on A, such that $\forall(a, b, c, d \in A)$:*

$$ab \succeq cd \Leftrightarrow \phi(a) - \phi(b) \geq \phi(c) - \phi(d)$$

Furthermore, ϕ is unique up to positive linear transformation i.e. if ϕ' is another such a function then $\exists(x, y \in \Re : x > 0,\ \phi = x.\phi + y)$.

Ramsey's basic strategy for proving his representation theorem is to use preference orderings of prospects to define an algebraic difference structure and then to invoke Theorem 8. We will now reconstruct his proof on that basis.

Theorem 9 *Let $\succeq$ be the relation on $\bar{\Gamma} \times \bar{\Gamma}$ induced by definition 3. Then $< \bar{\Gamma} \times \bar{\Gamma}, \succeq >$ is a difference algebra.*

Proof. We prove that $< \bar{\Gamma} \times \bar{\Gamma}, \succeq >$ satisfies the five conditions given in Definition 7.

(1) Follows directly from the completeness and transitivity of $\geq$.

(2) By Lemma 4(i).

(3) By R4*, if $\bar{\alpha} - \bar{\alpha}' \geq \bar{\beta} - \bar{\beta}'$ and $\bar{\beta} - \bar{\beta}' \geq \bar{\gamma} - \bar{\gamma}'$ then $\bar{\alpha} - \bar{\alpha}' \geq \bar{\gamma} - \bar{\gamma}'$. But by Corollary 4 (ii) $\bar{\alpha} - \bar{\alpha}' \geq \bar{\beta} - \bar{\beta}' \Leftrightarrow \bar{\alpha} - \bar{\beta} \geq \bar{\alpha}' - \bar{\beta}'$, $\bar{\beta} - \bar{\beta}' \geq \bar{\gamma} - \bar{\gamma}' \Leftrightarrow \bar{\beta} - \bar{\gamma} \geq \bar{\beta}' - \bar{\gamma}'$ and $\bar{\alpha} - \bar{\alpha}' \geq \bar{\gamma} - \bar{\gamma}' \Leftrightarrow \bar{\alpha} - \bar{\gamma} \geq \bar{\alpha}' - \bar{\gamma}'$. Hence if $\bar{\alpha} - \bar{\beta} \geq \bar{\alpha}' - \bar{\beta}'$ and $\bar{\beta} - \bar{\gamma} \geq \bar{\beta}' - \bar{\gamma}'$ then $\bar{\alpha} - \bar{\gamma} \geq \bar{\alpha}' - \bar{\gamma}'$.

(4) By R5, $\forall(\alpha, \beta, \gamma, \delta)[\exists(\epsilon, \epsilon')$: $(\alpha$ if $P)(\gamma$ if $\neg P) \approx (\epsilon$ if $P)(\beta$ if $\neg P)$ and $(\delta$ if $P)(\beta$ if $\neg P) \approx (\epsilon'$ if $P)(\gamma$ if $\neg P)]$. But by Lemma 4(i) $(\delta$ if $P)(\beta$ if $\neg P) \approx (\epsilon'$ if $P)(\gamma$ if $\neg P) \Leftrightarrow (\beta$ if $P)(\delta$ if $\neg P) \approx (\gamma$ if $P)(\epsilon'$ if $\neg P)$. Hence $\exists(\bar{\epsilon}, \bar{\epsilon}') : \bar{\alpha} - \bar{\epsilon} \approx \bar{\beta} - \bar{\gamma}$ and $\bar{\beta} - \bar{\gamma} \approx \bar{\epsilon}' - \bar{\delta}$.

(5) Follows from R8.

■

Theorems 5 and 6 clearly follow directly from Theorems 8 and 9. Ramsey does not seek to explicitly establish that the utility function, U, referred to in these theorems also represents the agent's preference ranking of possibilities in the sense that the utilities of prospects go by their position in the preference order. To establish this we would have to make a further, but unobjectionable, assumption. As it turns out, the assumption is presupposed in Ramsey's subsequent derivation of degrees of belief and so there good reason to make it explicit here.

R9 Let P be any proposition and α and β any worlds. Then:

$\alpha \geq \beta \Longleftrightarrow \alpha \geq (\alpha$ if $P)(\beta$ if $\neg P) \geq \beta$.

Corollary 10 $\alpha \approx (\alpha$ *if* $P)(\alpha$ *if* $\neg P)$

Theorem 11 *The utility function, U, on $\bar{\Gamma}$, is such that $\forall(\bar{\alpha}, \bar{\beta} \in \bar{\Gamma})$, $\bar{\alpha} \succeq \bar{\beta} \Leftrightarrow U(\bar{\alpha}) \geq U(\bar{\beta})$*

Proof. By Corollary 10 and Theorem 5 it follows that $\alpha \geq \beta \Leftrightarrow (\alpha$ if $P)(\alpha$ if $\neg P) \geq (\beta$ if $P)(\beta$ if $\neg P) \Leftrightarrow \bar{\alpha} - \bar{\beta} \succeq \bar{\beta} - \bar{\alpha} \Leftrightarrow U(\bar{\alpha}) - U(\bar{\beta}) \geq U(\bar{\beta}) - U(\bar{\alpha}) \Leftrightarrow U(\bar{\alpha}) \geq U(\bar{\beta})$. ■

2.3 Measuring Partial Belief

2.3.1 Defining Degrees of Belief

Recall from our informal presentation of his method that Ramsey's next move is to use the utility measure on worlds to determine the agent's degrees of belief in all propositions, including those that are not ethically neutral. The vehicle for doing so is the following definition.

Definition 12 (*Degrees of Belief*) *Suppose that worlds $\alpha \in P$, $\beta \in \neg P$ and ξ are such that $\alpha \not\approx \beta$ and $\xi \approx (\alpha \text{ if } P)(\beta \text{ if } \neg P)$. Then:*

$$\Pr(P) =_{defn} \frac{U(\xi) - U(\beta)}{U(\alpha) - U(\beta)}$$

The existence of the world ε in question is presumably secured by R7. As the Uniqueness Theorem establishes that ratios of utility difference are independent of choice of scale for the utility function, Definition 12 determines a unique measure of the agent's degrees of belief. Ramsey notes that in this definition the proposition P is not assumed to be ethically neutral, but that it is necessary to assume both that this definition is independent of the choice of worlds meeting the antecedent conditions and that

> "there is a world with any assigned value in which P is true, and one in which P is false"[18]

Why the latter assumption is necessary will only become clear once we look at Ramsey's proof that degrees of belief are probabilities. But if it is to be tenable it is patently necessary that P be neither logically true nor logically false. But this means that some separate treatment of such propositions is required e.g. by stipulating that $\Pr(P) = 1$, whenever P is logically true.

As regards the former assumption (of independence), Ramsey does not say how it might be formally expressed as a condition on preference or choice. But it must be possible to do so, as we know from the Uniqueness Theorem that the equality (or otherwise) of ratios of utility differences is determined by the preference ranking. One way to proceed would be to define the conditions under which the difference in values of α and β equals a particular fraction of the difference between the values of γ and δ. The definitions are cumbersome, so I will confine myself to illustrating the case of one-half. Suppose that Q is an ethically neutral proposition of probability one-half, that $\zeta \approx (\varepsilon \text{ if } Q)(\delta \text{ if } \neg Q)$ and that $\eta \approx (\varepsilon \text{ if } Q)(\gamma \text{ if } \neg Q)$. Then we can say that difference in value of α and β equals half the difference between the values of γ and δ iff $(\alpha \text{ if } Q)(\zeta \text{ if } \neg Q) \approx (\beta \text{ if } Q)(\eta \text{ if } \neg Q)$. And so on.

Theorem 13 *If $\alpha \in P$ and $\beta \in \neg P$, then:*
$U((\alpha \text{ if } P)(\beta \text{ if } \neg P)) = U(\alpha).\Pr(P) + U(\beta).(1 - \Pr(P))$

[18] Ramsey, ibid, p. 75

Proof. If $\alpha \approx \beta$, then it follows from axiom R9 that $\alpha \approx (\alpha \text{ if } P)(\beta \text{ if } \neg P) \approx \beta$. So the theorem follows immediately. If $\alpha \not\approx \beta$, then suppose that ξ is such that $\xi \approx (\alpha \text{ if } P)(\beta \text{ if } \neg P)$. Then by the definition of $\Pr(P)$, $U(\xi) = U(\alpha).\Pr(P) - U(\beta).(1 - \Pr(P)) = U((\alpha \text{ if } P)(\beta \text{ if } \neg P))$. ■

Definition 14 (*Conditional Degrees of Belief*) *Suppose that (α if Q)(β if $\neg Q$) $\approx$ (γ if PQ)(δ if $\neg PQ$)(β if $\neg Q$). Then the degree of belief in P given Q;*

$$\Pr(P \mid Q) =_{defn} \frac{U(\alpha) - U(\delta)}{U(\gamma) - U(\delta)}$$

As with the definition of degrees of belief it must be supposed (though Ramsey does not explicitly say so) that $\gamma \not\approx \delta$, $\alpha \in Q$, $\beta \in \neg Q$, $\gamma \in PQ$ and $\delta \in \neg PQ$, that the definition is independent of particular choices of worlds satisfying the antecedent conditions and that there is a world with any assigned value in which PQ, $P\neg Q$ and $\neg Q$ are true. It would also appear that the existence of equally ranked conditional prospects of kind referred to in the definition is not guaranteed by Ramsey's assumptions. There are a number of ways of filling in this gap. The most conceptually satisfactory would involve the postulation of compounded conditional prospects and a generalisation of R5. But somewhat more economically, we could simply add the following axiom to Ramsey's.

R10 Let $\{P_1, P_2, ..., P_n\}$ be a partition of propositions. Then $\forall(\gamma, \delta, ..., \beta), \exists(\alpha :$ $(\gamma \text{ if } P_1)(\delta \text{ if } P_2)...(\beta \text{ if } P_n) \approx (\alpha \text{ if } P_1 \cup P_2)...(\beta \text{ if } P_n))$.

2.3.2 Proving Coherence

Ramsey must still demonstrate that the degree of belief function $\Pr(\cdot)$ is truly a probability function. This is done in Theorems 15 and 17. Ramsey's proof of Theorem 17 requires a further assumption, not made explicit by him, but which is quite reasonable if one accepts his framework.

R11 Suppose that P and Q are inconsistent, $\alpha \in P$, $\beta \in Q$, and $\gamma \in P \cup Q$. If $\alpha \approx \beta \approx \gamma$, then $(\alpha \text{ if } P)(\beta \text{ if } Q)(\delta \text{ if } \neg P\neg Q) \approx (\gamma \text{ if } P \cup Q)(\delta \text{ if } \neg P\neg Q)$

Ramsey's own proof makes no use of R10, for the obvious reason that he does not explicitly postulate it. But we have seen that it is required elsewhere and by making use of it here a much simpler alternative proof of Theorem 17 can be given which does not require R11. Both proofs follow below.

Theorem 15 *Let P be any proposition. Then:*
(i) $\Pr(P) \geq 0$
(ii) $\Pr(P) + \Pr(\neg P) = 1$
(iii) $\Pr(P \mid Q) + \Pr(\neg P \mid Q) = 1$

Proof. Suppose that ξ is such that $\xi \approx (\alpha \text{ if } P)(\beta \text{ if } \neg P)$. Then:

(i) By R9, $\alpha \geq \beta \iff \alpha \geq \xi \geq \beta$. So $U(\xi) - U(\beta) \leq U(\alpha) - U(\beta)$, and it then follows from the definition of Pr that it never takes negative values.

(ii) By definition, $\Pr(P) + \Pr(\neg P) = \frac{U(\xi)-U(\beta)}{U(\alpha)-U(\beta)} + \frac{U(\xi)-U(\alpha)}{U(\beta)-U(\alpha)} = 1$.

(ii) Suppose that $\gamma \in PQ$ and $\delta \in \neg PQ$ are such that $(\alpha \text{ if } Q)(\beta \text{ if } \neg Q) \approx (\gamma \text{ if } PQ)(\delta \text{ if } \neg PQ)(\beta \text{ if } \neg Q)$. Then by definition of conditional degrees of belief, $\Pr(P \mid Q) = \frac{U(\alpha)-U(\delta)}{U(\gamma)-U(\delta)}$ and $\Pr(\neg P \mid Q) = \frac{U(\alpha)-U(\gamma)}{U(\delta)-U(\gamma)} = \frac{U(\gamma)-U(\alpha)}{U(\gamma)-U(\delta)}$. So $\Pr(P \mid Q) + \Pr(\neg P \mid Q) = \frac{U(\alpha)-U(\delta)-U(\alpha)+U(\gamma)}{U(\gamma)-U(\delta)} = 1$. ■

Lemma 16 *Suppose that $\beta \in \neg Q$, $\gamma \in PQ$ and $\delta \in \neg PQ$. Then $U((\gamma \text{ if } PQ)(\delta \text{ if } \neg PQ)(\beta \text{ if } \neg Q)) = (U(\gamma).\Pr(P \mid Q) + U(\delta)(1 - \Pr(P \mid Q))).\Pr(Q) + U(\beta).\Pr(\neg Q)$*

Proof. Let $\alpha \in Q$ be such that $(\alpha \text{ if } Q)(\beta \text{ if } \neg Q) \approx (\gamma \text{ if } PQ)(\delta \text{ if } \neg PQ)(\beta \text{ if } \neg Q)$. Then by Theorem 13, $U(\gamma \text{ if } PQ)(\delta \text{ if } \neg PQ)(\beta \text{ if } \neg Q) = U(\alpha).\Pr(Q) - U(\beta).(1 - \Pr(Q))$. But by Definition 14, $U(\alpha) = U(\gamma).\Pr(P \mid Q) + U(\delta)(1 - \Pr(P \mid Q))$. So $U((\gamma \text{ if } PQ)(\delta \text{ if } \neg PQ)(\beta \text{ if } \neg Q)) = (U(\gamma).\Pr(P \mid Q) + U(\delta)(1 - \Pr(P \mid Q))).\Pr(Q) + U(\beta).\Pr(\neg Q)$. ■

Theorem 17 $\Pr(P \mid Q) = \frac{\Pr(PQ)}{\Pr(Q)}$

Ramsey's Proof. Let $\Pr(Q) = x$ and $\Pr(P \mid Q) = y$. Then we need to show that $\Pr(PQ) = xy$. Let α and β be any worlds in Q and $\neg Q$ respectively, such that, for some real number t, $U(\alpha) = U(\xi) + (1-x)t$ and $U(\beta) = U(\alpha) - t = U(\xi) - xt$. By assumption such worlds α and β exist. Now $U(\xi) = U(\xi).x + U(\xi).(1-x) = U(\alpha).x + U(\beta).(1-x) = U((\alpha \text{ if } Q)(\beta \text{ if } \neg Q))$. Then by definition, $x = \frac{U(\xi)-U(\beta)}{U(\alpha)-U(\beta)}$.

Now let worlds $\gamma \in PQ$, $\delta \in \neg PQ$ be such that $U(\gamma) = U(\alpha) + \frac{t}{y} - t$ and $U(\delta) = U(\beta) = U(\alpha) - t$. Again by assumption such worlds γ and δ exist. Then by Lemma 16, $U((\gamma \text{ if } PQ)(\delta \text{ if } P\neg Q)(\beta \text{ if } \neg Q)) = (U(\gamma).y + U(\delta)(1-y)).x + U(\beta)(1-x) = U(\alpha).x + U(\beta).(1-x) = U((\alpha \text{ if } Q)(\beta \text{ if } \neg Q))$. So by Definition 14, $y = \frac{U(\alpha)-U(\delta)}{U(\gamma)-U(\delta)} = \frac{U(\alpha)-U(\beta)}{U(\gamma)-U(\beta)}$. So $xy = \frac{U(\xi)-U(\beta)}{U(\alpha)-U(\beta)}.\frac{U(\alpha)-U(\beta)}{U(\gamma)-U(\beta)} = \frac{U(\xi)-U(\beta)}{U(\gamma)-U(\beta)}$. It also follows from Axiom R11, that $U((\gamma \text{ if } PQ)(\delta \text{ if } P\neg Q)(\beta \text{ if } \neg Q)) = U((\gamma \text{ if } PQ)(\beta \text{ if } (\neg P \cup \neg Q)))$. Hence $\xi \approx (\gamma \text{ if } PQ)(\beta \text{ if } (\neg P \cup \neg Q))$ But then by definition $\Pr(PQ) = \frac{U(\xi)-U(\beta)}{U(\gamma)-U(\beta)} = xy$. ■

Alternative Proof. Let worlds $\delta \in \neg PQ$ and $\beta \in \neg Q$ be such that $\delta \approx \beta$. Now by R10, there exists worlds α and ϵ such that $\xi \approx (\alpha \text{ if } Q)(\beta \text{ if } \neg Q) \approx (\gamma \text{ if } PQ)(\delta \text{ if } \neg PQ)(\beta \text{ if } \neg Q) \approx (\gamma \text{ if } PQ)(\epsilon \text{ if } (\neg P \cup \neg Q))$. Then by definition, $\Pr(Q) = \frac{U(\xi)-U(\beta)}{U(\alpha)-U(\beta)}$, $\Pr(PQ) = \frac{U(\xi)-U(\beta)}{U(\gamma)-U(\beta)}$ and $\Pr(P \mid Q) = \frac{U(\alpha)-U(\delta)}{U(\gamma)-U(\delta)} = \frac{U(\alpha)-U(\beta)}{U(\gamma)-U(\delta)}$. So $\Pr(PQ) = \Pr(Q).\Pr(P \mid Q)$. ■

Corollary 18 $\Pr(PQ) + \Pr(\neg PQ) = \Pr(Q)$

Proof. By Theorem 17, $\Pr(PQ) = \Pr(P \mid Q).\Pr(Q)$ and $\Pr(\neg PQ) = \Pr(\neg P \mid Q).\Pr(Q)$. Therefore $\Pr(PQ) + \Pr(\neg PQ) = (\Pr(P \mid Q) + \Pr(\neg P \mid Q)).\Pr(Q) = \Pr(Q)$ by 15(iii). ■

2.3.3 Conditionalism

The importance of Ramsey's assumption that propositions contain worlds of every utility value should now be clear - it is what allows the assumption of the existence of the worlds α, β, γ and δ referred to in his proof of Theorem 17 and of worlds β and δ referred to in the alternative proof. Ramsey seems to think of this as a purely technical condition. But one might derive it from conditionalist premises, the relevant one in this context being as follows.

Proposition 19 *(Ethical Conditionalism) For any propositions P and Q, there are worlds $\alpha \in P$ and $\beta \in Q$ such that $\alpha \approx \beta$.*

Given the 1-1 correspondence between worlds and real numbers, Ethical Conditionalism implies the assumption of whose necessity Ramsey speaks. Namely, that whatever the range of utility values taken by prospects, for every value in that range and every proposition, there is a world which implies the truth of that proposition and which has the utility value in question.

Conditionalism is to my mind an imminently defensible doctrine. Essentially the conditionalist's claim is that however good (or bad) some possibility might be on average, there are imaginable circumstances in which it is not so. No prospect is good or bad in itself, but is only so relative to the conditions under which it is expected to be realised. Suppose, for instance, that P identifies some good prospect like winning the National Lottery and Q some dreadful one, like the death of a relative. The conditionalist claim is that even winning the lottery can be a bad thing, such as when it alienates one from one's friends or causes one to stop activities that gives one's life a sense of purpose. Likewise even the death of a relative can be a good thing, such as when it pre-empts a period of great suffering for them.

Defensible though it may be, Ethical Conditionalism is not consistent with Ramsey's atomistic framework. For consider worlds α and β such that $\alpha \not\approx \beta$ and the proposition - call it A - that α is the actual world. Then since worlds are (nearly) maximally specific it follows that any world in which A is true is ranked with α. But then there is no world in which A is true which is equally preferred to β.

The only way I can see of blocking this argument, is to deny that worlds imply the existence of propositions stating that they are the actual world. But this would be to argue, in effect, that worlds could not be represented propositionally and that contradicts the requirement that the agent be able to choose amongst them (which, I take it, presupposes that she can distinguish them propositionally). So this is not a response open to Ramsey. And though the issue may in some sense be 'merely technical', some modification to his framework will be required to deal with the problem.

3 The Evaluation

3.1 Ethical Neutrality versus State-Independence

> Ramsey's essays, though now much appreciated, seem to have had relatively little influence."[19]

Savage's remark applies equally well today and mainstream decision theory descends from Savage and not Ramsey. There are, I think, two reasons for his lack of influence. One is that Ramsey's style is so elliptical, and his writings so lacking in detail, that decision theorists have been unsure as to what exactly he has or has not achieved.[20] The second is that the distinction between the problem of justifying the claims of decision theory regarding the properties of rational belief and desire and the problem of the measurement of the decision theoretic variables - degrees of belief and desire - has not been properly recognised. Due, as I suggested before, to the different possible roles played by representation theorems with respect to the two problems.

This is important, because from the point of view of the problem of justifying the decision theory he employs, Ramsey's representation theorem is not particularly helpful. For one is very unlikely to accept his axioms as definitive of rational preference for conditional prospects unless one accepts the theory of expected utility that motivates them. This is particularly true of axiom R4, which seems to have no justification other than that it secures the meaningfulness of utility differences. Taken as axioms of measurement, however, they do much better for they specify in a precise way the conditions under which a measure of the agent's degrees of desire, unique up to a choice of scale, is determined by her choices amongst prospects.

With respect to problem of justification, on the other hand, a theory like Savage's is a good deal more impressive. Savage chooses his axioms of preference with an eye to their independent plausibility as rationality conditions. Independent, that is, of the quantitative theory of belief and desire that he will derive from them. Such a claim can justifiably be made for the Sure-Thing principle, for instance. One need not grant much plausibility to expected utility theory to grant that of two actions that yield the same outcomes when C is the case, one should choose the one with the preferred outcome when C is not. But Savage builds a very strong and implausible assumption into the very framework of his decision theory. He assumes that the desirability of any possible outcome of an action is independent of the state of the world in which it is realised.

Let us start by getting a general idea of the problem. It is a banal fact about our attitudes to many things that they depend on all sorts of contextual factors. Hot chocolate can be delightful on a cold evening, but sickly in the heat of a summer's day. Swimming on the other hand is best reserved for those hot days. I shall say, somewhat barbarically, that the swimming or drinking hot chocolate

[19] Savage [12, p. 96]

[20] For instance, Fishburn [5] rejects Ramsey-type theories in favour of Savage-type ones on the grounds of his 'restricted act space'. In fact, however, his set of conditional prospects is roughly equivalent to Savage's set of acts.

is desirabilistically dependent on the weather. Many things, on the other hand, are to all practical purposes desirabilistically independent, certainly swimming and the temperature on the moon are for me. Any reasonable theory of rational agency ought to recognise these banal facts.

How does Savage's theory violate them? Savage uses observations of choices amongst actions to determine agents' attitudes. Actions, on his account, are functions from states of the world to possible outcomes: when you choose an action you choose to make it true that if the world is in state s_1 then outcome o_1 will be realised, if it is in state s_2 then outcome o_2 will be realised, and so on. Now if we are to recognise that the desirability of the outcomes of actions depend on the state of affairs in which they are realised, then either the utilities we derive for them must be state-dependent i.e. of the form $U(o_1|s_1)$, or the outcome o_1 must include the fact that s_1 prevails (as outcomes of Ramsey's conditional prospects do). But Savage both assumes that any combination of state and outcome is possible and assigns state-independent utilities to outcomes.

On Ramsey's account, outcomes (worlds) are maximally specific with regard to things that matter to the agent, but not all outcomes are achievable in any given state. So his theory requires no violation of the banal facts concerning the interdependence of our attitudes. Instead of building desirabilistic independence into his framework, he postulates the existence of only a very limited class of possibilities - those represented by ethically neutral propositions - which are desirabilistically independent of all others. One may question whether there are any propositions that are truly ethically neutral, but there are clearly some that are good approximations. The postulation of their existence is not a heavy burden for such an idealised account to bare.

This is not, of course, the end of the matter. There have been numerous attempts to solve the problem of state-dependent utilities (as it has become known) within Savage's framework.[21] Many of the proposed solutions are ingenious, but they always come at the cost of greater complexity and more burdensome assumptions. This is not the appropriate place to review the literature, but anyone who has ploughed through it will have little difficulty in recognising the merits of the elegantly simple method that Ramsey invented. Indeed, despite the problems in the details that we discovered, there is nothing that matches it as an answer to the problem of measurement.

3.2 Jeffrey's Objection

In motivating his own method of measuring belief, Ramsey argues that the established method of offering bets with monetary rewards to someone to elicit their degrees of belief is 'fundamentally sound' but suffers from being both insufficiently general and necessarily inexact. Inexact partly because the marginal utility of money need not be constant, partly because people may be especially averse (or otherwise) to betting because of the excitement involved and partly because "the proposal of a bet may ... alter his state of opinion"[22]

[21] For a summary see Schervish et al [13]

[22] Ramsey [10, p. 68]

Ramsey seems to think that his own theory is not vulnerable to these problems, even though his method is similar in many ways to the one he is criticising. Not everyone would agree. Richard Jeffrey, for instance, has argued that just such a problem plagues Ramsey's own account.[23] In order to measure agents' partial beliefs, Ramsey requires that they treat possibilities like it being α if P and β if not as real prospects i.e. things that can be brought about by choice. But to persuade someone of our power to bring it about at will that it will be sunny tomorrow if the coin lands heads and snowy if it lands tails is to cause them to entertain possibilities which at present they do not. That is one must modify their beliefs in order that one may better measure them! There is of course no guarantee then that the measurements so effected are not, at least partially, artifacts of the measurement process itself.

It is worth noting that such an objection, if sustainable, can be directed with equal force at Savage. For when Savage invites agents to make choices amongst actions, he supposes that they know exactly what consequences the action has in every possible state of the world and hence what they are committing themselves to. This makes the choice of a Savage-type action rather more like a choice amongst Ramsey-type conditional prospects than amongst the sorts of things we normally think of as actions. Indeed, formally, they are just the same thing: functions from events to outcomes. What is of the essence, in any case, is that agents are invited to choose amongst causal mechanisms of some kind whose effects in each possible circumstance are advertised in advance. And the essence of Jeffrey's objection is that agents may legitimately doubt the efficacy of such mechanisms, and make their choice in the light of these doubts. If they do, their choices will reflect not their evaluation of the advertised outcomes of the chosen prospect, but their evaluations of the outcomes that they actually expect. Even in pure gambling situations agents will factor in such possibilities as the casino closing before paying up.

How might Ramsey respond to this problem? Sobel argues that Ramsey must require that the probability of a proposition P be measured only by means of conditional prospects which are such that P's probability is evidentially and causally independent of the conditional prospect being offered (by, for instance, addition of a further restriction in Definition 12).[24] But there is no obvious way of expressing this independence condition in terms of agents' preferences and so no way of applying it until the probability of P has already been measured.

A natural response to Jeffrey's objection would be to say that Ramsey does not, in fact, require that agents really believe in such fanciful causal possibilities. All that he requires is that they choose amongst gambles *as if* they believed that they would truly yield the advertised consequences under the relevant conditions. To be sure, such a response will not satisfy the behaviourist, for introspection on the part of agents must then play a crucial role in the production of their choices. For when we ask Mary to choose between an prospect which yields sunny weather if Labour wins the next election and rainy weather

[23] Jeffrey [6, chapter 10]
[24] Sobel [14, p.256]

if they do not, and one which yields rainy weather if Labour wins the next election and sunny weather if they do not, we are in effect asking her to determine *for herself* what she would prefer in the event that such gambles were reliable. But then we may as well just ask Mary what she would prefer and forget about the observation of choices altogether.

And indeed why not? Let us see what such a reconstrual of Ramsey's method would amount to in the context of the experimental determination of a subject's degrees of belief and desire, by comparing the following measurement schemes:

1. *Scheme A*: The subject introspects her degrees of belief and desire and then reports them to the observer.

2. *Scheme B*: The observer presents the subject with a number of options and her choice is recorded. The set of options offered is varied until a ranking over all of them has been constructed from the observations of her choices. This ranking is then used to construct a quantitative representation of her degrees of belief and desire.

3. *Scheme C*: The observer questions the subject as to which of various possibilities she would prefer were the true one. Her answers are then used to construct a ranking of all possibilities and this in turn determines a quantitative representation of her mental attitudes.

Scheme A is the method criticised by behaviourists and Ramsey alike for its naive dependence on introspection. Scheme B summarises the behaviourist's method, Scheme C the alternative interpretation of Ramsey's method. Both are underwritten by the representation theorems of Decision Theory. In Scheme C introspection plays an essential role: to provide answers to the experimenter's questions the subject must reflect upon and judge her own preferences. In Scheme B, on the other hand, though it is conceivable that the subject arrives at a choice via an introspection on her preferences, she need not do so. She may simply choose without reflection, indeed without even having the concept of preference. Scheme C is a method intimately tied to the possibility of linguistic communication and the kind of self-consciousness that typically accompanies it; Scheme B is just as applicable to earthworms as to philosophers.

I see no reason why Ramsey should be resistant to this interpretation of his method as a version of Scheme C. Although it requires him to disavow the behaviourist pretension that introspection can be completely eliminated in favour of rich observations of behaviour, it does not commit him to the view instantiated in Scheme A(and which he clearly rejects) that partial beliefs and desires can be directly introspected. In this sense this interpretation does not conflict with anything that he says. And it has the crucial advantage of extricating him from Jeffrey's objection.

3.3 Ramsey à la Jeffrey?

In filling in the details of Ramsey's theory of measurement we have had reason to raise a number of questions and to make a number of supplementary assumptions. But only the incompatibility of Conditionalism with his framework seems to raise a serious problem for Ramsey. In fact, however, this problem is largely a technical one and can be solved by modifications to Ramsey's framework that are not contrary to the 'spirit' of his account. I will content myself with sketching the essentials.

The basic move is to take (non-contradictory) propositions rather than worlds to be the elementary objects of preference. One immediate positive spin-off is that the notion of ethical neutrality can then be formulated in a manner less dependent on the peculiarities of Wittgenstein's theory of atomic propositions.[25]

Definition 20 *Suppose P and Q are mutually consistent propositions. Then P is* ***neutral*** *with respect to Q iff $PQ \approx Q \approx P\neg Q$.*

Definition 21 *P is* ***ethically neutral*** *iff P is neutral with respect to all propositions Q consistent with P.*

Conditional prospects must now be defined as functions from partitions of propositions to (non-contradictory) propositions, with the constraint that for any conditional prospect Φ and proposition X, $\Phi(X)$ implies X. But little else need change, since most of Ramsey's formal argument is carried out at the level of values. Of course, utilities as well as probabilities will now be defined on propositions. Finally the relevant definition of Conditionalism is as follows.

Proposition 22 *For any propositions P and Q there exists propositions P' and Q' such that P' implies P, Q' implies Q and $P' \approx Q'$.*

Proposition 22 can be satisfied only if there are no propositions X such that, for all propositions Y, X implies Y or X implies $\neg Y$. In other words, Conditionalism requires that the domain of the preference relation be atomless.

All of this takes Ramsey's framework quite a bit closer to that underlying Richard Jeffrey's decision theory and Ethan Bolker's representation theorem for it.[26] So too did the contention that Ramsey's work should be interpreted in such a way as to rid it of any dependence on dubious causal devices such as gambles. But I do not to propose to go much further in their direction, because from the perspective of the problem of the measurement of belief, the Jeffrey-Bolker theory suffers from a crucial weakness by comparison to Ramsey's.[27] For Bolker's representation theorem does not establish the existence of a unique measure of an agent's beliefs or a measure of her degrees of desire unique up to a choice of scale. In particular it allows for the possibility that two probability measures of an agent's degrees of belief, P_1 and P_2, both be consistent with her expressed preferences yet differ to the extent that there are propositions A and B such that $P_1(A) > P_1(B)$ but $P_2(B) > P_2(A)$.[28]

[25] See Bradley [2] for a more detailed development of these ideas.

[26] See Jeffrey [6] and Bolker [3].

[27] On the other hand, with respect to the problem of normative justification the Jeffrey-Bolker theory is much better than Ramsey's.

[28] For further discussion of this problem, see Bradley [2] and Joyce [7].

The essential difference, in this regard, between Ramsey's theory and that of Jeffrey and Bolker is that the latter make do without any conditional prospects of the kind postulated by Ramsey, working only with agents' attitudes to propositions. The price of this ontological economy would seem to be the underdetermination of agents' degrees of belief and desire by the evidence of their expressed preferences. If the price is too high, we have reason to favour a Ramsey-type theory when addressing the problem of measurement.

But this should not obscure the fact that this discussion raises a difficult question concerning the status of Ramsey's conditional prospects. For if conditional prospects could be given propositional expression then it should be possible to strengthen the Jeffrey's theory by simply adding to it suitably translated versions of Ramsey's postulates concerning preferences for conditional prospects. But the evidence is that this cannot be done without leading to some unpalatable consequences. But if Ramsey's conditional prospects have no adequate propositional correlates, as has already been suggested by his definition of ethical neutral propositions, what exactly is their nature?[29]

References

[1] Bradley, R. W. (1998) 'A Representation Theorem for a Decision Theory with Conditionals', *Synthese* 116, 187-229

[2] Bradley, R. W. (1997) '*The Representation of Beliefs and Desires within Decision Theory*', PhD dissertation, University of Chicago

[3] Bolker, E. D. (1967) 'A Simultaneous Axiomatisation of Utility and Subjective Probability', *Philosophy of Science* 34, 292-312

[4] Davidson, D. and Suppes, P. (1956) 'A Finitistic Axiomatisation of Subjective Probability and Utility', *Econometrica* 24, 264-275

[5] Fishburn, P. C. (1981) 'Subjective Expected Utility: A Review of Normative Theories', *Theory and Decision* 13, 139-199

[6] Jeffrey, R.C. (1983) *The Logic of Decision*, 2nd edn, Chicago: University of Chicago Press

[7] Joyce, J. (1999) *The Foundations of Causal Decision Theory*, Cambridge University Press

[8] Krantz, D. H., R. Duncan Luce, P. Suppes and A. Tversky (1971) *Foundations of Measurement*, Volume 1, Academic Press

[9] Pfanzagl, J. (1968) *Theory of Measurement*, New York: Wiley

[10] Ramsey, F. P. (1926) 'Truth and Probability' in *Philosophical Papers*, ed. D. H. Mellor, 1990, Cambridge: Cambridge University Press.

[29] An attempt to give their logical properties is to be found in Bradley [1] and [2].

[11] Sahlin, N.-E. (1990) *The Philosophy of F. P. Ramsey*, Cambridge: Cambridge University Press.

[12] Savage, L. J. (1954) *The Foundations of Statistics*, 2nd edition, New York: Dover, 1972

[13] Schervish, M. J., T. Seidenfeld, J.B. Kadane (1990) 'State-dependent Utilities', *Journal of the American Mathematical Association* 85, 840-847

[14] Sobel, J. H. (1998) 'Ramsey's Foundations Extended to Desirabilities', *Theory and Decision* 44, 231-278