
ECOLE POLYTECHNIQUE
CENTRE NATIONAL DE LA RECHERCHE SCIENTIFIQUE

Axiomatic Bayesian Utilitarianism

Richard Bradley

Mai 2003

Cahier n° 2003-009

LABORATOIRE D'ECONOMETRIE

1 rue Descartes F-75005 Paris

(33) 1 55558215

<http://ceco.polytechnique.fr/>
<mailto:labecox@poly.polytechnique.fr>

Axiomatic Bayesian Utilitarianism¹

Richard Bradley²

Mai 2003

Cahier n° 2003-009

Résumé: L'article examine les relations qu'entretiennent les conditions suivantes, formulées dans la théorie axiomatique de la décision de Jeffrey-Bolker. (1) La condition utilitariste voulant que l'utilité espérée de la société se représente comme une somme pondérée des utilités espérées individuelles. (2) Des conditions d'uniformité imposées aux fonctions d'utilité et de probabilité des individus. On montre en particulier que l'identité des probabilités individuelles est nécessaire et suffisante pour que l'on obtienne l'utilitarisme, et qu'elle dérive d'une condition antérieure de type parétien reliant les préférences individuelles et sociales, pour autant que ces préférences soient séparables en un sens convenable.

Abstract: This paper investigates, within the axiomatic framework of Jeffrey-Bolker decision theory, the relation between two kinds of conditions: (1) The Utilitarian condition that social expected utility function is a weighted sum of the individual expected utility functions; and (2) Homogeneity conditions on the probability and utility functions of individuals. In particular, we show that identity of individuals' probabilities is necessary and sufficient for the Utilitarian condition to hold and that it can be derived from an earlier Pareto condition on the relation between individual and social rankings, provided that these rankings are separable in a particular sense.

Mots clés : Utilitarisme, bayésianisme collectif, Bolker-Jeffrey, Savage, principe de Pareto, théorèmes d'identité

Key Words : Utilitarianism, Collective Bayesianism, Bolker-Jeffrey, Savage, Pareto Principle, Homogeneity Theorems

Classification JEL: D63, B21, D81

¹ This paper has benefited to a considerable extent from comments by Isaac Levi, John Broome and especially Philippe Mongin. The author thanks the Laboratoire d'économétrie for their invitation in February 2003.

² Department of Philosophy, Logic and Scientific Method, The London School of Economics

1 Introduction

Suppose that a group of rational individuals must construct a collective or joint evaluation of a number of prospects. Suppose also that they wish the group's evaluations to be both rational and positively sensitive to those of the individuals making it up. Then what sorts of evaluations of the options are open to the group? In general, of course, the answer will depend on what kind of prospects need to be evaluated and what is meant by rationality and positive sensitivity. So for definiteness, suppose that the rationality required of both the individuals' evaluations and those of the group be of the Bayesian variety and that positive sensitivity minimally amounts to satisfaction of a Pareto condition that unanimous individual rankings of one prospect over another (or co-ranking of the two) entails a collective ranking of the former over the latter (or a collective co-ranking of the two). These assumptions constrain the structure of both the cognitive and conative evaluations of the group to a surprising extent. This paper discusses two kinds of formal results about this structure: those that show

that under the conditions of Bayesian rationality and Paretian sensitivity to individual preferences, group preferences can and/or must take a Utilitarian form - the Utilitarian theorems - and those that show that under these conditions the beliefs and/or desires of individuals must be the same - the Homogeneity theorems.

The first Utilitarian theorem of this kind is due to Harsanyi [12], who showed that if we require of both the group and the individuals making it up that they rank the prospects - in this case lotteries - in a manner consistent with von Neumann-Morgenstern expected utility theory, then the Pareto condition entails that group preferences can be represented by a utility function that is a weighted sum of the expected utility functions representing the preferences of individuals. Harsanyi's work assumes that the probabilities of events are part of the identifying descriptions of the prospects to be evaluated and, hence, known to all. So homogeneity of belief is built into his formulation of the problem. But others have studied versions of the Utilitarian theorem in environments in which individuals may have different beliefs and in which homogeneity results are non-trivial. Particularly important in the context of this paper are the versions of the Utilitarian and Homogeneity theorems proven by Philippe Mongin [13], within the axiomatic environment of Savage's subjective expected utility theory, and those investigated by John Broome [6], within the axiomatic environment of Bolker-Jeffrey decision theory.¹

The primary aim of this paper is to review and build upon Broome's exploration of the conditions for, and consequences of, the Utilitarian and Homogeneity theorems. Since Jeffrey's decision theory remains relatively unknown and the differences between it and other Bayesian decision theories ill understood, the first section will both rehearse some of the reasons for favouring it over Savage's and give a brief presentation of Bolker's representation theorem. In the second section the relationship between Utilitarianism and individual belief and preference is examined. In particular, a proof is given of what, from a conceptual point of view, I take to be most interesting of the results derivable in the Jeffrey-Bolker framework: that it is necessary and sufficient for social preferences to be Utilitarian that individuals' beliefs be identical. The third section is devoted to a philosophical discussion of the implications of this theorem for Utilitarianism, including Broome's claim that it shows that it is goodness, not preference, that should be the object of the Utilitarian's concern. Those who prefer their meat raw can pass swiftly on to the fourth and final section where we investigate the technically rather complex question of the conditions for probability homogeneity in the Jeffrey-Bolker framework as a preliminary to the reformulation and generalisation of Broome's Utilitarian theorem that we undertake in the ultimate section.

¹Similar results, but using different axiom sets, are to be found in Hammond (1981) and Seidenfeld et al (1989).

2 Jeffrey-Bolker Decision Theory

2.1 Jeffrey versus Savage

Bayesian decision theories are formal theories of rational agency that tell us both what attitudes can consistently be taken to a given set of prospects (the theory of rational mind) and what action it is rational for an agent to perform, given her state of mind (the theory of choice). Bayesian theories differ with respect to the kinds of objects they postulate and the corresponding constraints they place on rational preference. In von Neumann and Morgenstern's theory, for instance, the basic prospects are lotteries over possible states of affairs, in Savage's theory they are actions, and in Jeffrey's they are propositions. But they also share a number of characteristic features. Rationality of mind is typically construed as internal consistency of belief and desire - formally expressed by modelling degrees of rational belief as probabilities and degrees of rational desire as utilities. Rationality of preference, on the other hand, is construed as a matter of the agent's preferring the option, or one of the options, which has the best expected consequences, given her partial beliefs and desires (an option being a prospect that can be realised at will). Typically the two parts are linked by a representation theorem showing under what conditions an agent whose beliefs and desires are rational must have rational choice-determining preferences and vice versa.

Savage's theory is perhaps the best known of the Bayesian decision theories, but it does suffer from a number of apparent limitations. The first concerns his notion of an action. Savage models actions as total functions from a set of possible states of the world to consequences or outcomes, with states and consequences being treated as independent of one another. For the purposes of his representation theorem, Savage supposes that every function from states of the world to consequences belongs to the set of options available to the agent. This results in actions being specified which, given their beliefs about the causal structure of the world, agents would not regard as being real options, and with respect to which they may have difficulty taking a sensible attitude. Consider, for instance, Savage's example of omelette making in which you are deciding whether to break a sixth egg into a bowl containing five broken eggs or throw it away, in light of the possibility that the sixth egg is rotten. The decision problem can be represented thus:

	STATES	
ACTS	<i>Good</i>	<i>Rotten</i>
1. <i>Break egg</i>	6-egg omelette, none wasted	no omelette, 5 eggs wasted
2. <i>Throw away</i>	5-egg omelette, 1 wasted	5-egg omelette, none wasted

Then the requirement that all functions from states to consequences belong to the set of options, implies that the following nameless acts (and others) are available to you.

ACTS	<i>Good</i>	<i>Rotten</i>
3.	6-egg omelette, none wasted	6-egg omelette, none wasted
4.	no omelette, five eggs wasted	6-egg omelette, none wasted

You may, however, legitimately doubt that the actions so described exist. Indeed, producing a six-egg (edible) omelette with five good eggs and one rotten one is nothing short of miraculous! To take such options seriously an agent would have to adjust their beliefs about what was physically possible. But then the beliefs that go into the determination of the expected utility of the acts will not be the actual beliefs of the agents, but the beliefs she must entertain in order to make the kinds of choices required of her.

Jeffrey regards the fact that his theory refrains from postulating the existence of devices (lotteries, Savage-type acts, etc.) which alter casual relations in ways that are potentially inconsistent with agents' beliefs as its principle virtue. To my mind, however, this is the wrong place to emphasise the distinctiveness of his theory. For one can go some way to addressing this difficulty for Savage's theory by interpreting actions not as real options, but as hypothetical possibilities of the kind identified by ordinary indicative conditional sentences such as 'if the egg is good, you will have an omelette, if it's bad you will have wasted the egg'. To be sure, this interpretation makes it impossible to operationalise the theory in purely behavioural terms and will still leave the agent with the question of what attitude to take to 'miraculous' possibilities. But in this regard Jeffrey's theory is in much the same position.

The main problem with Savage's theory lies elsewhere. On his account the preference-value of an action is equated with the expected value of the mapping from the domain of the action (the states of the world) to the utilities of its values (the consequences). This method for determining the action's choice-value has two noteworthy features (i) it uses the unconditional utilities of consequences i.e. the utilities of consequences are taken to be independent of the states of the world in which they are realised, and (ii) it uses the unconditional probabilities of the states as the weights i.e. it treats states as probabilistically independent of actions. Neither feature is especially attractive. Intuitively it often does matter greatly in what state of the world some outcome is realised (e.g. hot chocolate is better on a cold day than a hot one). Intuitively also actions are more or less preferable precisely because they affect the likelihood that certain states of the world will prevail.

These shortcomings of Savage's theory may rarely be crippling in practice, because when modelling specific 'small-world' decision problems one can choose descriptions of actions, events and consequences that ensure that the independence conditions are met. Sometimes it is natural to do so. But often it is not, and doing so can exacerbate the aforementioned problem of unrealistic actions. In the omelette example, for instance, it would be more natural to describe the consequences as six-egg omelette, five-egg omelette and no omelette and recognise that their desirability depends on whether or not eggs were wasted in their preparation. Savage cannot do this and so must built the wastage property into the description of the consequences. But then he is committed to the existence

of actions which produce no wastage even when the sixth egg is rotten. To be fair, it *is* possible to address these points by modifying Savage's framework in such a way as to allow the utilities of consequences to vary with the state of the world. This is not the place to review the field of state-dependent utility theory.² Suffice to note that when the implicit assumption of state-independence is dropped from the Savage framework, probabilities can no longer be uniquely identified from preferences. Indeed they are not even determined to the extent that they are in the Bolker-Jeffrey framework (see the next section).

In Jeffrey's framework actions, states of the world and consequences are all represented by propositions. Consequently both probabilistic and desirabilistic relations between the three are automatically incorporated in representations of such relations between propositions. Let X and Y be any propositions, $\neg X$ the negation of X and XY the conjunction of X and Y . Let P and V respectively be measures of the probability and desirability of propositions. Then Jeffrey's axiom of desirability relates as follows the desirability of the prospect that X to the possibility that Y :

Axiom 1 (*Desirability*) If $P(X) \neq 0$:

$$V(X) = V(XY).P(Y|X) + V(X\neg Y).P(\neg Y|X)$$

In contrast with Savage's theory (i) the desirability of the prospect that X depends on whether Y is the case or not, because $V(XY)$ will not generally equal $V(X\neg Y)$, and (ii) the desirability of X depends on the conditional probabilities of Y or $\neg Y$, given that X , and not the unconditional probabilities of Y or $\neg Y$. Jeffrey's theory makes no hard and fast contrast between events and consequences and Axiom 1 can be used to express the desirability of a prospect either in terms of its possible consequences or in terms of the possible states prevailing at its realisation, or indeed in terms of possibilities that are both desirabilistically and probabilistically relevant. It is noteworthy then that desirability values are attached to prospects independently of how the event or proposition space is partitioned. As a consequence of this partition-independence, Jeffrey faces no problem of reconciliation between small and grand world decisions and, unlike in Savage's theory, there is no need to find the 'right' partition for evaluating an action.

Causal decision theorists argue that Jeffrey has still not got things quite right and that the choice-worthiness of an action depends not on the conditional probability of its potential consequences, given that the action is performed, but on the probability that each *would* be a consequence of the action *were* it to be performed.³ In essence this is because what matters in decision making is not the evidential significance of the action but its causal potency. I believe there to be some force to this objection. But it is not of much importance in the context of this paper for, as Broome argues, the Utilitarian theorem addresses a claim

²For a survey see Drèze and Rustichini [8]. Strangely there is no mention of the Jeffrey-Bolker theory in this survey.

³See, instance, Joyce (1999).

about social valuation rather than social decision making, as to what it would be best to happen, rather than what we should make happen. On the issue of evaluation, causal decision theorists have no quarrel with Jeffrey.

2.2 Bolker's Representation Theorem

A representation theorem for Jeffrey's decision theory was first proved by Ethan Bolker [1]. Bolker assumes that the set of prospects, the objects with respect to which agents have preferences, forms a complete atomless Boolean algebra with the zero removed. We take as the canonical case a set of propositions, Ω , closed under the operations of conjunction, disjunction and negation, denoted here by \wedge , \vee and \neg , and containing F and T , respectively the logically false and true propositions. An algebra of propositions is atomless if for every proposition $X \neq F$ there exists a proposition $Y \neq F$ that implies X , but which is not implied by it. It is complete if it closed under disjunctions of arbitrary sets of mutually inconsistent propositions.

Let \geq be a complete and transitive relation on the set of propositions Ω . The expression $X \geq Y$ is standardly interpreted as saying that the agent does not prefer Y to X . Let $>$ and \approx be strict preference and indifference relations derived from \geq in the usual way. Assume that the relation \geq is continuous in the sense that if $\{Y_n\}$ is a chain (i.e. a countable, increasing sequence) in Ω , $Y = \vee\{Y_n\}$ and $X \geq Y \geq Z$, then $X \geq Y_n \geq Z$ for all large n . Assume also that is non-trivial in the sense that there exist $X, Y \in \Omega$ such that $X > Y$. Bolker proposes two further conditions on preference:

Axiom 2 (*Averaging*) If $XY = F$, then:

$$X > Y \Leftrightarrow X > X \vee Y > Y$$

$$X \approx Y \Leftrightarrow X \approx X \vee Y \approx Y$$

Axiom 3 (*Impartiality*) If $XZ = YZ = F$, $X \approx Y$ and $X \vee Z \approx Y \vee Z$, then for all $Z' \in \Omega$ such that $XZ' = YZ' = F$, it is the case that $X \vee Z' \approx Y \vee Z'$.

The axiom of Averaging is akin to, but weaker than, Savage's Sure-Thing principle in requiring that if one prefers the prospect of X to that of Y , then one should prefer the prospect that X is the case 'for sure' to the prospect that either X or Y is the case. The axiom of Impartiality is akin to Savage's Postulate 4 and plays a similar role in the representation theorem of ensuring the coherence of the procedure used to determine probabilities (see section 2.2.1 below). Neither the axiom of Impartiality nor Savage's Postulate 4 have much, if any, plausibility as principles of rational preference independently of the theory of subjective expected utility. They are there because the representation theorems require them.

Theorem 4 (*Bolker*) Let $\beta = \langle \Omega, \preceq \rangle$ be an atomless Boolean algebra of propositions and suppose that \geq is a complete, transitive and continuous relation on $\Omega - \{F\}$ that respects the axioms of Averaging and Impartiality. Then there

exists a probability measure P and signed measure U on Ω such that, $\forall (X, Y \in \Omega - \{F\})$, $P(X) \neq 0$ and:

$$X \geq Y \Leftrightarrow \frac{U(X)}{P(X)} \geq \frac{U(Y)}{P(Y)} \quad (1)$$

Furthermore P' and U' are another such pair of measures on Ω iff there exists real numbers a, b, c and d such that:

$$ad - bc > 0 \quad (2)$$

$$cU(T) + d = 1 \quad (3)$$

$$cU + dP > 0 \quad (4)$$

and:

$$P' = cU + dP \quad (5)$$

$$U' = aU + bP \quad (6)$$

The function $V =_{def} U/P$ both satisfies Axiom 1, Jeffrey's axiom of desirability, and representing preferences in the sense that $V(X) \geq V(Y) \Leftrightarrow X \geq Y$, for all prospects X . By Bolker's theorem V is unique up to positive fractional linear transformation i.e. if V' is a another such measure then $V' = \frac{aV+b}{cV+d}$ for real numbers a, b, c and d satisfying the equations above. The probability function P is standardly interpreted as measure of the credibility (for some agent) of the prospects in Ω . It is less a matter of consensus whether the signed measure U , what Bolker refers to as a measure of 'total utility', has an interpretation other than a derivative technical one. Comparative value judgements of the kind 'It is more important that X than that Y ' do seem, however, to cohere with the total utilities of prospects.

2.2.1 The Significance of the Uniqueness Theorem

A notable feature of Bolker's representation theorem is that his axioms do not sufficiently constrain preference so as to uniquely determine a representation of the agent's degrees of belief. This, it would seem, is the price you pay for a theory that makes no use of dubious causal devices and which does not assume that utilities are state-independent. There are some surprising consequences of this lack of uniqueness. For instance, two probability functions P and P' may both represent an agent's preferences and yet for any X and Y that are not co-ranked, it can be the case that $P(X) = P(Y)$ but $P'(X) \neq P'(Y)$.⁴ So within the Jeffrey-Bolker framework whether or not an agent believes one thing

⁴**Proof:** Suppose that $P(X) = P(Y)$, but $V(X) \neq V(Y)$. Then by equation 5, $P'(X) = cU(X) + dP(X) = P(X).(cV(X) + d)$ and $P'(Y) = P(Y).(cV(Y) + d) = P(X).(cV(Y) + d)$. Hence $P'(X) \neq P'(Y)$.

to the same degree as another is not completely determined by, or revealed in, her rankings of prospects. What is determined, on the other hand, is the equi-probability of co-ranked propositions and more generally ratios of probabilities of co-ranked propositions. The test for the equi-probability of two co-ranked propositions, X and Y , consists in finding some third proposition Z mutually exclusive with both X and Y and ranked differently from them. Then if $X \vee Z$ is ranked with $Y \vee Z$, the probabilities of X and Y must be equal. This follows from a simple lemma of Bolker's. As we shall need to use it later on and it is of some importance to the Bolker's uniqueness theorem, we display the proof here.

Lemma 5 *Suppose that $X \approx_i Y$. Let Z be any proposition inconsistent with X and Y and such that $X \not\approx_i Z$. Then $P_i(X) = P_i(Y) \Leftrightarrow X \vee Z \approx_i Y \vee Z$.*

Proof. By assumption, $V_i(X) = V_i(Y) \neq V_i(Z)$. Note that $P_i(Z) > 0$. So from the desirability axiom it follows that $X \vee Z \approx_i Y \vee Z$

$$\begin{aligned} \Leftrightarrow & \frac{V_i(X).P_i(X) + V_i(Z).P_i(Z)}{P_i(X) + P_i(Z)} = \frac{V_i(Y).P_i(Y) + V_i(Z).P_i(Z)}{P_i(Y) + P_i(Z)} \\ \Leftrightarrow & V_i(X).[P_i(X).P_i(Z) - P_i(Y).P_i(Z)] = V_i(Z).P_i(Z).[P_i(X) - P_i(Y)] \\ \Leftrightarrow & V_i(X).[P_i(X) - P_i(Y)] = V_i(Z).[P_i(X) - P_i(Y)] \\ \Leftrightarrow & P_i(X) = P_i(Y) \end{aligned}$$

■

There is a strong Behaviourist strain in decision theory that views facts about preference, as revealed in inter-subjectively observable choice behaviour, as the determinant of the empirical and/or semantic content of ascriptions of belief and desire to agents and which sees decision theoretic representation theorems as formal demonstrations of what is or is not empirically significant in such mentalistic talk. Behaviourists will infer from Bolker's uniqueness results that there is no empirical foundation for claims such as 'the agent believes X to the same degree as Y ' and hence no (scientific) sense to them either. Such a view would, because of the Homogeneity theorems that we discuss later, have very serious consequences for the intelligibility of Utilitarianism in this framework. In my opinion, however, Behaviourism both exaggerates the extent to which preference is unproblematically revealed in observable behaviour and ignores the existence of other routes, such as introspection, by which the content of mental ascriptions can be fixed and claims about mental properties validated.⁵ If there are independent grounds for treating as meaningful ascriptions of mental states, even quantitative ones like degrees of belief or desire, then it is more natural to regard Bolker's representation theorem as only partially relating the concept of rational preference to the richer ones of rational belief and desire. (This leaves open, of course, the question as to whether the gap between them can be bridged by further constraints on preference, or whether the relative

⁵This claim is developed in much greater detail in Bradley (2001)

independence of belief and desire from preference is in the ‘natural’ order of things).

We shall return to these issues later on. For now it is important to note that there is an exception to the general rule of non-unicity; namely when an agent’s preferences admit only of representation by unbounded utility measures. In this case, it follows from equation (5) that $c = 0$ and hence from (3) that $d = 1$. It is possible to characterise this case axiomatically: Jeffrey [9, p. 142], for instance, gives a necessary and sufficient condition for it. But the condition Jeffrey gives is intuitively opaque and it is difficult to say whether it is reasonable or not that an agent’s preferences should satisfy it. So we will follow Broome in simply characterising this case as follows:

Definition 6 (*Unboundedness*) *A preference ranking \geq_i will be said to be unbounded iff all signed measures U_i that represent \geq_i are unbounded from both above and below.*

2.3 Social Preference in the Jeffrey-Bolker Framework

Let $\beta = \langle \Omega, \preceq \rangle$ be an atomless Boolean algebra of propositions minus the contradiction F . Let G be a set of h individuals and $F = \{\geq_1, \geq_2, \dots, \geq_h\}$ be a corresponding set of h transitive, complete and continuous two-place relations on Ω . Let $G^+ = G \cup \{G\}$ and $F^+ = F \cup \{\geq_G\}$. The \geq_i s and \geq_G are respectively the individuals’ and group’s ordering of the prospects in Ω . The following axioms are assumed throughout the paper:

Axiom 7 (*Individual Rationality*) *All the \geq_i s in P satisfy Bolker’s axioms of rational preference*

Axiom 8 (*Group Rationality*) *\geq_G satisfies Bolker’s axioms of rational preference*

Axiom 9 (*Pareto Indifference*) *If $\forall(i \in G), X \approx_i Y$ then $X \approx_G Y$*

Axiom 10 (*Strict Pareto*) *If $\forall(i \in G), X \geq_i Y$ and $\exists(j \in G : X >_j Y)$ then $X >_G Y$*

Pareto Indifference and Strict Pareto are jointly called Strong Pareto. From the axioms of individual and group rationality and Bolker’s Representation Theorem, it follows that $\forall(i \in G^+)$ there exists a probability measure P_i and signed utility measure U_i that jointly represent \geq_i in the sense that $\forall(X, Y \in \Omega)$:

$$X \geq_i Y \Leftrightarrow \frac{U_i(X)}{P_i(X)} \geq \frac{U_i(Y)}{P_i(Y)} \Leftrightarrow V_i(X) \geq V_i(Y)$$

and which are unique up to the transformations identified by equations (1) to (6).

It is worthy noting at this point that Bolker’s axiom of Averaging in effect requires individuals to assign non-zero probability to all contingent propositions

in the domain of the preference relation. This may not be too much of a limitation in individual decision theory as propositions of zero probability can be identified and either removed from the algebra or treated separately - see Bolker (1967, p.337) and Jeffrey (1983, p.113-4). But it does present a more substantial issue in the context of group decision making as we should not assume that individuals assign probability one to the same propositions. The problem is, however, largely tangential to the concerns of this paper and can be finessed by assuming that all individual's preference relations are defined on the same domain; in effect, the largest Boolean set containing no proposition assigned probability zero by any individual.

In addition to the four basic axioms listed above we shall also have recourse at points in the discussion to the following technical conditions on preference.

Condition 11 (*Minimum Agreement*) $\exists(C \in \Omega)$ such that $\forall(i \in G), C >_i \neg C$

Condition 12 (*Weak Independence*) $\forall(i \in G), \exists(Z_i \in \Omega)$ such that $Z_i >_i T$ but $\forall(j \in G - i), Z_i \approx_j T$

Condition 13 (*Independence*) $\forall(i \in G, Z \in \Omega), \exists(Z_i \in \Omega)$ such that $\forall(j \in G - i), Z_i \approx_i Z$ but $Z_i \not\approx_i Z$

Condition 14 (*Strong Independence*) $\forall(i \in G, Z \in \Omega), \exists(Z_i \in \Omega)$ such that $Z_i \approx_i Z$ but $\forall(j \in G - i), Z_i \approx_i T$

The Minimum Agreement condition is a rather innocuous assumption that could probably be dispensed with, but which considerably simplifies the proofs that follow. The independence conditions are more contentious, though how much so depends on the precise interpretation of the ordering relation on prospects. More on this later.

2.3.1 Representations of the Group's Preferences

A vector valued measure $\langle P_1, P_2, \dots, P_h, P_G, U_1, U_2, \dots, U_h, U_G \rangle$ is said to be a representation of the group's preferences iff for all $i \in G^+$ the pair $\langle P_i, U_i \rangle$ meets the conditions of Bolker's theorem for representing the corresponding ranking \geq_i . The following conditions on representations of group preferences will play a role in our discussion.

Condition 15 (*Utilitarian Condition*) $\exists(\alpha_1, \alpha_2, \dots, \alpha_h \in \mathfrak{R})$ such that $\alpha_i > 0$, $\sum_{i=1}^h \alpha_i = 1$ and $V_G = \sum_{i=1}^h \alpha_i V_i$

Condition 16 (*Probability Identity*) $\forall i, P_i = P_G$

Condition 17 (*Utility Identity*) $\forall i, \exists(\alpha_i > 0, \beta_i \in \mathfrak{R})$ such that $U_G = \alpha_i U_i + \beta_i$.

If the Utility Identity condition holds for one representation of the group's preferences then it holds for all of them. This follows immediately from the fact that Utility Identity implies that $X \geq_G Y \Leftrightarrow X \geq_i Y$. But because of the non-uniqueness of probabilities in Jeffrey-Bolker framework, the fact that either the Utilitarian or the Probability Identity condition does or does not hold for one representation does not in itself imply that it holds for all of them. It is important therefore to distinguish weak and strong versions of the thesis that Utilitarian representations of social evaluations exist.

Strong Utilitarian Representation Thesis: Every representation of the group's preferences satisfies the Utilitarian condition.

Weak Utilitarian Representation Thesis: There exists a representation of the group's preferences that satisfies the Utilitarian condition.

It is not clear which of these two theses Broome seeks to establish, since the difference between the two plays no role in his paper. This is a consequence of the fact that in order to establish the Utilitarian thesis he assumes that preferences are unbounded which, as we noted before, suffices to determine unique representations, up to choice of scale, of individual's preference rankings. Given the assumption of unboundedness, the Weak and Strong Representation theses are equivalent.

2.3.2 Broome's Axiomatic Utilitarianism

We are now in a position to state more formally Broome's Utilitarian and Homogeneity theorems for social preference in the Bolker-Jeffrey framework.

Theorem 18 (*Probability Identity*) *Assume the axioms of Individual and Group Rationality and that the Utilitarian and Weak Independence conditions hold. Then the Probability Identity condition holds.*

Theorem 19 (*Utilitarian*) *Assume the axioms of Individual and Group Rationality, that individuals' preferences are unbounded and that the Strong Independence condition holds. Then Strong Pareto implies that the Utilitarian condition holds.*

Both of these results will be at the centre of our discussions. Broome's 1990 paper [6] contains a proof of the Utilitarian theorem, but he did not publish his proof of the Probability Identity theorem. We provide a different and shorter proof of this necessity claim in Section 3, which assumes the Minimum Agreement condition rather than Weak independence. This section also states and proves a corresponding Utility Identity theorem to the effect that the Strong Utilitarian Representation thesis implies the identity of the utilities of individuals, up to a choice of common scale, in any case in which preferences are bounded. The philosophical analysis of these Homogeneity theorems and an evaluation of Broome's response to them is hived off into Section 4. In Section 5 we turn to the examination of the Utilitarian theorem. Here we extend

Broome’s work in two ways. First we investigate in a more general way the conditions for homogeneity of belief in the Jeffrey-Bolker framework in the light of Bolker’s uniqueness result. Secondly, we partially recast Broome’s Utilitarian theorem in a way which separates the role played by the assumptions of unbounded preferences Strong Independence, and which replaces the latter with a more plausible condition. Finally we give sufficient conditions on (potentially bounded) preference for the truth of the Weak Utilitarian Representation thesis.

3 The Homogeneity Theorems

3.1 A Simple Illustration

In this section we prove, under a weak assumption, that the Utilitarian condition entails the identity of the probabilities of the individuals making up the group. Let us begin with a simple illustration of the result. Suppose we have two individuals, Ann and Bob, and a single action, A, of going for a walk that needs to be evaluated with respect to the two relevant (and exhaustive) possibilities, that of it being a hot day and of being a cool day. Suppose that Ann likes to walk when it’s hot, but not when it’s cool, and that Bob’s preferences are just the reverse. Suppose also that Ann believes that it is most likely to be a hot day and that Bob believes that it will be cool. Then they may both agree that it is a good idea for them to go walking, even though they disagree about everything that this choice depends on. To see this suppose that Ann and Bob’s degrees of belief and preference are as given below. (Implicitly the default option of doing nothing has zero utility for both).

Degrees of Belief

	Hot Day	Cool Day
Ann	0.9	0.1
Bob	0.1	0.9

Degrees of Preference

	Hot Day	Cool Day	Overall
Ann	5	-15	3
Bob	-15	5	3
Jointly	-5	-5	-5 or 3?

Ann and Bob want to make a decision on an impartial Utilitarian basis and so decide to give equal weighting to the utilities of each in determining their joint utility for taking a walk. Calculating their joint utility separately in the event of a hot day and in the event of a cool one, they find that in both cases it equals $(5 * 0.5) + (-15 * 0.5) = -5$. Since for them as a couple walking is worse than doing nothing in either event, they should decide to stay at home

irrespective of the likelihood of a hot or a cool day. But when they calculate the overall utility for walking separately for each of them, they find that both favour walking overall since for both it has a utility of $(5 * 0.9) + (-15 * 0.1) = 3$. So thinking about it this way they should jointly favour walking, indeed irrespective of whether they decide to make their decision impartially or not. So Utilitarian aggregation in this situation of uncertainty has not yielded a consistent decision.

The symmetry of the situation suggests that Ann and Bob can be sure that the value of their joint utility for walking will be independent of the order in which the utilities are aggregated only if they give the same probability to the events of it being a hot day and of it being a cold one. In fact, order independence can also be maintained in this example if Ann and Bob's preference ranking of the possible outcomes are the same, but only on the condition that they assign a joint probability for a hot day equal to a weighted average of their individual probabilities for a hot day, with the ratio of the weight on Ann's probability to Bob's being equal to the ratio of the difference between Ann's utility for walking on a hot day and her utility for walking on a cold day to the difference between Bob's utilities for these outcomes. This is clearly a rather special case and the formal results that follow show, in effect, that it is ruled out when larger sets of events and consequences are considered.

Our simple example of Ann and Bob's decision problem illustrates the difficulties faced by any attempt to aggregate a single quantity over two different dimensions (in this case persons and states of the world). It also serves to demonstrate a problem that afflicts any theory holding to the Pareto condition, Utilitarian or otherwise. Suppose that the utilities cited in the example have interpersonal significance, so that it can be said that they show Bob to dislike walking on a hot day more than Ann likes it (we are often in a position to make judgements of this kind). Then the reasonable thing to conclude would be that they should not go walking because whether it's hot or cold, one of them is going to suffer more than the other benefits. On the other hand both judge walking to be a good option overall because they believe it unlikely that the day will be such as to make walking unpleasant for them. Hence the Pareto condition requires that walking be considered a good option for them jointly, despite the fact that we know that is not. The problem here is that the fact that they agree that walking is a good option is misleading, since they arrived at this judgement on the basis of completely conflicting beliefs and preferences. Their agreement is spurious.

3.2 The Necessity of Probability Identity

We begin with a formal statement of the central claim of this section; namely that satisfaction of the Utilitarian condition requires that all individuals' probabilities be the same.

Theorem 20 (*Probability Identity*) *Assume the axioms of Individual and Group Rationality and that the Utilitarian and Minimum Agreement conditions hold. Then the Probability Identity condition holds.*

Theorem 20 follows immediately from Lemmas 22 and 23 which we prove below. To simplify the proofs we will assume a common choice of the tautology T as the zero of every individual's utility function and of the groups. No loss of generality is incurred by this simplification since, by Bolker's uniqueness theorem, any pair of probability and signed utility measures, $\langle P_i, U_i \rangle$, that represents i 's preferences, can be transformed into another such a representative pair of measures, $\langle P'_i, U'_i \rangle$ by defining $U'_i = U_i - U_i(T)$ and $P'_i = P_i$. By so doing we normalise the utilities in required manner without affecting the probabilities (or any conclusion we reach about them).

From our choice of zero⁶ and Bolker's uniqueness theorem it follows that there exists for all i another representative pair of measures, P'_i and U'_i , such that $U'_i(T) = 0$ iff there exists real numbers $a > 0$ and c such that $cU_i + P_i > 0$ and:

$$U'_i = aU_i \quad (7)$$

$$P'_i = P_i + cU_i \quad (8)$$

Note also that on this normalisation it follows from the axiom of Averaging that:

$$X >_i \neg X \iff V_i(X) > V_i(T) = 0 > V_i(\neg X) \quad (9)$$

Now let $\gamma = \langle P_1, \dots, P_h, P_G, U_1, \dots, U_h, U_G \rangle$ be any representative vector valued measure such that $\forall(i \in G^+), U_i(T) = 0$, and let Γ be its range on Ω . To prove the lemmas that follow, we make use of the following theorem proved in Bolker (1966, p.295).⁷

Theorem 21 (*Liapounoff*) *The range of a vector valued measure on an atomless measure algebra is convex.*

Lemma 22 *Assume Individual and Group Rationality and that the Utilitarian condition holds for γ . Then if proposition C is such that $\forall(i \in G), C >_i \neg C$, then $\forall(i \in G), P_i(C) = P_G(C)$*

Proof. By the Utilitarian Condition $\exists(\alpha_1, \alpha_2, \dots, \alpha_h)$ such that $\sum_{i=1}^h \alpha_i = 1$ and $V_G = \sum_{i=1}^h \alpha_i V_i$. Hence,

$$U_G(C) = \sum_{i=1}^h \alpha_i U_i(C) \cdot \frac{P_G(C)}{P_i(C)} \quad (10)$$

Note that since the U_i are measures,

$$U_i(T) = U_i(C \vee \neg C) = U_i(C) + U_i(\neg C) = 0 \quad (11)$$

⁶Strictly it follows from our choice of zero and the fact that any two pairs of representative probability and utility measures are related by the fractional linear transformations (5) and (6), for conditions (2) - (4) on the variables a, b, c and d . I am grateful to Philippe Mongin for this point.

⁷It is a version of a theorem first proved by Liapounoff for σ -additive measures on σ -algebras. Hence the name.

So from 10 and 11,

$$U_G(T) = \sum_{i=1}^h \alpha_i U_i(C) \frac{P_G(C)}{P_i(C)} + \sum_{i=1}^h \alpha_i U_i(\neg C) \frac{P_G(\neg C)}{P_i(\neg C)} = 0 \quad (12)$$

But from 11, $U_i(C) = -U_i(\neg C)$. So from 12:

$$\sum_{i=1}^h \alpha_i U_i(C) \left(\frac{P_G(C)}{P_i(C)} - \frac{P_G(\neg C)}{P_i(\neg C)} \right) = 0$$

Now suppose that $\forall(i \in G), C >_i \neg C$. Then all the $U_i(C)$ are positive, as are all the α_i . It follows that:

$$\frac{P_G(C)}{P_i(C)} - \frac{P_G(\neg C)}{P_i(\neg C)} = 0$$

So $P_i(C) = P_G(C)$. ■

Lemma 23 *Assume Individual and Group Rationality and that the Minimum Agreement condition holds. Suppose that if $\forall i, C >_i \neg C_i$ then $\forall i, P_i(C) = P_G(C)$. Then $\forall(i, X), P_i(X) = P_G(X)$*

Proof. Take any proposition X . Let $G1$ be set of individuals $i \in G$, such that $X \geq_i \neg X$, and $G2$ be $G - G1$. If either $G1$ or $G2$ are empty then by assumption $\forall(i \in G), P_i(X) = P_G(X)$. So suppose that neither are empty. Then let C be any proposition such that $\forall(i \in G), C >_i \neg C$. By the Minimum Agreement condition such a proposition exists and by assumption, $\forall(i \in G), P_i(C) = P_G(C)$. Now let i^* be any member of $G2$ such that:

$$\frac{-U_{i^*}(C)}{U_{i^*}(X) - U_{i^*}(C)} = \min_{i \in G2} \frac{-U_i(C)}{U_i(X) - U_i(C)}$$

Then define $\gamma(D) = \beta \cdot \gamma(X) + (1 - \beta) \cdot \gamma(C)$, where:

$$0 < \beta < \frac{-U_{i^*}(C)}{U_{i^*}(X) - U_{i^*}(C)} \leq 1$$

Note that since $U_{i^*}(C) > 0 \geq U_{i^*}(X)$, β must exist. It then follows by Liapounoff's Theorem that proposition D exists. Now clearly $\forall(i \in G1), U_i(D) > 0$ since $U_i(X) \geq 0$ and $U_i(C) > 0$. Hence $\forall(i \in G1), D >_i \neg D$. Equally since $\forall(i \in G2), C >_i \top \geq_i X$ it follows that $\forall(i \in G2)$:

$$\begin{aligned} U_i(D) &> \frac{-U_{i^*}(C)}{U_{i^*}(X) - U_{i^*}(C)} \cdot U_i(X) + \frac{U_{i^*}(X)}{U_{i^*}(X) - U_{i^*}(C)} \cdot U_i(C) \\ &\geq \frac{-U_i(C)}{U_i(X) - U_i(C)} \cdot U_i(X) + \frac{U_i(X)}{U_i(X) - U_i(C)} \cdot U_i(C) \\ &= 0 \end{aligned}$$

Hence $\forall(i \in G), D \succ_i \neg D$. So by assumption, $\forall(i \in G), P_i(D) = P_G(D)$. But by definition, $P_i(D) = \beta.P_i(X) + (1 - \beta).P_i(C)$. So

$$\begin{aligned} P_G(D) &= \beta.P_i(X) + (1 - \beta).P_G(C) \\ &= \beta.P_G(X) + (1 - \beta).P_G(C) \end{aligned}$$

Hence, $\forall(i \in G), P_i(X) = P_G(X)$. ■

3.2.1 Utility Identity

Except in the special case when individuals' preferences are unbounded, the Strong Utilitarian Representation thesis implies that all individuals rank the set of prospects in an identical manner. This claim is rather startling at first sight but, as the Utility Identity theorem that follows shows, it is in fact a rather straight-forward consequence of the Probability Identity theorem and the non-uniqueness of probability representations in Jeffrey-Bolker framework.

Theorem 24 (*Utility Identity*) *Let $\bar{\gamma} = \langle \bar{P}_1, \dots, \bar{P}_h, \bar{P}_G, \bar{U}_1, \dots, \bar{U}_h, \bar{U}_G \rangle$ be a representation of the group's preferences that is distinct from γ at least to the extent that, for some $i \in G$, $P'_i \neq P_i$. If $\bar{\gamma}$ also satisfies the Utilitarian condition, then the Utility Identity condition holds.*

Proof. By Bolker's uniqueness theorem, it follows from the existence of $\bar{\gamma}$ that there exists another representation of the group's preferences $\gamma' = \langle P'_1, \dots, P'_h, P'_G, U'_1, \dots, U'_h, U'_G \rangle$ such that $\forall(i \in G^+), P'_i = \bar{P}_i$ and $U'_i = \bar{U}_i - \bar{U}_i(T)$. Then for such i , $U'_i(T) = 0$ and γ' satisfies the Utilitarian condition. So by Theorem 23, $\forall(i \in G), P_i = P_G$ and $P'_i = P'_G$. Then by Bolker's Uniqueness Theorem, specifically equation 8, it follows that $\forall(i \in G) \exists(c_i, c \in \mathfrak{R})$ such that:

$$P'_i = P_i + c_i U_i = P'_G = P_G + c U_G$$

It follows that $c U_G = c_i U_i$. By hypothesis $\gamma \neq \gamma'$, so it is not the case that $\forall i, c_i = 0$. Hence $c \neq 0$ and $U_G = (\frac{c_i}{c}).U_i$. And it follows from the Minimum Agreement condition, that $\frac{c_i}{c} > 0$. So the Utility Identity condition holds. ■

3.3 The Sufficiency of Probability Identity

In this final part of section 3, we show that the identity of individuals' probabilities is not just necessary for the truth of the Utilitarian condition, but that it is also sufficient for it. The proof improves on Broome's [6, p. 490] in that it makes no use of the assumption of Weak Independence.⁸

Theorem 25 *If the Probability Identity condition holds then so too does the Utilitarian condition.*

⁸The argument I use in the proof is essentially that employed by Mongin to prove his Proposition 4 [13, p. 334].

Proof. Let $\mathbf{V} =_{def} \langle V_1, V_2, \dots, V_h \rangle$ and $\bar{\mathbf{V}}$ be its range on Ω . From Pareto Indifference it follows that there exists a strictly increasing real valued function f on $\bar{\mathbf{V}}$ such that $V_G(X) = f(\mathbf{V}(X))$, for all $X \in \Omega$. We prove the theorem by showing that $\bar{\mathbf{V}}$ is convex and that f is mixture preserving on $\bar{\mathbf{V}}$. For any $X, Y \in \Omega$ and $\alpha \in [0, 1]$, let $\lambda = \frac{\alpha P_G(Y)}{\alpha P_G(Y) + (1-\alpha)P_G(X)}$. Since $0 \leq \lambda \leq 1$, it follows by Liapounoff's Theorem that there exists $Z \in \Omega$, such that $\gamma(Z) = \lambda\gamma(X) + (1-\lambda)\gamma(Y)$. Then for all $i \in G^+$:

$$\begin{aligned} V_i(Z) &= \frac{U_i(Z)}{P_i(Z)} \\ &= \frac{\lambda U_i(X) + (1-\lambda)U_i(Y)}{\lambda P_i(X) + (1-\lambda)P_i(Y)} \\ &= \frac{\alpha P_G(Y).U_i(X) + (1-\alpha)P_G(X).U_i(Y)}{\alpha P_G(Y).P_i(X) + (1-\alpha)P_G(X).P_i(Y)} \end{aligned}$$

But by assumption $P_i = P_G$, so:

$$\begin{aligned} V_i(Z) &= \frac{\alpha P_i(Y).U_i(X) + (1-\alpha)P_i(X).U_i(Y)}{P_i(X).P_i(Y)} \\ &= \alpha V_i(X) + (1-\alpha)V_i(Y) \end{aligned}$$

So $\bar{\mathbf{V}}$ is convex. It also follows that $V_G(Z) = f(\mathbf{V}(Z)) = f(\alpha\mathbf{V}(X) + (1-\alpha)\mathbf{V}(Y))$ and that:

$$\begin{aligned} V_G(Z) &= \alpha V_G(X) + (1-\alpha)V_G(Y) \\ &= \alpha f(\mathbf{V}(X)) + (1-\alpha)f(\mathbf{V}(Y)) \end{aligned}$$

Hence f is mixture preserving. ■

4 Utilitarianism and Probability Identity

4.1 The Impossibility of Preference Utilitarianism

The Homogeneity theorems are in effect triviality theorems for Preference Utilitarianism: the view that the strength of society's preference for any prospect should be a weighted average of the strength of its members' preferences for it. Preference Utilitarianism fails because it combines two mutually inconsistent principles: that social preferences should be consistent and that they should systematically depend on the preferences of the individuals making up the society. Since the preferences of individuals depend on their beliefs, the second principle requires a group's preferences to depend on its member's beliefs. But when the beliefs of individuals differ there may be no way of retaining this dependency and at the same time for the group's preferences to reflect a consistent set of group beliefs. The simple example of Ann and Bob's difficulty in reaching a consistent judgement on the desirability of taking a walk together illustrated just how this can be a problem.

The Ann and Bob example also served to demonstrate that any theory holding to the Pareto condition, Utilitarian or otherwise, was likely to run into problems caused by spurious preference agreement. This is somewhat more manifest in the Homogeneity results that Mongin [13] proves for the Savage framework. He shows that if the preference relation on the set of prospects (acts, in this context) satisfies Savage's postulates, the Pareto Indifference axiom and the Minimum Agreement condition then:

1. If the individuals' utilities are affinely independent then the Probability Identity condition holds.
2. If the individuals' probabilities are linearly independent then the Utility Identity condition holds.

Results 1. and 2. may be regarded as the analogues of our Theorems 20 and 24. (Mongin also proves an exact analogue of our Theorem 25 for the Savage framework). To see the connection between 1. and Theorem 20, notice that in the Jeffrey-Bolker framework and given the chosen normalisation, Weak Independence implies that the U_i are affinely independent. The connection between 2. and Theorem 24 is somewhat looser. This reflects the fact that in Savage's framework probability and utility play symmetric roles, which is not the case in the Jeffrey-Bolker one.

Mongin's results are slightly weaker than ours in that they respectively need to suppose that individuals' probabilities are linearly independent and their utilities affinely independent. This is to be explained by the fact in the Savage framework agents' evaluations of consequences and beliefs about the states of the world are separately defined, while in the Jeffrey-Bolker framework desirabilities and probabilities are entangled to a much greater degree. On the other hand, Mongin's results are much stronger in that homogeneity of belief and desire follows from the assumption of Strong Pareto alone, without the assumption of the Utilitarian condition (which, of course, implies Strong Pareto). As we shall see in section 5.1, in the Jeffrey-Bolker framework, homogeneity of belief follows from Strong Pareto and an independence condition only for prospects that are ranked together by all members of the group. This difference can be explained by the fact that Savage's theory admits arbitrary functions from states to consequences into the domain of the ordering relation, making it is much easier to exploit spurious agreement.

Mongin's results for the Savage framework strongly suggests that the impossibility of Preference Utilitarianism derives from a conflict between its Bayesian and its Paretian elements. Mongin's view is that the fact that Pareto condition treats spurious agreements in the same way as genuine ones is good grounds for doubting its validity as an aggregative principle.⁹ This conclusion seems inescapable at least as long as we stay within the bounds of the standard interpretation of the ordering on prospects as preference relations. In this context, moreover, the rationality assumptions look secure, as preference is an intrinsically normative concept and it is part of the meaning of someone holding a

⁹See Mongin [14].

particular preference that it should stand in certain consistency relations to other preferences.

Under other interpretations of the ordering relations the Pareto principle has a stronger rationale. In social choice theory, the Pareto principle is not infrequently viewed as a minimal condition for democratic decision-making. This again involves interpreting the individuals' ordering relations as preferences, but the group's ordering relation as nothing more than a choice mechanism. Although the concept of choice itself may invite consistency conditions of a minimal kind (e.g. transitivity), there need be no implication that the choices are outcomes of full-blooded social agency and, hence, no requirement that choices cohere with respect to something like a consistent set of beliefs. Consequently the Democrat has nothing to fear from the Probability Identity theorem. Conditions of diversity of belief may imply that society's choices will lack Bayesian rationalisation, but the Democrat can argue that the rationalisability or otherwise of society's choices are simply tangential to the question of their rightness.

An alternative rehabilitation of the Pareto Principle is developed by Broome in his book '*Weighing Goods*' [7]. Broome argues that what Utilitarians are interested in is the social good. Preference Utilitarians think that what is good for a group depends on the preferences of its members. But, Broome claims, this is false; what is good for a group depends on what is *good* for its members. A Utilitarian should therefore be committed to what Broome terms the Principle of Personal Good (PPG): that if one prospect is better than another (or equally good) for everybody in the group, then the former prospect is better than the latter (or equally good) for the group as a whole. PPG is of course just the Strong Pareto principle with both the individual and group ordering relations on prospects interpreted as betterness relations. Nonetheless, on this interpretation too, the Utilitarian has nothing to fear from the Probability Identity theorem. Since what is good for each individual depends on the *objective* probabilities of events, not on the individuals' subjective beliefs, probability identity is to be expected, indeed embraced.

4.2 The Betterness Interpretation

The democratic interpretation evades the Probability Identity theorem by denying that social choice is constrained by Bayesian consistency conditions. The Utilitarian, on the other hand, embraces probability identity and so need not refuse the consistency conditions. Broome argues that, to the contrary, both individual and social betterness are subject to coherence constraints that are formally identical to those to which preference is subject. This opens the way to application of the Bayesian rationality axioms to betterness relations to obtain, in conjunction with the Principle of Personal Good, a Utilitarian theorem for representations of a group's betterness relation. Broome's view is that it is precisely in this context that Utilitarian theorems have their primary significance, as demonstrations of the structure of the social good and its relation to the goods of individuals.

Broome makes the case for this interpretation in great detail in his book

'*Weighing Goods*' and there is no need to reproduce it here. What is striking, however, is that he does so in the framework of Savage's decision theory. This is unfortunate because the aforementioned weaknesses of Savage's theory do not disappear with its re-interpretation as a theory of coherent betterness. The objections to Savage's postulation of dubious acts, for instance, carry over without qualification. More importantly, the argument against state-independent utility still holds. The relative goodness of say taking a bus versus walking to my destination depends on such factors as the level of traffic, the likelihood of rain and my physiological condition. So it should certainly not be built into a theory of the good that the goodness of a prospect is independent of the state of the world in which it is realised.

This suggests that there might be much to be gained by pursuing Broome's re-interpretation of the Utilitarian Theorem within Jeffrey-Bolker decision theory. This is the thesis that would result. Suppose that both the individuals' and the group's betterness rankings of prospects satisfy Bolker's axioms. Then from Bolker's representation theorem it follows that the goodness of any prospect for an individual or for the group is a probability weighted average of the ways in which the prospect can be realised. Call this derived principle Ethical Bayesianism. By Theorem 25, furthermore, if the same probabilities determine the good of each of the individuals making up the group, then the Principle of Personal Good is sufficient for the truth of the Utilitarian Representation Thesis. But the good of every individual depends on the actual or objective probabilities of prospects, so they must depend on the same probabilities. Given this, Ethical Bayesianism plus the Principle of Personal Good implies that the social good has a Utilitarian representation.

This is, I believe, a very important claim. But if it is to be defended an important question needs to be addressed first. We have seen that in the Jeffrey-Bolker framework the probabilities of individuals are not uniquely determined by their preference rankings of propositions. But if it is not determined what an individual's probabilities are, then it is not determined when they are the same as any other individual's. So it is not altogether clear what it means in this framework to say that individuals have identical beliefs, or on the betterness interpretation, that the goods of every individual depend on the same probabilities. But since the truth of the Utilitarian Representation Thesis presupposes that this condition holds we must find some way of interpreting it. It turns out that this problem is closely connected to issues raised by the Utility Identity theorem and it is to consideration of the latter that we now turn.

4.3 Utility Identity and the Uniqueness of Probabilities

Can the Utility Identity theorem also be deflated by Broome's proposed re-interpretation of ordering relations on prospects as betterness rankings? One immediate implication of this theorem is that if any one individual's utilities are bounded, then the Strong Utilitarian Representation thesis is sustainable only if all individuals have the same preferences, or on the proposed re-interpretation, the same betterness rankings. This is clearly implausible: one thing may pre-

ferred by you to another, but be less preferred by me. Equally it may be better for you, but worse for me. So the Utility Agreement theorem shows that the strong thesis must be abandoned in favour of the weak. To do so, however, considerably diminishes the extent to which the Utilitarian theorem can be said to reveal anything about the structure of goodness. For the weak thesis says only that it is *possible* to find Utilitarian representations of individual and social betterness rankings, not that they *have* to be so represented. For someone that already believes that betterness has a Utilitarian structure, such a possibility result is no doubt of some comfort. But the choice of whether to so represent things must now be made on grounds external to the axiomatic environment defined by Bayesian decision theory plus the Principle of Personal Good. This would diminish the interest of the Utilitarian theorems to a considerable extent.

This conclusion is perhaps a little too hasty, however. The distinction between a weak and strong version of the Utilitarian Representation thesis derived from the non-unicity of probability representations in the Jeffrey-Bolker framework. But it does not follow from the fact that probabilities are not uniquely *determined* by rankings of propositions that they are not, as a matter of fact, *determinate*. To assure this implication it must be established that the concept of comparative betterness has priority of a semantic and/or epistemic kind over that of goodness; so that what can be said and/or known about the goodness of prospects is constrained by what can be said and/or known about their comparative betterness. From the truth of this priority thesis, and on the assumption that the concept of betterness is exhausted by Bolker's axioms, it would follow that the non-unicity of the probability factor in the goodness of prospects is intrinsic or ineliminable. And then the Utility Identity theorem would bite.

The simplest way of denying this inference, and one which is consistent with the claim that goodness depends on betterness, is to argue that Bolker's axioms do not exhaustively characterise coherent betterness and that supplementation will open the way to the unique determination of probabilities by betterness relations. I suspect, however, that this cannot be done in a natural manner within the Jeffrey-Bolker framework. It is, of course, possible to obtain uniqueness by assuming that individuals' preferences are unbounded. But, as with preference relations, it is hard to say whether betterness relations can be unbounded, let alone whether they should be. In many decision-theoretic frameworks, including Savage's, it is possible, given an unbounded sequence of prospects, to construct a gamble or act with infinite utility. And the St. Petersburg paradox shows that infinite utilities cannot be admitted into Bayesian decision theory without unpalatable implications. But such gambles cannot be constructed within Jeffrey-Bolker framework and so the threat of infinite utilities cannot be used to support the claim that preferences must be bounded (but see below). On the other hand, the requirement that they not be bounded lacks any clear justification. Certainly it is not obvious that the *concept* of betterness requires it. So it is at best an open question as to whether preferences should be unbounded.

The alternative is to reject the priority thesis and appeal to considerations outside the scope of an axiomatic theory of comparative betterness: for instance, to direct arguments for the existence of a unique goodness-determining probabil-

ity. There is a cost to this approach, however, for it necessarily reduces the role that the Utilitarian theorem plays in justifying the Utilitarian Representation thesis. The theory of comparative betterness, as expressed by Bolker's axioms, has much greater initial plausibility than Ethical Bayesianism. So it would be of no small significance if the former implied the latter. But if Bolker's axioms do in fact exhaustively characterise comparative betterness relations between propositions, then the support that considerations of coherent comparative betterness affords the theory of goodness is less than total. In particular, it simply would not follow from the theory of betterness that there was a 'natural' probability factor that determines the goodness of prospects in the manner claimed by Ethical Bayesianism.

What this suggests, I think, is that the Jeffrey-Bolker framework is not the ideal vehicle for Broome's project. The solution is not to go back to Savage's - it's weaknesses are too profound - but to strengthen the Jeffrey-Bolker framework so as to allow the formulation of stronger conditions on rankings of prospects. A project of just this kind is undertaken in Bradley [2] and [3] where the set of prospects is supplemented with conditional ones, such as the prospect that if the train leaves promptly, you will be home in time for dinner and the prospect that if inflation rises the government will fall. In this extended framework probabilities *are* uniquely determined by the ordering of prospects and so it makes possible the pursuit of Broome's project without fear of any analogue of the Utility Identity theorem. We shall not pursue this possibility here, save to remark that the introduction of conditional prospects does allow the construction of something like a St. Petersburg game. Since the way it extends the Jeffrey-Bolker framework is a rather natural one, this amounts to an indirect argument against the assumption that preferences are unbounded.

5 The Utilitarian Theorems

We now turn our attention to the problem of determining sufficient conditions on rankings of prospects for the truth of the Utilitarian Representation Thesis. From Theorems 20 and 25, however, we know that it is both necessary and sufficient for this that we find conditions on rankings of prospects sufficient for individuals' probabilities to be identical. But the non-uniqueness of probabilities in the Jeffrey-Bolker framework means that this cannot be done outside of the special case of unbounded rankings. In the first section we look at types of probability homogeneity that are less demanding than probability identity but which *can* be characterised by properties of rankings of prospects. We then identify combinations of Paretian and independence conditions sufficient for the presence of probability homogeneity of this kind. These results are applied in the second section where we attempt to reformulate and generalise Broome's Utilitarian theorem.

5.1 Conditions for Homogeneity of Probabilities

In the Jeffrey-Bolker framework what is determined by rankings of prospects is the equiprobability of co-ranked propositions and more generally the ratio of the probabilities of co-ranked propositions. It is possible thus to characterise in terms of shared properties of rankings a kind of restricted homogeneity in the probabilities of individuals. In particular we will say that probabilities of some group of individuals are **JB-homogeneous** if its the case that whenever every individual in the group ranks two propositions X and Y together, then the ratio of their probabilities for X and Y are the same.

We show below, in Theorem 28, that the JB-homogeneity of a group's probabilities is derivable from Pareto Indifference and the Independence condition. This result generalises Broome's Theorem 2 [6, p. 91] and provides the promised counterpart to Mongin's result for the Savage framework (cited in section 4.1); that Pareto Indifference, Minimum Agreement and the affine independence of individuals' utilities are sufficient for Probability Identity. Though the probability homogeneity established by Theorem 28 is considerably weaker, its conceptual implications are similar; namely that the Pareto conditions severely restrict the extent to which individuals' probabilities can differ. The restriction simply manifests itself with greater ease in Savage's framework because the latter affords greater flexibility in the construction of actions.

JB-homogeneity of a group is not the maximum specifiable in the Jeffrey-Bolker framework in terms of properties of rankings. This is given by its natural strengthening: the condition that JB-homogeneity extends to the probabilities of every subset of the group. To give axiomatic conditions on rankings for such sub-group homogeneity, it would be necessary to modify the axiom of Pareto Indifference in such a way that it applied to arbitrary subsets of G . The generalisation is a natural one in the context of the betterness interpretation and has implications for some of the results in the next section as well, so it is worth our while stating it formally.

Axiom 26 (*Sub-group Pareto Indifference*) $\forall(g \subseteq G)$, if $\forall(i \in g)$, $X \approx_i Y$ then $X \approx_g Y$

Lemma 27¹⁰ Let γ be a vector valued measure on Ω and X, Y and Z be any propositions. Then there exists mutually contradictory propositions $X^*, Z^*, Y^* \in \Omega$ such that $\gamma(X^*) = \frac{1}{8}\gamma(X)$, $\gamma(Y^*) = \frac{1}{8}\gamma(Y)$ and $\gamma(Z^*) = \frac{1}{8}\gamma(Z)$.

Theorem 28 Assume Pareto Indifference and that the Independence condition holds. Then whenever $\forall i$, $X \approx_i Y$, it is the case that (a) $P_G(X) = P_G(Y) \Rightarrow P_i(X) = P_i(Y)$ and, hence, more generally that (b):

$$\frac{P_i(X)}{P_i(Y)} = \frac{P_G(X)}{P_G(Y)}$$

¹⁰This Lemma is a direct consequence of Bolker's Lemma 3.1 (1965, p.300)

Proof. Assume that $\forall i, X \approx_i Y$. Then by Pareto Indifference, $X \approx_G Y$. (a) Suppose that $P_G(X) = P_G(Y)$ and let X_i be such that $\forall(j \in G - \{i\}), X_i \approx_j X$, but $X_i \not\approx_i X$ (by the Independence condition such a proposition X_i exists). Then by Lemma 27 there exists mutually contradictory propositions X^*, Y^* and X_i^* , such that, $X^* \approx_i Y^*$, $X^* \not\approx_i X_i^*$, $X^* \approx_G Y^*$, $P_G(X^*) = P_G(Y^*)$ and $\forall(j \in G - \{i\}), X^* \approx_j Y^* \approx_j X_i^*$. So by the axiom of Averaging, $\forall(j \in G - \{i\}), X^* \vee X_i^* \approx_j Y^* \vee X_i^*$ and by Lemma 5, $X^* \vee X_i^* \approx_G Y^* \vee X_i^*$. So by Pareto Indifference, $X^* \vee X_i^* \approx_i Y^* \vee X_i^*$. But since $X^* \not\approx_i X_i^*$, it follows from Lemma 5 that $P_i(X^*) = P_i(Y^*)$. Hence $P_i(X) = P_i(Y)$. (b) Suppose now that $P_G(X) < P_G(Y)$. By Liapounoff's Theorem there exists a proposition W such that $\gamma(W) = \frac{P_G(X)}{P_G(Y)} \cdot \gamma(Y) + (1 - \frac{P_G(X)}{P_G(Y)}) \cdot \gamma(F)$. Then $P_G(W) = P_G(X)$, $V_G(W) = \frac{U_G(W)}{P_G(W)} = V_G(Y)$ and for all individuals i , $V_i(W) = V_i(Y)$. Hence $\forall i, X \approx_{i,G} W$ and by (a) above $P_i(W) = P_i(X)$. But $P_i(W) = \frac{P_G(X)}{P_G(Y)} \cdot P_i(Y)$. Hence $P_i(X)/P_i(Y) = P_G(X)/P_G(Y)$. ■

Corollary 29 *Assume Pareto Indifference and that the Independence condition holds. Then $\forall(i, j \in G)$, if $X \approx_{i,j} Y$, then:*

$$\frac{P_i(X)}{P_i(Y)} = \frac{P_j(X)}{P_j(Y)}$$

Theorem 30 *Assume Pareto Indifference and that the Weak Independence condition holds. If $\forall i, X \approx_i T$ then $P_i(X) = P_G(X)$.*

Proof. By the same argument as in Theorem 28, substituting T for Y throughout and noting that only Weak Independence is required, we prove that $\forall i, X \approx_i T$ then $P_i(X)/P_i(T) = P_G(X)/P_G(T)$. But $P_i(T) = P_G(T) = 1$. So $P_i(X) = P_G(X)$. ■

5.2 Broome's Utilitarian Theorem

We turn now to examination of Broome's Utilitarian theorem. Broome proves that if individuals' preferences are both strongly independent of one another and unbounded then there exists a representation of individual and group preferences that satisfies the Utilitarian condition. (Recall that under the assumption of unbounded preferences, the Weak and Strong Utilitarian Representation theses are equivalent). This is an immediate consequence of Theorem 25 and the following:

Theorem 31 (Broome) *Assume the Rationality and Strong Pareto axioms, that the \geq_i s are unbounded and that the Strong Independence condition holds. Then the Probability Identity condition holds.*

Proving this theorem is the central technical achievement of Broome's paper. Its conceptual significance is more difficult to assess. For a start neither the Strong Independence nor the Unboundedness condition are strictly necessary for

Probability Identity. Secondly, both assumptions are very strong, perhaps implausibly so. Broome himself concludes that the Utilitarian Theorem, "though it can be proven within the Jeffrey-Bolker theory, needs to be treated with caution" [6, p. 494].

Let us take a closer look at the assumptions of the theorem. We have already argued that the assumption that either preferences or goodness relations are unbounded lacks any convincing justification. But given that probability identity is a necessary condition for the truth of the Utilitarian Representation Thesis and that probability identity cannot be characterised in the Bolker-Jeffrey framework other than in the special case where preferences are unbounded, there is a sense in which the assumption is indispensable. It does, however, raise the question as to whether a natural generalisation of Broome's theorem holds; namely that, given the Rationality and Strong Pareto axioms, Strong Independence is a sufficient condition for the truth of the Weak Utilitarian Representation thesis. The short answer, as Broome notes, is that it does not. But there is a slightly longer answer, to the effect that a different generalisation of his theorem does hold true, which we investigate in the next section.

Let us turn now to the Strong Independence condition. Not only is this condition *not* necessary for probability identity, but is positively in tension with the interpretation of the \succeq_i s as preference relations. For it is frequently the case that the preferences of one individual will depend in a systematic way on those of others. Indeed dependencies of one kind or another mark the existence of social relationships such as those holding within a group of friends or a family. In such groups it is typical that one member's very strong preference for a certain state of affairs in itself gives the others reason to view it positively. Only a group that is 'a sack of potatoes' is at all likely to have preferences that come close to satisfying the Strong Independence condition - and even then it would be surprising if there were not some correlations.

It is more plausible that individual betterness relations are strongly independent. Although family members' preferences may be sensitive to the judgements of what is for the good of other members, it might still be the case that their goods are independent. Suppose, for instance, that the only thing that mattered for someone's good was the size of their allocation of some set of material goods. Then it might be possible to make small increases (or, more plausibly, decreases) to one person's allocation without affecting anyone else's. In which case both the Weak Independence and the Independence conditions would have some justification *qua* claims about the possibility of efficiency gains or losses in, as it were, the economy of goodness. But this scenario has less and less plausibility the larger the increase in the allocation to some individual. Realistically, *someone* has to pay for the increase - goods don't simply materialise. So even under the betterness interpretation, and especially in combination with the unboundedness assumption, Strong Independence looks highly implausible. In any case it is surely no part of the concept of 'good' that what is good for any individual can be so radically independent of what is good for others. Making this assumption robs the derivation of the Utilitarian condition from the Pareto axioms of much of its interest.

5.2.1 The Separation Condition

The reservations expressed above are not without their force, but we can go some way towards mollifying them by recognising that essentially what Broome's Utilitarian theorem requires is that the following condition holds.

Condition 32 (*Separation*) $\forall i, \exists X_i$ such that:

- (i) $\forall j, P_j(X_i) = P_G(X_i)$
- (ii) $X_i >_i T$, but $\forall j \neq i, X_i \approx_j T$

The Separation condition requires that, for each individual i , there exists a prospect of known or agreed probability, whose truth is a matter of indifference to all individuals except i . Such propositions play a rather similar role in our argument to that played by ethically neutral propositions of probability one-half in Ramsey's derivation of subjective probability ???. Indeed, the canonical example of an ethically neutral proposition of probability one-half - that of a fair coin landing heads - also provides a good example of a separating proposition. The sole difference is that we must suppose that the coin's landing heads brings some good to i , but to nobody else, so that it is ethically neutral for everyone except i .

The Separation condition is stronger than the Weak Independence condition, which it implies, though it is noteworthy that is derivable from Weak Independence plus Sub-group Pareto Indifference. More importantly, it is much easier to swallow than the Strong Independence condition and, unlike the latter, assumes very little about the structure of goodness. Despite this it is possible to substitute Separation for Strong Independence in Broome's Utilitarian Theorem, to give a result that is conceptually a good deal more satisfactory. (We omit the proof, which essentially follows the vital part of Broome's proof of his Theorem 3 [6, p. 498-500]).

Theorem 33 *Assume the Strong Pareto axioms. If the Separation condition holds for some representation of the group's preferences, then the Probability Identity condition holds for that representation.*

5.2.2 Unboundedness and the Possibility of Utilitarianism

Theorem 33 makes no use of the assumption of that preferences are unbounded, but without that assumption the Separation condition can only be partially expressed in terms of properties of the ordering relation on prospects. This does raise the question of whether, in the absence of the assumption that preferences are unbounded, it is possible to establish the validity of the Weak Utilitarianism Representation thesis. The natural strategy to follow here is to work with something weaker than the Separation condition but which can be expressed, using the results obtained in Section 5.1, in terms of rankings of prospects.

Definition 34 *A proposition X_i is said to **separate** i from group G iff:*

- (i) $X_i \not\approx_i T$
- (ii) $\forall (j, k \in G - i), X_i \approx_j T$ and $P_j(X_i) = P_k(X_i)$.

The question that now arises whether, in the absence of the assumption that preferences are unbounded, it is possible to establish the Weak Utilitarian Representation Thesis using only the assumption that for each individual i there exists a proposition that separates i from the others and without assuming that i 's probabilities agree with the others. The answer, however, appears to be negative. On the other hand, it is possible to establish the Weak thesis from the assumption that for every proposition in the ranking there exists a separating proposition ranked with it. (We omit the proof which is long and tedious. A version is to be found in [5].)

Theorem 35 *Assume the Rationality and Strong Pareto axioms and that for all individuals $i \in G$ and propositions $X \in \Omega$, there exists a proposition $X_i \approx_i X$ that separates i from G . Then the Weak Utilitarian Representation thesis holds.*

Theorem 35 does little to support the Weak Utilitarian Representation thesis: the conditions it postulates are just too strong. And we have, in any case, argued that the Weak Representation thesis is of considerably less conceptual interest than the Strong. What Theorem 35 does do, however, is show that case where preferences are unbounded is not completely discontinuous from the case where they are not. For the critical assumption - that for every proposition in the ranking there exists a separating proposition ranked with it - is derivable from the Strong Independence condition plus Sub-group Pareto Indifference. Given the strengthened Pareto condition, Broome's theorem is derivable as a special case of Theorem 35. This conclusion is broadly what one would expect, though perhaps not the degree to which it is sensitive to the precise assumptions invoked.

5.2.3 Conclusion

The implausibility of the assumptions of both Theorem 35 and Broome's Utilitarian theorem contrast with the relatively innocuous nature of the Separation condition. This strongly suggests that the strength of the assumptions required to prove a Utilitarian theorem within the Jeffrey-Bolker framework derive more from the peculiarities of the framework - in particular the non-unicity of probabilities - than from the Utilitarian Representation thesis itself. The latter emerges from our discussion in relatively strong shape. It has proven to be derivable from a set of quite plausible conditions and although one of them - the Separation condition - cannot be expressed solely in terms of properties of rankings of prospects within the Jeffrey-Bolker framework, we know that this can be done by strengthening the framework in the way discussed at the end of the last section (i.e. by introducing conditional prospects). But that is, in a sense, just a technical matter. I think we have done enough to be considerably less cautious about the Utilitarian theorem than Broome declares himself to be.

References

- [1] Bolker, E. (1966) "Bolker-Jeffrey Expected Utility Theory and Axiomatic Utilitarianism", *Review of Economic Studies* **57**, 477-502
- [2] Bradley, R. (1998) "A Representation Theorem for a Decision Theory with Conditionals", *Synthese* 116: 187-229, Kluwer Academic Press
- [3] Bradley, R. (2000) "Conditionals and the Logic of Decision", *Philosophy of Science* 67 (Proceedings): S18-S32, Philosophy of Science Association
- [4] Bradley, R. (2001) "Ramsey and the Measurement of Belief" in D. Corfield and J. Williamson (eds) *Foundations of Bayesianism*, 273-299, Kluwer Academic Press
- [5] Bradley, R. (forthcoming) "Broome's Axiomatic Utilitarianism", *CPNSS Discussion Papers*
- [6] Broome, J. (1990) "Bolker-Jeffrey Expected Utility Theory and Axiomatic Utilitarianism", *Review of Economic Studies* **57**, 477-502
- [7] Broome, J. (1991) *Weighing Goods*, Oxford, Basil Blackwell
- [8] Dreze, J and Rustichini, A. (forthcoming) "State-dependent Utility Theory" in S. Barbera, P. Hammond and C. Seidl (eds) *Handbook of Utility Theory*, Chapter 16, Kluwer Academic Press
- [9] Jeffrey, R. C. (1983) *The Logic of Decision*, 2nd ed, Chicago, University of Chicago Press
- [10] Joyce, J. (1999) *The Foundations of Causal Decision Theory*, Cambridge University Press
- [11] Hammond, P.J. (1981) "Ex ante and ex post welfare optimality under uncertainty" *Economica* **48**: 235-250
- [12] Harsanyi, J. (1955) "Cardinal Welfare, individualistic ethics, and interpersonal comparisons of utility", *Journal of Political Economy* **61**, 434-435
- [13] Mongin, P. (1995) "Consistent Bayesian Aggregation." *Journal of Economic Theory* **66** (2): 313-51
- [14] Mongin, P. (1997) "Spurious Unanimity and the Pareto Principle." *THEMA working papers*
- [15] Ramsey, F. P. (1926) 'Truth and Probability' in D. H. Mellor (ed) *Philosophical Papers*, Cambridge: Cambridge University Press, 1990.

- [16] Savage, L. J. (1972). *The foundations of statistics*, 2nd ed, New York, Dover.
- [17] Seidenfeld, T., Kadane, J.B., and Schervish (1989) “On the shared preferences of two Bayesian decision makers”, *Journal of Philosophy* **86**: 225-244