

“The Connection between Prudential Goodness and Moral Permissibility”, *Journal of Social Philosophy* 24 (1993): 105-28.

Peter Vallentyne

## 1. Introduction

In his book Weighing Goods John Broome (1991) gives an extremely elegant presentation, discussion, and qualified defense, of a mathematically formulated argument in favor of utilitarianism. In a nutshell, the argument says that if (1) for each individual, the ranking of alternatives in terms of prudential goodness satisfies certain standard coherence conditions from decision theory, (2) the moral ranking of alternatives satisfies these same coherence conditions, and (3) the moral ranking is positively responsive in a specified manner to the individual rankings, then, given a few additional assumptions, utilitarianism is the correct moral theory.

The basic idea of the theorem is not very new: it is a slight generalization of a theorem proved by John Harsanyi in the 1950s.<sup>1</sup> The power of the book comes from his interpretation of the theorem, and from his strikingly clear and insightful discussion of the various conditions.

The theorem establishes that if certain conditions hold, then utilitarianism is the correct moral theory. So, the key question is whether these conditions are plausible. I shall identify a number of assumptions that are doubtful.

## 2. Broome's Theorems

To introduce Broome's theorems, we need some terminology. For each action that an agent might perform there corresponds a prospect, which represents the different ways that things might go if the action is performed. If probabilities are assigned to the various possible outcomes (ways things

could go), then a prospect is a probability distribution over possible outcomes (e.g.,  $\langle p_1, x_1; p_2, x_2; \dots p_n, x_n \rangle$ , where the  $x_i$  are outcomes). If probabilities are not assigned, then a prospect identifies an outcome for each of the possible states of the world (e.g.,  $\langle s_1, x_1; s_2, x_2; \dots s_n, x_n \rangle$ ).

Broome assumes that for each person and for society there is a betterness relation that ranks actions in terms of how good they are for the individual (prudential goodness), or for society (moral goodness).<sup>2</sup> In addition, Broome makes the following assumption for both prudential and moral goodness (my terminology):

**Goodness Act Consequentialism:** The goodness of actions is based solely on features of their prospects.

Given this assumption one can, and Broome does, equate the goodness of actions with the goodness of their prospects.

Central to Broome's argument are assumptions that the betterness relations over prospects (and actions) have the following properties. A betterness relation is complete if and only for any two prospects,  $x$  and  $y$ , either  $x$  is at least as good as  $y$ , or vice-versa. A betterness relation is transitive if and only if when  $x$  is at least as good as  $y$ , and  $y$  is at least as good as  $z$ , then  $x$  is at least as good as  $z$ . A betterness relation is continuous roughly if and only if for any three prospects,  $x$ ,  $y$ , and  $z$ , such that  $x$  is at least as good as  $y$ , and  $y$  is at least as good as  $z$ , there is a probability,  $p$ , such that  $y$  is equally good as the compound prospect  $\langle p, x; 1-p, z \rangle$ .<sup>3</sup>

A betterness relation is strongly separable if and only if no matter how states are numbered,

for any  $i$ , and any outcomes,  $x_{i+1}-x_n$  and  $y_{i+1}-y_n$ , the relative ranking of prospects with the following form does not depend on which specific outcomes the  $O_1-O_i$  variables (over outcomes) stand for:

$\langle s_1, O_1; \dots; s_i, O_i; s_{i+1}, x_{i+1}; \dots s_n, x_n \rangle$

$\langle s_1, O_1; \dots; s_i, O_i; s_{i+1}, y_{i+1}; \dots s_n, y_n \rangle$

Prospects of this form differ only with respect to what happens under states  $s_{i+1}$  to  $s_n$ . They have the same, but unspecified, outcomes under  $s_1$  to  $s_i$ . Strong separability requires that the ranking of two options having this form be independent of the particular outcomes (within their respective ranges) that are assigned to the  $O_1-O_i$ .<sup>4</sup>

It is well known that any betterness relation that is complete, transitive, and continuous can be represented by a utility function. If in addition it is strongly separable, it is called coherent, and is representable by an expectational utility function, i.e., by one for which the utility of a prospect is the expected (probability weighted) utility of its components.

The only remaining condition needed to formulate Broome's main theorem is:

**The Principle of Personal Good:** (1) If two alternatives are prudentially equally good for each individual, then they are morally equally good. (2) If for every individual one alternative is prudentially at least as good as a second, and for at least one individual it is prudentially better, then the first alternative is morally better than the second.

This principle is just the familiar strong Pareto principle formulated in terms of betterness instead of preferences. It imposes a plausible connection between prudential goodness and moral goodness.

Broome's first theorem (Harsanyi's theorem with preferences replaced by goodness) is:

**Interpersonal Addition Theorem:** If (1) each person's prudential betterness relation is coherent, (2) moral betterness is coherent, and (3) the Principle of Personal Good is satisfied, then the moral betterness relation can be (ordinally) represented by an expectational utility function that is the sum of expectational utility functions (ordinally) representing the prudential betterness relations.

This does not yet give us total act utilitarianism (understood throughout as a principle of moral goodness), which is the following principle:

**Total Act Utilitarianism:** An action is morally at least as good as a second if and only if it produces at least as great a sum of prudential goodness as the second.

Utilitarianism presupposes that individuals have -- not merely betterness relations, which only specify the order of goodness, but also -- cardinal goodness functions, which specify intensities (amounts) of goodness (e.g., that the gain in goodness of x over y is twice that of the gain of w over z). Utilitarianism ranks prospects on the basis of the sum of the values of these goodness functions. The above theorem does not, however, guarantee that the individual expectational utility functions

that are summed represent the individuals' goodness functions cardinally (i.e., represent goodness intensities as opposed to mere order). To ensure this, we need to assume:

**Bernoulli's Hypothesis** (weak version): Any expectational utility function that (ordinally) represents prudential goodness of prospects also represents it cardinally.

According to this hypothesis, if prudential goodness is expectationally representable, then any such representation captures the intensity of goodness -- and not merely the order of goodness. One of the many insights of Broome's book is that expectational utility functions (ordinally) representing betterness do not necessary cardinally represent goodness, and so something like the above assumption is needed in the present context to get a standard sort of utilitarianism.

The Interpersonal Addition Theorem conjoined with Bernoulli's Hypothesis ensures that under the specified conditions moral betterness relation can be (ordinally) represented by an expectational utility function that is the sum of expectational utility functions cardinally representing prudential goodness. This still does not quite give us total act utilitarianism, for it does not ensure that the utility functions representing each person's good are on the same scale. Utilitarianism presupposes that interpersonal comparisons of good are possible, and requires that agents maximize the sum of prudential goodness when everyone's goodness is measured on the same scale. The Interpersonal Addition Theorem conjoined with Bernoulli's Hypothesis ensures that agents are required to maximize some sum of utilities cardinally representing prudential goodness, but it does not require that the relevant sum be of utilities on the same scale. It allows,

for example, that my utility might be measured on a hypersensitive scale (so that one unit of utility represents only a trivial difference in goodness), whereas your utility might be measured on a gross scale (so that one unit of utility represents a major difference in goodness). If this is so, then the sum of our utilities in effect gives much greater weight to my good than to yours (since a trivial difference in my good will be counted as equivalent to a major difference in yours). The above theorem does not rule out this possibility.

In effect, then, the Interpersonal Addition Theorem conjoined with Bernoulli's Hypothesis tells us that moral goodness is a weighted sum of utility functions cardinally representing, on the same scale, prudential goodness. (The weighting factors are simply whatever is required to convert all utility functions to the same scale.) In order to obtain total act utilitarianism, a further assumption is needed:

**Equal Weight** (Anonymity): When all utility functions measuring people's prudential goodness are on the same scale, moral goodness is unaffected by which people are associated with which utility function (everyone counts equally).<sup>5</sup>

Broome is not very clear in the book about the need for this additional assumption.<sup>6</sup> Originally he thought it unnecessary on the grounds that there was no difference between the amount of prudential goodness and how much it counts morally. Consequently, he denied the possibility that the same amount of goodness could count (be weighted) differently for different individuals. In personal correspondence, however, he has acknowledged that something like the

above assumption is necessary.

The Equal Weight assumption presupposes that there is a uniquely valid way of making interpersonal comparisons of goodness (to determine whether utility functions are on the same scale). So in order to ensure that the Equal Weight assumption is operative we need to add one final assumption:

**Interpersonal Comparisons of Goodness Units:** If a given utility function cardinally represents the prudential goodness of some individual, then for any second individual there is a unique, up to the zero points, utility function having the same unit that cardinally represents his/her prudential goodness.<sup>7</sup>

Given the above assumptions, it follows logically that utilitarianism is the correct theory of moral goodness. Below I shall question a number of these assumptions. First, however, I should mention an aspect of Broome's book that I will not discuss.

The Interpersonal Addition Theorem connects moral goodness with prudential goodness. The theorem remains valid if one reinterprets it by (1) replacing all references to prudential betterness by references to dated prudential betterness (that ranks prospects on the basis of what is good from the perspective of the specified time for that individual), and (2) by replacing all references to moral betterness by references to prudential betterness. The reinterpreted theorem (The Intertemporal Addition Theorem) establishes that under certain conditions prudential goodness must be the sum of dated prudential goodness. Unlike the Interpersonal Addition

Theorem, this result is intrapersonal, not interpersonal. Broome offers an ingenious defense of the conditions underlying the reinterpreted theorem by appealing to a Parfitian disuniting metaphysics of personhood. (Although, ultimately, he rejects this metaphysics, and so this defense is not one he endorses.)

The defense of the theorem under the intertemporal interpretation is an important and interesting part of the book -- especially in its demonstration that issues of separability (independence) over uncertainty, time, and people are intimately related. Nonetheless, due to space limitations, I shall focus solely on the interpersonal interpretation. Many of my criticisms, however, are equally applicable to the Intertemporal Addition Theorem.

### 3. The Nature of Goodness

All the assumptions of Broome's theorem are assumptions about prudential goodness, moral goodness, and the relation between the two. Thus, in order to assess the plausibility of these assumptions, we need to know something about the nature of goodness. Broome, however, does not specify any substantive account of the good. His theorem is valid for all conceptions of prudential and moral goodness. It is thus genuinely more general than Harsanyi's theorem, which was based on the preference satisfaction conception of prudential good.

How, then, are we to assess the plausibility of Broome's premisses? Broome does provide some guidance. For although he offers no substantive conceptions of prudential and moral goodness, he does tell us something about what he means by "goodness". He is not appealing to a notion of goodness that is "an externally determined objective", independent of permissibility, but



rather to one that is determined by the notion of permissibility (pp.13-17). Now, this is, I think, a non-standard way of understanding goodness. For goodness is often thought of as something that is independent of permissibility -- there being a substantive issue of whether there is a connection. In what follows, however, I will not address the issue of whether there is a relevant notion of goodness that is conceptually independent of permissibility. For I want to confront Broome on his own terrain in this matter, and so, like him, I shall understand the goodness of actions as intimately connected with permissibility.

But exactly how is goodness, so understood, connected with permissibility? Broome does not tell us. But in order to assess the plausibility of his assumptions about goodness, we need to know more about what he means by "goodness" and "betterness". I shall assume that weak betterness -- i.e., better-or-indifferentness (e.g., at-least-as-goodness) -- in his permissibility-connected sense, has the following connection with permissibility:

**Weak Betterness:** If one action is weakly better than a second then, in any choice situation in which both are feasible, if the second is permissible, then so is the first.

This assumption ensures that weak betterness is something that permissibility promotes: all and only those actions that are weakly better than some acceptable level are permissible. It allows, but does not require, that permissibility maximize weak betterness. (It would require maximization if it were a biconditional claim). Nor does it say anything about how goodness of actions is assessed. All it says is that, however the goodness of actions is assessed (based on their consequences, on the

agent's intentions, or whatever), if a given action is permissible, then so are all weakly better actions.<sup>8</sup>

This condition on weak betterness is fairly weak. It does not entail, as we shall see below, that weak betterness is complete or transitive. Of course, if permissibility is a maximizing notion (judging an action permissible iff its f-value is at least as great as that of any alternative, for some fixed function f), then the value maximized represents weak betterness. And in that case, weak betterness is complete and transitive. But that is true only on certain conceptions of permissibility. It is not a necessary part of the notion of weak betterness.<sup>9</sup>

In personal correspondence Broome has expressed serious reservations about the above assumption as applied to non-teleological theories. For he thinks that for such theories weak betterness might be connected to permissibility in other ways. (I, on the other hand, remain skeptical that there is an interesting notion of permissibility-determined weak betterness that violates the Weak Betterness condition.) For teleological theories, however, Broome is willing to grant this assumption. So, no great harm will be done by assuming it in the context of a defense of utilitarianism.

The above assumption will be made both for prudential weak betterness (as it is connected to prudential permissibility) and for moral weak betterness (as it is connected to moral permissibility).

With these assumptions we know something about both prudential and moral goodness, and we can start to assess the various properties that Broome asserts they have.

#### 4. Completeness and Transitivity

Broome assumes, for his theorem, that moral and prudential weak betterness are complete in the sense that for any two actions at least one of them is weakly better than the other. Conceptually, it is possible that weak betterness is not only incomplete (there being at least two actions that are incomparable) but radically incomplete (such that no actions are comparable). Suppose, for example, that there are just four actions,  $x_1$ ,  $x_2$ ,  $x_3$ , and  $x_4$ . Then it is conceptually possible that when the choice is from just two actions, only the action with the lowest subscript is permissible; when the choice is from just three actions, then all are judged permissible except the action with the lowest subscript; and when the choice is from all four actions, then only  $x_4$  is permissible. If this is the case, then nothing is weakly better than anything else. (For example,  $x_2$  is not weakly better than  $x_1$ , because  $x_1$ , but not  $x_2$ , is a permissible choice from  $\{x_1, x_2\}$ . And  $x_1$  is not weakly better than  $x_2$ , because  $x_2$ , but not  $x_1$ , is a permissible choice from  $\{x_1, x_2, x_3\}$ .)

Broome acknowledges that he has genuine misgivings about assuming that prudential and moral weak betterness are each complete. And I share his misgivings -- in particular about the assumption that moral weak betterness is complete. For recall Broome's Principle of Personal Good: (1) if for each person two actions are prudentially indifferent, then they are also morally indifferent; and (2) if (a) for each person one action is prudentially weakly better than a second, and (b) for some person it is better than the second, then the first action is morally better than the second.<sup>10</sup> Not only do I accept this principle, I also hold that it exhausts the content of moral weak betterness. I hold, that is, that weak moral betterness (i.e., better-or-indifferent-ness) is instantiated when -- and only when -- one of these two conditions is instantiated.<sup>11</sup> It is well-known, however,

that this relation (Pareto weak domination) is incomplete. If there are just two people, and x is better than y for one person and worse for the other, then neither of the two conditions is satisfied. Neither action Pareto weakly dominates the other, and so, on my view, neither is morally weakly better than the other.<sup>12</sup>

Of course, those who hold that objectively valid interpersonal comparisons of prudential goodness are both possible and relevant for moral permissibility will reject the above conception of moral weak betterness. They may accept the Principle of Personal Good, but deny that it exhausts the content of moral betterness (since interpersonal comparisons provide further content). So, the issue is controversial. My point here is that people who hold that objective interpersonal comparisons of prudential weak betterness are impossible or irrelevant should be very skeptical of the assumption of completeness.

Broome also assumes that prudential and moral weak betterness are transitive. Broome (pp. 11, 136) asserts that betterness is transitive as a matter of logic. But things are not quite that simple. We may grant that, as he claims, the relation "   is at least as F as   " (for any specified F) is, as a matter of conceptual necessity, transitive. But unlike Broome, we are not equating weak betterness with this relation. For we are starting with a weaker notion of betterness, and examining what properties it has. (The more that is defined into the notion of betterness, the more doubtful it is that such a relation is complete, or is even instantiated.) So, we must examine whether weak betterness is transitive. If it's not, then Broome's "   is at least as F as   " doesn't even exist (in the sense of being instantiated).

Unlike the issue of completeness our semantic assumption about betterness does not suffice

to establish that it is conceptually possible for it to fail to be transitive. But the following case will suffice to show that the assumptions made so far do not guarantee transitivity. For the following case is compatible with our assumptions: (1)  $x$  is weakly better than  $y$  (and so, for all sets containing both  $x$  and  $y$ ,  $x$  is permissible whenever  $y$  is); (2)  $y$  is weakly better than  $z$  (and so, for all sets containing both  $y$  and  $z$ ,  $y$  is permissible whenever  $z$  is); but (3)  $z$  alone is a permissible choice from  $\{x,z\}$ , and so it is not the case that  $x$  is weakly better than  $z$ . This example does not establish the conceptual possibility of transitivity failing, because weak betterness may have more formal properties than we have assumed, and these further properties may rule out the co-possibility of the above three cases. But if, for example, the notion of weak betterness is identical with -- as opposed to merely entailing -- the notion of a first action always being permissible whenever a second action is and both are feasible, then the above is a clear case in which transitivity fails.<sup>13</sup>

The above example suggests weak betterness is not necessarily transitive. Nonetheless, I agree with Broome, that as a matter of fact prudential and moral weak betterness are transitive on the correct account of prudential and moral permissibility. Prudential weak betterness, I will simply grant without argument, is transitive. Moral weak betterness is transitive, since moral weak betterness is, I hold, coextensive with the Pareto weak domination relation, and the latter is transitive. So we may readily grant Broome's assumptions of transitivity.<sup>14</sup>

In summary, Broome's theorem rests on the plausible assumption that prudential and moral weak betterness are transitive. It also rests on the assumption that these relations are complete. He is quite right to be worried about the completeness assumption -- especially with respect to moral goodness.

## 5. Assumptions about Betterness over Prospects: Continuity

So far we have been concerned exclusively with weak betterness over actions. The assumptions of Broome's theorem, however, are formulated in terms of weak betterness over prospects (how things might turn out if the action is performed). For the moment let us assume that the only relevant prospects are those that are prospects of some action, i.e., that there are no free-floating prospects. (The significance and plausibility of this assumption will be addressed below.) Given this assumption, and Broome's assumption that the betterness of actions is based solely on their prospects (Goodness Act Consequentialism), weak betterness over actions fully determines weak betterness over prospects. Prospects are ranked in exactly the same order as the actions producing them.<sup>15</sup> And thus, the ranking of prospects is complete and transitive if and only if the ranking of actions is.

So with respect to weak betterness over prospects, we may grant his assumptions that prudential weak betterness is complete and transitive, and agree that moral weak betterness is transitive. But the assumption that moral weak betterness is complete should be rejected.

The remaining conditions of Broome's theorem -- continuity, strong separability, Bernoulli's Hypothesis, and Rectangular Field -- can be formulated only with explicit reference to prospects. Let us start by considering continuity.

Broome assumes that prudential and moral weak betterness are each continuous. A betterness relation is continuous, recall, if and only if for any three prospects,  $x$ ,  $y$ , and  $z$ , such that  $x$  is at least as good as  $y$ , and  $y$  is at least as good as  $z$ , there is a probability,  $p$ , such that  $y$  is equally

good as the compound prospect  $\langle p, x; 1-p, z \rangle$ . The idea is roughly that "small" changes in the probabilities of prospects, produces "small" changes in their goodness. This may seem like an innocuous assumption, but it requires that if -- as is surely the case -- increasing one's wealth by one dollar over the status quo (SQ+1) is better than the status quo (SQ), and that is better than a slow painful death (SPD), then for some probability,  $p$ , the status quo is indifferent to  $\langle p, \text{SQ}+1; 1-p, \text{SPD} \rangle$ . It entails, that is, that for some  $p$ , the goodness of a  $p$ -chance of gaining a dollar is equal to the goodness  $1-p$  chance of a slow painful death instead of the status quo.

The assumption of continuity rules out the possibility that goodness is some sort of lexical ordering -- i.e., with some sorts of considerations (e.g., life and death considerations) always taking priority over others (e.g., financial). I see no reason to think prudential, or moral weak betterness, rules out this possibility. Indeed, I think that prudential weak betterness for a given individual is sensitive to the sorts of tradeoffs that the individual is willing, upon suitable reflection, to make. If a person is unwilling to make the above sorts of tradeoffs, then his/her prudential weak betterness relation is not continuous. And since, I hold, moral weak betterness is determined by prudential weak betterness, it too may fail to be continuous.

## 6. Strong Separability

There is no compelling reason to suppose that prudential and moral weak betterness must be continuous. For the sake of argument, suppose that they are. If they are also complete and transitive, then they are representable by a utility function. Nothing guarantees, however, that such a utility function is expectational (ranks on the basis of expected utility). In order to ensure that

weak betterness is representable by an expectational utility function, Broome makes the standard assumption that it is strongly separable (satisfies the sure thing principle).

A betterness relation is strongly separable, recall, if and only if no matter how states are numbered, for any  $i$ , and any outcomes,  $x_{i+1}$ - $x_n$  and  $y_{i+1}$ - $y_n$ , the relative ranking of prospects with the following form does not depend on which specific outcomes the  $O_1$ - $O_i$  variables (over outcomes) stand for:

$$\langle s_1, O_1; \dots; s_i, O_i; s_{i+1}, x_{i+1}; \dots; s_n, x_n \rangle$$

$$\langle s_1, O_1; \dots; s_i, O_i; s_{i+1}, y_{i+1}; \dots; s_n, y_n \rangle$$

Prospects of this form differ only with respect to what happens under states  $s_{i+1}$  to  $s_n$ . They have the same, but unspecified, outcomes under  $s_1$  to  $s_i$ . Strong separability requires that the ranking of two options having this form be independent of the particular outcomes (within their respective ranges) that are assigned to the  $O_1$ - $O_i$ .

Strong separability requires that weak betterness not be holistic over possible outcomes. It requires that the impact on betterness of the outcomes under any given set of states be independent of what the outcomes are under the remaining states. But it's not clear why we should think this to be so. Prospects have components (state-relative outcomes). Why can't weak betterness be sensitive to how these components fit together?

To see the force of the assumption of strong separability, consider (as Broome does) the following well-known example, introduced by Peter Diamond, purporting to show that moral weak



betterness can violate strong separability. Suppose that there two people each of whom will continue living if and only she obtains a certain medical operation, but resources permit the operation for only one of them. For simplicity, suppose that the effects on other people are constant, and can be ignored. Let "L-D" designate the outcome that the first person lives and the second person dies, and similarly for "L-L", "D-L", and "D-D". Suppose also that a coin will be flipped with a 50% probability that it will come up heads (H) and a 50% chance that it will come up tails (T).

Now consider two prospects having the following forms:

x: <H,O;T,D-L>

y: <H,O;T,L-D>

where O is one of D-L, D-D, L-L, L-D.

Is x morally weakly better than y? If tails comes up, then x would have an outcome of the first person dying and the second person living (D-L), whereas y would have the "opposite" outcome (i.e., L-D: first person lives, second person dies). If heads comes up from the coin flip, then x and y would have the same outcome as each other (O), but we don't know what it is. Now, Diamond argues that x and y cannot be morally ranked until the outcome under Heads (O) is specified. For if that outcome is L-D, then x is better than y, since x would give each person an equal chance of surviving; whereas y would guarantee the first person's survival and the second person's death. On

the other hand, if the outcome under Heads is D-L, then the ranking would be reversed: y would be better than x. For in that case y, but not x, would give each person an equal chance of survival.

If, as just suggested, the moral ranking of x and y (which differ only with respect to what happens under Tails) depends on what happens under Heads, then moral weak betterness violates strong separability. For strong separability requires that if two prospects have the same outcome under Heads, and differ only with respect to the outcomes under Tails, then their ranking in no way depends on what the outcome under Heads is. If x is better than y when the outcome under Heads is L-D (where only x would give them an equal chance of survival), then strong separability requires that x is also better than y when the outcome under Heads is D-L (where only y would give them an equal chance of survival).

So this example purports to show that moral weak betterness can violate strong separability. And a similar argument (with a slightly modified example) can be run for prudential weak betterness. Broome, however, claims that these examples are unsuccessful. For he claims to have both a general argument for why betterness (prudential or moral) satisfies strong separability, and a corresponding diagnosis of why the above example does not support the view that weak betterness violates strong separability.

There are two steps in the general argument. One connects goodness with reasons:

**Reasons Thesis:** A prospect A is weakly better than prospect B if and only if the goodness-directed reasons for preferring A to B are at least as strong as the goodness-directed reasons for preferring B to A. (pp. 95-96, 137-39).

We may grant this step as long as the nature of reasons is left open. In particular, it should not be assumed that reasons are complete (i.e., that for any A and B, either the reasons for preferring A to B are at least as strong as those for preferring B to A, or vice-versa).

Broome makes, however, the following further assumption about reasons:

**Outcomism:** All goodness-directed reasons for preferring one prospect to another emanate from features specified in their outcomes. (p. 138)<sup>16</sup>

This assumption is simply an assumption that reasons are strongly separable! For it says that reasons are always based solely on what the nature of the possible outcomes in isolation from each other, and never on the global features of how the different possible outcomes fit together as a package with their corresponding states. Anyone who denies that weak betterness satisfies strong separability will also deny that reasons satisfy Outcomism. So, Outcomism is going to be little help in establishing that weak betterness is strongly separable.

To see exactly how Outcomism works and why it is problematic, consider how Broome uses it to argue that Diamond's life and death case does not show the moral weak betterness violates strong separability. There it was suggested that if moral betterness is sensitive to equality considerations, then it must violate strong separability. Broome denies this: moral betterness, he claims, can be sensitive to equality considerations without violating strong separability.

The claim, recall, was that (because of equal chances of survival) <H,D-L;T,L-D> was

better than  $\langle H, D-L; T, D-L \rangle$ , but that  $\langle H, L-D; T, D-L \rangle$  was better than  $\langle H, L-D; T, L-D \rangle$ . This appears to violate strong separability, but Broome says that we have been too hasty. For, according to the Reasons Thesis if one prospect is better than a second, then the good-directed reasons for the former must be stronger than those for the latter; and by Outcomism the reason-giving features must be features of the outcomes (i.e., the D-L, D-D, etc.). So if  $\langle H, D-L; T, L-D \rangle$  is better than  $\langle H, D-L; T, D-L \rangle$ , it must be because the reasons for L-D under T are stronger than the reasons for D-L under T (since the reasons for D-L under H apply to both). And if that is the case, then  $\langle H, L-D; T, D-L \rangle$  must be worse -- not better -- than  $\langle H, L-D; T, L-D \rangle$  (for the reasons are the same here). So, if all the relevant facts about the case have been given, then the purported violation of strong separability cannot, he claims, take place.

Broome's argument fails, however, because it relies on Outcomism, the doubtful assumption that reasons are strongly separable.

Now, Broome allows that the ranking of prospects may be sensitive to considerations such as equality of chances of survival (as in the above example); but he insists that, if it is, then some relevant information is missing above in the outcomes. If, for example, equality of chances of survival is relevant for the moral ranking then, he claims, facts about such chances need to be explicitly added to the outcomes. If we let "EC" stand for "everyone had equal chance of survival", then we would have to redescribe the four possible outcomes as follows.

$x_1: \langle H, \sim EC \ \& \ D-L; T, \sim EC \ \& \ D-L \rangle$

$y_1: \langle H, EC \ \& \ D-L; T, EC \ \& \ L-D \rangle$

$x_2$ : <H, EC & L-D;T, EC & D-L>

$y_2$ : <H,~EC & L-D;T,~EC & L-D>

If this is how the case is described, Broome rightly points out, there is no violation of strong separability. For strong separability has bite only when there is a given state (here: H or T) such that two prospects have the same outcome in that state. But once facts about equality of chances of survival are factored in, no two of the above prospects have the same outcome under any of the states. So there is no violation of separability.

This specific diagnosis relies on the dubious assumption of Outcomism. For it assumes that relevant facts must be in the outcome components of prospects, and not in their specification of possibilities, or probabilities, over the outcomes, nor in their specification of the states of nature. And again, I see now reason why this must be so. As long as all relevant facts are specified by the prospect, there is no reason to suppose that this specification must be internal to individual outcomes -- as opposed to in the specification of the states of nature (i.e., the  $s_i$ ) or of the possibilities, or probabilities, over outcomes. For example, in a prospect of the form  $\langle s_1, o_1; s_2, o_2; \dots; s_n, o_n \rangle$ , then some of the relevant facts could be in the  $s_1, \dots, s_n$  as they attach to the outcomes, and in the global distribution of the outcomes. In order for prospects to be assessed all the relevant information must be specified in the prospects, but it need not be in their outcomes.<sup>17</sup>

The issue here concerns how outcomes are individuated for the purposes of prudential and moral assessment. I have just argued that outcomes need not be individuated so that all relevant

reasons are in the outcomes (as required by Broome's Outcomism). And I have suggested that if outcomes are not so individuated, then prudential and moral weak betterness may violate strong separability by being sensitive to how the possible outcomes and their probabilities, or conditioning states, fit together.<sup>18</sup>

Outcomes need not be individuated in accordance with Broome's Outcomism, but they can be. And so Broome can legitimately to stipulate how outcomes are to be individuated for the purposes of his theorem. Note, however, that, although stipulating that Outcomism holds does guarantee that strong separability is satisfied, it does this by rendering strong separability vacuous (i.e., true, but only because its antecedent clause is never satisfied). For according to Outcomism, whenever two prospects appear to violate strong separability, the outcomes of these prospects have not been fully specified. Because Outcomism requires that all reasons for ranking to be in the outcomes, it makes it impossible for strong separability to fail. If Outcomism holds, strong separability is empty (and vacuously true).<sup>19</sup>

So, Broome has not provided a sound argument for the claim that Outcomism must be accepted. If we grant Outcomism as a stipulation about how outcomes are to be individuated for the purposes of his theorem, then prudential and moral weak betterness are each transitive and strongly separable. Their continuity, and the completeness of moral weak betterness, remains doubtful, but if we grant these features, then each weak betterness relation is representable by an expectational utility function (i.e., one for which prospects are ranked on the basis of their expected utility). But when outcomes are individuated in accordance with Outcomism, the existence of an expectational representation (in the context of the other assumptions) has no substantive

significance. For Outcomism requires that outcomes be individuated as finely as is necessary to ensure (by leaving enough leeway to utility functions in how they assign utility to outcomes) the existence of an expectational representation.<sup>20</sup>

## 7. The Principle of Personal Good and Bernoulli's Hypothesis

A weak betterness relation is coherent just in case it is complete, transitive, continuous, and strongly separable. Coherent weak betterness relations are representable by expectational utility functions. If prudential and moral betterness are coherent, then the only remaining condition needed to formulate Broome's main theorem is:

**The Principle of Personal Good:** (1) If two alternatives are prudentially equally good for each individual, then they are morally equally good. (2) If for every individual one alternative is prudentially at least as good as a second, and for at least one individual it is prudentially better, then the first alternative is morally better than the second.

As I have indicated already a few times, this is an extremely plausible principle. Of course, those who hold that moral goodness depends in part on features other than prudential goodness may well reject both parts of the principle. And those who hold that moral goodness is determined by the pattern (or distribution) of prudential goodness (such as strict egalitarians who hold that strict equality is better than inequality, even if everyone is worse off under strict equality) may well reject the second part of the principle.<sup>21</sup> Broome, however, effectively argues against those views, and I

join him in accepting this assumption.

Broome's main theorem is:

**Interpersonal Addition Theorem:** If (1) each person's prudential weak betterness relation is coherent, (2) moral weak betterness is coherent, and (3) the Principle of Personal Good is satisfied, then the moral weak betterness relation can be (ordinally) represented by an expectational utility function that is the sum of expectational utility functions (ordinally) representing the prudential weak betterness relations.

As Broome makes wonderfully clear, this theorem does not guarantee that moral weak betterness ranks alternatives on the basis of the total amount of prudential good. All it establishes is that moral weak betterness ranks alternatives on the basis of the sum of some expectational utility functions that (ordinally) represent prudential weak betterness. There is no guarantee that these expectational utility functions cardinally represent the individual prudential goodness functions -- there is no guarantee, that is, they these utility functions represent the intensity of prudential weak betterness (e.g., that the gain in goodness of x over y is twice that of the gain of w over z).

In order to derive total utilitarianism from the above theorem a further assumption is needed. For total utilitarianism ranks prospects on the basis of the sum of cardinal prudential goodness. Broome therefore makes the following claim:

**Bernoulli's Hypothesis** (weak version): Any expectational utility function that (ordinally)



represents prudential goodness of prospects also represents it cardinally.<sup>22</sup>

According to this hypothesis, if prudential goodness is expectationally representable, then any such representation captures the intensity of goodness -- and not merely the order of goodness.

Broome offers a tentative defense of Bernoulli's Hypothesis by assuming that prudential weak betterness has cardinal features of intensity, and then arguing that these cardinal features are expectationally based on the goodness of outcomes. Both claims are questionable.

Is prudential weak betterness cardinal? Remember that prudential weak betterness, as we are understanding it, is something at least partly determined (since permissibility promotes betterness) by the nature of prudential permissibility. And so the question is whether prudential permissibility requires prudential weak betterness to be cardinal. And the answer, I believe, is that it is far from clear that it does. Prudential permissibility might simply require that an action's prospect be maximally (or sufficiently) good, and that is a purely ordinal matter (it in no way depends on intensities). Prudential permissibility might, that is, might be sensitive to configurations in prospects of probabilities/possibilities of outcomes without in any way attaching any cardinal significance to the ranking of prospects.

Assuming prudential betterness is cardinal, is the goodness of a prospect based on the expected goodness of its outcomes? Broome argues for the affirmative. He claims that there is no difference between how much the goodness of an outcome (under a given state of nature) counts in the determination of the goodness of a prospect of which it is part, and the amount of that goodness. He claims, that is, that other than differences reflecting different probabilities of realization, the

goodness of all outcomes must receive the same weight in determining the goodness of a prospect. If this claim is granted, then Broome has indeed established that if weak betterness is cardinal, then it is expectationally so.

But why should we grant this claim? It precludes, for example, the possibility of there being decreasing marginal weights -- with the difference between 10 and 20 units of goodness of the outcomes under a given state of nature counting for less (for the determination of the goodness of prospects) than the difference between 110 and 120 units of goodness of the outcomes under a different state of nature. More generally, it precludes the possibility of prudence allowing prospects to be ranked on the basis of the worst possible outcome, the best possible outcome, or in various other ways that do not require betterness to be risk neutral (as expectational cardinality does). I see no reason to think that prudence precludes such possibilities.

Now this does not establish that prudence doesn't require all quantities of goodness of outcomes to count the same. It only establishes that this is not obviously so. It all depends on the nature of prudence. Broome is quite aware of this, but in the book, at least, he remains skeptical of the possibility of a quantitative measure of goodness that does not equate quantity with how much it counts.

Without Bernoulli's Hypothesis, Broome's Interpersonal Addition theorem does not yield total act utilitarianism. The implausibility of the hypothesis does not, however, negate the importance of Broome's work. For the Interpersonal Addition theorem is an interesting theorem in its own right. Furthermore, Broome's identification of the need for Bernoulli's Hypothesis is an advance in itself. This assumption has almost always been invisible in the past.<sup>23</sup>

## 8. Rectangular Field Assumption

Let us return to the Interpersonal Addition Theorem. This theorem rests on the following assumption, which I have so far left implicit.

**Rectangular Field of Outcomes:** For all outcomes  $o_1$ - $o_n$ , if  $o_1$ - $o_n$  are possible outcomes under states  $s_1$ - $s_n$  respectively, then  $\langle s_1, o_1; \dots; s_n, o_n \rangle$  is a prospect.

This assumption says that any arbitrary combination of possible outcomes, one for each state, yields a prospect. This requires two things: (1) Outcomes under each state are conceptually compatible with the outcomes under any other state. (If they weren't conceptually independent, then some combinations would be impossible.) (2) Prospects cover all the conceptually possible combinations of outcomes.

Consider the second requirement first. This requires that there be enough prospects to cover all the possibilities. One possibility is simply to stipulate that the domain of prospects consists of all conceptually possible combinations of outcomes under the states. Doing this, however, has the result that some prospects may be "free-floating" in the sense that there is no action that leads to them. We have been assuming that there are no such free-floating prospects, and this assumption has enabled us to move from the ranking of actions to the ranking of prospects (since each prospect can be ranked in accordance with its associated action). If prospects have no necessary connection with actions, however, then the ranking of some prospects may be underdetermined by the ranking

of actions. As a result, the completeness and transitivity of the ranking of actions would not guarantee the completeness and transitivity of the ranking of prospects.

Although formally it is open to Broome to simply stipulate that prospects cover all conceptually possible combinations of outcomes, it's unlikely that he would want to. For doing so breaks the intimate connection between the goodness of actions and that of prospects, and I suspect that he would not want to break this connection. More promisingly, I think, he could assert (as Savage does) that actions are individuated solely in terms of what outcomes they produce under each state of nature<sup>24</sup>, and that the ranking of actions ranks all conceptually possible actions. For if that is so, then the ranking of actions ranks all conceptually possible combinations of outcomes under states. And so every conceptually possible combination of outcomes corresponds to the prospect of some action. Because this approach seems plausible, Broome should have little problem with the second part of the Rectangular Field assumption.<sup>25</sup>

More problematic, however, is the first requirement above of the Rectangular Field assumption noted above. This requires that outcomes under each state be conceptually compatible with the outcomes under any other state. If they are not conceptually compatible, then some combinations of outcomes, one for each state, will be conceptually impossible, and so there will be no corresponding prospect. Broome acknowledges very clearly that his approach to the individuation of outcomes (requiring that all relevant facts be in the outcomes) ensures that this requirement is violated. For his approach requires that modal facts (facts about what happens in alternative outcomes), if they are relevant, be built into each outcome, and clearly not just any combination of modal facts are compatible. For example, in the Diamond example above where

only one person can be given the necessary operation to survive, EC&D-L under Heads (each person has equal chance of survival and the first dies and the second lives) is incompatible, for a given prospect, with  $\sim$ EC&D-L under Tails (the chances of survival are unequal and the first dies and the second lives). For if the chances of survival are equal (a modal fact) under Heads, then they must be so under Tails as well.<sup>26</sup> If such modal facts are built into outcomes, then the Rectangular Field assumption is sure to be false.

Broome states this very clearly in his book. He goes on to express optimism that the Rectangular Field assumption may not be necessary for his proof. His hope is that a weaker, and more plausible, assumption will suffice. In correspondence, however, he has stated that he is now less optimistic about being able to do without Rectangular Field than he was initially. Given his way of individuating outcomes, the assumption is demonstrably false, and so until his theorem can be proven without this assumption, the significance of the theorem is diminished.

## 9. Conclusion

Throughout we have been assuming that one action is weakly better than a second only if, in any choice situation in which both are feasible, if the second action is permissible, then so is the first (Weak Betterness). This assumption was made because Broome made clear that he was concerned with a notion of betterness that was intimately connected with permissibility, but he did not tell us how they are connected. This assumption is a fairly weak and natural way of connecting the two. Of course, if Broome rejects it, then some of my claims may miss their target (such as the incompleteness of moral impermissibility).

We have also been assuming, following Broome, that the relative weak betterness of actions is based solely on features of their prospects (Act Goodness Consequentialism). This may sound like a controversial assumption, but it's not. For it does not say that the goodness of actions is based on some independent notion of goodness of prospects. It only says that the goodness of actions is based solely on what their prospects are like. And since the prospects may include anything that is relevant for assessing the goodness and permissibility of actions (including the performance of actions, agents's intentions, etc.), no relevant fact is ruled out in principle. The assumption really just ensures that the prospects include whatever is relevant.

With respect to prudential weak betterness, I have granted that it is complete and transitive, but I have denied that it must be continuous (e.g., not involve any lexical orderings), be weakly separable (be insensitive to how the parts of prospect fit together) in any substantive sense, satisfy Bernoulli's Hypothesis (be expectationally cardinal), allow for objective interpersonal comparisons, or satisfy the Rectangular Field assumption (for the modally-packed individuation of outcomes defended by Broome). Even if we grant all these assumptions, however, the implications of Broome's theorem would depend on some problematic assumptions about moral weak betterness.

Consider first the moral assumptions that I have endorsed. Moral weak betterness is transitive and determined by, and positively sensitive to, prudential goodness (Principle of Personal Good). Furthermore, we may grant that when utilities are expressed on the same scale, then each person's utility gets an equal weight in determining moral weak betterness (Equal Weight).<sup>27</sup> Finally, if prudential weak betterness is weakly separable, we may even grant that moral weak betterness is as well.<sup>28</sup>

Two assumptions, however, should not be granted. These are completeness and continuity. Moral weak betterness is, I hold, exhausted by the Principle of Personal Good (moral weak betterness just is Pareto weak domination with respect to prudential weak betterness). It is thus incomplete (e.g., where x is better than y for one person, but worse for a second person). And because it is incomplete, it also fails to be continuous (even if prudential weak betterness is).<sup>29</sup>

The seemingly innocuous assumptions of completeness and continuity of moral goodness are in fact rather powerful. For in the context of the other assumptions of Broome's theory, they entail that moral goodness is based on interpersonal comparisons of prudential goodness. Their rejection opens the door to a theory of moral weak betterness and permissibility that is significantly different than utilitarianism.<sup>30</sup> Although I am critical of a number of Broome's assumptions (many of which he himself identifies as problematic), I stand in admiration of the beauty and insightfulness of the book. Broome has opened up significant new approaches to examining moral goodness, the individuation of actions and their outcomes, strong separability, and Bernoulli's Hypothesis. That alone makes it an important book.<sup>31</sup>

## Notes

1. See, for example, Harsanyi (1977).
2. Throughout I use "moral goodness" to designate what Broome calls "general goodness". And I use "prudential goodness" to designate what Broome calls "personal goodness" or "individual goodness". (Thus, prudential goodness is used in its broad sense for which it is not necessarily limited to prudential concerns with respect to future events).
3. Broome leaves implicit the need for a continuity assumption, but he needs some such assumption, and the one given in the text is the one that makes sense.
4. Expressed with reference to probabilities the assumption of strong separability is: A betterness relation is independent if and only if for any probability,  $p$ , and any prospects,  $x$  and  $y$ , the relative ranking of  $\langle p, P; 1-p, x \rangle$  and  $\langle p, P; 1-p, y \rangle$  does not depend on what particular prospect is assigned to  $P$ .
5. Actually, because utilitarianism is sensitive only to the unit of the scale, and not to the origin (zero), scales here need only have the same unit.
6. Harsanyi (1977), p.69, however, is quite explicit about the need for the assumption.
7. In the book, Broome (p.220) holds that this assumption is redundant in that it is entailed by the other assumptions. He claims that if the Principle of Personal Good is complete and the Principle of Personal Good holds (so that moral goodness is determined solely by individual goodness), then interpersonal comparisons of utility must be possible. For without them moral goodness could not be complete. The force of Broome's reasoning depended, however, on the assumption that the amount of individual goodness is identical to how much it counts morally. Now that he has dropped this assumption, he shows that the possibility of interpersonal comparisons does not follow from his other assumptions. An analogy will make this obvious. Consider two "utility" functions from two "individual" perspectives: one cardinally represents the temperature of objects in Celsius and the other in Fahrenheit.



their weight in kilograms. We could also have a "general" utility function that is the sum of these two utility functions. This "general" utility function is complete and satisfies the counterpart of the Principle of Personal Good. Nonetheless, there is no basis for determining whether the general utility function is on the same scale as the weight utility function. Of course, the general utility function treats 1 unit of kilograms as equivalent to 1 (degree) of Celsius, but that hardly establishes that kilograms and Celsius are the same scale.

8. Note that the defined notion of weak betterness is context-free, that is, it is not relativized to a choice situation (with a given set of actions). Thus, completeness, transitivity, etc. are stronger conditions than they would be if weak betterness were context-relative. If weak betterness is context-relative, or if actions are individuated in a manner that they are part of only one feasible set, then completeness is trivial. For discussion of context-relative orderings see McClennen (1990). For discussion of how theories of permissibility promote complete and transitive context-free orderings, see Sen (1977b), Vallentyne (1988d), and Oddie & Milne (1993).

9. Thus, we are not assuming that weak betterness is identical with the notion of "being at least as good as". For, as Broome notes, "being at least as good as" is not necessarily transitive (by the logic of "is at least as F as", for any F). But the transitivity of weak betterness needs to be established.

10. Broome's official formulation of the Principle of Personal Good uses the notion of "being at least as good as". For reasons (namely to avoid building transitivity into the definition of betterness), I here formulate it in terms of weak betterness. Through the following standard definitions: x is indifferent to y if and only if x is weakly better than y and y is weakly better than x; x is better than y if x is weakly better than y but y is not weakly better than x.

11. My conception of moral weak betterness, as coextensive with Pareto weak domination, does not rule out the possibility that rights considerations are also relevant for determining what is permissible. Indeed, as I indicated below, I hold that rights considerations are relevant in a manner that sometimes allows the violation of rights).

12. The conception of moral weak betterness as Pareto weak domination involves a conception of permissibility that violates both Condition Alpha and Condition Beta from social choice theory. Alpha says that if a given action is an impermissible choice from a given set, then it is also impermissible from all subsets containing it. Beta says that if two actions are each permissible choices from a given set, then they are either both permissible or both impermissible, from all larger sets containing them. My conception of morality violates these two fairly standard conditions from social choice theory because it makes permissibility depend on contextual features (what other actions are feasible). Fortunately, I am not alone. The following theories violate Alpha: %Gauthier (1986) the theory of rational choice of %Levi (1986), minimax risk, and a satisficing theory that judges all actions to be permissible (with respect to some measure) all violate Alpha. Levi's theory of rational choice and the satisficing theory just described also violate Beta. (For example, %McClennen (1990), Ch.2, for discussion.)

13. Here I assume that for the purposes of prudential and moral assessment of actions, actions are individuated in a manner such that an action that is feasible in a given choice situation is also feasible in a choice situation that differs only with respect to what other actions are feasible. Broome is not committed to this assumption. He is committed to the assumptions of completeness, etc. without a substantive account of how actions are individuated for the purposes of the assessment. Broome uses this very point to argue against Humeans who endorse formal criteria of rationality (such as completeness and transitivity) and substantive criteria (requiring a specific preference or indifference). For, Broome cogently argues that specifying how actions are individuated for the purposes of prudential or moral assessment is equivalent to imposing conditions of indifference among actions. That is, requiring that two actions be treated as equally good (in terms of weak betterness or permissibility) if they are alike in a specified respect is equivalent to treating the two actions as equally good for the purposes of assessment. This is an extremely important insight.

14. Here I ignore some very real problems of how weak betterness for contingently possible future people should be dealt with.

15. By definition, for each action there corresponds exactly one prospect. Given the just mentioned assumption, for each prospect

least one action. It is possible that two distinct actions might correspond to a given prospect, but given Goodness Act Consequences will be equally good, so there will be no conflict in their implications for how the prospect is ranked.

16. In his discussion of the purported intransitivity of Maurice's preferences (pp. 100-107) Broome cogently argues that actions, and should be individuated on the basis reason-relevant features (justifiers). This corresponds to Reasons. It does not show (as Broome does) that all the reason-relevant features must be in the outcomes (as opposed specification of the events under, or probabilities with which) not show, that is, that Outcomism is true.

17. Here I assume that if there are probabilities, then prospects are defined in terms of them (e.g.,  $\langle p_1, O_1; \dots; p_n, O_n \rangle$ ). If the probabilities are then probability ranges, rather than points, will be assigned to each outcome.

18. Of course, the rejection of strong separability does not entail the rejection of all separability conditions. There may be weaker conditions that are plausible. Here I mention two variants.

Strong separability requires that (no matter how states are numbered), for any number i, the relative ranking of prospects  $\langle s_1, O_1; \dots; s_i, O_i; s_{i+1}, x_{i+1}; \dots; s_n, x_n \rangle$  and  $\langle s_1, O_1; \dots; s_i, O_i; s_{i+1}, y_{i+1}; \dots; s_n, y_n \rangle$  does not depend on what those common outcomes are. That is, the relative ranking of  $\langle s_1, O_1; \dots; s_i, O_i; s_{i+1}, x_{i+1}; \dots; s_n, x_n \rangle$  and  $\langle s_1, O_1; \dots; s_i, O_i; s_{i+1}, y_{i+1}; \dots; s_n, y_n \rangle$  does not depend on which specific outcomes the  $O_1$ - $O_i$  vari- stand for. A weaker version, weak separability, would require this independence only for only for cases where prospects agree up to  $O_{i-1}$ . would be the above principle restricted to the case where  $i=n-1$ .

Expressed with reference to probabilities a weaker assumption is: A betterness relation is weakly independent if and only if for any probability, and any outcomes  $x_{i+1}$ - $x_n$  and  $y_{i+1}$ - $y_n$  [as opposed to any prospects], and any values of  $O_1$ - $O_i$  ranging over sure outcomes, the relative ranking of  $\langle p_1, O_1; \dots; p_i, O_i; p_{i+1}, x_{i+1}; \dots; p_n, x_n \rangle$  and  $\langle p_1, O_1; \dots; p_i, O_i; p_{i+1}, y_{i+1}; \dots; p_n, y_n \rangle$  not depend on what outcome values  $O_1$ - $O_i$  take on.

For further criticism of strong separability and discussion of such weakenings, see, Broome's Ch.4, %McClennen (1990), and

19. Furthermore, as I shall discuss below, Outcomism creates problems for Broome's Rectangular Field assumption.

20. As I was finishing the first draft of this paper I discovered that the main points made in this section (e.g., about the vacuity of Outcomism is accepted) were made long ago by %Samuelson (1952), and more recently and more fully by %Machina (1991). The same points were also made by %Bacharach & Hurley (1991).

21. Note that those who accept a Rawlsian leximin principle (with no constraints) will accept the Principle of Personal Good.

22. This is weaker than the official version of the hypothesis given by Broome. His official version is: One alternative is at least as good as another if and only if it gives the person at least as great an expectation of goodness. This trivially entails that the person's good is representable by an expectational utility function. But since the General Addition Theorem already entails this, it is not needed for the derivation of the official version. The official version is weaker, since it says only that if goodness is (ordinally) representable by an expectational utility function, then each such alternative represents the good cardinally.

23. There have been a few precursors to Broome regarding the need for something like Bernoulli's Hypothesis. The need is clearly stated in %Sen (1976a), %Sen (1977a), %Sen (1986), %McClennen (1981), and %Weymark (1991).

24. If actions are individuated by their outcomes under states of nature, and there is objective chance or uncertainty as to what will occur under each state of nature, then the outcomes will need to include a specification of such chances or uncertainty.

25. See %Resnik (1987), Ch.6.4 for discussion, in the context of Harsanyi's theorem, for the need for the Rectangular Field assumption ("The Distributable Goods Assumption").

26. And it's not just modal facts that create problems for the Rectangular Field assumption when Outcomism is accepted. Facts about

the same problems. We are always ignorant of some aspects of the past. So some states of nature will involve different pasts. If then Outcomism requires them to be part of the outcomes. But outcomes with different pasts are conceptually incompatible. So, be violated.

27.This does not commit us to accepting the relevance of interpersonal comparisons of utility. For we may also hold (as I do) that has an equal weight no matter what scale it is on. This ensures that interpersonal comparisons are irrelevant, while guaranteed ordinal prudential weak betterness relation has an equal role (anonymity).

28.If, as I hold, the Principle of Personal Good exhausts the notion of moral good (and so it's conditionals also hold in the opposite person's prudential weak betterness is strong separable, then moral weak betterness will be as well (since it depends on nothing else).

29.Suppose, for example, that, with respect to prudential goodness, for Smith  $x = y > z$ , and for Jones  $x > y = z$ . Then (assuming people) the Principle of Personal Good requires that, with respect for moral weak betterness,  $x > y > z$ . Continuity then requires that  $p, y$  be morally indifferent to  $\langle p, x; 1-p, z \rangle$ . But given that moral weak betterness is exhausted by the Principle of Personal Good both individuals are prudentially indifferent between  $y$  and  $\langle p, x; 1-p, y \rangle$ . And that is not so. For any  $p$  less than 1,  $\langle p, x; 1-p, z \rangle$  is prudentially better than  $y$  for Smith; and for  $p=1$ ,  $\langle p, x; 1-p, z \rangle$  is prudentially better than  $y$  for Jones. So there is no  $p$  for which  $y$  is morally indifferent. Continuity is violated.

30.See %Vallentyne (1988a), %Vallentyne (1989a), and %Vallentyne (1991) for an elaboration and defense of a theory of moral goodness that satisfies all the assumptions of Broome's theory except the completeness and continuity of moral goodness.

31.I'm deeply indebted to John Broome for taking the time in correspondence to straighten me out on a number of matters. I've also received helpful comments from Brad Hooker, Shelly Kagan, Gene Mills and Tony Ellis.