

Review of *Ethics out of Economics* by John Broome

Nicolas Gravel

February 3rd 2000

The work of John Broome hardly needs an introduction to the readership of *Social Choice and Welfare*. It is a rare example of an intellectual effort of a high quality which has managed to remain in the difficult boundary zone between economics and philosophy. *Ethics out of Economics* is a collection of papers which, with one exception,¹ have all been published by the author in the last ten years. The papers are consistently presented and organized in such a way as to complete the book “*Weighing Goods*” ([3]) written by John Broome some ten years ago. In this latter book, Broome develops a wide encompassing theory of ethics that is based on the notion of “goodness”. Roughly putted, Broome’s theory says that alternative social states should be ethically compared only on the basis of the “goodness” of the consequences that they entail for the individuals. Comparing alternative social states requires therefore an assessment (weighing) of the different consequences that these social states may have on the “good” of the different individuals involved and to aggregate these different consequences into a ranking of the social states. To that extent, and at least when examined from the view point of an economist, Broome’s ethical theory is welfarist in the sense given to that word by Amartya Sen. It stands on the premise that the only information that is relevant to compare alternative states of affairs is the well-being (the “good”) achieved by individuals in those states. It is obviously outside the scope of this review to discuss the merits and limitations of this view. Broome’s himself has remained remarkably unresponsive to the various challenges which have been launched against welfarist principles of ethics from various spheres (notably Rawls [15], Sen ([17],[18],[19], [20]), Dworkin [5] and Roemer [16]). Taken these principles for granted, Broome tries instead to apply them to specific and delicate issues involving uncertainty, time and variable populations.

One should not be misled by the meaning of the word “apply” in the last sentence. For Broome’s work could certainly not be considered as “applied” in any ordinary account of the word. Quite to the contrary his work is abstract and philosophical in nature, based on discussions and argumentations rather than the building of specific models. This feature of Broome’s work, which is more apparent in *Ethics out of Economics* than it was in *Weighing goods*,

¹The exception is chapter 11 which has been published in the seventies.

has *pros* and *cons* when considered from the view point of someone trained as an economist. One bad thing is that the reading of the book may leave on occasion the reader with a sensation of vagueness or indefiniteness which is bit frustrating for someone used to definite and specific conclusion obtained out of no less definite and specific set of axioms and assumptions. Another weakness is the semantic nature of some of the discussions which tend sometimes to hidden the more substantive issues involved. But the *pros* largely outweigh the *cons* (after all are we not talking about weighing ?). It is always beneficial for economists to be recalled the limits and the specificity of assumptions and concepts which they are used to take for granted. And Broome's excels in doing just that. But the economist who read Broome also gets a somewhat less noble, but no less intense, satisfaction of having contributed, by the methods and the way of thinking they have developed, to philosophical ethics. For economists are not used to receiving this kind of compliments from philosophical spheres. Economists are more than often accused by philosophers to make "controversial methodological statements" ([9] p. 3) or to write "undisciplined musings, based on outmoded or ill-digested philosophical views" ([9]; p. 4). If would ever an economist suffer from such an inferiority complex with respect to philosophy ², then surely reading Broome would cure him or her all right.

The fourteen papers which constitute the body of "*Ethics out of Economics*" are grouped into three categories. They are preceded by an introduction which is an edited version the inaugural lecture delivered by the author at the University of Bristol in February 1995. For this reason this introduction may look a bit disconnected from the chapters it is supposed to introduce. The partition of the book into the three sections is also somewhat artificial. I will therefore ignore this partition and will briefly discuss the content of the book chapter by chapter.

The second chapter of the book (the first one is the introduction) is devoted to the meaning of the word "utility" as used by economists. John Broome first recalls the historical (and interesting) "twisting" in the meaning of this word which was used originally (and is still used in the daily language) to denote a property of an object (utility as usefulness) and which is now used (in economics) to refer to the pleasure obtained by the user of the object. According to Broome, this semantic twisting has been the source of a confusion about the actual meaning of the word "utility". In standard economic textbooks, the utility of an object is number assigned to this object by a function that represents the ordinal preferences of an agent. According to this view, utility has an *ordinal* meaning. But one finds also in economics many *cardinal* views of utility according to which the utility of an object is an exact (up to an affine monotonic transform) measure of the pleasure or well-being that this object provides to the agent. Broome argues that the word "utility" should *not* be used in the latter context and should be restricted to its ordinal meaning. He might be right although I feel inclined to believe that what really matters is the

²Roemer [16] (p. 3) seems to be one of them when we writes "I do not (...) believe that that the economist's way of thinking has produced, or will ever produced, important new insights into what distributive justice is. The key concepts in the last thirty years in the theory of distributive justice (...) have all come from the philosophical way of thinking".

meaning of the object to which the word applied rather than the word itself.

The third chapter concerns the important issue of whether the “extended preference” approach examined by Arrow ([1],[2]), Kolm [11] and Suppes [21] among others can serve as a basis for making interpersonal comparisons of utility. Recall that if individuals have “extended preferences”, they are able to make ordinal comparisons of the type “I would prefer to be individual i with consumption bundle x rather than being individual j with consumption bundle y ”. Many authors have argued that if individuals had preferences of this kind, then, provided that we take into account *all* characteristics that distinguish individuals from each others, they should have the *same* extended preference. Pursuing ideas developed in this journal (see [4]), he first explains with great care why this argument is faulty within an ordinalist framework in which utility only represents preferences. He then goes on to propose a way to salvage the argument if utility is taken from the beginning as an index of the “true” well-being (the “good”) of an agent. But the “extended preferences” approach certainly loses much of its appeal if it can only be applied to a framework in which utility is assumed to be cardinally measured from the outset.

The fourth chapter examines an important distinction concerning the issue of whether or not we should discount the future when making cost-benefit analysis. Broome argues that while discounting future goods (one potato tomorrow does not have the same social value as does one potato today) has some merit (it also has some flaws that Broome recalls to our attention), discounting future well-being is indefensible. Discounting future goods is sensible because time is productive and individuals themselves prefers current consumption to future one. Discounting future goods then merely amounts to respecting individuals (current) preferences and taking proper account of the production possibilities of the economy. However discounting future well-being is objectionable because it involves treating future well-being recipients differently from current ones. It may involve treating future generations differently from current ones on an arbitrary basis. Although straightforward, the distinction is well worth recalling because it is often ignored in many applied discussions.

The next chapter is more philosophical and, I feel inclined to say, more semantic. It addresses the question of whether there is a room for a “moderate” Humean viewpoint with respect to individual preferences. The Humean viewpoint states that any individual preference is rational or, analogously, that there can not be any rational discussion about individual preferences (*de gustibus et coloribus non disputandum*). An extreme Humean view would claim that any preference pattern whatsoever is rational. A moderate Humean would weaken such a claim by limiting its scope to preferences that satisfy a consistency requirement such as transitivity. Broome argues that the moderate view point is untenable and that the only choice is that between the extreme view point or a non Humean view point (according to which preference judgements can be the object of rational discussion). More specifically, he argues that a consistency requirement such as transitivity does not constrain in any way an individual’s *practical* preferences (although it does constrain what Broome calls her *non-practical* preferences). Broome’s argument goes as follow. Suppose we observe

an individual revealing, by a sequence of choices, a preference for A over B , for B over C and for C over A . According to Broome, such an observation should not be taken as evidence against the transitivity of the preferences of the individual because it could well be that the description of the alternatives A , B and C involved does not coincide with their description as they matter for the individual. For instance, it might be that the individual, in fact, cares not only about the objects A , B and C but also about the alternative objects with respect to which they are compared in the choice process. If that is the case, one could interpret the preceding sequence of choice as implying that alternative A when B is available (call this s) is preferred to B when A is available (call this t), that B when C is available (call this u) is preferred to C when B is available (call this v) and that C when A is available (call this w) is preferred to A when C is available (x). This new relabeling of the object makes the preceding description of choice perfectly consistent with transitivity. I must confess that I am not entirely convinced by Broome's argument. It seems to be based on at least two doubtful procedures. First, it confuses preferences as they may be revealed by actual choices with preferences as they may exist in the mind of the individual (without necessary implications for actual choices). If Broome wants to make the two notions synonymous (as I have the impression he does), I think he should mention it explicitly. After all the whole theory of revealed preferences suggests that there is a matter for distinction. Second, Broome's argument works by changing the definition of the set of alternative *after* the definition of a preference relation over this set. This is slightly dishonest, at least if we are playing seriously the game of modeling decisions by set theoretic notions. If a preference relation is defined with respect to a specific set of objects, then any of its property must be defined with respect to this very set of objects. If a preference relation is not transitive on this set of objects, then it is not! It does not seem convincing to me to question this conclusion by redefining the objects compared and the binary relation defined over them.

Chapter 6 contains a discussion of a theory of choice in uncertainty that is not well-known to most economists: That of Bolker-Jeffrey (see e.g. [10]) It also uses Bolker-Jeffrey theory of choice in uncertainty to prove a version of Harsanyi's ([7], [8]) utilitarian theorem. The Bolker-Jeffrey theory of choice under uncertainty differs from the usual theories developed in the Savagian framework by ignoring the distinction between actions and consequences which are crucial in this latter setting. In the Bolker-Jeffrey setting, uncertain prospects are simply *propositions* which are interpreted as set of all states of the world in which the proposition is true. Propositions in this theory can always be subdivided into subevents whose conjunction gives the proposition. While Bolker-Jeffrey's theory imposes axioms on preferences over such propositions that lead to a numerical representation which takes the form of the expectation of utility numbers, these numbers do not receive the same interpretation as they do in standard expected utility theory. In particular *none* of these utility numbers can be interpreted as the utility of a certain consequence defined independently from its probability of occurrence. The utility of any proposition depends upon both the probability and the utility of all its subevents. This theory has *pros*

(it does not force one to believe in the existence of consequences which are free from any remaining uncertainty) and *cons* (it does not work well in game theoretic settings where an individual's strategy takes naturally the form of a function from the strategy sets of this individual's opponents to the set of pay-off). Both are carefully recalled to our attention by Broome. The only weakness of this chapter is the insufficient motivation, at least in my view, for using the Bolker-Jeffrey theory to prove (and therefore somehow endorse) Harsanyi's famous expected utility justification for utilitarianism. Insofar as I can understand it, utilitarianism is, like all welfarist theories, a *consequentialist* theory. It compares alternative social states on the basis of the ultimate (and I suppose well-defined) consequences that they entail for the individual. It seems therefore somewhat surprising that such a theory could find its justification into a framework which dismisses the existence of well-defined consequences. Moreover, and as is well-known, there has been serious criticisms addressed to Harsanyi's classical justification for utilitarianism (see Weymark [22] for a thorough account of what has become to be known as the Harsanyi-Sen debate). For sure, these criticisms apply to the Bolker-Jeffrey version of Harsanyi's theorem provided by Broome. Yet the chapter contains no mention of the Harsanyi-Sen debate, and more generally of the relevance of Harsanyi's theorems for social choice.

In chapter 7, Broome examines the role of random lotteries as a "fair" device for allocating indivisible goods between many individuals. A typical example where the use of a random lottery may seem at first glance to be a fair method to allocate an indivisible good is organ transplants. Suppose one and only one kidney is available and that there are four patients with identical relevant characteristics (age, severity of illness, etc.) waiting for a kidney. Clearly, there is no reason to give the kidney to a particular person even though there is a good reason to give the kidney to someone). In such a case, deciding randomly who should get the kidney (by using a device that gives an equal probability to each of the four individuals) seems indeed to be a "fair" way of making the decision. It is not "fair" because it is random. It is fair because it allows to break the tie among four equally good solutions of a social choice problems that leaves no room to the discretionary power of the person in charge of allocating the kidney. But Broome makes the more controversial claim that lotteries are *intrinsically fair* methods for allocating indivisible goods, and that they are so even in circumstances where the individuals who compete for having the indivisible good are not identical in the relevant characteristics. He gives the good example of a dangerous mission, which, if successful, would lead to very good consequences to a group of people but that exposes the individual who performs it to very high risk of death. The example also supposes that individuals differ in their ability to make the mission successful (the benefit of the mission depends on the person who performs it) but not in their probability of dying. He argues that choosing at random the person who should perform the mission would be fair in that setting (even though some could have been inclined to believe that assigning the mission to the more able person would also have been reasonable here)³. It

³Given any rule for redistributing the benefit of the mission, this latter solution would be

would be nice to see the formal argument underlying Broome's reasoning.

Chapters 8 and 9 are devoted to the notion of incommensurability, which unavoidably arises when aggregation over individuals of some notion of "goodness" is involved. After all, few people would be willing to consider two different individuals' well-being as perfectly commensurable magnitudes. Neither are the various and possibly conflicting values that individuals may wish to satisfy in the conduct of their lives. The conceptual tool used by Broome to address the problem of commensurability is what he calls a 'standard configuration'. Suppose we have an individual (Bob) whose happiness can be measured on a continuous scale. Take a level of happiness very high on Bob's scale and another very low on the same scale. Take now a given level of happiness of someone else (Natacha). One could plausibly argue that Natacha's level of happiness is comparable to (and lower than) the high level of happiness of Bob, that Natacha's level of happiness is higher than Bob's low level of happiness but that Natacha's level of happiness is not comparable to (incommensurable with) a spectrum of levels of Bob's happiness comprised somewhere between the high and the low level. This describes a standard configuration (in which the specific level of happiness selected for Natacha serves as the 'standard'). In chapter 8, Broome relates the notion of incommensurability with the notion of vagueness taken from formal logic. A proposition is vague in this sense if its truth value (in the formal logical sense of this concept) is allowed to differ from 0 and 1. Broome goes on in distinguishing carefully statements of non-comparability of the type "Natacha's level of happiness is not comparable to Bob's level" from vague statements of the type "it is neither true or false that Natacha is happier than Bob and it is neither true or false that Bob is happier than Natacha". I am uncertain about the importance of the distinction but this may be due to my lack of expertise in formal logic. In chapter 9, Broome discusses, within the 'standard configuration' framework, the more substantial issue of existence incommensurable values. More specifically, he addresses two questions: 1) Are options available to a decision maker (which may be a social planner) always ordered by their "goodness" and, if not, 2) how important is this fact for ethics? The answer given is a plain (and unsurprising) yes to the first question. The answer given to the second question is also affirmative and stands on a careful examination made by Broome of the difference between *equivalence* and *non-comparability* of objects.

In chapter 10, Broome argues in favor of defining "goodness" in "relative" (or comparative) rather than "absolute" terms as often done in the philosophical literature. Some properties can be expressed in terms of comparative statements and other not. The property of "correctness" for instance, as applied to calculations, is not a property that can be expressed in terms of comparative statements. A calculation is correct or is not correct. A calculation can not be said to be more correct than another. Another property that can not easily be expressed in terms of comparative statement is the notion of envy-freeness as applied to allocation of goods across individuals with selfish preferences. An al-

selected by the maximisation of the sum of the expected utility of the individuals for instance.

location of good is envy-free or involves envy. But an allocation can not be more envy-free than another. Broome argues convincingly in this chapter that a plausible notion of “goodness” for ethics would not be absolute like “correctness” or envy-freeness. It would instead express itself in the form of comparative statements of “betterness” relating alternative options. Broome uses this relative notion of “betterness” to appraise a philosophical examination performed by Nagel [12] (ch. 1) of a statement made by Epicurus about the fear of dying.

The last five chapters of the book are more consistently connected together since they all address the delicate issue of valuing life. Chapter 11 is a criticism of an approach, proposed by Mishan and others, to evaluate projects (like building nuclear plants) that affect the risk of death faced by individuals. The approach criticized by Broome is based on the potential Pareto criterion: Accepts a project if it is possible for the gainers of the projects to compensate the loser while still remaining (after paying the compensation) gainers. Clearly, such an ethic can not be applied to projects involving the death of someone since no amount of money could compensate someone for her death. Yet very few projects involve the killing of a certain known person. What most projects do is to modify the *probability of dying* faced by individuals. Given this, it has been suggested by Mishan to apply the potential Pareto argument to the probability of dying rather than to the final deaths. A project is worth doing if the compensations required *ex ante* to make people willing to accept the increase in the probability of dying brought about by the project could be paid by the beneficiaries of the project. Broome shows the flaw of the argument which rides on the doubtful idea that making individuals ignorant of the identity of those who will die can help in rendering a project normatively acceptable.

Chapter 12 discusses the difference between a “structured” and an “unstructured” approach to cost-benefit analysis. By the “unstructured approach”, Broome means an approach that does not take a full account of the formal structure of the social states compared by the cost-benefit analysis as opposed to the structured approach which do take account of this structure. The “formal structure” to which Broome is referring is that underlying time and uncertainty. When a project involves time (and in particular when it affects the well-being of yet unborn individuals), the unstructured approach would base the project appraisal on the current point of view of those who are alive at the moment where the project is decided. A structured approach by contrast would use information on the future well-being of those who alive at the present instant (which does not necessarily coincide with their current evaluation of it) as well as the well-being of future generations. When a project involves uncertainty, the unstructured approach would amount to taking an *ex ante* point of view rather than an *ex post* one as would the structured approach (using Hammond’s [6] distinction). Unsurprisingly, Broome argues in favor of the structured approach. He does so by proposing his own methodological assessment of what cost-benefit analysis should be: An exercise aiming at discovering what ought to be done rather than a prescription of action

Chapter 13 discusses the use of QALYs (Quality Adjusted Life Years) in health economics. Defining a QALY requires a numerical indicator of the good-

ness of one year of life. As such, it is an exercise that is bound to meet the same kind of difficulties as any other exercise of welfare measurement (viz. cardinality, interpersonal comparability, etc.). Broome also raises a few problems of QALY analysis when applied to decisions involving yet unborn (or almost unborn) individuals. In standard QALY analysis, the benefits obtained out of applying a medical treatment to someone is the number of quality adjusted years of life that this treatment will produced. In doing this analysis, no account is made for the QALYs of the descendants of the person concerned with the treatment. People who are yet unborn do not count in QALY analysis. But as pointed out by Broome, this may lead to an analysis that is too sensitive to the exact definition of the boundary of life. Is saving the life of a fetus as beneficial as saving that of a recently born baby ? What about a decision of having a baby as compared to the same decision of saving the life of a recently born baby ? An analysis only based on the number of quality adjusted remaining years of life would answer these two questions by a definite yes, even though, for many of us, they would seem to have rather different answers.

Economics and philosophers have only started recently to address questions of this kind into a consistent framework.⁴ The two last chapters of the book provides the reader with a sample of Broomes's views on these matters. Both are devoted to the analysis of policies that affect both the quality of the life of those alive and the well-being of yet unborn individuals. They are addressed in the welfarist framework favored by Broome which views states of nature, individuals and time periods as different possible locations of a single magnitude - the good or well-being - which needs to be allocated among these. Chapter 14 examines two difficulties in applying welfarism to such a setting. The first concerns the basis tenet of *separability of time* which is often accepted as a sensible property to impose on social welfare functions when there is a fixed amount of individuals. Broomes argues that separability of time is not defensible as a normative principle for analysis variable populations problems because it forces one to consider the prolongation of one's life and the creation of a new life as being equivalent socially (provided that the life created has the same length and the same per period utility as does the extension of the actual life). Yet, according to Broome, there is strong intuition according to which prolonging existing lives has more social value than creating new ones. Critical level utilitarianism characterized in Blackorby, Bossert and Donaldson [?] captures this intuition to some extent since it considers the addition of a new life to be a social improvement only if the new life has a life-time utility above some exogenous critical level. On the other hand it does not require the prolonging of one's existing life to have a (remaining) life-time utility above the critical level to consider it as a social improvement.

The problem with critical levels utilitarian principles lies of course in the choice of the critical level. If the critical level is chosen to be too low, then it would be possible to improve matters simply by producing a very large amount

⁴See Parfit [14] on the philosophical side and Blackorby, Bossert and Donaldson [?] on the economics one.

of new babies and by giving each of them an arbitrary low (but above the critical level) life time utility (a state of affairs referred by Parfit [14] as to the “repugnant conclusion”). If on the other hand the critical level is too high (and correspond to a decent standard of living), then one would be led to the no less “repugnant conclusion” that it would be worthwhile socially to spend actual resources to prevent to birth of individuals who would enjoy reasonable levels of lifetime utility.

Broome examines also (especially in chapter 15) a principle which would support his “intuition” that a social “priority” should be given to existing people over yet unborn ones. The principle examined is the one proposed by Narveson [13] and referred to as the “constituency principle”. This principle says that, in comparing two social states involving different number of individuals, we should only consider the view point of those who are alive in the two states. Put bluntly, the well-being of someone who is not yet borne should not count in deciding whether or not we should bring this person into existence. It is easy to see that such a principle can lead to intransitive rankings of social states. It does not seem that easy however to show that this principle could lead to a ranking that violates *quasi-transitivity* in a framework that prevents the possibility of killing individuals. Broome proposes a violation of quasi-transitivity in his example 3 of chapter 15 but the example stands heavily of the possibility of ranking a state in which someone has been killed above one in which this person was alive. But it could easily be argued (while staying in tune with Narveson’s intuition) that the constituency principle should only be applied when dealing with *addition* of new persons. It should therefore not be applied to situations involving the *subtraction* (killing) of existing persons. After all the view that only the preferences of people who remain alive should count in deciding whether or not to kill someone is not at all appealing. Broome’s example 3 fails to make this distinction and one may find this a bit surprising. I conjecture that there can be no violation of quasi-transitivity by the sole constituency principle if one restricts its application to decisions involving the creation of new individuals.

It is my hope that this lengthy review will give to the readership of *Social choice and Welfare* a foretaste of the quality as well as the importance of the contributions gathered in *Ethics out of Economics*, despite what the minor critical sound of some of the comments may suggest.

References

- [1] K. J. Arrow. *Social Choice and Individual Values*. Cowles Foundation Monograph. Yale University Press, New Heaven, 1963.
- [2] K. J. Arrow. Extended sympathy and the possibility of social choice. *American Economic Review: Papers and Proceedings*, 67:219–225, 1977.
- [3] J. Broome. *Weighing Goods*. Blackwell, Oxford, 1991.

- [4] J. Broome. A cause of preference is not an object of preference. *Social Choice and Welfare*, 10:57–68, 1993.
- [5] R. Dworkin. What is equality ? part 1: Equality of welfare, part 2: Equality of resources. *Philosophy and Public Affairs*, 10:185–246, 1981.
- [6] P. J. Hammond. Ex ante and ex post welfare optimality under uncertainty. *Economica*, 48:235–250, 1981.
- [7] J. Harsanyi. Cardinal utility in welfare economics and in the theory of risk taking. *Journal of Political Economy*, 61:434–435, 1953.
- [8] J. Harsanyi. Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility. *Journal of Political Economy*, 63:309–321, 1955.
- [9] D. M. Hausman and M. S. McPherson. Editors’ introduction. *Economics and Philosophy*, 1:1–6, 1985.
- [10] R. Jeffrey. *The Logic of Decision*. University of Chicago Press, Chicago, 1965. reprint 1983.
- [11] S. C. Kolm. *Justice et Équité*. Editions du CNRS, Paris, 1972.
- [12] T. Nagel. *Mortal Questions*. Cambridge University Press, Cambridge UK, 1979.
- [13] J. Narveson. Moral problems of population. *Mind*, 76:62–72, 1967.
- [14] D. Parfit. *Reasons and Persons*. Oxford University Press, Oxford, UK, 1984.
- [15] J. Rawls. *A Theory of Justice*. Berknap Press of Harvard University Press, Cambridge, Massachusetts, 1971.
- [16] J. E. Roemer. *Theories of Justice*. Harvard University Press, Cambridge MA, 1996.
- [17] A. K. Sen. Utilitarianism and welfarism. *Journal of Philosophy*, 76:463–469, 1979.
- [18] A. K. Sen. Well-being, agency and freedom: The dewey lectures 1984. *Journal of Philosophy*, 82:169–221, 1985.
- [19] A. K. Sen. *The Standard of Living*. Cambridge University Press, Cambridge, UK, 1987.
- [20] A. K. Sen. Welfare, preferences and freedom. *Journal of Econometrics*, 50:15–29, 1991.
- [21] P. Suppes. Some formal models of grading principles. *Synthese*, 16:284–306, 1966.

- [22] J. A. Weymark. A reconsideration of the harsanyi-sen debate on utilitarianism. In J. Elster and J. Roemer, editors, *Interpersonal Comparisons of Well-Being*, pages 255–320. Cambridge University Press, Cambridge, U. K., 1991.