

Assessing risky social situations*

Marc Fleurbaey[†]

May 2006; this version Jan. 2007

Abstract

This paper analyzes social situations in the context of risk. A new argument is proposed in order to defend the ex-post approach against Diamond's (1967) famous critique of Harsanyi's (1955) utilitarian theorem. It leads to a characterization of the criterion consisting in computing the expected value of the "equal-equivalent". Characterizations of the ex-post maximin and leximin criteria, as well as a variant of Harsanyi's theorem, are also obtained. It is examined how to take account of concerns for ex-ante fairness within this ex-post approach. Related issues are also addressed, such as the rescue problem or preference for catastrophe avoidance.

Keywords: risk, social welfare, ex ante, ex post, fairness, Harsanyi theorem.

JEL Classification: D63, D71, D81.

*This paper has benefited from comments by R. Bradley, J. Broome, A. Casella, N. Gravel, P. Mongin, A. Trannoy, from discussions with K. Arrow, P. Casal, T. Christiano, M. Clayton, G. Cohen, F. Dietrich, P. Hammond, R. Hockett, C. List, F. Maniquet, S. Orr, M. Otsuka, J. Roemer, A. Voorhoeve, A. Williams, J. Wolff, as well as from reactions from audiences at the LSE, Carlos III (Madrid) and GREQAM (Marseilles).

[†]CNRS-CERSES, University Paris 5, and IDEP. Email: marc.fleurbaey@univ-paris5.fr.

1 Introduction

This paper defends a family of criteria for the evaluation of risky social situations, namely, those which rely on the expected value of the equal-equivalent utility. The equal-equivalent utility of a given distribution of utilities is the utility level which, if equally enjoyed by the population, yields a distribution that is judged as good as the examined distribution. If the equal-equivalent utility is lower than the average utility, this criterion displays aversion to ex-post inequalities. The ex-post maximin criterion, which computes the expected value of the smallest utility, is one example.

The literature on risky social decisions has been strongly influenced by Harsanyi's (1955) defense of the utilitarian criterion. Harsanyi showed that this criterion has the unique virtue of being at the same time based on expected social utility – corresponding to rationality at the social level – and based on ex-ante individual preferences – i.e., respecting the Pareto principle. Hammond (1983) and Broome (1991) have offered additional arguments in favor of this criterion, and made a careful scrutiny of the scope of this result.¹ A well known drawback of the utilitarian criterion is its indifference to ex-ante and ex-post inequalities, while intuition suggests that ex-ante inequalities matter, as noted by Diamond (1967) and Keeney (1980a), and ex-post inequality aversion appears also rather pervasive (Meyer and Mookherjee 1987, Harel et al. 2005). For those who seek an inequality-averse criterion, the choice between ex-ante criteria which focus on individual expected utilities and ex-post criteria which compute the expected value of social welfare appears to take the form of a dilemma.

The problem can be illustrated as follows.² Consider two individuals, Ann and Bob,

¹See also Weymark (1991) for a careful examination of Harsanyi's arguments in favor of utilitarianism.

²This is inspired by Broome (1991) and Ben Porath et al. (1997). See also Fishburn (1984) for an early axiomatic analysis of these issues.

and two equally probable states of natures, Heads and Tails (as with coin flipping). There are three lotteries to compare, L_1 , L_2 and L_3 , which are described in the following tables by their utility consequences for Ann and Bob.

L_1	H	T
Ann	1	1
Bob	0	0

L_2	H	T
Ann	1	0
Bob	0	1

L_3	H	T
Ann	1	0
Bob	1	0

Harsanyi's additive criterion is indifferent between the three lotteries. An ex-post approach is indifferent between L_1 and L_2 (provided it is impartial between Ann and Bob), because the ex-post distribution is $(0, 1)$ for sure for both lotteries, while an ex-ante approach is indifferent between L_2 and L_3 , because the individuals face the same individual lotteries (0 or 1 with equal probability) in the two lotteries.

But, intuitively, L_1 is worse than L_2 which is worse than L_3 . Indeed, with L_1 Ann is advantaged for sure whereas L_2 randomizes over the two individuals, as with a coin tossing. This is Diamond's (1967) observation. With L_3 , there is the same ex-ante prospect for each individual as with L_2 , but with less inequality ex post, as observed in Broome (1991).

Facing this problem, Ben Porath et al. (1997) suggest taking a weighted sum of an ex-ante and an ex-post egalitarian criterion.³ Broome (1991) proposes to introduce a measure of ex-ante fairness in the measurement of individual utilities. In a similar vein, Hammond (1983) alludes to the possibility of recording intermediate consequences jointly with terminal consequences.

This paper takes sides with the ex-post approach and against the ex-ante approach, on the ground that risky situations are fundamentally situations of imperfect information. In this light the Pareto principle applied to ex-ante prospects does not appear compelling, whereas the ex-post approach is justified by the argument that one should try to mimic an

³A similar proposal is made in Gajdos and Maurin (2004).

“omniscient” evaluator whenever this is possible. It is further argued that this argument is fully compatible with ex-post inequality aversion. The next section introduces the “omniscience principle” which is the key analytical element of this paper, and Section 3 formalizes it and derives its consequences over the definition of an ordering of risky social alternatives. In particular, it offers characterizations of the ex-post equal-equivalent, maximin and leximin criteria. Section 4 provides a variant of Harsanyi’s (1955) theorem based on the omniscience principle instead of the expected social utility assumption, and discusses the conflict between aversion to ex-post inequality and the Pareto principle. Section 5 discusses the related literature, and in particular examines how ex-ante fairness, which has been a concern for many authors after Diamond (1967), can be accommodated by the ex-post criteria proposed here. Section 6 concludes.

2 Risk as incomplete information

One of the greatest achievements of economic theory is the Arrow-Debreu-Mackenzie general equilibrium model, and especially its extension to the case of contingent goods (see e.g. Debreu 1959), which allows one to analyze risky situations with the standard analytical concepts of consumer demand. The formal analogy between consumer choice over, say, carrots and tomatoes and consumer choice over money if it rains and money if it is sunny is a very powerful insight which underlies the economics of risk and insurance and enables us to understand insurance demand and risk-taking behavior in a very convenient way. This analogy is generally carried over from positive analysis to normative prescriptions. The economics of insurance typically analyzes the performance of insurance markets in terms of efficiency and equity by referring to the satisfaction of the agents’ ex-ante preferences, just as one is used to catering to consumer preferences over carrots and tomatoes. The analysis of public policies reducing risk is also generally made from

an ex-ante perspective.⁴

This extension of the ex-ante viewpoint from positive economics to normative issues is unwarranted. Consider the following example. A corn-flakes producer discovers that a small number b of boxes have been contaminated with potentially lethal chemicals. The boxes are now in the market and, for the sake of the example, let us imagine that the producer has no way of tracking where these boxes are. Suppose that b is so small that the risk to which consumers are submitted is not greater than the risk of accident they daily endure in other activities, and that, even knowing about the risk, consumers would still be willing to buy corn-flakes at the posted price.

From the ex-ante perspective, the contamination is just like a small reduction in the average quality of corn-flakes, and therefore, since there is nothing problematic in letting consumers accept a small reduction in quality, it is no more problematic to let them accept the risk of dying from a contaminated box. By taking the ex-ante perspective, one therefore concludes that the reasons for concern are neither smaller nor greater in any of the two cases, quality reduction or hazard. But something is going awry here. There is a key difference between the consumer who buys a box knowing that the quality is lower, and the consumer who buys a box without knowing that this particular box is contaminated. The risky situation involves a lack of information that is absent from the other situation. The ex-ante perspective blurs this difference.

Does this difference have any practical implication? After all, nobody knows if the particular box bought by this consumer is contaminated. It seems that this consumer makes his decision with the best available information, so that nobody can warn him against buying it, pretending to know better. The central idea which underlies this paper is that, actually, it sometimes happens that an external observer is in a situation that is

⁴See e.g. Viscusi (1992), Pratt and Zeckhauser (1996). (The latter even take a veil-of-ignorance perspective, which introduces artificial risk in order to achieve an impartial judgment.)

equivalent to knowing better than the individual agent.

Imagine an omniscient evaluator who knows where the contaminated boxes are. This evaluator sees that b consumers are buying a box at a price they would not accept if they knew the content of the box. Let us ask this evaluator to make a judgment about two situations. One is the *laissez-faire* situation in which the boxes are freely sold. The other is the withdrawal of all the boxes, contaminated and non-contaminated, from the market. Since the former implies that b consumers are contaminated, one can imagine that for reasonable social preferences, the evaluator prefers the withdrawal policy. Let us assume that these are indeed the social preferences in the case at hand.

The interesting fact here is that this omniscient evaluator's judgment presumably relies only on information that we already have. Suppose that we prefer the distribution of consumer satisfaction that results from the withdrawal policy to the distribution that results from the contamination of b consumers. Then any omniscient evaluator who shares our view will also prefer the withdrawal policy. The additional information that this evaluator has, namely, who actually gets contaminated, is not relevant for impartial social preferences. Therefore, when the withdrawal policy is implemented, we can say to the consumer: "We do not know whether you would be contaminated, but even if we knew, we would prefer the withdrawal policy to the *laissez-faire* policy."⁵ If the consumer agrees with the ethical principles which underlie the preference for this policy, he will not object, even though, as a consumer, he was willing to buy the box.

In more general terms, the argument is the following. Let us imagine an omniscient evaluator who has exactly the same social preferences as ours for cases involving no risk. By mere deduction, we can also guess a subset of her preference ordering over cases involving risk for us (for her there is no risk because she knows the true state of nature).

⁵Of course, the first-best policy for the omniscient evaluator would be to withdraw only the contaminated boxes, but we cannot implement this policy with the imperfect information that we have.

Because she has better information, we should construct our own social preferences in such a way that they never contradict what we know about hers. This is what I propose to call the “omniscience principle”: *Never oppose what you can guess about your omniscient alter ego.*

The next section will formalize this principle and examine what actually can be guessed about an omniscient evaluator’s preferences over risky situations. Before that, a related point must be made. Interpreting risk as a lack of information renders the Pareto principle suspect when it is applied to ex-ante prospects. It is widely recognized that, when individuals have different beliefs, their unanimity over prospects may be spurious and carries little normative weight.⁶ But even if all individuals have the same beliefs about the probabilities of the various possible states of nature, their beliefs are “wrong” insofar as there is only one true state of nature. Therefore, their preferences over prospects do not call for the same respect as fully informed preferences.

Consider the three lotteries described in the following tables. There are two equiprobable states of nature (Heads or Tails), and the figures in the table are the agents’ von Neumann-Morgenstern (VNM) utilities.

M_1	H	T
Ann	1	1
Bob	1	1

M_2	H	T
Ann	0	3
Bob	3	0

M_3	H	T
Ann	0	3
Bob	0	3

Both agents prefer M_2 to M_1 , but their assessment is based on the equiprobability of the two states of nature. In reality, only one of the states is realized, and one agent’s better informed preferences, ex post, will oppose his or her ex-ante preferences. Respecting the agents’ ex-ante preferences therefore goes against the fully informed preferences of one of them. Since the latter preferences are more worthy of respect than the former, the Pareto

⁶See e.g. Broome (1999, p. 95), or Mongin (2005).

principle applied to ex-ante preferences is not compelling.

Suppose that for cases involving no risk we actually prefer the distribution of utilities $(1, 1)$ to $(0, 3)$ and to $(3, 0)$. Then we know that our omniscient alter ego prefers M_1 to M_2 , because no matter whether Heads or Tails is the true state, she prefers M_1 to M_2 . Therefore, the omniscience principle, combined to these preferences over non-risky cases, forces us to go against the agents' ex-ante preferences. This can be intuitively defended as follows. In M_2 , the worse-off agent will regret his or her ex-ante preferences and realize that M_1 was actually better for him or her. If we give enough priority to the fully informed preferences of the worse-off agent, it makes sense to prefer M_1 .

The situation is quite different when one looks at M_3 and compares it to M_1 . In this case we do not know if the fully informed agents will agree or disagree with their ex-ante preferences, and there is no way to guess what our omniscient alter ego thinks. In this case, therefore, the Pareto principle applied to ex-ante preferences regains some appeal.

The difference between M_2 and M_3 can be described in terms of micro versus macro risk. In M_2 the negative correlation between agents produces a final distribution of utilities which does not depend on the true state of nature, and this enables us to guess how our omniscient alter ego compares it to M_1 . In M_3 , in contrast, the correlation is positive, one has a macro risk and the final distribution is unknown. Our omniscient alter ego either likes or dislikes M_3 , compared to M_1 , but there is no way to guess, and therefore individual ex-ante preferences provide a good criterion. In a one-agent case, in particular, all risks are macro risks, and there is no reason not to accept the agent's own evaluation of his prospects.

3 Ex-post criteria

In this section we formalize the argument of the previous section and derive its consequences.

The framework is as simple as possible. The population is finite and fixed, $N = \{1, \dots, n\}$.⁷ The set of states of nature is finite, $S = \{1, \dots, m\}$, and the evaluator has a fixed probability vector $\pi = (\pi^s)_{s \in S}$, with $\sum_{s \in S} \pi^s = 1$. It does not matter whether the probabilities are objective or subjective, but if they are subjective they belong to the evaluator.⁸ Since what happens in zero-probability states can be disregarded, we simply assume that $\pi^s > 0$ for all $s \in S$.

The evaluator's problem is to rank lotteries U , where $U = (U_i^s)_{i \in N, s \in S} \in \mathbb{R}^{nm}$ describes the utility attained by every i in every state s . Let \mathcal{L} denote the set of such lotteries. Let U_i denote $(U_i^s)_{s \in S}$ and U^s denote $(U_i^s)_{i \in N}$. The social ordering over the set \mathcal{L} is denoted R (with strict preference P and indifference I).

The utility numbers U_i^s are assumed to be fully interpersonally comparable, and they may measure any subjective or objective notion of advantage which the evaluator considers relevant for social evaluation. In particular, it is not necessary for the analysis that they derive from individual VNM⁹ utility functions. However, it is assumed that,

⁷This means that we consider the possibility of individuals being killed but not risky decisions that may affect their mere existence. This is because issues of variable population trigger a whole set of different considerations that are better kept aside here. On variable population issues, see Hammond (1996), Blackorby et al. (2005) and Broome (2004).

⁸There is no problem of aggregation of beliefs here, since the evaluator uses his own probability vector in order to compute expected values. This way of proceeding can be justified by the argument that one should always use one's best estimates of the probabilities when evaluating a lottery. The evaluator's beliefs may be influenced by the population beliefs, but we work here with his final probability vector.

⁹Some authors make a distinction between the VNM utility function bearing over lotteries or prospects and the Bernoulli function bearing over consequences. For simplicity, both are called VNM functions here.

for one-person evaluations, the evaluator considers that the expected value $\sum_{s \in S} \pi^s U_i^s$ correctly measures agent i 's ex-ante interests.

In Appendix B, it is shown how the analysis can be extended to a more complex framework in which individual situations are described by vectors (representing consumption bundles for instance), and individual VNM utility functions appear distinctly.

We now introduce requirements that R should satisfy. The first, and the main one, formalizes the omniscience principle. The social ordering R over lotteries induces an ordering over non-risky cases, i.e., degenerate lotteries in which all vectors U^s are the same. Note that the omniscience principle, as such, imposes no restriction on this ordering. It is a purely formal principle which can be applied to any kind of social preferences, e.g., egalitarian, anti-egalitarian or inequality-neutral. Let $[U^s]$ denote the degenerate lottery in which vector U^s occurs in all states of nature. Our omniscient alter ego, ex hypothesi, agrees with R on degenerate lotteries. Her ranking of the whole set of lotteries, which can be denoted R^* , is therefore defined as follows. Let s^* be the true state of nature. For all $U, U' \in \mathcal{L}$,

$$U R^* U' \text{ if and only if } [U^{s^*}] R [U'^{s^*}].$$

In other words, our omniscient alter ego is only interested in what happens in state s^* , and evaluates options as we would do if what happens in s^* were to happen for sure. When can we guess what our omniscient alter ego thinks? Since we do not know s^* , we can only be sure about $U R^* U'$ when $[U^s] R [U'^s]$ for all $s \in S$. That is the most we can guess about our omniscient alter ego, and the omniscience principle is then formalized by an axiom guaranteeing that we always agree with this guessing.

Axiom 1 (Omniscience Principle) *For all $U, U' \in \mathcal{L}$, one has $U R U'$ if for all $s \in S$, $[U^s] R [U'^s]$.*

In other words, this means that a lottery is at least as good as another if we are

sure that it produces an at least as good situation, whatever the state of nature. This is, in essence, a basic consequentialist principle which the reader might find attractive independently of any invocation of an omniscient alter ego. But, as will be discussed in the last section, an important subset of the literature implicitly or explicitly rejects it, especially in relation to Diamond's (1967) critique of Harsanyi, and we will see that the analysis of risk as imperfect information, which underlies the omniscience principle, is useful in order to defend it against such views.

This kind of condition, if it were applied to individual decision-making under risk, would seem totally uncontroversial. A prospect is obviously better for an individual if it dominates another in all states of nature. Viewed in this light, it is somewhat amazing that one has to go to great lengths to defend such a basic condition in the context of social preferences.

The second axiom expresses the idea, also introduced in the previous section, that, for pure macro risks, the evaluation should coincide with the ex-ante appreciation of individual situations, since the evaluator cannot take advantage of his knowledge of the statistical distribution of luck in this case.¹⁰ A pure macro risk is described here as a situation in which in all states of nature, the distribution of utilities is equal.

Axiom 2 (Pareto for Macro Risks) *For all $U, U' \in \mathcal{L}$ such that for all $i, j \in N$, $U_i = U_j$ and $U'_i = U'_j$, one has $U R U'$ if for all $i \in N$, $\sum_{s \in S} \pi^s U_i^s \geq \sum_{s \in S} \pi^s U_i'^s$.*

¹⁰As already explained, it is not assumed here that π , used by the evaluator, coincides with the agents' beliefs, since an evaluation simply has to be made with the best knowledge of the evaluator himself. Moreover, the utility figures need not be individual von Neumann-Morgenstern utilities. Therefore the Pareto axioms in this paper do not necessarily reflect a respect for individual ex-ante preferences. More basically they describe a social evaluation that respects the ex-ante evaluation of individual situations by the same evaluator: "If the evaluator thinks that nobody is ex-ante worse-off, he also thinks the lottery is not worse".

The next two axioms are standard monotonicity and continuity conditions with respect to U . Vector inequalities are denoted $\geq, >, \gg$.

Axiom 3 (Weak Monotonicity) For all $U, U' \in \mathcal{L}$, one has $U P U'$ if $U \gg U'$.

Axiom 4 (Continuity) Let $U, U' \in \mathcal{L}$ and $(U(t))_{t \in \mathbb{N}} \in \mathcal{L}^{\mathbb{N}}$ be such that $U(t) \rightarrow U$. If $U(t) R U'$ for all $t \in \mathbb{N}$, then $U R U'$. If $U' R U(t)$ for all $t \in \mathbb{N}$, then $U' R U$.

We can now state a first result, which describes how R must be computed in order to satisfy these axioms.

Theorem 1 R satisfies the four axioms if and only if, for all $U, U' \in \mathcal{L}$,

$$U R U' \Leftrightarrow \sum_{s \in S} \pi^s e(U^s) \geq \sum_{s \in S} \pi^s e(U'^s),$$

where $e(U^s)$ is defined by the condition

$$[U^s] I [(e(U^s), \dots, e(U^s))].$$

The function $e(U^s)$ is the “equal-equivalent” value of U^s , i.e. the utility level which, if enjoyed uniformly by all agents, would yield a distribution that is equivalent to U^s . The theorem says that the expected value of the equal-equivalent should be the criterion. The proofs of the results (and examination of the independence of axioms) are in the appendix, but the proof of this one is particularly simple. Weak Monotonicity and Continuity guarantee the existence of the function $e(\cdot)$. Omniscience Principle implies that, for all lotteries U , one can replace each vector U^s by $(e(U^s), \dots, e(U^s))$, because our omniscient alter ego is, like us, indifferent between $[U^s]$ and $[(e(U^s), \dots, e(U^s))]$. Once this is done for all s , one obtains a lottery which involves a pure macro risk. One therefore simply has to know how to rank lotteries with pure macro risks, and the rest of the ordering

is automatically derived from this. Pareto for Macro Risks then intervenes in order to require relying on expected utilities for such lotteries.

The result is compatible with e being partial in favor of some agents, and it is easy to exclude biased criteria by imposing the following anonymity requirement.

Axiom 5 (Anonymity) *For all $U, U' \in \mathcal{L}$, one has $U I U'$ if U' differs from U only by permuting the vectors U_i .*

With this axiom added to the list, the function e must be symmetrical in its arguments.

The classical utilitarian criterion is a special case of this, and corresponds to the situation in which $e(U^s)$ equals the average utility in U^s . But the above result is also compatible with incorporating inequality aversion into R . Let us introduce an adaptation of the Pigou-Dalton principle to this setting. The original Pigou-Dalton principle is about inequality reduction between two agents with unequal utilities (or incomes, in the standard framework of inequality studies). We apply the same idea when one agent has greater utility in all states of nature, and inequality reduction is performed in every state of nature.

Axiom 6 (Pigou-Dalton Equity) *For all $U, U' \in \mathcal{L}$ such that for some $i, j \in N$, some $\delta \gg 0$,*

$$U'_i - \delta = U_i \geq U_j = U'_j + \delta$$

while $U_k = U'_k$ for all $k \neq i, j$, one has $U R U'$.

Contrary to the classical Pigou-Dalton principle, it is not required here that the transfer always strictly improves the lottery. This weaker version is more reasonable.¹¹

¹¹It is also more similar to the Hammond Equity condition which is introduced later, and which was formulated by Hammond (1976, 1979) in terms of weak preference.

Indeed, it is questionable that among very rich people, a Pigou-Dalton transfer strictly improves the situation. In particular, the maximin criterion which focuses on the worst-off only does not satisfy the strict version of the Pigou-Dalton principle because it is indifferent to transfers among the better-off, but it satisfies this weaker version.

Corollary 1 *R satisfies the six axioms if and only if, for all $U, U' \in \mathcal{L}$,*

$$U R U' \Leftrightarrow \sum_{s \in S} \pi^s e(U^s) \geq \sum_{s \in S} \pi^s e(U'^s),$$

where $e(U^s)$ is defined as above and is Schur-concave.¹²

This result, however, is compatible with an arbitrarily weak, even zero, aversion to inequality (with the strict version of the axiom one would obtain that e must be *strictly* Schur-concave but that is still compatible with an arbitrarily weak positive aversion to inequality). One may want to require a greater concern for equality. In particular, one can adapt the axiom of Hammond equity (Hammond 1976) to this case.

Axiom 7 (Hammond Equity) *For all $U, U' \in \mathcal{L}$ such that for some $i, j \in N$,*

$$U'_i \gg U_i \geq U_j \gg U'_j$$

while $U_k = U'_k$ for all $k \neq i, j$, one has $U R U'$.

This axiom is logically stronger than Pigou-Dalton Equity and requires an infinite inequality aversion in two-person situations. With this axiom substituted to Pigou-Dalton Equity, the ex-post maximin criterion, which computes the expected value of the smallest utility, is singled out.

¹²The function e is Schur-concave if $e(AU^s) \geq e(U^s)$ for all $n \times n$ non-negative matrices A such that $\sum_{k=1}^n A_{kj} = \sum_{k=1}^n A_{ik} = 1$ for all i, j (i.e. bistochastic matrices). It is strictly Schur-concave if $e(AU^s) > e(U^s)$ for all matrices A satisfying the above conditions and additionally $A_{ij} \neq 0, 1$ for some i, j (i.e. bistochastic matrices which are not permutation matrices).

Corollary 2 *R satisfies the first four axioms and Hammond Equity if and only if, for all $U, U' \in \mathcal{L}$,*

$$URU' \Leftrightarrow \sum_{s \in S} \pi^s \min_{i \in N} U_i^s \geq \sum_{s \in S} \pi^s \min_{i \in N} U_i'^s.$$

A drawback of the maximin criterion is that it does not satisfy the following stronger monotonicity condition.

Axiom 8 (Strong Monotonicity) *For all $U, U' \in \mathcal{L}$, one has UPU' if $U \geq U'$ and $U_i \gg U'_i$ for some $i \in N$.*

This axiom is satisfied, along with the others except Continuity, by the ex-post leximin criterion, defined as follows. Let $U_{(i)}^s$ denote the utility of i th rank (by increasing order) in vector U^s . The symbol \geq_{lex} denotes the ordinary leximin criterion, which compares two vectors by comparing the smallest component, and if they are equal it compares the second smallest component, and so on. The ex-post leximin criterion weakly prefers U to U' if

$$\left(\sum_{s \in S} \pi^s U_{(i)}^s \right)_{i \in N} \geq_{lex} \left(\sum_{s \in S} \pi^s U'_{(i)}^s \right)_{i \in N}$$

In other words, this criterion computes the expected value of the utility of i th rank in U , for all $i = 1, \dots, n$, and applies the standard leximin criterion to such vectors.

The ex-post leximin criterion is not the only one which satisfies Omniscience Principle, Pareto for Macro Risks, Hammond Equity and Strong Monotonicity. This list is also satisfied, for instance, by the criterion which is defined by

$$(z_i(U))_{i \in N} \geq_{lex} (z_i(U'))_{i \in N},$$

where

$$\begin{aligned} z_1(U) &= \sum_{s \in S} \pi^s U_{(1)}^s, \\ z_{i+1}(U) &= z_i(U) + \min_{s \in S} (U_{(i+1)}^s - U_{(i)}^s) \text{ for } i = 1, \dots, n-1. \end{aligned}$$

One sees that this criterion does not rely on π in order to evaluate the situation of the agents who are not the worst-off. This is rather questionable in situations where the evaluator has no better data than π , on the basis of the statistical distribution of utilities, in order to make his judgment. It would make sense to extend the application of the Pareto principle to situations in which the agents may have unequal utilities but are always ranked in the same way whatever the state of nature. Such situations may contain micro risks which are insurable, but they do not contain reversals of utility rankings in different states of nature. In absence of such reversals, one can never be sure of what an omniscient and impartial evaluator would say, and relying on the agents' expected utilities in these cases is fully compatible with the Omniscience Principle axiom and with inequality aversion.

Let us say that U is comonotonic if for all $s, s' \in S$ and all $i, j \in N$, $(U_i^s - U_j^s)(U_i^{s'} - U_j^{s'}) \geq 0$.

Axiom 9 (Pareto for Comonotonic Risks) *For all $U, U' \in \mathcal{L}$ such that U and U' are comonotonic, one has $U R U'$ if for all $i \in N$, $\sum_{s \in S} \pi^s U_i^s \geq \sum_{s \in S} \pi^s U_i'^s$.*

This axiom is satisfied by the ex-post maximin and leximin criteria, as well as by many other ex-post criteria. In combination with the other axioms, however, strenghtening Pareto for Macro Risks into Pareto for Comonotonic Risks forces us to adopt the ex-post leximin criterion.

Theorem 2 *R satisfies Omniscience Principle, Pareto for Comonotonic Risks, Hammond Equity, Strong Monotonicity and Anonymity if and only if it is the ex-post leximin criterion.*

4 A variant of Harsanyi's theorem

In his famous aggregation theorem, Harsanyi (1955) directly assumed that the evaluator seeks to maximize the expected value of social welfare. This assumption implies satisfaction of the Omniscience Principle axiom, but is much stronger. For instance, the double leximin criterion which simply applies the leximin criterion to the matrix $(U_i^s)_{i \in N, s \in S}$ (viewed as a nm -vector), satisfies Omniscience Principle without involving the expected value of social welfare, and without even being continuous. In this section we show that, in the framework of this paper, Harsanyi's result still holds if his expected social welfare assumption is weakened into Omniscience Principle.

Let us extend the Pareto requirement further in order to capture Harsanyi's requirement of respecting agents' ex-ante preferences. Moreover, in order to espouse Harsanyi's reasoning more closely, we retain the strong version of the Pareto condition, such that if the expected utility of at least one agent increases (and no one else's decreases), the resulting lottery is strictly better.

Axiom 10 (Strong Pareto) *For all $U, U' \in \mathcal{L}$, one has URU' if for all $i \in N$, $\sum_{s \in S} \pi^s U_i^s \geq \sum_{s \in S} \pi^s U_i'^s$; one has UPU' if, in addition, for some $i \in N$, $\sum_{s \in S} \pi^s U_i^s > \sum_{s \in S} \pi^s U_i'^s$.*

We have the following result, which characterizes the classical utilitarian criterion.

Theorem 3 *R satisfies Omniscience Principle, Strong Pareto and Anonymity if and only if, for all $U, U' \in \mathcal{L}$,*

$$URU' \Leftrightarrow \sum_{\substack{s \in S \\ i \in N}} \pi^s U_i^s \geq \sum_{\substack{s \in S \\ i \in N}} \pi^s U_i'^s.$$

The proof of this result takes a very different route from Harsanyi's theorem since there is no continuity and one cannot even use the intermediate argument that the ordering

must be an ex-post criterion (which, by Theorem 1, would be true under Continuity). While Harsanyi's result relies on separability arguments (such as Gorman's theorem),¹³ the above theorem relies on the observation that the conjunction of Omniscience Principle and Anonymity makes it acceptable, by permuting utilities across agents in certain states of nature, to transfer expected utility across them at will. This is what triggers the absence of inequality aversion. As a simple illustration of this reasoning, consider the following lotteries, in a two-agent, two-state context, with equiprobable states (each agent has a row, each state has a column):

$$U = \begin{pmatrix} 2 & 0 \\ 0 & 2 \end{pmatrix}, V = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}, W = \begin{pmatrix} 2 & 2 \\ 0 & 0 \end{pmatrix}, X = \begin{pmatrix} 0 & 0 \\ 2 & 2 \end{pmatrix}.$$

By Strong Pareto, $U I V$. By Anonymity, $W I X$, so that by Omniscience Principle, one can replace the second column of W , $\begin{pmatrix} 2 \\ 0 \end{pmatrix}$, by $\begin{pmatrix} 0 \\ 2 \end{pmatrix}$, and obtain that $W I U$. By transitivity, $W I V$, which shows that no inequality aversion is possible.

This theorem is problematic for those who would also like R to display aversion to inequality. It rules out inequality aversion with respect to utilities (although it is of course compatible with inequality aversion with respect to the arguments of individual utility functions, if they can be concave). But the weak pillar of this result is the Strong Pareto axiom. As explained in Section 2, the Pareto requirement is safe when applied to macro risks, possibly comonotonic risks, but much more questionable when applied across the board as in this last result. Further elaboration on this point is made in the next section, which discusses the literature.

¹³See in particular Hammond (1981) for a version of the theorem in the framework with state-contingent alternatives. In Harsanyi (1955) and Weymark (1991), the theorem is formulated in the von Neumann-Morgenstern framework where lotteries are instead described by probabilities over a fixed vector of gains.

5 Discussion of the literature

The ex-post egalitarian approach defended in this paper goes against the branch of the literature which accepts Diamond’s (1967) critique of Harsanyi as well as against the branch which accepts the Pareto principle applied to ex-ante prospects. Let us examine where disagreements occur exactly. Afterwards, we will briefly discuss other problems raised in the literature in relation to the ex-ante-ex-post divide.

Diamond’s critique revisited. Diamond argues that, since it is obviously more fair to toss a coin than to give a prize directly to one individual, one should prefer L_2 to L_1 (recalled from the introduction):

L_1	H	T	L_2	H	T
Ann	1	1	Ann	1	0
Bob	0	0	Bob	0	1

This argument implies that the social criterion should not maximize the expected value of an anonymous social welfare function.

Hammond (1983) provided a defense of the expected utility approach for the social criterion, in terms of dynamic consistency. Diamond’s argument, indeed, seems likely to produce dynamically inconsistent decisions. After a coin is tossed, one would still be in the same situation when considering another coin tossing: It would always be “more fair” to toss an additional coin, and that seems absurd.¹⁴ However, Hammond defended the ex-post approach by combining the requirement of dynamic consistency with an independence axiom which requires decisions to depend only on future possibilities. Epstein and Segal (1992) showed that one can reconcile ex-ante egalitarianism and dynamic consistency by letting decisions depend on the past, i.e. by dropping Hammond’s independence axiom. For instance, whether flipping a coin appears necessary or not may depend on whether a

¹⁴See also Myerson (1981), Broome (1984, 1991) and Hammond (1988, 1989).

previous coin flipping has already taken place. In summary, Hammond's defense of the ex-post approach does not appear fully successful.¹⁵

But Diamond's reasoning is not only incompatible with expected social welfare, it is actually incompatible with the much weaker Omniscience Principle axiom (under Anonymity), and this suggests another way to rebut Diamond's critique. Would it make sense to consider that, even if a lottery produces a worse situation in every state of nature, it can be better overall? This is what Diamond's reasoning would entail, under Anonymity and Continuity. Indeed, if we strictly prefer L_2 to L_1 , then we still prefer it to L'_1 for a sufficiently small ε :

L'_1	H	T
Ann	$1+\varepsilon$	$1+\varepsilon$
Bob	ε	ε

L_2	H	T
Ann	1	0
Bob	0	1

This appears to be a much worse violation of rationality than a failure to maximize expected social welfare. All ex-ante egalitarian criteria suffer from this problem. For instance, Epstein and Segal's approach, which involves quadratic social welfare functions applied to individual expected utilities, also implies this violation of rationality.

These considerations still leave us with the overwhelming intuition that it is more fair to toss a coin. Observe again, however, that if, in L_1 , Ann has been chosen by an anterior procedure of coin tossing, L_1 is no less fair than L_2 , and it becomes very plausible to argue that the two lotteries are equivalent.¹⁶ Therefore, if we remove any difference in

¹⁵As explained in Hammond (1989), the independence axiom can be saved by enriching the description of consequences, incorporating information about the history from which they emerge. But when this is done the axiom loses its bite, a classical dilemma for independence axioms in decision theory.

¹⁶As noted in Epstein and Segal (1992, p. 704), it would be unfair to toss again only when Ann wins in the first round. But there is no unfairness in tossing the coin twice and disregarding the outcome of the first round. It appears obviously equivalent to tossing the coin only once, even when the evaluation is made after the first tossing, which justifies indifference between L_1 and L_2 .

the degree of fairness of the two lotteries, Diamond's critique falls apart.

Diamond's critique appears compelling only when one ignores that there are ways in which the fairness of tossing a coin can be registered in the ex-post description of a social situation. There are at least four ways to do it. First, fairness can be an external feature of a lottery. The ex-post criterion would then measure social welfare ex-post as a function of individual utilities and of the fairness of the situation. Viewed in this light, the analysis of the previous sections is valid only for lotteries which are equivalent at the bar of fairness. An extension of the framework would be needed in order to compare lotteries with different degrees of fairness. A second way of recording fairness consists, as proposed by Broome (1991), in measuring unfairness as a harm inflicted to the victims. In L_1 , Bob's utility might then be below zero due to the unfairness he suffers. Any reasonable ex-post criterion would prefer L_2 to this corrected L_1 . A third way in which tossing a coin can be valued is when it has welfare effects over the period which precedes the outcome. If individual utility in the first period depends on the expected utility of the second period, then lifetime utilities are more equal when a coin is tossed, because the utilities in the first period are more equal. Even a purely ex-post criterion would register this reduced inequality.¹⁷ A fourth way in which fairness can be registered by an ex-post approach is simply by the fact that if a systematic bias in favor of Ann and against Bob is enacted throughout their lives, this will typically entail inequality between them ex post. If no inequality ex post follows, it means that Ann is naturally handicapped and the bias in its favor has simply restored fairness at a broader level. In this case, tossing a coin is actually less fair. At any rate, an ex-post criterion does not seem to be wanting in this respect.

Ex-ante Pareto. The literature contains surprisingly few explicit arguments sup-

¹⁷See Deschamps and Gevers (1979).

porting the application of the Pareto principle to individual ex-ante preferences over prospects. Harsanyi (1955, 1977) simply considers it a “very weak” individualistic postulate embodying the principle of “consumer sovereignty”. Hammond (1981) clearly warns that the context of risk makes the Pareto principle more demanding, and identifies three problems: 1) individuals may misperceive probabilities; 2) their attitude to risk may be questionable; 3) one may be worried about ex-post inequalities. However, he later excludes the third problem on the ground that “there is no reason to take account of any possible correlation between different individuals’ personal consequences.” (Hammond 1996, p. 107). Broome (1991) devotes more attention to the conflict between the Pareto principle and ex-post egalitarianism, and ends up accepting the former on the ground that it is *prima facie* appealing and that the best way of recording ex-post inequalities is to measure the unfairness they entail as an individual harm. In this way, even ex-post inequalities do appear in individual ex-ante prospects, and are taken into account by an ordinary utilitarian criterion applied to individual expected utilities. This solution appears to salvage utilitarianism and the ex-ante Pareto principle in a purely formal way, while surrendering to the substance of the ex-post egalitarian challenge. Moreover, it appears very artificial to measure individual welfare as a function of social welfare, and in particular it is questionable that an improvement in the well-off’s welfare should be viewed as a direct harm to the less well-off.

In Section 2 we have already argued that, because individual ex-ante preferences may clash with ex-post preferences, which are, by definition, better informed, the Pareto principle is not so appealing. Let us hammer this point here, and explain further why the Pareto principle applied to ex-ante prospects is, contrary to what is said in the literature, *prima facie unappealing* and actually betrays the principle of consumer sovereignty.

As recalled above, Hammond (1981) correctly notes that misperceptions of probabil-

ities undermine the Pareto principle. If we consider that we live in a deterministic world, this problem is very serious. In such a world, all probabilistic beliefs are mistaken beliefs. The example of contaminated boxes of corn flakes has already illustrated it. When the consumer buys a contaminated box, he is wrong to believe that this box could be safe. By relying on his ex-ante preferences, we are betraying his consumer sovereignty, because his informed preferences would disagree with his decision. As another example, imagine that we rely on coin flipping in order to allocate a prize. Suppose the coin has been tossed long before the problem of allocating the prize arose. It is then clear that the loser is wrong to believe that he has a chance of getting the prize. Now, whether the coin is tossed before or after the allocation procedure is launched clearly does not make a difference if it is a deterministic coin.

One might think that the situation is different in an indeterministic world. Suppose the coin is truly random and there is no way to determine its outcome, even in full knowledge of the involved mechanics. Then the agents can have “correct” probabilities. But this does not make a difference for social evaluations. First, the fairness of the procedure is not affected by it. When we toss a coin, do we worry about the metaphysics of determinism? The only thing that matters for the fairness of the procedure is that the operating agents all have a subjective probability of one half, not that the true probability is one half. Second, our ignorance about the true state of nature is no less an ignorance when the world is indeterministic. We can still imagine an omniscient alter ego who already knows the outcome (she reads not only the mechanics of the world but also the future), and observe that, like her, we already know that there will be only one winner. Knowing that there will be only one winner is the only relevant information, and the fact that the mechanism picking the winner is truly or only epistemically random does not make a difference. (It is interesting to note that, even if the world is truly indeterministic,

we are sometimes in a position that is equivalent to our omniscient alter ego's, because like her we have all the relevant information that is needed.)

Another line of defense of the ex-ante Pareto principle must be discussed. As mentioned in Hammond (1983) and Kolm (1998), one may associate the principle of respecting individual preferences over risk with the idea that individuals have the right to take risks and should be able to assume the consequences. But this defense appears very weak when one sees risk as a situation of ignorance. Why should individuals bear the consequences of a lack of information for which they are not responsible?

In conclusion, understanding risk as an instance of ignorance makes the ex-ante Pareto principle radically suspect. It is not a respect of consumer sovereignty to cater to individual ignorance.

Risk equity versus catastrophe avoidance. A problem related to the ex-ante-ex-post divide was raised by Keeney (1980a), who argued that for public decisions which submit individuals $1, \dots, n$ to independent risks of death (p_1, \dots, p_n) , “risk equity” requires a preference for a more equal over a less equal distribution of probabilities (for a fixed sum of probabilities $p_1 + \dots + p_n$, which is equal to the expected number of fatalities). He also claimed that many people express a preference for “catastrophe avoidance”, meaning that for a given expected number of fatalities, one prefers a smaller number of actual fatalities (even if they must be, logically, more probable).

Now, in this simple context, an impartial ex-post criterion simply has to define preferences over lotteries $(\pi_f)_{f=1, \dots, n}$ where π_f is the probability of having f fatalities. Keeney (1980a) then shows that, if the social criterion is an expected utility $\sum_{f=1}^n \pi_f u(f)$, the utility function must be strictly convex¹⁸ (risk prone) under the equity requirement, but

¹⁸Strict convexity for a function taking only n values is understood here as

$$u(f) < \frac{1}{2} [u(f-1) + u(f+1)]$$

strictly concave (risk averse) under catastrophe avoidance.¹⁹

Which one is correct? Risk equity over independent risks is obviously satisfied by an ex-ante egalitarian criterion, but also, less obviously, by an ex-post egalitarian criterion, because a more equal probability of death increases the expected value of social welfare. For instance, the ex-post maximin criterion maximizes the probability that nobody dies (because all other situations are identical: the worst-off dies), and this probability equals $(1 - p_1) \dots (1 - p_n)$, which, viewed as a criterion over (p_1, \dots, p_n) , is inequality averse.

What about preference for catastrophe avoidance? It implies a puzzling preference for ex-post inequalities. It may be that the intuition for catastrophe avoidance relies on an illusion. Fearing large numbers of fatalities appears in fact somewhat paradoxical if one would not similarly want to avoid joint losses of smaller magnitude, or joint gains. It appears obvious, by equity, that a $1/n$ chance that n individuals all gain \$10 is preferable to the certainty of one individual alone gaining \$10. By symmetry, one should prefer a $1/n$ risk that n individuals lose \$10 to the certainty of one individual losing \$10. Now, dying prematurely is just a loss of bigger magnitude, so that one should presumably uphold the same judgment in the case of fatalities.

Nonetheless, there are two ways of making sense of catastrophe avoidance. In the case of a big loss, one may abandon inequality aversion and adopt “triage” preferences that seek to minimize the number of individuals suffering losses.²⁰ For instance, one might consider it less bad if one individual dies painfully than if two individuals die without pain. With such social preferences, it is indeed possible to combine preference for catastrophe

for all $f = 2, \dots, n - 1$. The reverse inequality defines strict concavity.

¹⁹The related literature includes Keeney (1980b), Broome (1982), Fishburn (1984), Keeney and Winkler (1985) and Fishburn and Straffin (1989), Bommier and Zuber (2006).

²⁰An example of such social preferences is the poverty head count, which accepts worsening the situation of the remaining poor if this reduces the number of poor. This kind of social preferences is examined in Roemer (2004).

avoidance in the case of large losses such as fatalities and aversion to ex-post inequality in the case of gains and small losses.

Another way of putting catastrophe avoidance in positive light is to take account of the fact that fatalities may not only consist in the premature death of some individuals but may also prevent the existence of other individuals (their descendants). Suppose that the death of one individual prevents the existence of q other individuals, but that the death of n individuals prevents many more than qn individuals from coming into existence. For instance, if n is the size of the whole population it may be that human life is wiped out entirely. In such a context it may be reasonable to prefer one fatality for sure to a $1/n$ risk of n fatalities.

In summary, preference for catastrophe avoidance should not be adopted as an axiom, and can only appear as the result of special “triage” preferences in the case of big losses, or come up naturally as the result of taking account of exponential externalities in the case of certain hazards.

Rescue versus prevention. Schelling (1968) has made the striking observation that people are generally much more willing to rescue identified people currently at great risk than to invest in the prevention of future uncertain fatalities, even if the expected number of lives saved is greater with the prevention policy. This apparently irrational preference has been analyzed in many ways, such as the difference between known victims and “statistical lives”, the symbolic expression of our commitment to life by undertaking rescue operations, or preferences about how to die.²¹

Although all of these considerations are relevant in order to understand the full force of the appeal of rescue operations in concrete examples, it seems possible to understand this as, more than anything else, another instance of the ex-ante vs. ex-post dilemma.

²¹For a penetrating and synthetic analysis of this issue, see Fried (1969).

An ex-post approach would minimize the number of fatalities and would favor the prevention policy if, as generally assumed in these examples, this policy is more efficient in this respect. But ex-ante, the rescue operation looks preferable because it reduces the inequality of expected utility between the victims and the rest of the population.

However, even if one sticks to a pure ex-post approach, there is a way to make sense of the apparent duty of rescue that many people feel compelling. There is a difference between victims facing a sure prospect of death if no rescue is undertaken and other people facing a slightly greater risk in absence of prevention, and this difference can be described in terms of ex-post distributions. With the rescue operation there is still a positive probability that everybody will survive, whereas the prevention policy condemns the victims for sure. The ex-post maximin criterion, if it is applied by maximizing the probability of having no fatality, therefore favors the rescue operation. This may also be true when the victims are not condemned to death in absence of rescue.

In order to illustrate this, let us consider an example with independent risks, with m victims facing a greater probability of death \bar{p} than the rest of the n individuals, for whom it is $\underline{p} < \bar{p}$. Suppose that rescue and prevention reduce the expected number of fatalities by the same number Δ .²² We have the choice between the rescue lottery

$$\left(\underbrace{\bar{p} - \Delta/m}_m, \underbrace{\underline{p}}_{n-m} \right)$$

and the prevention lottery

$$\left(\underbrace{\bar{p}}_m, \underbrace{\underline{p} - \Delta/(n-m)}_{n-m} \right),$$

with $\bar{p} - \Delta/m \geq \underline{p}$. The probability of no fatality is, respectively,

$$\left(1 - \bar{p} + \frac{\Delta}{m} \right)^m (1 - \underline{p})^{n-m} \quad \text{and} \quad (1 - \bar{p})^m \left(1 - \underline{p} + \frac{\Delta}{n-m} \right)^{n-m}.$$

²²If we obtain strict preference for rescue when the two policies are equally efficient, the preference will remain when prevention is more efficient at reducing expected fatalities and the gap is not too large.

If one maximizes this probability, one prefers the rescue lottery if $\bar{p} = 1$ (the victims are condemned in absence of rescue), or if $\bar{p} < 1$ and

$$\left(1 + \frac{\Delta}{(1 - \bar{p})m}\right)^m > \left(1 + \frac{\Delta}{(1 - \underline{p})(n - m)}\right)^{n-m}.$$

This is satisfied for all n such that $n - m > m$ if

$$1 + \frac{\Delta}{(1 - \bar{p})m} > e^{\frac{\Delta}{(1 - \underline{p})m}}.$$

To fix ideas, let us assume that victims are saved with the rescue operation, in the sense that $\bar{p} - \Delta/m = \underline{p}$ (after rescue, they face the same risks as the rest of the population).

The above inequality then reads

$$1 + \frac{\bar{p} - \underline{p}}{(1 - \bar{p})} > e^{\frac{\bar{p} - \underline{p}}{1 - \underline{p}}},$$

which holds true whenever $\bar{p} > \underline{p}$. This shows that even a pure ex-post approach, if it is egalitarian, can rationalize the preference for rescue. This is not, in essence, different from Keeney's observation that an ex-post egalitarian criterion can obey "risk equity" when the risks are independent.

6 Conclusion

The expected value of the equal-equivalent utility yields a criterion which satisfies the omniscience principle and the Pareto principle restricted to macro risks. The omniscience principle excludes ex-ante egalitarian social welfare functions, but we have seen that concerns for ex-ante fairness can be accommodated to a substantial, probably sufficient, extent by an ex-post criterion. The utilitarian criterion does obey the omniscience principle, and combining this principle with a full-fledged version of the Pareto principle imposes embracing utilitarianism (under mild monotonicity and anonymity conditions). But we

have seen that the omniscience principle actually weakens the case for the Pareto principle because ex-ante preferences at the individual level ignore correlations, a crucial information for anyone wanting to mimic an omniscient evaluator. Therefore we may conclude that ex-post egalitarianism emerges from this analysis as the most plausible approach to evaluating risky social alternatives.

References

- Ben-Porath E., I. Gilboa, D. Schmeidler 1997, "On the measurement of inequality under uncertainty", *Journal of Economic Theory* 75: 194-204.
- Blackorby C., W. Bossert, D. Donaldson 2005, *Population Issues in Social-Choice Theory, Welfare Economics and Ethics*, Cambridge: Cambridge University Press.
- Bommier A., S. Zuber 2006, "Can preferences for catastrophe avoidance reconcile social discounting with intergenerational equity?", mimeo.
- Broome J. 1982, "Equity in risk bearing", *Operations Research* 30: 412-414.
- Broome J. 1984, "Uncertainty and fairness", *Economic Journal* 94: 624-632.
- Broome J. 1991, *Weighing Goods. Equality, Uncertainty and Time*, Oxford: Blackwell.
- Broome J. 1999, *Ethics Out of Economics*, Cambridge: Cambridge University Press.
- Broome J. 2004, *Weighing Lives*, Oxford: Oxford University Press.
- Debreu G. 1959, *Theory of Value*, New York: Wiley.
- Deschamps R., L. Gevers 1979, "Separability, risk-bearing and social welfare judgments", in J.J. Laffont ed., *Aggregation and Revelation of Preferences*, Amsterdam: North-Holland.
- Diamond P.A. 1967, "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility: Comment", *Journal of Political Economy* 75: 765-766.
- Epstein L.G., U. Segal 1992, "Quadratic social welfare functions", *Journal of Political*

- Economy* 100: 691-712.
- Fishburn P.C. 1984, "Equity axioms for public risks", *Operations Research* 32: 901-908.
- Fishburn P.C., P.D. Straffin 1989, "Equity considerations in public risks evaluation", *Operations Research* 37: 229-239.
- Fried C. 1969, "The value of life", *Harvard Law Review* 82: 1415-1437.
- Gajdos T., E. Maurin 2004, "Unequal uncertainties and uncertain inequalities: An axiomatic approach", *Journal of Economic Theory* 116: 93-118.
- Hammond P.J. 1976, "Equity, Arrow's conditions, and Rawls' difference principle", *Econometrica* 44: 793-804.
- Hammond P.J. 1979, "Equity in two-person situations", *Econometrica* 47: 1127-1136.
- Hammond P.J. 1981, "Ex-ante and ex-post welfare optimality under uncertainty", *Economica* 48: 235-250.
- Hammond P.J. 1983, "Ex post optimality as a dynamically consistent objective for collective choice under uncertainty", in P.K. Pattanaik, M. Salles eds., *Social Choice and Welfare*, Amsterdam: North-Holland.
- Hammond P.J. 1988, "Consequentialist foundations for expected utility", *Theory and Decision* 25: 25-78.
- Hammond P.J. 1989, "Consistent plans, consequentialism, and expected utility", *Econometrica* 57: 1445-1449.
- Hammond P.J. 1996, "Consequentialist decision theory and utilitarian ethics", in F. Farina, F. Hahn, S. Vannucci eds, *Ethics, Rationality, and Economic Behaviour*, Oxford: Clarendon Press.
- Harel A., Z. Safra, U. Segal 2005, "Ex-post egalitarianism and legal justice", *Journal of Law, Economics and Organization* 21: 57-75.
- Hardy G.H., Littlewood, J.E. and Pólya, G. 1952, *Inequalities*, Cambridge: Cambridge

University Press.

Harsanyi J.C. 1955, "Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility", *Journal of Political Economy* 63: 309-321.

Keeney R.L. 1980a, "Equity and public risk", *Operations Research* 28: 527-534.

Keeney R.L. 1980b, "Utility functions for equity and public risk", *Management Science* 26: 345-353.

Keeney R.L., R.L. Winkler 1985, "Equity and public risk", *Operations Research* 33: 955-970.

Kolm S.C. 1998, "Chance and justice: Social policies and the Harsanyi-Vickrey-Rawls problem", *European Economic Review* 42: 1393-1416.

Marshall A., I. Olkin 1979, *Inequalities : Theory of Majorization and Its Applications*, New York: Academic Press.

Meyer M.A., D. Mookherjee 1987, "Incentives, compensation, and social welfare", *Review of Economic Studies* 54: 209-226.

Mongin P. 2005, "Spurious unanimity and the Pareto principle", LSE-CPNSS Working Paper, vol. 1, no. 5.

Myerson R.B. 1981, "Utilitarianism, egalitarianism, and the timing effect in social choice problems", *Econometrica* 49: 883-897.

Pratt J.W., R.J. Zeckhauser 1996, "Willingness to pay and the distribution of risk and wealth", *Journal of Political Economy* 104: 747-763.

Roemer J.E. 2004, "Eclectic social preferences", *Politics, Philosophy and Economics* 3: 267-281.

Schelling T.C. 1968, "The life you save may be your own", in S.B. Chase, Jr ed., *Problems in Public Expenditure Analysis*, Washington, D.C.: Brookings Inst.

Viscusi W.K. 1992, *Fatal Tradeoffs: Private and Public Responsibilities for Risk*, Oxford:

Oxford University Press.

Weymark J.A. 1991, “A reconsideration of the Harsanyi-Sen debate on utilitarianism”, in J. Elster, J.E. Roemer ed., *Interpersonal Comparisons of Well-Being*, Cambridge: Cambridge University Press.

Appendix A: Proofs

Proof of Theorem 1: By Weak Monotonicity and Continuity, there is always a solution e to the equation

$$[U^s] I [(e, \dots, e)],$$

and therefore the function $e(U^s)$ is well-defined. By Weak Monotonicity, $e(U^s) > e(U'^s)$ whenever $U^s \gg U'^s$. By Continuity, e is continuous.

Take any lotteries $U, U' \in \mathcal{L}$ such that

$$\sum_{s \in S} \pi^s e(U^s) = \sum_{s \in S} \pi^s e(U'^s).$$

By application of Omniscience Principle, and the fact that by definition,

$$[U^s] I [(e(U^s), \dots, e(U^s))],$$

one can replace each U^s by $(e(U^s), \dots, e(U^s))$. Let \bar{U} denote the resulting lottery in which, by construction, $\bar{U}_i^s = \bar{U}_j^s$ for all $s \in S$ and all $i, j \in N$. One has $U I \bar{U}$. Similarly one can replace each U'^s by $(e(U'^s), \dots, e(U'^s))$ and construct \bar{U}' such that $U' I \bar{U}'$. By Pareto for Macro Risks, $\bar{U} I \bar{U}'$. By transitivity, $U I U'$.

Consider now the case

$$\sum_{s \in S} \pi^s e(U^s) > \sum_{s \in S} \pi^s e(U'^s).$$

By the continuity and monotonicity properties of e proved above, there is $V \in \mathcal{L}$ such that $V \gg U'$ and

$$\sum_{s \in S} \pi^s e(U^s) = \sum_{s \in S} \pi^s e(V^s).$$

By the previous step, $U I V$ and by Weak Monotonicity, $V P U'$. By transitivity, $U P U'$.

We now check that no axiom is redundant.

Drop Omniscience Principle. Take the maximin criterion applied to the vector of expected utilities.

Drop Pareto for Macro Risks. Take a criterion based on $\sum_{s \in S} \hat{\pi}^s e(U^s)$ for some arbitrary $\hat{\pi}$ that differs from π .

Drop Weak Monotonicity. Take universal indifference.

Drop Continuity. Take the ex-post leximin criterion.

Proof of Corollary 1: By Theorem 1, we know that for all $U, U' \in \mathcal{L}$,

$$U R U' \Leftrightarrow \sum_{s \in S} \pi^s e(U^s) \geq \sum_{s \in S} \pi^s e(U'^s).$$

Pigou-Dalton Equity and Anonymity imply that $e(\cdot)$ must satisfy Pigou-Dalton Equity and Anonymity on U^s vectors. This implies that it must be Schur-concave (see Marshall and Olkin 1979, pp. 22, 54). Conversely, if $e(\cdot)$ is Schur-concave (and therefore symmetric), then R satisfies Pigou-Dalton Equity and Anonymity.

Let us check that no axiom is redundant. For the first four axioms, the counterexamples are the same as for Theorem 1.

Drop Anonymity. Take a function $e(\cdot)$ that satisfies Pigou-Dalton Equity but is not symmetric, such as:

$$e(u_1, \dots, u_n) = \frac{1}{|\{i \in N \mid u_i \leq u_1\}|} \sum_{\substack{i \in N \\ u_i \leq u_1}} u_i.$$

Drop Pigou-Dalton Equity. Take a strictly Schur-convex $e(\cdot)$.

Proof of Corollary 2: By Theorem 1, we know that for all $U, U' \in \mathcal{L}$,

$$U R U' \Leftrightarrow \sum_{s \in S} \pi^s e(U^s) \geq \sum_{s \in S} \pi^s e(U'^s).$$

Hammond Equity implies that $e(\cdot)$ must satisfy Hammond Equity on vectors U^s . Take two vectors U^s and V^s such that $\min_{i \in N} U^s < \min_{i \in N} V^s$. Construct $U'^s \gg U^s$ such that for some $i_0 \in N$ and for all $i \neq i_0$,

$$\begin{aligned} \min_{i \in N} U^s &< U'_{i_0} < U'_{i^s}, \\ U'_{i_0} &< \min_{i \in N} V^s, \end{aligned}$$

and $V'^s \ll V^s$ such that for all $i \in N$,

$$U'_{i_0} < V'_i < \min_{i \in N} V^s.$$

By repeated applications of Hammond Equity between i_0 and the other agents, one can construct U''^s from U'^s , such that for some $i_0 \in N$ and for all $i \neq i_0$,

$$U'_{i_0} < U''_{i_0} < U''_i < V'_i.$$

By Hammond Equity, $e(U''^s) \geq e(U'^s)$. By Weak Monotonicity, $e(U^s) < e(U'^s)$, $e(U''^s) < e(V'^s) < e(V^s)$. Therefore $e(U^s) < e(V^s)$.

By Continuity, it then follows that for all U^s, V^s , $e(U^s) \leq e(V^s)$ if and only if $\min_{i \in N} U^s \leq \min_{i \in N} V^s$.

Take any U^s and define $V^s = (\min_{i \in N} U^s, \dots, \min_{i \in N} U^s)$. One has $e(U^s) = e(V^s)$, and since V^s is egalitarian, $e(V^s) = V_1^s = \dots = V_n^s = \min_{i \in N} U^s$. Therefore $e(U^s) = \min_{i \in N} U^s$.

Let us check that no axiom is redundant. For the first four axioms, the counterexamples are the same as for Theorem 1.

Drop Hammond Equity. Take classical utilitarianism.

Proof of Theorem 2: Consider any lottery U . Let V be defined by $V_i^s = U_{(i)}^s$ for all $s \in S$ and all $i \in N$. By Anonymity, for all $s \in S$, $[U^s] I [V^s]$. Therefore, by Omniscience Principle, $U I V$.

Notice that V is comonotonic. The above argument implies that we can restrict attention to comonotonic lotteries and we do so hereafter. By Pareto for Comonotonic Risks, there is an ordering \tilde{R} over \mathbb{R}^n such that $U R U'$ if and only if

$$\left(\sum_{s \in S} \pi^s U_i^s \right)_{i \in N} \tilde{R} \left(\sum_{s \in S} \pi^s U_i'^s \right)_{i \in N}.$$

By Hammond Equity, Anonymity and Strong Monotonicity, this ordering satisfies the following property: For all $u, v \in \mathbb{R}^n$, if there are $i, j \in N$ such that for all $k \neq i, j$, $u_k = v_k$, then $u \tilde{P} v$ if and only if

$$\begin{aligned} \min \{u_i, u_j\} &> \min \{v_i, v_j\} \text{ or:} \\ \min \{u_i, u_j\} &= \min \{v_i, v_j\} \text{ and } \max \{u_i, u_j\} > \max \{v_i, v_j\}. \end{aligned}$$

(See e.g. the proof of Th. 5 in Hammond 1979 for details about this step.) By Hammond (1979, Th. 4), \tilde{R} is then the leximin ordering.

Let us check that no axiom is redundant.

Drop Omniscience Principle. Take the leximin criterion applied to the vector of expected utilities.

Drop Pareto for Comonotonic Risks. Take the ex-post leximin criterion based on some arbitrary $\hat{\pi}$ that differs from π .

Drop Hammond Equity. Take classical utilitarianism.

Drop Strong Monotonicity. Take R based on the computation of

$$\sum_{s \in S} \pi^s \left(\min_{i \in N} U_i^s - \max_{i \in N} U_i^s \right).$$

Drop Anonymity. Let \succeq_{lexico} denote the lexicographic criterion that compares the first component of vectors, and then the second component, and so on. Take R defined

by: URU' if

$$\begin{aligned} \left(\sum_{s \in S} \pi^s U_{(i)}^s \right)_{i \in N} &>_{lex} \left(\sum_{s \in S} \pi^s U_{(i)}'^s \right)_{i \in N} \quad \text{or:} \\ \left(\sum_{s \in S} \pi^s U_{(i)}^s \right)_{i \in N} &=_{lex} \left(\sum_{s \in S} \pi^s U_{(i)}'^s \right)_{i \in N} \quad \text{and} \quad \left(\sum_{s \in S} \pi^s U_i^s \right)_{i \in N} \geq_{lexico} \left(\sum_{s \in S} \pi^s U_i'^s \right)_{i \in N}. \end{aligned}$$

Proof of Theorem 3: By Strong Pareto, there is an ordering \tilde{R} over \mathbb{R}^n such that URU' if and only if

$$\left(\sum_{s \in S} \pi^s U_i^s \right)_{i \in N} \tilde{R} \left(\sum_{s \in S} \pi^s U_i'^s \right)_{i \in N}.$$

By Omniscience Principle and Anonymity, one can permute U_i^s and U_j^s without changing the value of a lottery. Such a permutation has the effect of adding $\pi^s (U_j^s - U_i^s)$ to $\sum_{s \in S} \pi^s U_i^s$ and of subtracting the same quantity from $\sum_{s \in S} \pi^s U_j^s$. By the universal domain \mathcal{L} , this implies that for all $u, u' \in \mathbb{R}^n$ and for all $a \in \mathbb{R}$, $u \tilde{I} u'$ if u' is obtained from u by transferring a from i to j , for any $i, j \in N$.

By a finite number of such transfers (as proved in Hardy, Littlewood and Polya 1952, pp. 47-48), one can obtain the vector $(\bar{u}, \dots, \bar{u})$, where $\bar{u} = \frac{1}{n} \sum_{i \in N} u_i$, from any $u \in \mathbb{R}^n$. Therefore, for every $u \in \mathbb{R}^n$, $u \tilde{I} (\bar{u}, \dots, \bar{u})$.

By Strong Pareto, one then obtains that for all $u, v \in \mathbb{R}^n$,

$$u \tilde{R} v \Leftrightarrow \frac{1}{n} \sum_{i \in N} u_i \geq \frac{1}{n} \sum_{i \in N} v_i.$$

Let us check that no axiom is redundant.

Drop Omniscience Principle. Take the maximin criterion applied to the vector of expected utilities.

Drop Pareto. Take the ex-post maximin criterion.

Drop Anonymity. Take the criterion based on the expected utility of one particular agent.

Appendix B: Extension

The framework of Section 3 appears in a part of the literature (e.g. Broome 1991, Ben Porath et al. 1997),²³ but another part (e.g. Weymark 1991) distinguishes the physical alternatives and the individual utility functions bearing over such alternatives. We briefly examine here how the above analysis can be adapted to such a framework.

Let a lottery now be described as a matrix $x = (x_i^s)_{i \in N, s \in S} \in \mathcal{X}$, where $x_i^s \in X \subset \mathbb{R}^\ell$ is a vector such as a consumption bundle. Individual i has a VNM function $u_i : X \rightarrow \mathbb{R}$ which appears in the computation of expected utility. This function and its domain are assumed to be such that $u_i(X)$ is an interval of \mathbb{R} .

The axiom of Omniscience Principle is immediately adapted to this new framework.

Axiom 11 (Omniscience Principle) *For all $x, x' \in \mathcal{X}$, one has $x R x'$ if for all $s \in S$, $[x^s] R [x'^s]$.*

The axiom of Pareto for Macro Risks is not so easily adapted. What is a macro risk in this context? Even if $x_i^s = x_j^s$ for all $i, j \in N$, the individuals are not necessarily in an equal situation. This may depend in particular on their preferences over the dimensions of \mathbb{R}^ℓ . The same can be said if $u_i(x_i^s) = u_j(x_j^s)$ for all $i, j \in N$, since these utility functions may not capture the correct measure of well-being for interpersonal comparisons.²⁴ We will assume here that the equal-equivalent exists and can serve to check whether the individuals are in an equal situation. It is defined by a collection of functions $e_i : X^n \rightarrow X$ such that $e_i(x^s)$ is a vector of \mathbb{R}^ℓ which, if given to individual i , for all $i \in N$, would yield a situation

²³Harsanyi (1955) adopts a similar framework but describes prospects as probability distributions over possible vectors of individual utilities.

²⁴Whether equality is defined in terms of equality of utilities or in terms of equality of bundles has to do with whether the approach is welfarist or not. This analysis covers both possibilities.

that is really equal and that is socially equivalent to x^s . In any arbitrary allocation, one can then be sure that individuals are in really equal situations if $x_i^s = e_i(x^s)$ for all $i \in N$.

Another issue is raised by the multidimensionality of x_i^s . When there is no risk, it makes sense to rely on individual preferences in order to trade-off the various dimensions of x_i^s , since, in absence of risk, such preferences are well informed. We can therefore extend the notion of “macro risk” to cover situations with no risk. With these two modifications, we obtain the following adaptation of Pareto for Macro Risks.

Axiom 12 (Pareto for Macro Risks) *For all $x, x' \in \mathcal{X}$ such that either for all $s \in S$, for all $i \in N$, $x_i^s = e_i(x^s)$ or for all $s, s' \in S$, $x^s = x'^s$, and similarly for x' , one has $x R x'$ if for all $i \in N$, $\sum_{s \in S} \pi^s u_i(x_i^s) \geq \sum_{s \in S} \pi^s u_i(x_i'^s)$.*

We then have the following adaption of Theorem 1, in which Weak Monotonicity and Continuity are no longer needed because the existence of e has been assumed.

Theorem 4 *R satisfies the two axioms if and only if, for all $x, x' \in \mathcal{X}$,*

$$x R x' \Leftrightarrow \left[\left(u_i^{-1} \left(\sum_{s \in S} \pi^s u_i(e_i(x^s)) \right) \right)_{i \in N} \right] R \left[\left(u_i^{-1} \left(\sum_{s \in S} \pi^s u_i(e_i(x'^s)) \right) \right)_{i \in N} \right],$$

where $u_i^{-1} : u_i(X) \rightarrow X$ is any function satisfying the property: for all $q \in u_i(X)$,

$$q = u_i(u_i^{-1}(q)).$$

Proof. Take any lotteries $x, x' \in \mathcal{X}$. By application of Omniscience Principle, and the fact that by definition,

$$[x^s] I [(e_1(x^s), \dots, e_n(x^s))],$$

one can replace each x^s by $(e_1(x^s), \dots, e_n(x^s))$. Let \bar{x} denote the resulting lottery. One has $x I \bar{x}$. Similarly one can replace each x'^s by $(e_1(x'^s), \dots, e_n(x'^s))$ and construct \bar{x}' such that $x' I \bar{x}'$.

By Pareto for Macro Risks, one can replace all the components of i 's row $(e_i(x^1), \dots, e_i(x^m))$ of \bar{x} by a certain-equivalent $u_i^{-1}(\sum_{s \in S} \pi^s u_i(e_i(x^s)))$. (Since $u_i(X)$ is an interval of \mathbb{R} , $\sum_{s \in S} \pi^s u_i(e_i(x^s)) \in u_i(X)$, which means that there exists $q \in X$ such that $u_i(q) = \sum_{s \in S} \pi^s u_i(e_i(x^s))$.) The resulting lottery, denoted \hat{x} , is no longer egalitarian but is non-risky. Pareto for Macro Risks then implies that $\hat{x} I \bar{x}$. By transitivity, $\hat{x} I x$. One can proceed similarly for x' and obtain a non-risky lottery \hat{x}' such that $\hat{x}' I x'$.

One has $\hat{x} R \hat{x}'$ if and only if $x R x'$, which completes the proof. ■

This is a criterion which can be informally described as focused on the individual certain-equivalent of the equal-equivalent, where by “certain-equivalent” of a given prospect we mean a non-risky prospect which gives an individual the same utility ex-ante as the prospect under consideration.

An alternative formulation of the result is possible, assuming Continuity. A continuous criterion is representable by a function. The Pareto for Macro Risks axiom, applied to non-risky situations, implies that any function that represents the criterion for non-risky situations can be decomposed in the Bergson-Samuelson fashion:

$$[x^s] R [x'^s] \Leftrightarrow W((u_i(x_i^s))_{i \in N}) \geq W((u_i(x_i'^s))_{i \in N}).$$

One can then rewrite the above criterion simply as

$$x R x' \Leftrightarrow W\left(\left(\sum_{s \in S} \pi^s u_i(e_i(x^s))\right)_{i \in N}\right) \geq W\left(\left(\sum_{s \in S} \pi^s u_i(e_i(x'^s))\right)_{i \in N}\right).$$

Besides, note that if the equal-equivalent is defined such that one always has $u_i(e_i(x^s)) = u_j(e_j(x^s))$ for all $i, j \in N$, one can define a function $E(x^s) = u_i(e_i(x^s))$ (for any arbitrary $i \in N$), and the above formulae can be simplified into

$$x R x' \Leftrightarrow \sum_{s \in S} \pi^s E(x^s) \geq \sum_{s \in S} \pi^s E(x'^s),$$

which is reminiscent of Theorem 1.

Analyzing how inequality aversion can be introduced in this framework would require more precision about the description of interpersonal comparisons. Depending on whether the approach is welfarist or not, comparisons would be made in terms of utilities or in terms of bundles (or indifference curves). This analysis need not be developed here, since it is clear that, as above, there is no obstacle to ex-post egalitarianism in this framework. A detailed axiomatic analysis of inequality aversion in this setting will await another occasion.