Utilitarian preferences with multiple priors*

Shiri Alon† and Gabi Gayer‡

February 22, 2015

Abstract

This paper proposes a social preference that aggregates individual preferences in the context of uncertainty. Individuals are assumed to abide by Savage's model of Subjective Expected Utility, each having their own utility and subjective probability. Disagreement on probabilities among individuals' gives rise to uncertainty in the social level, thus society may entertain a set of probabilities instead of just one., and the social preference admits a Maxmin Expected Utility representation. Employing two Pareto-type conditions is then shown to be equivalent to the social utility function being a weighted average of the individuals' utility functions, and the set of social priors containing only weighted averages of the individuals' priors. This conveys an important aspect of the model, whereby no ambiguity is perceived by society when individuals are in full agreement about beliefs.

Keywords: Preference aggregation, Decision under uncertainty, Multiple priors

JEL classification: D71, D81

---

*We are grateful to Itzhak Gilboa and David Schmeidler for many helpful comments, and to the participants of the D-TEA Workshop in Paris, 2014, for their useful feedback. We thank the Israeli Science Foundation, grant number 1188/14.

†Bar-Ilan University, Ramat-Gan 5290002, Israel. e-mail: Shiri.Alon-Eron@biu.ac.il

‡Bar-Ilan University, Ramat-Gan 5290002, Israel. e-mail: Gabriella.Gayer@biu.ac.il

1 Introduction

In the process of trying to create policies intended to benefit society government decision makers need to consider several alternatives. Clearly, alternatives may be evaluated differently across individuals as a result of personal interests, personal values and so on. However evaluations may still be subjective even when no apparent conflict of interest exists, as the alternatives being evaluated depend upon unknown factors that can be differently assessed. This is true for many governmental decisions, concerning the budget, environmental issues, public regulations and the like.

In this paper we propose a social welfare function that aggregates individual preferences when the alternatives are uncertain. Uncertainty is modeled by an abstract set of states of nature as formulated by Savage (1954). Individual preferences are assumed to abide by Savage's theory of Subjective Expected Utility (SEU), the most prominent model of decision making under uncertainty. Each individual preference is therefore assumed to admit an expected utility representation, where both the utility function (over outcomes) and the prior probability (over events) are subjective.

In welfare economics and social choice theory, the Pareto criterion is a widely accepted principle used link the social planner's preferences to those of individuals. It states that any unanimous preference that prevails in society must be obeyed by the social planner, thus formalizing the idea of autonomy of individuals. When alternatives are certain or risky (involving exogenous probabilities), conforming to the Pareto criterion is expected to increase the welfare of all members of society. Harsanyi's celebrated result [23] shows that under an expected utility hypothesis, when probabilities are known, the Pareto principle is tantamount to the social utility being a weighted sum of the individual utilities. Averaging of individual utilities represents a compromise between personal interests when individuals disagree, while accepting the force of their combined authority when individuals unanimously agree.

However, the implication of the Pareto criterion in the context of risk does not carry over to uncertain environments. Mongin in [30] and [32] proved that Harsanyi's result cannot be reconstructed in Savage's framework.[1] His results show that it may be impossible to aggregate SEU individual preferences into a non-dictatorial social SEU preference, if society is required to respect every unanimous preference of individuals.

The crux of this impossibility result is there being two components that determine preferences, a tastes component (a utility function in the SEU model) and a belief component (a subjective probability in the SEU model), instead of just one as in certain or risky environments. With two components individuals can agree on the ranking of two alternatives, but still hold very different tastes regarding their possible outcomes and very different beliefs on the likelihood of these outcomes. Such unanimous preferences are not based on genuine agreement, but on conflicts that cancel out. Adhering to such a unanimity, cannot be expected to increase the welfare of all members of society, and therefore may not be compelling.

The impossibility results for preference aggregation are akin to the problem of judgment aggregation that is demonstrated in the Discursive Dilemma (also known as the Doctrinal Paradox, see for example Kornhauser and Sager [27] Kornhauser [26] and List and Pettit [28]). The problem is caused by the tension that resides between judgment procedures that are premise-based and those that are conclusion-based. A conclusion-based approach revolves around the final judgments of individuals, which in our context are their rankings of alternatives, whereas a premise-based approach examines their premises or reasons for their judgments, which in our context are their tastes and beliefs. The paradox arises when these two procedures lead to opposite judgments. However when the two procedures coincide there is no doubt as to what should be the aggregated judgment, making those instances that are based on agreement both on premises and on conclusions much more sound than those that are based on agreement on conclusions alone.

In accord with this rationale Gilboa, Samet and Schmeidler [20] narrowed the scope of the Pareto criterion to apply only to rankings that agree on both premises and conclusions. In that model, both individuals and society are assumed to have

[1] Mongin's result builds on ideas found in Raiffa [37] and a related result in Hylland and Zeckhauser [24].

SEU preferences, but society is required to respect only those unanimous preferences that are based on agreed-upon probabilities. By employing this weaker version of Pareto these authors were able to circumvent the impossibility result of Mongin and obtain a social utility and social probability that are weighted averages of individuals' utilities and subjective probabilities, respectively.

The present work adopts the same approach, compelling the social planner to accept any unanimity in society that concerns alternatives in which one decision component is agreed-upon across individuals. Keeping one component constant across individuals ensures that any exhibited unanimous preference is generated by a unanimous ranking of the second component. Contrary to that, when a consensus in society is not both conclusion-based and premise-based, namely it is not supported by agreement on both tastes and beliefs, the social planner is allowed to disregard it.

Specifically, the social planner is required to conform to two Pareto criteria. One is Lottery Pareto, which states that for two acts that involve events with agreed-upon probabilities across individuals, a unanimous ranking of these acts will have to be respected by the social planner. This requirement reflects the understanding that such unanimity must be the result of a real agreement over the ranking of tastes. Analogously, Likelihood Pareto compels the social planner to accept any unanimous preference over acts that are contingent on the same pair of agreed-upon outcomes, amounting to preferences that express the same ranking of likelihoods.

These two conditions determine the ground rules that bind the social planner's choices in some, straightforward decisions. To complete the decisions for the general case, a framework for the social preferences needs to be designated. A common assumption in models of the type developed here is that society conforms to the same decision rule as do the individuals. In our context this would mean that the social planner adopts a single prior probability over events. In the simple case where all individuals in society share the same prior, for instance when there are irrefutable scientific data supporting this prior, this is also the obvious choice of the social planner. But if individual priors differ, which prior should the social planner adopt?

The fundamental message of our model is that failure of individuals to agree on probabilistic assessments introduces ambiguity into the social decision. The mere fact that individuals hold conflicting evaluations at the very least suggests that there is no obvious prescribed prior probability. Therefore the requirement that the social planner is always able to form a single prior when facing multiple priors appears too stringent. We hence relax the assumption of a single social prior by allowing the social planner to entertain a set of priors. Furthermore, considering that social decisions potentially affect the welfare of many individuals, and bearing in mind that the social planner is responsible their wellbeing, a conservative approach suggests itself. This is true especially when the alternatives under consideration may inflict considerable or irreversible harm.

A manifestation of this approach is the precautionary principle that is a widely accepted standard for policy concerning threats to public health and the environment. It essentially states that when an action or policy may cause harm to the public or to the environment, in the absence of scientific consensus that the action or policy is not harmful the burden of proof that it is not harmful falls on those taking an action. This principle was first endorsed internationally in 1982 by the United Nations that adopted the World Charter of Nature. Soon afterwards it was incorporated into many other environmental treaties, such as the Montreal protocol in 1987, the Rio declaration in 1992, and the Kyoto protocol in 1997. Since then the principle has become even more prominent, serving as a guideline for many international and domestic legislations concerning genetic engineering products, cell-phone radiation, food safety, extinction of species, global warming and so forth. There is no single accepted interpretation of the precautionary principle and in its most extreme formulation it requires absolute proof of safety before allowing an action to be taken. But even in its weaker versions the precautionary principle expresses the need of decision makers to adopt a cautious approach in the face of uncertainty.

To incorporate this concept we assume that society abides by Gilboa and Schmeidler's [22] Maxmin Expected Utility (MEU) decision rule, entertaining a set of priors as its belief component and evaluating each alternative by its minimal expected utility with respect to this set of priors. Importantly, the distinction between individuals who are

SEU maximizers and society that is an MEU maximizer coveys the idea that there is ambiguity that is intrinsic to social decision making, that arises at a social level even if no ambiguity is perceived at an individual level. The social planner exhibits aversion to this ambiguity, exercising a precautionary approach.

The assumptions of SEU-maximizing individuals and of an MEU-maximizing society hold throughout the paper. Under these assumptions, our main theorem shows that the two Pareto-type conditions mentioned above are equivalent to the social planner having utilitarian tastes, and entertaining a set of prior probabilities that are all weighted averages of individual priors. An implication of this result it that all the priors of the social planner identify with the individual probabilities of events when these happen to agree. Accordingly, the social planner perceives no ambiguity when all individuals agree on the probabilities of events, and may perceive ambiguity only when disagreement on probabilities prevails. In other words, under our assumptions the Pareto criteria entail that the only source of ambiguity is conflicting probabilistic assessments of individuals. Furthermore, social priors being weighted averages of individual probabilities only, is a reflection of the principle that the social belief is based on, and composed exclusively of, individual beliefs.

In the main theorem of the paper the set of priors representing the social belief is not required to contain all the individual priors. The social planner is allowed the discretion to under-weigh or even completely disregard some of the individual priors, as he or she sees fit. In a complementary result a strong notion of ambiguity aversion on behalf of the social planner identifies those cases in which the social belief is composed of all the individuals' prior probabilities.

Finally, it should be noted that the motivation for the choice of a social decision rule as well as the interpretations of the Pareto criteria are based on utilities and subjective probabilities being meaningful representations of tastes and beliefs of individuals under Savage's theory. This issue is discussed further in Subsection 3.1.

1.1 More commentary on the literature

Our aggregation results may be viewed as an answer to the impossibility result of Chambers and Hayashi [7]. Chambers and Hayashi [7] showed that Mongin's [30] and [32] impossibility results extend to a broader class of preferences, determining that the Pareto condition is incompatible with society and individuals having distinct notions of tastes over outcomes, and beliefs over events. In the present paper possibility is restored for SEU agents and an MEU social planner by replacing the standard Pareto criterion with two weaker conditions, restoring the possibility to separate social tastes from belief.

Broadening the scope of Mongin's impossibility result in a different direction, Gajdos, Tallon and Vergnaud [16] proved for a wide class of preferences that the standard Pareto condition is incompatible with aggregating individual preferences into a social preference when non-neutral attitudes towards ambiguity prevail. Our paper only partially addresses this impossibility in the special case of SEU agents. A more direct response to [16] was proposed by Danan, Gajdos, Hill and Tallon [9] and by Qu [36], in which individuals as well as society may express a non-neutral attitude toward ambiguity. This is done in an Anscombe-Aumann setup, where similar to our work the Pareto condition is weakened to overcome the impossibility result in [16].[2]

As opposed to the papers just mentioned, in which individuals and society are assumed to abide by the same decision rule, a feature of our development is the distinction between the decision rules of individuals and society. It has been suggested before, notably by Diamond [12], that it may be normatively inappropriate to require the same decision rule for individuals and for society. Recently Danan, Gajdos and Tallon [10] investigated, in the context of risk, Paretian aggregation under an assumption of incomplete expected utility preferences for both individuals and society. They discuss the special case where individual preferences are complete while those of society are incomplete. Society's preferences are then represented by a unanimity rule over a set of convex combinations of the individual utilities. This set of convex

[2] Billot and Vergopoulos [4] in an alternative approach obtained a utilitarian social welfare function on an extended state space, imposing a standard Pareto condition.

combinations is a dual of the priors set in our model, that is analogously composed of convex combinations of individual priors.

An incomplete social welfare criterion was also provided by Brunnermeier, Simsek and Xiong [6], this time for uncertain alternatives. According to the criterion in [6], the social planner evaluates outcomes via a utilitarian aggregate of individuals' tastes. One alternative is then said to be superior to another if the expected social utility of the former is higher than that of the latter, according to every individual prior.[3] Whenever such a domination holds the social planner in our model, under any one of our two theorems, will follow the same ranking of alternatives. Conceptually speaking, the main difference between the paper of Brunnermeier, Simsek and Xiong and our paper is that while we axiomatically derive a complete social welfare function, these authors define efficiency criteria, investigate their properties and study their implications in various economic setups. Furthermore, Brunnermeier, Simsek and Xiong [6] is aimed for situations in which the reason agents hold different beliefs is that they distort them due to behavioral biases, this being the justification for weakening the standard Pareto criterion. Contrary to that, in the current model standard Pareto is weakened to avoid cases where agreement is based on conflicts, not presupposing that beliefs are 'wrong' or distorted in any manner.

A problem that is related to the aggregation of individual preferences is that of a decision maker who consults with several experts. Experts are assumed to share the decision maker's tastes but may differ with respect to their likelihood assessments, and their opinions are combined employing various likelihood-Pareto criteria. The aggregation of beliefs in our model resembles the aggregation of experts' beliefs in Cres, Gilboa, and Vielle [8] in that multiple weights are assigned to the individual beliefs (be they those of experts or those of members of society), and the minimum utility over these weights is used for evaluating alternatives. However, Cres, Gilboa and Vielle [8] axiomatized the case of aggregating MEU beliefs rather than SEU beliefs as in our paper. Another decision rule, in the form of a minimum over more

[3]These authors present another criterion of Pareto efficiency that is less related to the present work.

general opinions, is Nascimento's [34], allowing for experts with different perceptions of ambiguity and attitudes towards it.

Mongin and Pivato [33] take another approach for aggregating individual preferences. They investigate ex ante and ex post Pareto conditions in a multi-dimensional array framework, where uncertainty is modeled by a two-component state: one corresponding to a subjective source of uncertainty and the other to an objective source. The utilitarian-like representation for society that is derived agrees with individual priors with regard to the objective part of the state space, but makes no statement on the relationship between social and individual beliefs in the subjective dimension of the state space.

In light of the financial crisis of 2007, several recent papers have claimed that the Pareto principle may be undesirable in connection to unregulated trade (see for instance Weyl [41]; Posner and Weyl [35]; Blume, Easley, Sargent, and Tsyrennikov [5]; and Brunnermeier, Simsek, and Xiong [6]). These papers claimed that the standard Pareto criterion may not be suitable when agents hold heterogeneous beliefs. While agents whose beliefs are close to the truth may prosper, those agents holding 'wrong' beliefs may end up making poor decisions.

In the search for an appropriate Pareto criterion for uncertain environments Gilboa, Samuelson, and Schmeidler [21] presented a model of a regulator whose objective is to interfere as little as possible with voluntary trade. They suggested that this regulator act in accordance with a Pareto criterion called no-betting Pareto which requires that there be a single probability, on top of the unanimity of personal preferences, that can rationalize the preferences of all the SEU agents. This probability measure need not have any connection to the probabilities of the individuals nor to that of the regulator. This is a sensible condition when the sole objective of the regulator is to ban actions that are essentially betting. However, it is a less appealing criterion for making social decisions. In particular, there is no reason to expect that individual welfare increases when alternatives are evaluated using a probability measure that is unrelated to any of the individuals' conceptions of reality. Likelihood Pareto, in contrast to no-betting Pareto, disregards any domination among alternatives that cannot be justified by a probability measure that is a convex combination of individuals' beliefs. A related Pareto concept, which weakens standard Pareto but generalizes both Lottery Pareto and Likelihood Pareto, can be found in Gayer, Gilboa, Samuelson, and Schmeidler [17].

1.2 Outline of the paper

The next section contains the setup and the basic assumptions of the model, and Section 3 presents the Pareto criteria and our two aggregation theorems. A discussion of the identification of tastes and beliefs is given in Subsection 3.1. Finally, the proofs appear in Section 4.

2 Setup and basic assumptions

Let $N = \{1, \ldots, n\}$ be the set of individuals whose preferences are to be aggregated. Subjective uncertainty is represented by a nonempty set S of *states of nature* that is endowed with a sigma algebra Σ of *events*. The set of possible *outcomes* is X and the set of (simple) *acts* is $\mathcal{F} = \{f \in X^S \mid f \text{ obtains finitely many values and is measurable w.r.t. } \Sigma\}$. For acts f and g and an event E, fEg stands for the act which assigns the outcome $f(s)$ if $s \in E$ and $g(s)$ otherwise. The preferences of individual i are modeled by a binary relation $\succsim^i$ over $\mathcal{F}$, with $\succ^i$ and $\sim^i$ being its asymmetric and symmetric components. The social preferences are modeled by another binary relation over acts, $\succsim^0$, with asymmetric and symmetric components $\succ^0$ and $\sim^0$.

Throughout the paper it is assumed that individual preferences satisfy the axioms of the Subjective Expected Utility model of Savage [38], supplemented with a monotone continuity condition (see Villegas, [40]). A subjective expected utility representation with a sigma-additive probability measure ensues, as stated in the next assumption.

Assumption 1: SEU Individuals. For each individual i there exists a cardinal utility function over outcomes[4], u_i, and a unique non-atomic and sigma-additive

[4] That is, a utility function which is unique up to an increasing linear transformation.

subjective probability measure over events[5], P_i, such that the preference $\succsim^i$ is represented by the functional:

$$SEU_i(f) = E_{P_i}(u_i \cdot f) \quad , \quad \forall f \in \mathcal{F} \ .$$

A weak notion of agreement among agents is required for the aggregation, whereby the agents need only to agree upon the strict ranking of some pair of outcomes, as stipulated next.

Assumption 2: Minimal Agreement. There are outcomes x^* and x_* such that $x^* \succ^i x_*$ for all individuals i.

The notation x^*, x_* is henceforth reserved for a specific pair of outcomes that satisfy Minimal Agreement, and all individual utilities are calibrated so as to satisfy $u_i(x^*) = 1$ and $u_i(x_*) = 0$. Next a definition of unanimously agreed-upon half-events is given.

Definition 1. An event E is a unanimous half-event if for every individual i, $x^* E x_* \sim^i x_* E x^*$.

It is assumed in this paper that the social preference, $\succsim^0$, admits an MEU representation with a set of sigma-additive probabilities. It is furthermore supposed that the social preference may be represented by an MEU functional that agrees with the individuals on one 'fair coin', so to speak: there is one event which all individuals unambiguously perceive as a half-event, and which the social planner also takes to be an unambiguous half event. In a companion paper (see [1]) we derive rather than assume the MEU representation along with such an agreed-upon half event, based on the assumptions of SEU individuals (Assumption 1), Minimal Agreement (Assumption 2), and a set of behavioral axioms for the social preference, among them the two Pareto conditions that are given in the next section. Notably, these axioms express the view that the social planner behaves in an ambiguity neutral manner towards acts that involve

[5] A probability measure p is non-atomic if, for any event E with $p(E) > 0$ and any $\lambda \in (0,1)$, there exists an event $F \subset E$ such that $p(F) = \lambda p(E)$.

only probabilities on which all individuals agree. This neutrality is what essentially underlies the existence of a 'fair coin', agreed-upon by the individuals and the social planner alike. It should be noted that in an Anscombe-Aumann framework (see [15]) the fair coin assumption would be redundant, as objective probabilities, on which everybody is assumed to agree at the outset, are made available as part of the setup itself.

Assumption 3: MEU Social Preference. There exists a cardinal utility function u_0 and a unique non-empty, convex and closed set $\mathcal{C}$ of non-atomic and sigma-additive probabilities,[6] such that the relation $\succsim^0$ is represented by the functional:

$$MEU(f) = \min_{p \in \mathcal{C}} E_p(u_0 \cdot f) \quad , \quad \forall f \in \mathcal{F} \ ,$$

where for some unanimous half-event E, $MEU(x^* E x_*) = \frac{1}{2}u_0(x^*) + \frac{1}{2}u_0(x_*)$.

3 Aggregation results

Two Pareto criteria connect the preferences of individuals to those of society, by laying down the ground rules of how and when the social preference must obey the unanimous wishes of individuals. As alluded to in the introduction, agreement among individuals on the ranking of acts under the Bayesian paradigm might be the result of double disagreement, both on tastes and on beliefs. The Pareto criteria employed in this model regard as 'legitimate' only unanimous preferences with underlying consensus either on tastes or on likelihoods, relieving the social planner from the obligation to comply those unanimous rankings that result only from conflicts.

The first Pareto criterion pertains to acts with the same underlying likelihoods.

[6] We take the set $\mathcal{C}$ to be closed in the weak* topology over the space of finitely additive set functions. This is equivalent to the set $\mathcal{C}$ being closed under event-wise convergence, that is, if there exists a generalized sequence μ_b in $\mathcal{C}$ such that $\mu_b(A)$ converges to $\mu(A)$ for every event $A \in \Sigma$, then $\mu \in \mathcal{C}$. Consequently, $\mathcal{C}$ is compact (see Maccheroni and Marinacci [29]) the minimum operator is well defined.

It hinges on a definition of *socially unambiguous partitions* and *socially unambiguous acts.*

Definition 2. *A partition* $\{E_1, \ldots, E_m\}$ *is a* ***socially unambiguous partition*** *if for every individual* i*,* $x^* E_k x_* \sim^i x^* E_\ell x_*$*, for all* k *and* ℓ*.*

A ***socially unambiguous act*** *is an act which is measurable w.r.t. a socially unambiguous partition.*

In terms of probabilities, a socially unambiguous partition $\{E_1, \ldots, E_m\}$ satisfies $P_i(E_k) = \frac{1}{m}$ for every event E_k and every probability P_i. Technically speaking, Lyapunov's theorem guarantees that such partitions exist for every $m \in N$. A socially unambiguous act f satisfies $P_i(f = x) = P_j(f = x)$ for every outcome x and every pair of individual probabilities P_i and P_j, namely it induces the same distribution over outcomes according to all individual priors. Such an act is akin to a lottery with known probabilities in the sense that it creates no ambiguity in the social level .

The next Pareto condition, that was introduced in Gilboa, Samet and Schmeidler [20], applies only to socially unambiguous acts. It requires that the probability component of all individuals be equal, but allows tastes to vary. Whenever consensus obtains under such circumstances, this Pareto condition commands the social planner to respect it.

Lottery Pareto. For two socially unambiguous acts f and g, if $f \succ^i g$ for all individuals i, then $f \succ^0 g$.

An immediate implication of this condition is that the social preference agrees with the individual preferences on the strict ranking of the outcomes x^* and x_*, namely $x^* \succ^0 x_*$.

An agreement on the likelihoods of events, as depicted in the Lottery Pareto criterion, may prevail when individuals form beliefs using a large body of statistical data that can be freely accessed, or there being well established institutions that provide the information to the public. This is relevant to many important questions in the fields of finance, education, health and the like, where data are abound, and official statistics are available.[7].

One example is the problem of overweight and obesity in children and adolescents that has been recognized as a serious public health concern. There is an ample amount of statistics showing that the level of childhood overweight and obesity has more than doubled over the years, and its connection to the consumption of high calorie foods and lack of physical activity in well known. In this context there is a discussion on whether soft drink vending machines should be banned from schools. Opponents of the vending machines claim that schools are responsible for providing a healthy environment for children during school hours. Proponents of vending machines are well aware of the potential harms, but believe that children and their parents should be free to make their own decisions regarding what the children drink. Others mention the revenues generated by vending machines that help fund school activity. Given the debate on this topic, banning vending machines is clearly not in consensus. However, whatever the choice is, adding vending machines of mineral waters is an alternative that is not opposed by any party.

Another instance where probabilities are agreed upon is in the trade of corporate and government bonds. Credit rating agencies are responsible to publish ratings of corporate and government bonds that are based on assessing the likelihood that the debt will be repaid. These ratings serve as a common basis for investors, to either buy or sell bonds. Voluntary trade on such an asset for which there is much reliable information is considered Pareto improving compared to no trade, as it is based on different risk attitudes but on the same beliefs.

The second Pareto condition we employ is Likelihood Pareto, which is Lottery Pareto's dual (this dual condition was already mentioned in Gilboa, Samet and Schmeidler [20]). While Lottery Pareto applies only to acts that involve agreed upon probabilities, Likelihood Pareto applies only to acts that return agreed upon outcomes. More specifically, Likelihood Pareto is limited to acts that return the

[7] An example of this type of statistical data are those provided by the UN Statistical Commission, http://unstats.un.org/unsd/default.htm.

identically-ranked outcomes x^* and x_* and can thus be translated by all the individuals into bets on the same events. Likelihood Pareto states that should all individuals prefer to bet on one event over another, meaning that they find the former more likely than the latter, then so does society.

Likelihood Pareto. For two events E and F, if $x^*Ex_* \succ^i x^*Fx_*$ for all individuals i, then $x^*Ex_* \succ^0 x^*Fx_*$.

The only outcomes that are returned in the acts stated in the Likelihood Pareto condition are x^* and x_*, so there is not a question as to which is the preferred outcome. Individual preference is determined solely by the subjective evaluation of the likelihood of event E compared with that of event F. This condition resembles those found in connection to the aggregation of experts opinions, as in both cases the question concerns likelihoods rather than tastes (see for instance Cres, Gilboa and Vieille [8]).

An example of an issue in which there is disagreement over likelihoods but not over the ranking of outcomes may be found in connection to the financial crisis that began in 2007. This crisis led to a deep and prolonged global economic downturn causing a high rate of unemployment. In response the Federal Reserve took extraordinary actions to help stabilize the U.S. economy and financial system, including reducing the level of short-term interest rates to near zero. Today, given the improved economic conditions, there is question whether this policy is still appropriate. However, at the time, there was a consensus that this monetary policy would likely be effective to achieve maximum employment and stable prices, in line with Likelihood Pareto.

Another example may be found in connection with programs that provide training and supportive services to long-term unemployed individuals. These programs include job-search assistance, vocational and on-the-job training, as well as commitments from employers to interview and hire program participants. Once again in relation to Likelihood Pareto, for a given budget there is a common view that a program is more likely to be successful if it provides a combination of these services rather than only one.

Other examples where people share the same tastes, but may differ on assessments of probabilities include comparisons of methods to produce clean energy, designing auctions to sell public resources (such as spectrum auctions), and the process of approval of new developed drugs by the FDA. In all of these instances Likelihood Pareto entails that consensus in society regarding the preferred alternative should be respected by the social planner, as it must be the consequence of true agreement.

The main result of the paper states that when Lottery Pareto and Likelihood Pareto are used to aggregate SEU individual preferences into a social MEU preference, then (i) the social utility is a weighted average of the individual utilities and (ii) the social set of priors is composed only of probabilities that are weighted averages of the individual beliefs. Put differently, the set of priors of the social planner is contained in the convex hull of the individuals' probabilities.

Theorem 1. *Suppose* ***SEU Individuals*** *(Assumption 1),* ***Minimal Agreement*** *(Assumption 2) and* ***MEU Social Preference*** *(Assumption 3). Then Lottery Pareto and Likelihood Pareto hold if and only if the social utility is a nonnegative, nonzero combination of the individual utilities, and each prior in the social set of priors is a convex combination of individual priors.*

Theorem 1 characterizes a social MEU preference whose tastes (the utility) and beliefs (the set of priors) are both based on individual preferences. The social utility is utilitarian, being a convex combination of the utilities of the individuals. The social belief is composed of convex combinations of the individuals' beliefs. This is in line with the underlying view that the only factors that come into play in the social planner's likelihood evaluations are the likelihood evaluations of the individuals. Such an assumption is appropriate in situations where any information that is relevant for the formation of the social belief is already known by individuals in society (and will certainty apply when the social planner is one of the individuals in society, so that any extra information he or she might possess is also apprehended by one of the individuals).

Note that the priors set representing the social belief is not required to contain all individual probabilities. It is left to the social planner's discretion to determine whether to assign low weight to some individual beliefs or even ignore them altogether, for instance if he or she regards some of the beliefs as too extreme.

The full proof of Theorem 1 appears in the appendix. Essentially it is conducted in three steps. First it is shown that whenever all individual probabilities assign the same probability to an event, then so do all the probabilities in the social set of priors. The proof of this part is based on Likelihood Pareto and on the MEU assumption, which includes existence of a unanimous half-event that is perceived as unambiguous by the social planner. After showing that all the priors of the social planner identify with the individual probabilities when these agree, Lottery Pareto is employed to derive a Harsanyi-like result, by which the social utility is a nonnegative combination of the individual utilities. Finally, a separation theorem is used to establish that there can be no prior in the social set of priors that lies outside the convex hull of the individual probabilities, otherwise Likelihood Pareto would be contradicted.

To illustrate what may go wrong if the social planner's belief is not contained in the convex hull of the individuals' beliefs imagine a society contemplating which method of renewable energy to invest in. Suppose that all members of society believe that it is more likely to produce enough energy from solar technologies than from wind turbines. For the sake of illustration assume that the beliefs held in society are that the probability of sufficient sun is at least 0.6, while the probability of a sufficient amount of wind is no more than 0.5. It is therefore a consensus in society that it is preferable to invest in solar energy. Assume that the social planner agrees with the individuals that the probability of wind is 0.5, but maintains that the probability of sufficient sun is as low as 0.4, which is outside the convex hull of the individuals' probabilities. This social planner will advance an investment in wind turbines countering individuals' consensual preference, even though this is rooted in identical ranking of outcomes, based on real agreement over the likelihood assessments of the events under consideration.

To complement Theorem 1, the second theorem of the paper characterizes those cases where society does consider all individual beliefs as viable and therefore includes all of them in its priors set. Under the assumptions of Theorem 1 this is also the largest possible set of social priors. In the context of MEU models, where a larger set of probabilities corresponds to a preference that is more ambiguity averse, this specific social preference exhibits the strongest aversion to ambiguity possible.

To characterize this social preference a strong condition of social ambiguity aversion is presumed, named *Avoidance of Societal Ambiguity.* It states that whenever at least one individual strictly prefers to bet on an unambiguous event compared to some other event then so does society, thereby asserting that every individual likelihood judgment counts. We note that Avoidance of Societal Ambiguity is related to an axiom that appears in Gilboa et. al. [19] named Default to Certainty, that ties a decision maker's objective preference to his or her subjective one.

Avoidance of Societal Ambiguity. Let F be an event and E an unambiguous event. If $x^*Ex_* \succ^i x^*Fx_*$ for some individual i, then $x^*Ex_* \succ^0 x^*Fx_*$.

Avoidance of Societal Ambiguity implies that society can rank a *not* agreed-upon event above an agreed-upon event only when this ranking is consensual among all individuals (as already implied by Likelihood Pareto). In any other case society should opt for the bet whose rewarding event is unambiguous, which represents a well-understood risk. This condition is akin to a comparative notion of ambiguity aversion proposed by Epstein [14] and Ghirardato and Marinacci [18], whereby decision maker A is considered more ambiguity averse than B if whenever B prefers an unambiguous act then so does A.[8] While unambiguous acts in the two works mentioned are either exogenous (in the former paper) or constant (in the latter), unambiguous alternatives in our model are defined endogenously as those that involve events with probabilities that are agreed-upon by all members of society. This once again is an expression

[8] The two works differ in their notions of ambiguity neutrality, the former assuming it is equivalent to a probability sophisticated preference and the latter assuming it is an SEU preference. However when only two outcomes are concerned, as in the axiom discussed here, the two notions coincide.

of the ambiguity in the proposed model being a result of different beliefs held in society. By this token, society preferring the consensual event whenever any one of the individuals does, implies that the social planner is more ambiguity averse than each of the individuals in society.

The next result shows that adding Avoidance of Societal Ambiguity to the conditions of Theorem 1 pins down the social belief so that the set of prior probabilities $\mathcal{C}$ must equal the entire convex hull of individual probabilities. It thus implies that society evaluates each act by its minimum expected social utility, the minimum being taken over all the individual priors.

Theorem 2. *Suppose* ***SEU Individuals*** *(Assumption 1),* ***Minimal Agreement*** *(Assumption 2) and* ***MEU Social Preference*** *(Assumption 3). Then Lottery Pareto, Likelihood Pareto and Avoidance of Societal Ambiguity hold if and only if the social utility is a nonnegative, nonzero combination of the individual utilities, and for any two acts f and g,*

$$f \succsim^0 g \iff \min_{i \in N} E_{P_i}(u_0 \cdot f) \geq \min_{i \in N} E_{P_i}(u_0 \cdot g) \ .$$

To further compare this theorem with the previous one, note that the social planner's preferences under the conditions of Theorem 1 could exhibit a neutral attitude towards ambiguity, in the special case where the social belief reduces to a single prior probability. This possibility is excluded under the conditions of Theorem 2. When imposing Avoidance of Societal Ambiguity a corresponding impossibility arises, where society cannot be ambiguity neutral (i.e. admit an SEU representation) unless all individuals share a common prior. This is a manifestation of the strong form of ambiguity aversion described by Avoidance of Societal Ambiguity. If opting for *not* agreed upon events is allowed only under consensus then the only way to restore possibility is to allow the social planner to exhibit a non-neutral attitude towards ambiguity.

To prove the second theorem the only supplement to the previous proof that is needed is to show that all the individual priors must be contained in the social set of priors. This inclusion is sustained using a separation theorem, showing that Avoidance of Societal Ambiguity would be contradicted if an individual prior resides outside the social set of priors: there must be an unambiguous event and another arbitrary event such that the 'outsider' individual prior ranks the unambiguous event as more likely than the other event, whereas society's ranking is reversed.

3.1 Identification of tastes and beliefs

The normative appeal of the model we offer relies on the interpretation that tastes and beliefs of individuals are meaningful notions under the Savage theory. This can be seen in the distinct role each of these components plays in society's decision rule, and in the justifications of the two Pareto conditions employed in the model. A fundamental assumption that allows the identification of tastes and beliefs is that individual preferences are state-independent. Without this assumption, a unique extraction of tastes and beliefs is no longer guaranteed, so that the utilities and probabilities that appear in our axioms and social representation may not be faithful descriptions of the individuals' actual tastes and beliefs.

Whether or not this assumption applies will depend on the context. In situations where people care only about the outcome but not its causes, such as the realized value of a stock, state-independence will commonly prevail. On the other hand when there is a clear connection between the payoffs and the states of nature, as in health-related decision problems, this assumption is less appropriate.

For pure state-dependent problems where tastes cannot be distinguished from beliefs the impossibility results presented in the introduction do not apply, and there are results to show that aggregation is compatible with the standard form of Pareto (see Mongin [32], Gajdos, Tallon and Vergnaud [16] and Chambers and Hayashi [7]). When more structure on the sate-dependent preferences is imposed, specifically when subjective probabilities are uniquely identified (for example in the model of Karni, Schmeidler and Vind [25]), aggregation may once again become impossible or be very limited. In those cases a modification of our Pareto criteria so that they employ the 'correct' probabilities and utilities under the specific state-dependent model could be considered.

A general criticism that may come up in connection to our work is that utility and probability are just two technical artifacts, the result of a modeler's intellectual exercise rather than real concepts that influence people's choices. Under this paradigm our use of these artifacts and the motivation we provide lose normative appeal. This broad criticism is hardly specific to our work, but applies to many decision-theoretic models under uncertainty. We thus refrain from pursuing it in this paper, and the interested reader is referred to Gayer et al. [17] for an extensive discussion.

4 Proofs

4.1 Proof of Theorem 1

Let all the individual preferences be represented by SEU functionals, with utilities $\{u_i\}_{i=1}^n$ and sigma-additive probabilities $\{P_i\}_{i=1}^n$. By the assumption of Minimal Agreement all the individual utilities can be calibrated to satisfy $u_i(x_*) = 0$ and $u_i(x^*) = 1$. Suppose that the social preference $\succsim^0$ admits an MEU representation in accordance with Assumption 3, with a cardinal utility u_0 and a nonempty, closed and convex set of sigma-additive probabilities $\mathcal{C}$ (see footnote 2). Assume that Lottery Pareto and Likelihood Pareto are satisfied. An immediate consequence is that $x^* \succ^0 x_*$, therefore the social utility u_0 can be calibrated in the same manner as the individual utilities, assigning $u_0(x^*) = 1$ and $u_0(x_*) = 0$.

We proceed to define mixtures of acts.

Definition 3. *Let $f = [x_1, E_1; \ldots; x_m, E_m]$ and $g = [y_1, E_1; \ldots; y_m, E_m]$ be two acts, such that $\{E_1, \ldots, E_m\}$ is a partition with respect to which both acts are measurable. For $0 \leq \alpha \leq 1$, an $\alpha : (1-\alpha)$ mixture of f and g is an act $h = [x_1, G_1; y_1, E_1 \setminus G_1; \ldots; x_m, G_m; y_m, E_m \setminus G_m]$ for events G_k that satisfy $G_k \subseteq E_k$ and $P_i(G_k) = \alpha P_i(E_k)$ for every P_i.*

By Lyapunov's theorem such events G_k exist, thus the defined mixtures exist.

Consider the set of all events that have a unanimously agreed-upon probability (not necessarily rational). This set of events, denote it by $\mathcal{E}$, is precisely the sigma-algebra generated by all the socially unambiguous partitions of S. All the probabilities P_i identify on all the socially unambiguous partitions. All these probabilities are assumed to be sigma-additive, thus they identify on $\mathcal{E}$ (on account of their continuity). Denote the restriction of some (hence all) individual beliefs P_i to $\mathcal{E}$ by π.

Employing Lyapunov's theorem again, for any real number $\rho \in [0,1]$ there exists an event $E \in \mathcal{E}$ such that $\pi(E) = \rho$. The set of $\mathcal{E}$-measurable acts are henceforth referred to as *lotteries*. Note that the set of lotteries is closed under the mixtures defined above.

Claim 1. *For all p in $\mathcal{C}$ and E in $\mathcal{E}$, $p(E) = \pi(E)$.*

Proof. Let E be a unanimous half-event for which $MEU(x^* E x_*) = \frac{1}{2}u_0(x^*) + \frac{1}{2}u_0(x_*) = \frac{1}{2}$, as in Assumption 3. The MEU representation of $\succsim^0$ and Likelihood Pareto imply that $MEU(x^* E x_*) = \min_{p\in\mathcal{C}} p(E) = \min_{p\in\mathcal{C}} p(E^c) = 1 - \max_{p\in\mathcal{C}} p(E)$, thus $\min_{p\in\mathcal{C}} p(E) = \max_{p\in\mathcal{C}} p(E) = \frac{1}{2}$, implying $p(E) = \frac{1}{2}$ for every $p \in \mathcal{C}$. For any other unanimous half-event F, the implied rankings $x^* F x_* \sim^0 x_* F x^* \sim^0 x^* E x_*$ entail $p(F) = \frac{1}{2}$ for all $p \in \mathcal{C}$.

Now suppose a socially unambiguous partition $\{E_1, \ldots, E_4\}$. Any union of two partition elements is a unanimous half-event, therefore each such union is assigned a probability of half according to every prior in $\mathcal{C}$. In particular, for every $p \in \mathcal{C}$, $p(E_1) + p(E_2) = p(E_1) + p(E_3) = p(E_1) + p(E_4) = 0.5$, implying $p(E_1) = \frac{1}{4}$ for every $p \in \mathcal{C}$. The same follows for every event in a foursome socially unambiguous partition. In the same manner events in every dyadic socially unambiguous partition are assigned the corresponding dyadic probability by every $p \in \mathcal{C}$. Since all the probabilities in $\mathcal{C}$ and all the probabilities P_i are sigma additive by assumption it follows that $p(E) = \pi(E)$ for every $p \in \mathcal{C}$. ■

Claim 2. *The social preference $\succsim^0$ over lotteries is represented by a vNM utility function that is a convex combination of the individuals' utilities.*

Proof. For two socially unambiguous acts L and L' the condition in Lottery Pareto, when translated to the representation, states $E_\pi(u_i \cdot L) > E_\pi(u_i \cdot L')$ for all i. Employing the previous claim, the conclusion of the axiom translates to $E_\pi(u_0 \cdot L) > E_\pi(u_0 \cdot L')$.

Although the axiom addresses only acts that involve rational valued π-probabilities, the same implication in representation follows for irrational probabilities as well. Suppose on the contrary that for two lotteries L and L', which involve irrational π-probabilities, $E_\pi(u_i \cdot L) > E_\pi(u_i \cdot L')$ for all i, but $E_\pi(u_0 \cdot L') \geq E_\pi(u_0 \cdot L)$. Construct a new socially unambiguous act Q' from L' by assigning all the outcomes other than the socially most favorable outcome under L' to socially unambiguous events with slightly lower probabilities than in L', so that the new probabilities are rational, and assigning the most favorable outcome to the complementary event, which now has a higher, rational probability. Similarly construct a socially unambiguous act Q from L by amplifying the probability of the socially least favorable outcome under L, to obtain that outcomes are assigned to events with rational π-probabilities. The inequality for the social preference still holds, and for some small enough π-probabilities the strict inequalities for the individual preferences hold as well, obtaining a violation of Lottery Pareto.

The range of the vector-valued function $(E_\pi(u_0(\cdot)), E_\pi(u_1(\cdot)), \ldots, E_\pi(u_n(\cdot)))$ on the set of lotteries is convex (employing the mixtures from Definition 3), therefore according to De Meyer and Mongin ([11]; by employing also Minimal Agreement), $E_\pi(u_0(\cdot))$ is a nonnegative, non-zero linear combination of the functions $E_\pi(u_i(\cdot))$. It may be calibrated by setting $u_0(x_*) = 0$ and $u_0(x^*) = 1$, to obtain that $u_0 = \sum_{i=1}^{n} \theta_i u_i$ for non-negative weights θ_i that sum to one. ∎

It is next shown that the set of probabilities entertained by the social planner, $\mathcal{C}$, is contained in $conv\{P_1, \ldots, P_n\}$, the convex hull of the individuals' probabilities. This inclusion is proved using a separation theorem, for which further notation and definitions are required.

Denote by $B_0(S, \Sigma)$ the space of all Σ-measurable, finite-valued functions over

S (equivalently, this is the vector space generated by the indicator functions of the elements of Σ), endowed with the supremum norm. Denote by $ba(\Sigma)$ the space of all bounded and finitely additive functions from Σ to $\mathbb{R}$, endowed with the total variation norm. The space $ba(\Sigma)$ is isometrically isomorphic to the conjugate space of $B_0(S,\Sigma)$.

Consider an additional topology on $ba(\Sigma)$. For $\varphi \in B_0(S,\Sigma)$ and $m \in ba(S,\Sigma)$, let $\varphi(m) = \int_S \varphi dm$. Every φ defines a linear functional over $ba(\Sigma)$, and $B_0(S,\Sigma)$ is a total space of functionals on $ba(\Sigma)$.[9] The $B_0(S,\Sigma)$ topology of $ba(\Sigma)$, by its definition, makes a locally convex linear topological space, and the linear functionals on $ba(\Sigma)$ which are continuous in this topology are precisely the functionals defined by $\varphi \in B_0(S,\Sigma)$. This topology is called the weak* topology of $ba(\Sigma)$.

Lemma 3. $\mathcal{C} \subseteq conv\{P_1,\ldots,P_n\}$.

Proof. Suppose on the contrary that there exists $p' \in \mathcal{C}$ such that $p' \notin conv\{P_1,\ldots,P_n\}$. By a standard separation theorem (based on the topological properties detailed above; see for instance Corollary V.2.12 in Dunford and Schwartz [13]) there exists a separating non-zero vector $\varphi = [\varphi_1, E_1;\ldots;\varphi_m, E_m] \in B_0(S,\Sigma)$ and a scalar c such that $\varphi(q) \geq c > \varphi(p')$ for all $q \in conv\{P_1,\ldots,P_n\}$. Note that as q and p' are probabilities the separating vector φ cannot be constant.

In order to maintain that all coordinates of φ are between zero and one subtract $\min_k \varphi_k$ from all sides of the two inequalities and then divide all of them by the sum of the resulting coordinates (which cannot be zero as φ is non-constant). Denote the transformed vector φ by $\alpha = [\alpha_1, E_1;\ldots;\alpha_m, E_m]$, for which $0 \leq \alpha_k \leq 1$, and not all coordinates α_k are zero. The corresponding transformed scalar is also between zero and one, strictly larger than zero since $\alpha(p') \geq 0$. Denote by $\hat{c}$ the transformed scalar c, subtracted by some small enough ϵ, so that it satisfies: $\alpha(q) > \hat{c} > \alpha(p')$ for all $q \in conv\{P_1,\ldots,P_n\}$.

Once again by Lyapunov's theorem, for every event E_k there exists an event $G_k \subseteq E_k$ such that $p'(G_k) = \alpha_k p'(E_k)$ as well as $P_i(G_k) = \alpha_k P_i(E_k)$ for all i. The act $f = [x^*, G_1; x_*, E_1 \setminus G_1;\ldots;x^*, G_m; x_*, E_m \setminus G_m]$ satisfies $E_{P_i}(u_i \cdot f) = \alpha(P_i)$ for

[9] That is, $\varphi(m) = 0$ for every $\varphi \in B_0(S,\Sigma)$ implies that $m = 0$.

all i, and $E_{p'}(u_0 \cdot f) = \alpha(p') \geq \min_{p \in \mathcal{C}} E_p(u_0 \cdot f)$. In addition, let $F \in \mathcal{E}$ be such that $\pi(F) = \hat{c}$. As $F \in \mathcal{E}$, all the probabilities P_i and every $p \in \mathcal{C}$ agree that $P_i(F) = p(F) = \pi(F) = \hat{c}$, therefore $E_{P_i}(u_0 \cdot (x^* F x_*)) = \min_{p \in \mathcal{C}} E_p(u_0 \cdot (x^* F x_*)) = \hat{c}$, for all i. The separation result thus yields $E_{P_i}(u_i \cdot f) > E_{P_i}(u_i \cdot (x^* F x_*))$ for all i, and at the same time $\min_{p \in \mathcal{C}} E_p(u_0 \cdot (x^* F x_*)) > \min_{p \in \mathcal{C}} E_p(u_0 \cdot f)$. As the acts involved return only the outcomes x^* and x_* a contradiction to Likelihood Pareto results. It follows that $\mathcal{C} \subseteq conv\{P_1, \ldots, P_n\}$. ■

For the other direction suppose that f and g are two socially unambiguous acts. Then all the individual probabilities, hence all the priors in $\mathcal{C}$, agree on the distributions over outcomes induced by f and g. Each one of these acts is thus evaluated by the expected utility of outcomes, based on the same probabilities. If all these computations, for all the utilities u_i, yield that f is preferred to g, then the same will be true for any nonnegative combination of the u_i's, hence for u_0. The social planner will therefore strictly prefer f to g whenever strict preference holds for all individuals.

Now suppose that for two acts that involve only the outcomes x^* and x_*, one is unanimously preferred to the other (by all individuals). The two acts can be written $x^* E x_*$ and $x_* F x_*$, and the preference corresponds to $P_i(E) \geq P_i(F)$ for all individuals i. As each $p \in \mathcal{C}$ is a convex combination of P_i's it follows that $p(E) \geq p(F)$ for all $p \in \mathcal{C}$, therefore $x^* E x_* \succsim^0 x_* F x_*$ by the assumed representation of the social preference.

4.2 Proof of Theorem 2

Theorem 1 has confirmed that $\mathcal{C} \subseteq conv\{P_1, \ldots, P_n\}$. Thus the only part that must be verified for the proof of Theorem 2 is that Avoidance of Societal Ambiguity holds if and only if the opposite inclusion of the priors set holds, namely that the set of priors $\mathcal{C}$ in the social preference representation contains all the individual priors.

Lemma 4. *$conv\{P_1, \ldots, P_n\} \subseteq \mathcal{C}$.*

Proof. Suppose on the contrary that there exists P_i such that $P_i \notin \mathcal{C}$. By a standard separation theorem there exists a separating non-zero vector $\varphi = [\varphi_1, E_1; \ldots; \varphi_m, E_m] \in B_0(S, \Sigma)$ and a scalar c such that $\varphi(p) \geq c > \varphi(P_i)$ for all $p \in \mathcal{C}$, where this vector cannot be constant. Once again φ is subtracted $\min_k \varphi_k$ from all its coordinates and divided by the sum of the (nonnegative) coordinates, so that all its coordinates become between zero and one. The resulting vector is denoted by $\alpha = [\alpha_1, E_1; \ldots; \alpha_m, E_m]$, with $0 \leq \alpha_k \leq 1$ for all k, not all coordinates being zero. The transformed scalar is again between zero and one and strictly larger than zero. Let $\hat{c}$ denote a separating scalar between zero and one that is rational (Such a separating scalar exists since any separating scalar as implied by the separation theorem may be approached arbitrarily close by rational numbers).

In the same manner as in the previous lemma for every E_k let G_k be events such that $G_k \subseteq E_k$ and $P_i(G_k) = \alpha_k P_i(E_k)$ for all i. As $\mathcal{C}$ is known by the previous lemma to be contained in the convex hull of the P_i's it follows that $p(G_k) = \alpha_k p(E_k)$ for all $p \in \mathcal{C}$. The act $f = [x^*, G_1; x_*, E_1 \setminus G_1; \ldots; x^*, G_m; x_*, E_m \setminus G_m] = x^*(\bigcup_{k=1}^m G_k)x_*$ satisfies $E_{P_i}(u_i \cdot f) = \alpha(P_i)$, and $E_p(u_0 \cdot f) = \alpha(p)$ for all $p \in \mathcal{C}$. Furthermore, there exists an event $F \in \mathcal{E}$ such that $\pi(F) = \hat{c}$, hence $E_{P_i}(u_i \cdot (x^* F x_*)) = \hat{c}$, and also $E_p(u_0 \cdot (x^* F x_*)) = \hat{c}$ for all $p \in \mathcal{C}$. The separation result thus yields $x^*(\bigcup_{k=1}^m G_k)x_* \succsim^0 x^* F x_*$, but at the same time $x^* F x_* \succ^i x^*(\bigcup_{k=1}^m G_k)x_*$, contradicting Avoidance of Societal Ambiguity. It follows that $conv\{P_1, \ldots, P_n\} \subseteq \mathcal{C}$. ■

It is concluded that $\mathcal{C} = conv\{P_1, \ldots, P_n\}$, and by the structure of the social priors set it follows that the social preference admits the representation, for every $f \in \mathcal{F}$,

$$V(f) = \min_{i \in N} E_{P_i}(u_0 \cdot f) \ .$$

In the other direction suppose that $\succsim^0$ is represented by a minimum expected utility over all individual priors, let F be some event and let E be a socially unambiguous event. If for some individual i, $P_i(E) = \pi(E) > P_i(F)$, then $\min_{j \in N} P_j(E) = \pi(E) > P_i(F) \geq \min_{j \in N} P_j(F)$, yielding $x^* E x_* \succsim^0 x^* F x_*$.

References

[1] Alon, S. and G. Gayer (2015), “Social maxmin expected utility”, working paper.

[2] Alon, S. and D. Schmeidler (2014), “Purely subjective Maxmin Expected Utility”, *Journal of Economic Theory*, 152, 382–412.

[3] Anscombe, F.J. and R.J. Aumann (1963), “A Definition of Subjective Probability”, *The Annals of Mathematical Statistics*, 34, 199-205.

[4] Billot, A., and V. Vergopoulos (2014), “Utilitarianism with Prior Heterogeneity", working paper.

[5] Blume, L. E., T. Cogley, D. A. Easley, T. J. Sargent, and V. Tsyrennikov (2013), “Welfare, Paternalism, and Market Incompleteness", working paper.

[6] Brunnermeier, M. K., A. Simsek, and W. Xiong (2014), “A Welfare Criterion for Models with Distorted Beliefs”, *The Quarterly Journal of Economics*, 129, 1753-1797.

[7] Chambers, C.P., and T. Hayashi (2006), “Preference aggregation under uncertainty: Savage vs. Pareto”, *Games and Economic Behavior*, 54, 430–440.

[8] Cres, H., I. Gilboa, and N. Vieille (2011), “Aggregation of multiple prior opinions”, *Journal of Economic Theory*, 146, 2563-2582.

[9] Danan, E., Gajdos, T., Hill, B., and J.-M. Tallon (2014), “Aggregating Tastes, Beliefs, and Attitudes under Uncertainty", working paper.

[10] Danan, E., T. Gajdos, and J.-M Tallon (2015), “Harsanyi's Aggregation Theorem with Incomplete Preferences", *American Economic Journal: Microeconomics*, 7, 61-69.

[11] De Meyer, B., and P. Mongin (1995), “A Note on Affine Aggregation”, *Econonic Letters*, 47, 177-83.

[12] Diamond, P.A. (1967), “Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility: a comment”, *Journal of Political Economy*, 75, 765–766.

[13] Dunford N., and J.T Schwartz (1957), Linear Operators, Part I. Interscience, New York.

[14] Epstein, L.G. (1999), “A definition of uncertainty aversion", *Review of Economic Studies*, 66, 579–608.

[15] Fishburn, Peter C. (1970), Utility Theory for Decision Making. Wiley, New York.

[16] Gajdos T., J.-M. Tallon, and J.-C. Vergnaud (2008), “Representation and aggregation of preferences under uncertainty”, *Journal of Economic Theory*, 141, 68-99.

[17] Gayer. G., I. Gilboa, L. Samuelson, and D. Schmeidler (2014), “Pareto Efficiency with Difference Beliefs”, *Journal of Legal Studies*, forthcoming.

[18] Ghirardato P. and M. Marinacci (2002), “Ambiguity Made Precise: A Comparative Foundation”, *Journal of Economic Theory*, 102, 251-289.

[19] Gilboa, I., F. Maccheroni, M. Marinacci and D. Schmeidler, “Objective and Subjective Rationality in a Multiple Prior Model”, *Econometrica*, 78, 755-770.

[20] I. Gilboa, D. Samet, and D. Schmeidler (2004), “Utilitarian aggregation of beliefs and tastes ”, *Journal of Political Economy*, 112, 932-938.

[21] Gilboa, I., L. Samuelson, and D. Schmeidler (2014), “No-Betting Pareto Dominance”, *Econometrica*, 82, 1405-1442.

[22] Gilboa, I. and D. Schmeidler (1989), “Maxmin expected utility with nonunique prior“, *Journal of Mathematical Economics*, 18, 141-153.

[23] Harsanyi, J. C. (1955), “Cardinal welfare, individualistic ethics, and interpersonal comparisons of utility”, *Journal of Political Economy*, 63, 309-321.

[24] Hylland, A., and R. Zeckhauser (1979), “The impossibility of Bayesian group decision making with separate aggregation of beliefs and values“, *Econometrica*, 47, 1321–1336.

[25] E. Karni, D. Schmeidler, and K. Vind (1983), “On State Dependent Preferences and Subjective Probabilities“, *Econometrica*, 51, 1021-1031.

[26] Kornhauser, L. A. (1992), “Modeling collegial courts II. Legal doctrine”, *Journal of Law, Economics and Organization*, 8, 441-470.

[27] Kornhauser, L. A., and L. G. Sager (1986), “Unpacking the court”, Yale Law Journal, 96, 82-117.

[28] List, C., and P. Pettit (2002), “Aggregating sets of judgments. An impossibility result”, *Economics and Philosophy*, 18, 89-110.

[29] Maccheroni F., and M. Marinacci (2001), “A Heine-Borel theorem for $ba(\Sigma)$”, RISEC.

[30] Mongin, P. (1995), “Consistent Bayesian aggregation”, *Journal of Economic Theory*, 66, 313-351.

[31] Mongin, P. (1997), “Spurious Unanimity and the Pareto Principle”, *Economics and Philosophy*, forthcoming.

[32] Mongin, P. (1998), “The paradox of the Bayesian experts and state-dependent utility theory”, *Journal of Mathematical Economics*, 29, 331-361.

[33] Mongin, P., and M. Pivato (2014), “Ranking Multidimensional Alternatives and Uncertain Prospects”, working paper.

[34] Nascimento, L. (2012), “The ex ante aggregation of opinions under uncertainty”, *Theoretical Economics*, 7, 535-570.

[35] Posner, E. A. and G. Weyl (2012), “An FDA for Financial Innovation: Applying the Insurable Interest Doctrine to 21st-Century Financial Markets", working paper.

[36] Qu, X. (2014), “Separate Aggregation of Beliefs and Values under Ambiguity”, working paper.

[37] Raiffa, H. (1968), Decision Analysis. Addison–Wesley, Reading, MA.

[38] Savage, L.J. (1952, 1954), The Foundations of Statistics. Wiley, New York (2nd ed. 1972, Dover, New York).

[39] Schmeidler, D. (1989), "Subjective probability and expected utility without additivity", *Econometrica*, Vol. 57, No. 3, 571-587.

[40] Villegas, C. (1964), "On Qualitative Probability Sigma-Algebras", *Annals of Mathematical Statistics*, 35, 1787-1796.

[41] Weyl, G. (2007), "Is Arbitrage Socially Beneficial?", working paper.