



HAL
open science

Utilitarian Aggregation with Heterogeneous Beliefs *

Antoine Billot, Xiangyu Qu

► **To cite this version:**

Antoine Billot, Xiangyu Qu. Utilitarian Aggregation with Heterogeneous Beliefs *. American Economic Journal: Microeconomics, 2020, 13 (3), pp.112-123. 10.1257/mic.20180344 . hal-03034701

HAL Id: hal-03034701

<https://hal.science/hal-03034701v1>

Submitted on 17 Sep 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Utilitarian Aggregation with Heterogeneous Beliefs[†]

By ANTOINE BILLOT AND XIANGYU QU*

The utilitarian aggregation rule requires social utility and beliefs to be a convex combination of individual utilities and beliefs, respectively. Since, in the case of belief heterogeneity, the standard Pareto condition is incompatible with such a separate aggregation, a new condition, called the belief-proof Pareto condition, is proposed to alleviate occurrences of spurious agreement by restricting unanimity to beliefs that can be considered reasonable by society. Then, we show, in the Anscombe-Aumann and the Savage framework, that the belief-proof Pareto condition is equivalent to separate aggregation of individual beliefs and tastes. (JEL D11, D71, D83)

Harsanyi (1955) proposes an axiomatic justification of the utilitarian aggregation rule that is based, on the one hand, on Bayesian rationality and, on the other hand, on the Pareto condition. These two principles are well known to imply social utility to be a combination of individual utilities (Harsanyi's aggregation theorem). Nevertheless, these two requirements lead to an impossibility result in a framework à la Savage (1954), where individual beliefs are subjective and possibly heterogeneous (see, e.g., Hylland and Zeckhauser 1979 and Mongin 1995). In the case of belief heterogeneity, either Bayesian rationality, i.e., subjective expected utility, as imposed on society, or the standard Pareto condition should be rejected. At least two main approaches have been suggested to address this impossibility. On the one hand, the *non-Bayesian* approach calls for retaining the original Pareto condition in full strength and relaxing the assumption of social preferences consistent with the *subjective expected utility model* (SEU).¹ On the other hand, the *Bayesian* approach advised by Mongin (1995) calls for retaining the assumption of SEU social preferences and, consequently, criticizes the Pareto condition as a normative principle.

At first glance, the Pareto condition seems to be uncontroversial. However, Gilboa, Samet, and Schmeidler (2004)—henceforth, GSS—produce an example, a duel, in which both agents prefer fighting to not fighting, although based on conflicting beliefs. Such unanimity without unanimous reasons—called *spurious unanimity* by

*Billot: Université Panthéon-Assas, Lemma (email: billot@u-paris2.fr); Qu: CNRS, Centre d'Economie de la Sorbonne and Wuhan University of Technology, School of Management (email: xiangyuqu@gmail.com). Michael Ostrovsky was coeditor for this article. Billot is grateful to the project Labex MMEII (ANR 11-LABX-0033-01) for its financial support. Qu is grateful to the ANR CHOp project (ANR-17-CE26-0003) for its financial support. The authors are very grateful to Itzhak Gilboa, Marcus Pivato, Philippe Mongin, David Schmeidler, and the anonymous referees for very helpful comments and suggestions.

[†]Go to <https://doi.org/10.1257/mic.20180344> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

¹For instance, Crès, Gilboa, and Vieille (2011); Alon and Gayer (2016); and Qu (2017) propose that social preferences admit a *maxmin expected utility* representation, while Danan et al. (2016) propose that social preferences are Bewley.

Mongin (1995)—has indeed no true normative standing. The standard Pareto condition cannot be considered compelling in this situation. Specifically, the impossibility of preference aggregation in the case of belief heterogeneity reveals the existence of inadmissible beliefs. A natural way to adapt the Pareto condition is to remove these inadmissible profiles. GSS accordingly propose a solution that corresponds to a weakening of the Pareto condition such that preference unanimity can only be considered conclusive if unanimous preferences are based on agreed-upon beliefs.²

The purpose of this paper is to suggest a new criterion for belief restriction. A society that does not want to (or cannot) discriminate among individuals is led to adopt something of an “electoral” principle: one man, one prior. Consequently, we assume here that unanimity is restricted to all cases in which individual preferences remain unchanged with respect to any other individual prior. Formally, individuals and society have expected utility preferences. We then single out a new Pareto condition, called the *belief-proof Pareto condition*. This condition entails that society restricts unanimity only to those preferences that do not depend on individual beliefs, that is, which remain unchanged for all individual beliefs. Then, the belief-proof Pareto condition is proved to imply that social utility is a convex combination of individual utilities and social beliefs are a convex combination of individual beliefs.

I. Motivation

Let us revisit the famous GSS example of a duel, showing that the standard Pareto condition may not be very plausible for individuals with heterogeneous priors. Suppose that society \mathcal{I} consists of two duelists, Alice and Bob, equipped with SEU preferences u_a and u_b and priors π_a and π_b over set of states Ω . Suppose that Ω consists of two possible states $\{A, B\}$, where A corresponds to Alice’s win and B to Bob’s win. The set X of possible outcomes is defined by the triplet $\{x, y, z\}$, where x means that Alice is alive but Bob is dead, y means that Bob is alive but Alice is dead, and z means that there is no duel. Utilities and beliefs are described in the following matrices.

	X			
u	x	y	z	
u_a	1	-5	0	
u_b	-5	1	0	

and

	Ω		
π	A	B	
π_a	0.9	0.1	
π_b	0.1	0.9	

Now, compare the *duel act* $f = xAy$ with the constant *nonduel act* z . Since both individual preferences are assumed to be SEU, it is straightforward that both individuals prefer fighting to not fighting. Should society necessarily consider that it is compelling to adopt such a preference, that is, to prefer the duel? In agreement with many others, especially GSS, we do not believe this is necessarily the case. Observe that this apparent unanimity is based on conflicting individual utilities and beliefs.

²However, for being applied within SEU, this suggestion needs individual priors to conform to countable additivity. Recent versions of the Pareto condition are also defined as restrictions: while GSS restrict Pareto dominance to common-belief acts, Danan et al. (2016), for instance, restrict it to common-taste acts.

In fact, no single probability can be found such that both individuals prefer the duel to the nonduel. Although society would believe *ex ante* that there exists a reasonable prior supporting the duel, in this example, no hypothetical prior can rationalize such apparent unanimity. Such unanimity is therefore spurious, and indeed, it is not appealing for society to comply with it.

An essential question is the following: what are the persuasive reasons for society to comply with unanimity preferences? One might suggest that if there exists a probability distribution such that unanimity preferences can be rationalized, then it should be compelling for society to adopt these preferences. However, the choice of an arbitrary probability violates methodological individualism since methodological individualism requires society to be neutral and respectful of individual motives.³ According to this principle, it is natural for society to consider all individual beliefs reasonable and to accommodate them rather than choosing at random any discretionary prior. Consequently, for society, there is *a priori* no difference between the set of individual beliefs and the set of all reasonable priors. Hence, since, on the one hand, any individual prior is reasonable, and on the other hand, to be compelling, unanimity preferences are expected to preserve unanimity under reasonable beliefs, unanimity preferences should therefore preserve unanimity under any individual prior. This is consistent with the modern postulate whereby a democratic society should commit itself to not discriminate among individuals. Society is then supposed to restrict unanimity to the only situation where this principle can be active, that is, whenever individual preferences remain unchanged with respect to any individual prior.

In the original GSS duel example, unanimity is not preserved for both priors. Now, consider a slightly modified duel example where individual utilities and beliefs are given as follows.

u	X	x	y	z	
u_a		10	-1	0	
u_b		-1	10	0	

and

π	Ω	A	B
π_a		0.9	0.1
π_b		0.1	0.9

The two duelists improve their respective utilities for x and y , while they maintain the same ranking among x , y , and z . Both individuals still prefer the duel act $f = xAy$ to the nonduel act z . However, this unanimous preference can now be supported by any reasonable prior, i.e., π_a and π_b . If we switch individual beliefs from π_a to π_b for Alice and vice versa for Bob, both individual preferences for f and z remain unchanged and unanimous. Such unanimity is clearly independent of individual beliefs in the sense that all permutations between priors do not modify individual preferences. Then, it seems sensible for society to adopt it.

This situation is reminiscent of a strategy-proof game, namely, a game that is immune to information manipulation. By analogy, we called *belief-proof Pareto* a condition such that, every time unanimity preferences can be rationalized by all individual priors, society complies with it—that is, society is immune to belief

³ See Hayek (1948) for a detailed discussion of the epistemological issue of individualism when studying social phenomena and Arrow (1994) for a more recent discussion.

manipulation. In contrast to the standard Pareto condition where unanimity is required with respect to individual rationality, the belief-proof Pareto condition instead requires a kind of social rationality in the sense that none of the individual priors prevents unanimity preferences. To a certain extent, in our view, this constitutes a minimal democratic requirement for a society consisting of individuals with heterogeneous beliefs.

The paper proceeds as follows: in Section II, the normative background is described, and the belief-proof Pareto condition is displayed. Then, we prove (under a mild assumption of minimal agreement) that it implies separate aggregation in the Anscombe-Aumann model. In Section III, the belief-proof Pareto condition is updated for the fully subjective expected utility model. Then, we prove that it is also equivalent to separate aggregation. In Section IV, the assumption of minimal agreement is relaxed. Two slightly stronger conditions are introduced, and we show that these conditions are also equivalent to separate aggregation. In Section V, the way these conditions relate to alternative conditions is discussed. In particular, we speculate on the meaning of the choice between a “pure” Savage setting or a “Savage-Arrow” setting. This difference appears to be central to understanding the relation between our results and those of GSS.

II. Anscombe-Aumann Preferences

We use Fishburn and Fishburn’s (1970) version of the Anscombe-Aumann (AA) model. The set of *states of the world* is Ω , and it is endowed with an algebra \mathcal{A} of *events*. Let X be a set of *outcomes*. Let Δ be the set of all the *probability distributions* on X with a finite support. An *act* is a finite-outcome function $f: \Omega \rightarrow \Delta$.⁴ The set of all acts is denoted by $\mathcal{F}(\Delta)$. Society is a finite set of individuals $\mathcal{I} = \{1, \dots, I\}$. Individual $i \in \mathcal{I}$ has *preferences* $\succsim_i \subset \mathcal{F}(\Delta) \times \mathcal{F}(\Delta)$, whereas social preferences are denoted by $\succsim_0 \subset \mathcal{F}(\Delta) \times \mathcal{F}(\Delta)$. For $0 \leq i \leq I$, the relations \sim_i and \succ_i are defined as the symmetric and asymmetric parts of \succsim_i . We assume that each preference relation is represented by an *expected utility function* $\mathbf{E}_i: \mathcal{F}(\Delta) \rightarrow \mathbb{R}$: for $0 \leq i \leq I$, there exists a utility function u_i on X and a finitely additive probability measure π_i on \mathcal{A} such that every act f is ranked on the basis of $\mathbf{E}_i(f) = \int_{\Omega} u_i(f) d\pi_i$.⁵ Moreover, (u_i, π_i) is the unique pair that represents \succsim_i (up to an increasing affine transformation of u_i).

We introduce a mild restriction of minimal agreement, which will be relaxed in Section IV.

Minimal Agreement Outcome (MAO).—There are $z, z' \in X$ such that $z \succ_i z'$ for all $0 \leq i \leq I$.

Note that MAO implies immediately that each utility function is nonconstant. Owing to the affine utility function, we can normalize the utility function by assuming $u_i(z) = 1$ and $u_i(z') = 0$ for $0 \leq i \leq I$.

⁴ f is a finite-outcome function if its outcome set $f(\Omega) = \{f(\omega) | \omega \in \Omega\}$ is finite.

⁵Actually, $\mathbf{E}_i(f) = \int_{\Omega} (\sum_{x \in X} u_i(x) f(\omega)(x)) d\pi_i(\omega)$. By a slight abuse of notation, we also define $u_i: \Delta \rightarrow \mathbb{R}$ by $u_i(P) = \sum_{x \in X} u_i(x) P(x)$. Using this definition, \mathbf{E}_i can be written exactly as above.

In the face of heterogeneous beliefs, it is important for society to form an acceptable probability measure to evaluate every possible act. However, we assume that society does not possess any extra information. Therefore, social beliefs rely solely on individual probability estimations. Without taking a stance on which individual beliefs are correct, society could not rule out any individual prior by proving it wrong. Instead, it is natural for society to consider each individual prior as a *reasonable* probability estimation. The insight from the following Pareto condition is that the unanimity condition applies if this unanimity is invariant with respect to every individual prior. In that sense, it recalls strategy-proofness in game theory since this form of unanimity is immune to any belief manipulation. Specifically, we propose that if each individual prefers an act to another act under every reasonable belief, then so does society. For $f \in \mathcal{F}$, we denote by

$$\mathbf{E}_{ij}(f) = \int u_i(f) d\pi_j$$

the *virtual* expected utility of f with respect to i 's utility and j 's belief.

Belief-Proof Pareto Condition (AA).—For all acts f and g in $\mathcal{F}(\Delta)$, if, for every $i, j \in \mathcal{I}$, $\mathbf{E}_{ij}(f) \geq \mathbf{E}_{ij}(g)$, then $\mathbf{E}_0(f) \geq \mathbf{E}_0(g)$.

If individuals have common beliefs, as in Harsanyi (1975), that is, if $\pi_i = \pi$ for every $i \in \mathcal{I}$, then the belief-proof Pareto condition (BPPC) coincides with the standard Pareto condition. The standard Pareto condition states that if every individual prefers one act to another act, then so does society. However, to rule out spurious unanimities, our criterion requires that every individual should not change her choice between two options even if her prior is replaced by any other individual prior. This principle makes social choice rely not only on the direct preferences comparison but also on the belief-adjusted preferences comparison, which leaves no room for spurious unanimity.⁶

THEOREM 1: *Assume MAO. BPPC (AA) is satisfied if and only if π_0 is a convex combination of $\{\pi_i\}_{i=1}^I$ and u_0 is a convex combination of $\{u_i\}_{i=1}^I$.*

PROOF:

The necessity part is straightforward. We only prove that BPPC (AA) implies that u_0 and π_0 are convex combinations of $\{u_i\}_{i=1}^I$ and $\{\pi_i\}_{i=1}^I$, respectively. Note that for each $i, j \in \mathcal{I}$, every virtual expected utility \mathbf{E}_{ij} is an affine function on $\mathcal{F}(\Delta)$. Due to the convexity of set $\mathcal{F}(\Delta)$, Harsanyi's aggregation theorem applies.⁷ That is, there exist $\gamma_{ij} \geq 0$ with some γ_{ij} strictly positive such that, for all $f \in \mathcal{F}$,

$$\mathbf{E}_0(f) = \sum_{i,j \in \mathcal{I}} \gamma_{ij} \mathbf{E}_{ij}(f).$$

⁶By symmetry, we suggest that each individual taste can be a reasonable representation of social taste. As a result, society is supposed to restrict unanimity only to cases in which individual preferences remain unchanged with respect to any other individual taste. It is clear that such a modified condition also leads to a utilitarian representation of preferences. We thank a referee for highlighting this point.

⁷For a more detailed result, see De Meyer and Mongin (1995).

Let $\alpha_i = \sum_j \gamma_{ij}$. Then, for every $x \in X$

$$u_0(x) = \sum_{ij} \gamma_{ij} u_i(x) = \sum_i \alpha_i u_i(x).$$

Recall, by MAO, that there exist $z, z' \in X$ such that $u_i(z) = 1$ and $u_i(z') = 0$, for each $0 \leq i \leq I$. As a result, we have

$$u_0(z) = \sum_i \alpha_i u_i(z) \Rightarrow \sum_i \alpha_i = 1.$$

Therefore, social utility is a convex combination of individual utilities. Furthermore, we have

$$\sum_{ij} \gamma_{ij} = 1.$$

Let $\beta_j = \sum_i \gamma_{ij}$. Now select any event E in \mathcal{A} . Then, we have

$$\pi_0(E) = \mathbf{E}_0(zEz') = \sum_{ij} \gamma_{ij} \mathbf{E}_{ij}(zEz') = \sum_{ij} \gamma_{ij} \pi_j(E) = \sum_j \beta_j \pi_j(E).$$

It is then immediate that $\pi_0 = \sum_j \beta_j \pi_j$. Hence, the social belief is also a convex combination of individual beliefs. ■

Theorem 1 illustrates the relationship between BPPC (AA) and the separate aggregation rule. Note that BPPC (AA) is more general than the standard Pareto condition because the unanimity condition is applied to both I individuals and $I(I - 1)$ virtual individuals.⁸

III. Savage Preferences

In this section, we consider a Savage setup. Let Ω be a set of *states of the world*. The algebra of all subsets of Ω is denoted by \mathcal{A} . Let X be a set of *outcomes*. An *act* is a finite-outcome function $f: \Omega \rightarrow X$. The set of all acts is denoted by \mathcal{F} . Society is a finite set of individuals $\mathcal{I} = \{1, \dots, I\}$. Individual $i \in \mathcal{I}$ has preferences $\succsim_i \subset \mathcal{F} \times \mathcal{F}$, whereas social preferences are denoted by $\succsim_0 \subset \mathcal{F} \times \mathcal{F}$. For $0 \leq i \leq I$, the relations \sim_i and \succ_i are defined as the symmetric and asymmetric parts of \succsim_i , as usual. We assume that each preference relation is represented by a subjective expected utility SEU function $\mathbf{E}_i: \mathcal{F} \rightarrow \mathbb{R}$. For $0 \leq i \leq I$, there exists a utility function u_i on X and a finitely additive and nonatomic probability measure π_i on \mathcal{A} , such that every act f is evaluated by $\mathbf{E}_i(f) = \int_{\Omega} u_i(f) d\pi_i$.⁹ Moreover, (u_i, π_i) is the unique pair that represents \succsim_i (up to an increasing affine transformation of u_i). As in the preceding section, MAO is assumed. As above, we normalize

⁸In our framework, MAO is necessary to exclude a constant social utility function (Qu 2017). To achieve separate aggregation without MAO, the Pareto condition must be further strengthened. See Section IV.

⁹A probability measure π is said to be *nonatomic* if, for every event E with $\pi(E) > 0$ and any $\alpha \in (0, 1)$, there exists an event $F \subset E$ such that $\pi(F) = \alpha\pi(E)$.

the expected utility representation by assuming that $u_i(z) = 1$ and $u_i(z') = 0$ for each $0 \leq i \leq I$.

Note that since the set \mathcal{F} of acts is not convex, the Pareto condition and technique developed in the AA framework are no longer valid in a Savagian framework.

Fix $0 \leq i \leq I$. For $f \in \mathcal{F}$, $Y \subseteq X$, define a probability distribution λ_i^f on X by

$$\lambda_i^f(Y) = \pi_i(\{\omega \in \Omega : f(\omega) \in Y\}).$$

That is, $\lambda_i^f(Y)$ is the probability that the outcome is in Y if act f is taken. Let $\Delta(X)$ be the set of simple probabilities on X .¹⁰ Then, the set of all act-induced probability distributions over outcomes equals $\Delta(X)$. (See theorem 14.3 in Fishburn and Fishburn 1970 for a proof.)

LEMMA 1: For $0 \leq i \leq I$, $\{\lambda_i^f | f \in \mathcal{F}\} = \Delta(X)$.

Fix $0 \leq i \leq I$. We can use \succsim_i to define a binary relation \succsim_i^* on $\Delta(X)$. If $p \in \Delta(X)$, Lemma 1 allows us to construct an act $f \in \mathcal{F}$ such that $p = \lambda_i^f$. Therefore, for p and q in $\Delta(X)$, let f, g be such that $p = \lambda_i^f$ and $q = \lambda_i^g$, and define \succsim_i^* as follows:

$$p \succsim_i^* q \quad (\text{or, equivalently } \lambda_i^f \succsim_i^* \lambda_i^g) \quad \text{if } f \succsim_i g.$$

SEU representation of \succsim_i implies that \succsim_i^* is well defined for any f and g such that $\lambda_i^f = p$ and $\lambda_i^g = q$. The equivalence between the two relations refers to proposition (9.15) in Kreps (2018).

LEMMA 2: For $0 \leq i \leq I$, for $p, q \in \Delta(X)$, let f and g be such that $p = \lambda_i^f$ and $q = \lambda_i^g$. Then, $p \succsim_i^* q$ if and only if $\mathbf{E}_i(f) \geq \mathbf{E}_i(g)$.

Seek now a condition implying separate aggregation in a Savage framework. Assume that for three acts f, g , and g' , all individuals agree that a half-half mixture between the induced probabilities of g and g' cannot be better than f . That is, in terms of preferences over induced probability distributions, for each $i \in \mathcal{I}$,

$$\lambda_i^f \succsim_i^* \frac{1}{2}\lambda_i^g + \frac{1}{2}\lambda_i^{g'}.$$

In addition, we assume that every j also agrees with i about these induced probability distributions. That is, for each $i, j \in \mathcal{I}$,

$$\lambda_i^f \succsim_j^* \frac{1}{2}\lambda_i^g + \frac{1}{2}\lambda_i^{g'}.$$

¹⁰We say a probability is *simple* if it contains finite outcomes.

Then, we argue that it stands to reason that society would reach the same conclusion, namely,

$$\lambda_0^f \succsim_0^* \frac{1}{2} \lambda_0^g + \frac{1}{2} \lambda_0^{g'}$$

The following axiom generalizes this example. If all individuals prefer a mixture of a set of induced acts to a mixture of another set of induced acts with respect to every individual belief, then so does society.

Belief-Proof Pareto Condition (S).—For every act $f_n, g_m \in \mathcal{F}$, $n = 1, \dots, N$, $m = 1, \dots, M$, and every pair of numbers $\alpha_n \geq 0$ and $\beta_m \geq 0$ such that $\sum_n \alpha_n = 1$ and $\sum_m \beta_m = 1$, if $\sum_n \alpha_n \lambda_i^{f_n} \succsim_j^* \sum_m \beta_m \lambda_i^{g_m}$ for every $i, j \in \mathcal{I}$, then $\sum_n \alpha_n \lambda_0^{f_n} \succsim_0^* \sum_m \beta_m \lambda_0^{g_m}$.

Observe that since $\lambda_i^{f_n}, \lambda_i^{g_m} \in \Delta(X)$ for all n, m , their mixtures, $\sum_n \alpha_n \lambda_i^{f_n}$ and $\sum_m \beta_m \lambda_i^{g_m}$, are also in $\Delta(X)$. Thus, BPPC (S) means that if each individual prefers a certain mixing of act-induced probability distributions to another mixing, irrespective of the reasonable beliefs that are chosen, then society also prefers the former to latter mixture. However, the probability distributions involved will typically vary across individuals and society. What relates $\lambda_i^{f_n}$ ($\lambda_i^{g_m}$) across individuals is the fact that each of them is obtained by the same mixture α_n (β_m) of induced probability distributions of the same acts f_n (g_m). Since, however, individual priors in general do not coincide, these induced probability distributions will typically vary. Finally, note that BPPC (S) implies BPPC (AA).

THEOREM 2: *Assume MAO. BPPC (S) is satisfied if and only if π_0 is a convex combination of $\{\pi_i\}_{i=1}^I$ and u_0 is a convex combination of $\{u_i\}_{i=1}^I$.*

PROOF:

Since the necessity part is straightforward, we only prove the sufficiency part. Let

$$\Delta(\mathcal{F}) = \left\{ (\alpha_1, f_1; \dots; \alpha_N, f_N) \mid N \geq 1, \sum_{n=1}^N \alpha_n = 1, \alpha_n \geq 0, f_n \in \mathcal{F} \right\}$$

be the set of finite-outcome probability distributions over \mathcal{F} . Clearly, $\Delta(\mathcal{F})$ is a convex set. Now, define a function $\tilde{\mathbf{E}}_{ij}$ on $\Delta(\mathcal{F})$, for every $F = (\alpha_1, f_1; \dots; \alpha_N, f_N)$, as follows:

$$\tilde{\mathbf{E}}_{ij}(F) = \sum_n \alpha_n \mathbf{E}_{ij}(f_n).$$

We can define $\tilde{\mathbf{E}}_0$ in a similar way. Consider any pair of acts $F = (\alpha_1, f_1; \dots; \alpha_N, f_N)$ and $G = (\beta_1, g_1; \dots; \beta_M, g_M)$ in $\Delta(\mathcal{F})$, and suppose that, for every $i, j \in \mathcal{I}$,

$$\tilde{\mathbf{E}}_{ij}(F) \geq \tilde{\mathbf{E}}_{ij}(G).$$

By definition, it is equivalent to $\sum_n \alpha_n \mathbf{E}_{ij}(f_n) \geq \sum_m \beta_m \mathbf{E}_{ij}(g_m)$, that is,

$$\sum_n \alpha_n \int u_i(f_n) d\pi_j \geq \sum_m \beta_m \int u_i(g_m) d\pi_j.$$

By the linearity of the expected utility, it is also equivalent to

$$\int u_i(x) d\sum_n \alpha_n \lambda_j^{f_n} \geq \int u_i(x) d\sum_m \beta_m \lambda_j^{g_m},$$

which means that the necessity part of BPPC (S) is satisfied. Then, we must have

$$\int u_0(x) d\sum_n \alpha_n \lambda_0^{f_n} \geq \int u_0(x) d\sum_m \beta_m \lambda_0^{g_m},$$

which is equivalent to $\tilde{\mathbf{E}}_0(\sum_n \alpha_n f_n) \geq \tilde{\mathbf{E}}_0(\sum_m \beta_m g_m)$. We have already proven that the unanimity condition is satisfied with respect to $(\{\tilde{\mathbf{E}}_{ij}\}_{i,j \in \mathcal{I}}, \tilde{\mathbf{E}}_0)$ on the convex set $\Delta(\mathcal{F})$. Since each $\tilde{\mathbf{E}}_{ij}$ is an affine function, Harsanyi's aggregation theorem applies: there exist nonnegative numbers γ_{ij} for $i, j \in \mathcal{I}$ such that, for every $F \in \Delta(\mathcal{F})$, $\tilde{\mathbf{E}}_0(F) = \sum_{ij} \gamma_{ij} \tilde{\mathbf{E}}_{ij}(F)$. For acts $F = (1, f)$ on $\Delta(\mathcal{F})$, we have $\tilde{\mathbf{E}}_0(F) = \mathbf{E}_0(f)$ and $\tilde{\mathbf{E}}_{ij}(F) = \mathbf{E}_{ij}(f)$. As a result, we have, for every $f \in \mathcal{F}$,

$$\mathbf{E}_0(f) = \sum_{ij} \gamma_{ij} \mathbf{E}_{ij}(f).$$

Let $\alpha_i = \sum_j \gamma_{ij}$ and $\beta_j = \sum_i \gamma_{ij}$. Repeat a similar argument as in Theorem 1. It is then immediate that $u_0 = \sum_i \alpha_i u_i$ and $\pi_0 = \sum_j \beta_j \pi_j$. ■

IV. Disagreement

MAO was previously assumed for technical reasons. Now, we want to relax it since separate aggregation might be an appropriate rule for social contexts without any minimal agreement on outcomes. For that purpose, suppose, for simplicity, that individuals are AA expected utility maximizers as in Section II. If individual beliefs are known, it is possible for society to transform each act into a lottery. Therefore, individual preferences over acts can be automatically translated into individual preferences over lotteries. Whenever preferences over lotteries are provided, the Pareto condition over lotteries should be applied; that is, unanimous preference comparison for lotteries is compelling.

We define \succsim_i^* in a similar way as in Section III. Then, a so-called lottery Pareto condition can be formally formulated.

Lottery Pareto Condition (LPC).—For any $p, q \in \Delta(X)$, if $p \succsim_i^* q$ for every $i \in \mathcal{I}$, then $p \succsim_0^* q$.

Since LPC coincides with the original Pareto condition, the utilitarian result follows immediately.

PROPOSITION 1: *LPC is satisfied if and only if u_0 is an affine combination of $\{u_i\}_{i=1}^I$.*

The role of individuals in economic theory does not truly differ from the role of “these individualistic atoms of the rare gas in my balloon” (Samuelson 1966, 1411).¹¹ Consequently, it is not always necessary for society to fully respect individual beliefs. In this situation, LPC is sensible, and society would only be utilitarian with respect to individual tastes. The next axiom would be useful for the characterization of belief aggregation only.

Societal Pareto Condition (SPC).—For all $f, g \in \mathcal{F}(\Delta)$, if $\lambda_i^f \succsim_0^* \lambda_i^g$ for every $i \in \mathcal{I}$, then $f \succsim_0 g$.

SPC basically states that if society prefers an induced probability distribution of one act f to that of another act g with respect to every individual prior, then society should prefer f to g .

PROPOSITION 2: *SPC is satisfied if and only if π_0 is a convex combination of $\{\pi_i\}_{i=1}^I$.*

PROOF:

Denote the social *virtual* expected utility of f with respect to individual i 's prior by $\mathbf{E}_{0i}(f) = \int u_0(f) d\pi_i$. SPC can be rewritten as follows: for all $f, g \in \mathcal{F}(\Delta)$, if $\mathbf{E}_{0i}(f) \geq \mathbf{E}_{0i}(g)$, for all $i \in \mathcal{I}$, then $\mathbf{E}_0(f) \geq \mathbf{E}_0(g)$. Owing to the property of SEU, for $i \in \mathcal{I}$, \mathbf{E}_{0i} is an affine function on the convex set $\mathcal{F}(\Delta)$. Then, Harsanyi's aggregation theorem implies that there exist numbers $\beta_i \geq 0$ with $\sum_i \beta_i = 1$ such that $\mathbf{E}_0(f) = \sum_i \beta_i \mathbf{E}_{0i}(f)$. Hence, $\pi_0 = \sum_i \beta_i \pi_i$.

LPC and SPC imply BPPC (AA), but the opposite is not true. In other words, with a weaker condition, MAO is also a necessary condition for separate aggregation. ■

PROPOSITION 3: *LPC and SPC imply BPPC (AA).*

PROOF:

Let $f, g \in \mathcal{F}(\Delta)$. Suppose that for all $i, j \in \mathcal{I}$, $\mathbf{E}_{ij}(f) \geq \mathbf{E}_{ij}(g)$. Then, for all positive numbers α_i with $\sum_i \alpha_i = 1$, we have, for all j , $\sum_i \alpha_i \mathbf{E}_{ij}(f) \geq \sum_i \alpha_i \mathbf{E}_{ij}(g)$. Equivalently, we have

$$\int \left(\sum_i \alpha_i u_i(f) \right) d\pi_j \geq \int \left(\sum_i \alpha_i u_i(g) \right) d\pi_j.$$

According to LPC, this implies $\mathbf{E}_{0j}(f) \geq \mathbf{E}_{0j}(g)$ for all j . Hence, SPC directly implies $\mathbf{E}_0(f) \geq \mathbf{E}_0(g)$. ■

¹¹ See Arrow (1994) for a similar viewpoint.

V. Concluding Remarks

In a Bayesian environment, where both individuals and society are SEU, many alternative principles to guide society in its choices can be found in the literature. However, these principles do not all share the same methodological approach. To the best of our knowledge, the first alternative principle is suggested by Mongin (2005) in his working paper and later published in Mongin (2015), who proposes to consider that a unanimous preference comparison is relevant if and only if all individuals agree on both the probabilities and utility rankings underlying these preferences. Nevertheless, ordinal rankings do not convey any quantifiable information about the difference between evaluations. Therefore, unanimity preferences fail to be compelling insofar as homogeneous ordinal rankings do not constitute a persuasive basis for social decision. Alternatively, Billot and Vergopoulos (2016) propose defining the social space as the Cartesian product of individual spaces and assume a somewhat extended Pareto principle leading society to be utilitarian. Another approach is provided by Mongin and Pivato (2016). They distinguish between objective and subjective uncertainty and suggest possible solutions for situations where Pareto conditions can be translated through weighted additive utility representations.

In terms of motivation and content, this paper is close to GSS. Let us formally restate their *restricted Pareto condition*. An act f is said to be a *GSS lottery* if, for every $x \in X$, $\lambda_i^f(x) = \lambda_j^f(x)$ for each $i, j \in \mathcal{I}$.

Restricted Pareto Condition (RPC).—For every GSS lottery f and g , if $f \succsim_i g$ for every $i \in \mathcal{I}$, then $f \succsim_0 g$.

RPC requires that the Pareto condition only apply to GSS lotteries. It allows for separate aggregation only if identical events exist for every possible probability. This assumption is quite demanding and rather unrealistic even if when a GSS lottery is unanimously preferred to another one, then, according to BPPC, society should also adopt such a preference. In standard settings, Anscombe-Aumann and Savage, RPC is then too weak to imply the separate aggregation of individual utilities. In contrast, BPPC, though stronger than RPC, does not rely on the existence of these identical events. The economic meaning of the difference between BPPC and RPC can be expressed as follows: BPPC does not require the existence of events that are equally estimated by all individuals. Moreover, RPC seems to be too demanding for a reliable interpretation of the Pareto principle: for instance, there may not exist a GSS lottery required by RPC, even when an act unanimously first-order stochastically dominates another act. In contrast, under BPPC, in the case of unanimity preferences respectful of first-order stochastic dominance, society would evidently adopt these preferences. Indeed, this is not to dispute the usefulness of RPC. Gilboa, Samet, and Schmeidler (2004) show that RPC is a necessary and sufficient condition for separate aggregation in a “Savage-Arrow” setting, that is, a special Savage setting with a σ -additive probability measure over events. The extension of the Savage domain to the “Savage-Arrow” domain is particularly welcome in this context because σ -additivity guarantees the existence of identical events for all possible probabilities. However, both de Finetti (1974) and Savage (1954) point out that σ -additivity is not a natural requirement for

behavior.¹² Note that since BPPC is stronger than RPC, it is straightforward to see that the two principles are equivalent in a “Savage-Arrow” setting.

REFERENCES

- Alon, Shiri, and Gabi Gayer.** 2016. “Utilitarian Preferences with Multiple Priors.” *Econometrica* 84 (3): 1181–1201.
- Arrow, Kenneth J.** 1994. “Methodological Individualism and Social Knowledge.” *American Economic Review* 84 (2): 1–9.
- Billot, Antoine, and Vassili Vergopoulos.** 2016. “Aggregation of Paretian Preferences for Independent Individual Uncertainties.” *Social Choice and Welfare* 47 (4): 973–84.
- Crès, Hervé, Itzhak Gilboa, and Nicolas Vieille.** 2011. “Aggregation of Multiple Prior Opinions.” *Journal of Economic Theory* 146 (6): 2563–82.
- Danan, Eric, Thibault Gajdos, Brian Hill, and Jean-Marc Tallon.** 2016. “Robust Social Decisions.” *American Economic Review* 106 (9): 2407–25.
- de Finetti, Bruno.** 1974. *Theory of Probability: A Critical Introductory Treatment*. New York: Wiley.
- De Meyer, Bernard, and Philippe Mongin.** 1995. “A Note on Affine Aggregation.” *Economics Letters* 47 (2): 177–83.
- Fishburn, Peter C., and C. Simone Fishburn.** 1970. *Utility Theory for Decision Making*. Hoboken, NJ: Wiley.
- Gilboa, Itzhak, Dov Samet, and David Schmeidler.** 2004. “Utilitarian Aggregation of Beliefs and Tastes.” *Journal of Political Economy* 112 (4): 932–38.
- Harsanyi, John C.** 1955. “Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparisons of Utility.” *Journal of Political Economy* 63 (4): 309–21.
- Harsanyi, John C.** 1975. “Nonlinear Social Welfare Functions.” *Theory and Decision* 6 (3): 311–32.
- Hayek, Friedrich A.** 1948. *Individualism and Economic Order*. Chicago, IL: University of Chicago Press.
- Hylland, Aanund, and Richard Zeckhauser.** 1979. “The Impossibility of Bayesian Group Decision Making with Separate Aggregation of Beliefs and Values.” *Econometrica* 47 (6): 1321–36.
- Kreps, David M.** 2018. *Notes on the Theory of Choice*. New York: Routledge.
- Mongin, Philippe.** 1995. “Consistent Bayesian Aggregation.” *Journal of Economic Theory* 66 (2): 313–51.
- Mongin, Philippe.** 2005. “Spurious Unanimity and the Pareto Principle.” LSE Choice Group Working Paper Series vol. 1, no. 5.
- Mongin, Philippe.** 2015. “Spurious Unanimity and the Pareto Principle.” *Economics and Philosophy* 32 (3): 511–32.
- Mongin, Philippe, and Marcus Pivato.** 2016. “Social Preference under Twofold Uncertainty.” https://papers.ssrn.com/sol3/papers.cfm?abstract_id=2796560#.
- Qu, Xiangyu.** 2017. “Separate Aggregation of Beliefs and Values under Ambiguity.” *Economic Theory* 63 (2): 503–19.
- Samuelson, Paul.** 1966. *The Collected Scientific Papers of Paul A. Samuelson*, Vol. 2, edited by Joseph E. Stiglitz. Cambridge, MA: MIT Press.
- Savage, Leonard J.** 1954. *The Foundations of Statistics*. New York: Dover.
- Seidenfeld, T.** 2001. “Remarks on the Theory of Conditional Probability: Some Issues of Finite versus Countable Additivity.” In *Probability Theory: Philosophy, Recent History and Relations to Science*, edited by Vincent F. Hendricks, Stig Andur Pedersen, and Klaus Froyen Jørgensen, 167–78. New York: Springer.

¹²See Seidenfeld (2001) for a detailed discussion of σ -additivity.