No-Betting Pareto under Ambiguity*

Itzhak Gilboa† and Larry Samuelson‡

January 13, 2020

Abstract It has been argued that Pareto-improving trade is not as compelling under uncertainty as it is under certainty. The former may involve agents with *different* beliefs, who might wish to execute trades that are no more than betting. In response, the concept of *No-Betting Pareto* dominance was defined to capture Pareto improvements that can also be rationalized by *common* probabilities. In this paper we argue that this definition might be too narrow for use when agents are not Bayesian. Agents who face ambiguity might wish to trade in ways that can be justified by common ambiguity, though not necessarily by common probabilities. We accordingly extend the notion of No-Betting Pareto dominance to characterize trades than are "No-Betting Pareto" ranked according to the maxmin expected utility model.

Contents

*Gilboa gratefully acknowledges ISF Grant 1077/17, the Investissements d'Avenir ANR -11- IDEX-0003 / Labex ECODEC No. ANR - 11-LABX-0047, the AXA Chair for Decision Sciences at HEC and the Foerder Institute at Tel-Aviv University.

†Tel-Aviv University and HEC, Paris. tzachigilboa@gmail.com.

‡ Department of Economics, Yale University. Larry.Samuelson@yale.edu.

No-Betting Pareto with Ambiguity

January 13, 2020

1 Introduction

Voluntary trade is normally considered to be normatively appealing, reflecting reasons that range from ideological commitment to liberalism to the First Welfare Theorem. There are clearly exceptions having to do with the object traded, the circumstances of trade, or the agents involved—organ trading is generally considered "repugnant", insider trading is deceitful, and trade involving children may be unfair. Yet, when mentally healthy adults with symmetric information choose to trade monetary assets, many believe that they should be allowed to do so.

By contrast, betting is often frowned upon. Gambling can be addictive and become a psychological disorder. The gambling industry is often linked to criminal activities. Indeed, in most countries gambling is regulated, including those that boast freedom and liberalism. Betting is often viewed as a way of exploiting other people's mistaken views without creating value for society as a whole.

Where do we draw the line between trade and betting? There are many cases in which the distinction is simple. When an agent buys insurance for her car from an insurance company, we have voluntary trade that can easily be justified as Pareto improving among agents with common and relatively objective beliefs. One party starts off with considerable uncertainty and manages to reduce it by trading with another party who is either less risk averse or better able to diversify. On the other hand, when two people bet on the outcome of a sports match, they start out with no uncertainty in consumption, and both agree to introduce uncertainty to their bundle without any net production. These trades can typically be justified only if the agents entertain different beliefs.[1]

In between, there are cases that are harder to classify. Gilboa, Samuelson, and Schmeidler (2014) suggest the example of an entrepreneur who seeks the investment of a venture capitalist. This case is similar to the insurance example, as there are initially different levels of uncertainty, and the entrepreneur can be viewed as buying insurance against commercial failure. However, it is also similar to the betting example in that one of the parties involved might well be more optimistic than the other. Also unlike the car insurance problem, there are no probabilities that may be deemed reasonably objective.

Gilboa, Samuelson, and Schmeidler (2014) propose to refine Pareto domination in a way that rules out pure betting: beyond the agents' preference for trade, the definition requires that there be at least one probability according to which all agents concerned are better off with trade than without. In the simple case of objectively given probabilities, as in the car insurance problem, the shared probability can obviously serve as the probability according to which

[1]By "typically" we refer mostly to the common assumption of (weak) risk aversion.

all agents are better off with trade. In the equally simple case of pure betting, no such probability exists. Finally, in mixed cases the definition arguably captures the intuitive idea of ruling out bets, but *only* bets (Gilboa, Samuelson, and Schmeidler (2014, Theorem 2)). For example, in the entrepreneur-venture capitalist example, even if one is more optimistic than the other, one can typically point to a shared probability (in this case, the optimist's), according to which both parties will benefit from trade.

One may, however, take issue with the requirement that trade can be justifiable by a "shared belief" that is necessarily probabilistic. Suppose two agents own different properties whose values depend on the oceans' level in the future. Both agents are concerned about the value of their future holdings, but they do not exactly know the probability distribution of the oceans' level. They read studies, consult experts, and form estimates, but obtain answers that vary, concluding that they are in a state of ambiguity (aka Knightian uncertainty).[2] Let us also assume that this ambiguity matters, and that the need to split uncertainty goes beyond risk sharing that can be explained by a concave utility function.[3] A third agent, say an insurance company, offers an ambiguity-sharing market and suggests the two agents insure their properties for a fee.

It appears that a three-way trade, by which the insurance company allows each of the agents to hold a proportion of both assets, where these proportions add up to less than unity, is not a bet. The agents have precisely the same beliefs and agree on all predictions, but do not have enough hard data to calculate a single, additive probability measure that would describe the effects of global warming. No single shared probability would make both of them better off (by the expected utility criterion under Bayesian beliefs), but there are shared ambiguity-averse models of beliefs that can simultaneously justify trade for all agents. The proposed trade is then well within the realm of "shared risks." Thus, we may find that the no-betting criterion proposed by Gilboa, Samuelson, and Schmeidler (2014) is too strict for some purposes.

This paper accordingly offers a generalization of the no-betting Pareto dominance relation of Gilboa, Samuelson, and Schmeidler (2014). We weaken the shared-justification requirement to allow for justifications that are based on ambiguity. After a formal definition of the concept for maxmin expected utility, we present some examples that illustrate the new concept. We then move on to our main result, which is a characterization of No-Betting Pareto improvements for the maxmin expected utility model, and conclude with a discussion.

2 Definitions

Let there be a set of agents $A = \{1, ..., m\}$ and a set of states of the world, $S = \{1, ..., n\}$. Let k index agents and i index states. An allocation $x = \left(x_i^k\right)_{i \leq n, k \leq m} \in \mathbb{R}^{m \times n}$ is a utility profile in which agent k receives x_i^k in state

[2] See Millner, Dietz, and Heal (2013).

[3] That is, assume that the payoff in this example are already given in expected utility terms, as measured by equivalent lotteries with known distributions.

i. For $k \leq m$, x^k denotes the vector $\left(x_i^k\right)_{i\leq n} \in \mathbb{R}^n$ that is agent k's bundle in the allocation x. Agent k is characterized by a preference relation $\succsim_k$ over her payoff vectors.

For a probability vector $p \in \Delta(S) \simeq \Delta^{n-1}$ and a bundle x^k the inner product $p \cdot x^k = \sum_{i\leq n} p_i x_i^k$ will be agent k's expected payoff under x, according to p. The x values are considered to be observable. The two leading interpretations are that (i) x_i^k is monetary, and the agents are risk neutral; and (ii) x_i^k is given in utility terms (which would be vNM utilities under risk), and are observable.

We assume that each $\succsim_k$ is a maxmin expected utility (MEU) relation. Specifically, for $k \in A$ there exists a closed (nonempty) and convex set of probabilities $C_k \subseteq \Delta(S)$ such that, for all $x^k, y^k \in \mathbb{R}^n$,

$$y^k \succsim_k x^k \qquad \textit{iff} \qquad \min_{p\in C_k} p \cdot y^k \geq \min_{p\in C_k} p \cdot x^k.$$

Clearly, if a set C_k is a singleton, agent k is Bayesian.

We say that y *improves* upon x if for all $k \in A$, we have $y^k \succsim_k x^k$. Note that this allows for the possibility that all agents are indifferent between y and x, and it is indeed a reflexive relation. There are two reasons for using weak preferences in the definition of "improvement". First, this simplifies the discussion, obviating the need to distinguish between agents who are indifferent and others who are strictly better off. Second, from a conceptual viewpoint, we prefer to be cautious in excluding trades on the grounds that they are bets. Thus, if all agents are indifferent between y and x, it can still be an "improvement" in our sense. It stands to reason that, due to transaction costs, such a trade would anyway not take place, but we have no particular reason to object to it on the grounds of absence of strict preferences.

Next, we adapt the definition of no-betting Pareto dominance proposed by Gilboa, Samuelson, and Schmeidler (2014). For this purpose it is convenient (though not necessary) to first assume that all agents are Bayesian.

Definition 1 *For two allocations y, x, we say that y* Probability No-Betting improves upon *x if:*

(i) y improves upon x;

(ii) There exists a probability $p_0 \in \Delta(S)$ such that, for all $k \in A$,

$$p_0 \cdot y^k \geq p_0 \cdot x^k.$$

The main point of the present paper is to allow for more freedom in (ii), thus proposing the following:

Definition 2 *For two allocations y, x, we say that y* Ambiguity No-Betting Improves (ANBIs) *upon x if:*

(i) y improves upon x;

(ii) There exists a closed and convex set of probabilities $C \subseteq \Delta(S)$ such that, for all $k \in A$,

$$\min_{p\in C} p \cdot y^k \geq \min_{p\in C} p \cdot x^k.$$

In each case, part (i) of the definition is the requirement that y improves upon x, where each agent's bundles (under y and under x) are evaluated based on that agent's beliefs. Part (ii) in each definition imposes an additional constraint, namely, that the proposed trade can be justified also using "common beliefs". These beliefs are Bayesian in Definition 1 and ambiguity-averse in Definition 2. In both, the intuition behind the additional condition (ii) is the same: it asks the agents to prove that the trade involved is not a bet. Thus, it is perfectly legitimate for the agents to differ in their beliefs; but they are asked to prove that these differences are not the driving force behind their eagerness to trade. As long as there exist some beliefs (Bayesian in Definition 1 and non-Bayesian in Definition 2) that can explain, simultaneously, why no one is hurt by the proposed trade, the latter passes the test.

Both definitions could be used for agents that are Bayesian or ambiguity averse. For example, even if all the agents are ambiguity averse and follow the MEU model, an outside observer asked to judge whether a trade constitutes a bet might still insist on a Bayesian justification in order to rule that it does not. By contrast, the agents might all be Bayesian, but, especially if they differ in their Bayesian beliefs, the observer might wonder whether a set of probabilities—say, the (convex hull of the) union of the agents' subjective probabilities—can be used to simultaneously justify trade for all agents.[4]

Our main result is a characterization of condition (ii) in Definition 2 (as are the results of Gilboa, Samuelson, and Schmeidler (2014) for Definition 1). The agents' preferences will not play any role in this result, and it can thus apply to agents who make decisions according to the Bayesian model, to MEU, or to any other model. However, it appears more natural to pose the question where the external observer uses a model that is similar to the agents'. If the agents are all Bayesian, so that each seems to think that she knows what the "true" probability is, the inability to find a single probability according to which they all (weakly) benefit from trade is troublesome. The observer may be able to justify trade via ambiguous beliefs, but she would still have to admit that in the agents' own eyes the other agents are fundamentally wrong. By contrast, if each of the agents is ambiguity averse—say, along the lines of the MEU model—then the absence of a single probability that can simultaneously justify trade for all is less problematic. In this case the observer might reason along the lines of, "Well, I am dealing here with people who aren't sure that they know what the probability really is; they allow for sets of probabilities. Even a single such agent might make a collection of binary decisions that can't be justified by a single probability. Hence, the absence of such a probability is no indication of betting." As a result, we find it most intuitive to think of the results of Gilboa, Samuelson, and Schmeidler (2014) as applying to Bayesian agents, and the current analysis as applying to MEU agents.

[4]Alon and Gayer (2016) propose a related idea, according to which society is ambiguity averse while the agents are Bayesian. They show that Pareto-type conditions imply that society's set of probabilities is the convex hull of the agents' probabilities. Alon and Gayer assume that society has a complete ordering over alternatives, whereas in the current paper we only characterize a relation that need not be complete.

3 Examples

We present examples designed to explore our intuition concerning ambiguity aversion. The idea is that we conventionally accept risk aversion at face value, and allow agents to secure Pareto gains by insuring this risk aversion. Do we have the same attitude toward ambiguity aversion?

In the examples, we improve readability by using letters to denote agents and numbers to denote states. Thus, we think of a set of two agents as $A = \{a, b\}$ in the examples and $A = \{1, 2\}$ in the general analysis.[5]

Example 1 (Ambiguity Insurance) Let there be two agents, $A = \{a, b\}$ and two states, $S = \{1, 2\}$. Each gent $k = a, b$ believes that the probability of state 1 lies in the interval $[0, 1]$. The two agents have identical intervals of beliefs, and so condition (ii) of Definition 2 adds nothing to condition (i). That is, if each agent already finds an allocation y (weakly) better for her than x, then the shared set of beliefs can also be used to show that condition (ii) holds. We are thus examining only the implications of extending our notion of improvement from risk to ambiguity.

Let the allocation x provide utility $(0, 1)$ to agent a and $(1, 0)$ to agent b. The contract curve of efficient allocations is the diagonal. Moreover, every point on the diagonal except the two endpoints ANBIs upon x. These agents are extremely ambiguity averse, and so are willing to purchase insurance at extraordinarily unfavorable rates.

We might think it unusual that the agents should be so ambiguity averse. Further, some of the trades they might wish to engage in can seem rather unfair. For example, an allocation y that guarantees a $(1-\varepsilon, 1-\varepsilon)$ and leaves b with $(\varepsilon, \varepsilon)$ (for a small $\varepsilon > 0$) would improve, and therefore also ANBI, upon x. However, the type of moral objection one might have to such a trade is unrelated to betting: this can be a standard case in which Pareto improvement divides the surplus of trade rather unevenly, as can happen under certainty and with no role to be played by betting. ■

Example 2 (Ambiguity Insurance vs. Risk Insurance) Consider the utility profiles (with 1, 2 being states and a, b, c denoting the agents):

$$x = \begin{pmatrix} & 1 & 2 \\ a & 10 & 0 \\ b & 0 & 10 \\ c & 0 & 0 \end{pmatrix} \quad \text{and} \quad y = \begin{pmatrix} & 1 & 2 \\ a & 4 & 4 \\ b & 4 & 4 \\ c & 2 & 2 \end{pmatrix} \quad .$$

We can think of agent c as an insurance company and y as an allocation in which c insures agents a and b.

If these were material payoffs, then this would look like a standard insurance contract. If agents a and b are risk averse, then agent c can extract a risk premium while offering insurance. Given a shared probability distribution over

[5] Switching to an abstract set of agents A throughout complicates notation throughout the proofs without adding any insight.

states, agents a and b both prefer to purchase insurance from agent c. Allocation y probability no-betting improves x.

Now suppose we remove risk aversion from consideration either by assuming the agents are risk neutral or more generally by assuming that these are utility payoffs. Then y does not probability no-betting Pareto improve upon x, as there is no single probability under which agents a and b both receive higher payoffs under y than under x. However, we can specify a set of beliefs C for which y ANBIs upon x.

One might argue that this example raises no difficulty, and that we should allow the agents to insure ambiguity just as we allow them to insure risk. On the other hand, the maxmin expected utility calculation calls for agent a to evaluate the insurance contract according to a different probability distribution (selected from the common set C) than does agent b. Is "fair" for the insurance company to justify its trade with a and with b based on different probability distributions?

In the physical payoff interpretation, we are typically comfortable with the idea that the insurance company can extract a risk premium from the agents (which it can do with a single probability for all agents, exploiting the fact that different agents incur losses in different states). We might similarly allow the company to extract an ambiguity premium (with agents now holding intervals of probabilities, and the company exploiting the fact that different agents apply different probabilities in different states). ■

Example 3 (Reducing Ambiguity to Risk) This example shows that one could extract an ambiguity premium by converting a situation with ambiguity to a situation with risk, *without* eliminating the risk. The two agents start with an allocation that (because of their ambiguity aversion) is worth zero, while there is an alternative allocation exhibiting risk but no ambiguity that ANBIs upon the initial allocation.

Suppose that there are two individuals, a, b, and four states. We can think of the states as those in the two-urn Ellsberg example: there is a "known" urn, with 50 black and 50 red balls, and an "unknown" urn, with 100 balls that are each black or red; a state specifies the color of the ball drawn from each urn, should that urn be chosen. States $(1, 2, 3, 4)$ can be thought of as (BB, BR, RB, RR), that is, state 1 corresponds to a black ball being drawn from either urn, state 2 corresponds to a black ball drawn from the known urn (should it be chosen) and a red ball being drawn from the unknown urn (should it be chosen), and so on. Let the allocations be given by

$$\begin{array}{ccccc} & 1 & 2 & 3 & 4 \\ x & (2,0) & (0,2) & (2,0) & (0,2) \\ y & (1,0) & (1,0) & (0,1) & (0,1) \end{array} ,$$

where the first coordinate is the (utility) payoff to agent a and the second is the payoff to agent b. Thus, at allocation x the two agents have 2 units between them, but the allocation of these units depends on the outcome of a draw from the unknown urn. At y, by contrast, the total quantity they have is reduced to

1, but that unit is allocated to one of them in a way that involves no ambiguity: it depends only on the outcome of a draw from the known urn.

The trade of x for y can be justified by the set of probabilities given by

$$\{ p \in \Delta^3 \,|\, p_1 + p_2 = .5, p_3 + p_4 = .5 \}.$$

This is the set of *all* probabilities consistent with the description of the problem: they only reflect the knowledge of the distribution of balls in the known urn, leaving all possible distributions for the unknown urn. Both agents prefer y to x. In trading from x to y, they pay a unit for sure (between them), in order to switch the payoff between state 2 and 3. This switch eliminates the ambiguity for both agents, giving rise to their willingness to pay, but leaves them subject to risk. ■

4 Characterization Result

Suppose we are given a pair of allocations y and x, specified in terms of utilities, with y improving on x. When does there exist a set of probabilities such that y ANBIs upon x?

4.1 A Theorem

Theorem 1 *Given two allocations,* $x = \left(x_i^k\right)_{i \leq n, k \leq m}$ *and* $y = \left(y_i^k\right)_{i \leq n, k \leq m}$, *the following are equivalent:*

- *Condition (A): There exists a convex, closed set of probabilities on* S, $C \subseteq \Delta^{n-1}$ *such that, for each* $k \leq m$,

$$\min_{p \in C} p \cdot x^k \leq \min_{p \in C} p \cdot y^k.$$

- *Condition (B): For every* $\left(\lambda^k\right)_{k \leq m}, \left(\alpha^{kl}\right)_{k,l \leq m}, \left(\beta^{kl}\right)_{k,l \leq m}, \left(\gamma^{kl}\right)_{k,l \leq m}, \left(\delta^{kl}\right)_{k,l \leq m} \geq 0$, *for every* $k \leq m$ *there exist* $i_k, j_k \leq n$, *such that*

$$\sum_{k=1}^{m} \lambda^k \left(y_{j_k}^k - x_{i_k}^k\right) \geq \sum_{k=1}^{m} \left[\begin{array}{c} \sum_{l=1}^m \alpha^{kl} \left(x_{i_k}^k - x_{i_l}^k\right) \\ + \sum_{l=1}^m \beta^{kl} \left(x_{i_k}^k - x_{j_l}^k\right) \\ + \sum_{l=1}^m \gamma^{kl} \left(y_{j_k}^k - y_{i_l}^k\right) \\ + \sum_{l=1}^m \delta^{kl} \left(y_{j_k}^k - y_{j_l}^k\right) \end{array} \right]. \qquad (1)$$

4.2 Interpretation

4.2.1 Special Cases

To understand Condition (B), we first consider three simplified scenarios:

(i) There is but one agent ($m = 1$);
(ii) A single probability can justify trade;

(iii) The initial allocation x is a full-insurance allocation.

Scenario (i) is a rather trivial case of the model. In this set-up there are no agents who can disagree, and the question of justifying "trade" boils down to the question, are there any beliefs according to which y is at least as good as x. The natural answer we would expect to get is that this is the case if and only if x does not strictly dominate y (state-by-state). That is, if at one state ℓ we have $y_\ell \geq x_\ell$, this should suffice. It is reassuring to know that this is indeed the case:[6]

Remark 1 *If $m = 1$, then Condition B holds iff there exists a state ℓ with $y_\ell \geq x_\ell$.*

Scenario (ii) directs attention to trades that can be justified by a single probability. This is the focus of Gilboa, Samuelson, and Schmeidler (2014), but that paper uses strict inequalities (for the subsets of agents that are involved in trade). To facilitate the comparison we first present the (weak inequality) analogue of Gilboa, Samuelson, and Schmeidler's (2014) characterization:

Remark 2 *There exists a probability $p_0 \in \Delta^{n-1}$ such that, for each $k \leq m$,*

$$p_0 \cdot x^k \leq p_0 \cdot y^k$$

iff, for every nonnegative $(\lambda^k)_{k \leq m}$ there exists a state l such that

$$\sum_{k=1}^{m} \lambda^k (y_\ell^k - x_\ell^k) \geq 0.$$

We omit the proof which mimics that of Theorem 1 in Gilboa, Samuelson, and Schmeidler (2014) (the latter being stated with strict inequalities throughout).

We will therefore define

- Condition (GSS'): For every $(\lambda^k)_{k \leq m} \geq 0$ there exists a state ℓ such that

$$\sum_{k=1}^{m} \lambda^k (y_\ell^k - x_\ell^k) \geq 0.$$

Because the justification of trade by maxmin expected utility can also be, as a special case, justification by a single probability, it should better be the case that Condition (GSS') implies Condition (B).[7] Indeed,

Remark 3 *Condition B is implied by Condition (GSS').*

[6] The proofs of all Remarks in this subsection, provided in **Section 6, are direct** and do not rely on Theorem 1.

[7] Observe that here we do not provide a special case of the model, where an equivalence result should still hold as equivalence, but rather deal with a special case in which Condition (A) can be satisified, and wish to see that it indeed implies Condition (B).

Finally, in **Scenario (iii)**, we focus on the special case in which x is a full-insurance allocation, i.e., each agent k is guaranteed a vector x^k that is constant across states of nature. If for all k there exists $\xi^k \in \mathbb{R}$ such that $x_i^k = \xi^k$ for all i, then, for any probability p, $px^k = \xi^k$. Hence, for any set of probabilities C we also have $\min_{p \in C} px^k = \xi^k$.

Assume that another allocation, y, ANBIs upon x. That is, there is a set of probabilities C such that, for every agent k,

$$\min_{p \in C} py^k \geq \min_{p \in C} px^k = \xi^k$$

then for each p in C we have $py^k \geq \xi^k = px^k$. In other words, when x is a full insurance allocation, such a set C exists if and only if each element thereof is also such a set when considered as a singleton. Intuitively, for all agents to agree to switch from a full-insurance allocation x to another allocation y that is not fully insured, there has to be at least one probability p that makes them at least as well off under y than under x. The addition of probabilities to the set $C = \{p\}$ cannot make ambiguity-averse agents prefer y to x if they do not already prefer it for any single probability. (Note that this argument fails if the agents face uncertainty also under x.) This line of reasoning suggests that, in case x is a full insurance allocation, the converse of Remark 3 should also hold. Indeed,

Remark 4 *If $x_i^k = x_j^k$ for every pair of states i, j and every agent k, then Condition B is equivalent to the Condition (GSS').*

4.2.2 Interpretation

Beyond these special cases, Condition (B) is evidently more complicated. Let us start with the characterization of *Bayesian* justifications of trade in Gilboa, Samuelson, and Schmeidler (2014) or (essentially equivalently) with Condition (GSS'). This condition says that for every weighted average of the agents' utilities, there should be a state at which the new allocation is at least as good as the old one. That is, if we think of an "average agent", or some weighted-utilitarian aggregation of the agents, they should be able to point to a state in which trade does not make them worse off. The fact that the average agent can point to a *single* state at which trade makes her (weakly) better off has to do with the agents' agreement on beliefs: as they are supposed to share the same probability, they can find (at least) one state in which the trade makes their average better off.

When ambiguity is concerned, things are more involved. There is a set of probabilities that are considered possible, and different acts will be evaluated by different probabilities. This has two ramifications: the probabilities that different agents use to evaluate a given allocation vary, and the probabilities that the same agent uses to evaluate different allocations also vary. Thus, instead of a single state that is used in Condition (GSS') for a given weight vector, here we have $2m$ states: for each of the m agents, and for each of the two allocations under consideration there, might be a different state that explains trade.

More specifically, imagine that we ask each agent k why she prefers allocation y over x. She should be able to point to a state j_k at which y is very promising (for her), that is, with a high $y^k_{j_k}$, and to a state i_k at which x is threatening low, that is, with a low $x^k_{i_k}$. These forces are captured in the left side of (1). However, it is not enough to have one state at which y is attractive to a given agent and another at which x is worrisome to her. The agents might be wrong, and the expressions on the right side of (1) can be viewed as taking into account the various "mistakes" that may result from one agent failing to identify the good or the bad state.

To see this we can examine the implications of Condition B when x is not constant. In this case, there may exist convex combinations of agents with the property that there is no single state at which that combination prefers y to x, even though there may exist a set of beliefs under which all convex combinations of agents find y better than x. For an example, let there be two agents and two states, let y give both agents 0 in both states, and let $x^1 = (-1, 2)$ and $x^2 = (2, -1)$. Let $\lambda^1 = \lambda^2 = 1/2$. Then there is no choice of state j for which

$$\frac{1}{2}(y^1_j - x^1_j) + \frac{1}{2}(y^2_j - x^2_j) \geq 0,$$

but y is better for both agents when the set C contains a prior that puts probability 1 on state 1 and another prior that puts probability 1 on state 2.

Perhaps the difficulty in the preceding paragraph is that we forced ourselves to choose the same state for agents 1 and 2. Since our criterion for comparing y and x involves a minimization over a set of probability measures C, it may well be that different states are the crucial states for different agents. This leads to the hypothesis that the appropriate generalization of Condition (GSS') is that every set of weights $(\lambda^k)_k$, there exists a state j_k for every agent k such that

$$\sum_{k=1}^{m} \lambda^k (y^k_{j_k} - x^k_{j_k}) \geq 0.$$

Notice that this is just the condition that for every agent, there exist a state at which y is better than x. This is too weak to be sufficient. To see this, let there be two agents and two states, and let

$$\begin{array}{lll} x^1 = (4, 2) & \qquad & x^2 = (2, 4) \\ y^1 = (5, 0) & \qquad & y^2 = (0, 5). \end{array}$$

Each agent has a state at which y is better than x, and hence we can satisfy the candidate condition, but there is no set of probability measures that will make both agents prefer y to x. In order for agent 1 to prefer y to x, the set C must contain a prior putting probability at least 2/3 on state 1, but that ensures agent 2 will not prefer y to x.

We must then look for some strengthening. Before doing so, however, we note that our motivation for allowing different agents to appeal to different states in explaining why they like y better than x, namely that different elements of the set C are likely to be relevant for different agents, should be carried further. An individual agent is likely to appeal to a different element of C when examining y than when examining x. This suggests the requirement that for every $(\lambda_k)_k$, there is a state j_k and a state i_k for every k with the property that

$$\sum_{k=1}^{m} \lambda^k \left(y^k_{j_k} - x^k_{i_k}\right) \geq 0.$$

This gives us the left side of (1). But this condition is even more dramatically too weak to be sufficient. To see this, let there be two agents and two states, and let

$$\begin{aligned} x^1 &= (1,3) & x^2 &= (3,1) \\ y^1 &= (0,2) & y^2 &= (2,0). \end{aligned}$$

Then by choosing $i_1 = i_2 = 1$ and $j_1 = j_2 = 2$,

$$\begin{aligned} x^1_{i_1} &= x^1_1 = 1 & x^2_{i_2} &= x^2_2 = 1 \\ y^1_{j_1} &= y^1_2 = 2 & y^2_{j_2} &= y^1_1 = 2, \end{aligned}$$

we can ensure that the left side of (1) holds, for any choice of weights, but for each agent x strictly dominates y, and no set of probability measures can make either agent prefer y.

We thus need to couple the left side of (1) with the right side. The right side of (1) is more involved, and certainly more complicated than the corresponding right side in Condition (GSS), which would simply be 0. Here we have a weighted sum of all the payoffs that agent k might be giving up in the old allocation x, when looking at the state she considers dangerous (i_k) as compared to the states that others consider dangerous (i_l) or promising (j_l), and the payoff she is gaining in the new allocation, when looking at the state she considers promising (j_k) as compared to those the others consider dangerous (i_l) or promising (j_l).

5 Discussion

GSS (2014) is not the only paper that suggested a reconsideration of the concept of Pareto optimality in the presence of varying beliefs. Brunnermeier, Simsek, and Xiong (2014) suggested an alternative approach, according to which Pareto domination would hold only when it can be justified by any belief that is in the convex hull of the agents' beliefs. One may ask a similar question about their concept, namely, how does it generalize to non-Bayesian beliefs. Indeed, Danan, Gajdos, Hill and Tallon (2016) allow for non-Bayesian beliefs that vary across agents in the related context of preference aggregation.

The present paper is related to Gilboa and Samuelson (2020), which characterizes collections of evaluations of uncertain acts that can be simultaneously justified by a single belief—both in the Bayesian case, in which the belief is a probability, and in a non-Bayesian case where "belief" is modeled as a set of priors with respect to which minimal expected utility is computed. The two papers share a normative approach, dealing with collections of decisions that can or cannot be justified. However, they have rather different motivations, which lead to important differences in the models. Gilboa and Samuelson's (2020) focus is a single decision maker who wishes (or is committed) to justify her decisions. An important class of cases includes public officers who make decisions on behalf of a constituency, and who need or wish to have a reasonable account of the way their decisions were taken. By contrast, in the present paper the focus is on different agents who wish to trade, and a market regulator who needs to be convinced that their transactions do not amount to betting. The different setups give rise to different models of observable data. When dealing with the records that public officers need to keep, we may demand that they evaluate each possible act and document its expected utility. When potential regulations on trade are concerned, we do not require here that pre- and post-trade bundles be declared equivalent to risk-free ones. This implies that the two papers have rather different assumptions regarding observable data, and this, in turn, renders their results independent.

6 Appendix: Proofs

6.1 Proof of Theorem 1

We will show that the following eight statements, about two allocations, $x = \left(x_i^k\right)_{i\leq n,k\leq m}$ and $y = \left(y_i^k\right)_{i\leq n,k\leq m}$, are equivalent:

Statement 1: There exists a convex, closed set of probabilities on S, $C \subset \Delta^{n-1}$ such that, for each $k \leq m$,

$$\min_{p\in C} p \cdot x^k \leq \min_{p\in C} p \cdot y^k.$$

Statement 2: There are $2m$ probability vectors, $\left(p^k\right)_{k\leq m}, \left(q^k\right)_{k\leq m} \subset \Delta^{n-1}$ such that, for each $k \leq m$,

$$p^k \cdot x^k \leq q^k \cdot y^k \tag{2}$$

and, for all $k, l \leq m$

$$\begin{aligned} p^k \cdot x^k &\leq p^l \cdot x^k, q^l \cdot x^k \\ q^k \cdot y^k &\leq p^l \cdot y^k, q^l \cdot y^k. \end{aligned} \tag{3}$$

Statement 3: The following linear programming problem is feasible:

$$(P) \qquad Max_{\left(p_i^k\right)_{i\leq n,k\leq m},\left(q_i^k\right)_{i\leq n,k\leq m}} \sum_{k=1}^{m}\sum_{i=1}^{n} 0 \cdot p_i^k + \sum_{k=1}^{m}\sum_{i=1}^{n} 0 \cdot q_i^k$$

subject to

$$\sum_{i=1}^{n} x_i^k \cdot p_i^k - \sum_{i=1}^{n} y_i^k \cdot q_i^k \quad \leq \quad 0 \qquad \forall k \leq m \qquad \text{(P1)}$$

$$\sum_{i=1}^{n} x_i^k \cdot p_i^k - \sum_{i=1}^{n} x_i^k \cdot p_i^l \quad \leq \quad 0 \qquad \forall k, l \leq m, \quad k \neq l \qquad \text{(P2)}$$

$$\sum_{i=1}^{n} x_i^k \cdot p_i^k - \sum_{i=1}^{n} x_i^k \cdot q_i^l \quad \leq \quad 0 \qquad \forall k, l \leq m \qquad \text{(P3)}$$

$$\sum_{i=1}^{n} y_i^k \cdot q_i^k - \sum_{i=1}^{n} y_i^k \cdot p_i^l \quad \leq \quad 0 \qquad \forall k, l \leq m \qquad \text{(P4)}$$

$$\sum_{i=1}^{n} y_i^k \cdot q_i^k - \sum_{i=1}^{n} y_i^k \cdot q_i^l \quad \leq \quad 0 \qquad \forall k, l \leq m, \quad k \neq l \qquad \text{(P5)}$$

$$\sum_{i=1}^{n} p_i^k \quad = \quad 1 \qquad \forall k \leq m \qquad \text{(P6)}$$

$$\sum_{i=1}^{n} q_i^k \quad = \quad 1 \qquad \forall k \leq m \qquad \text{(P7)}$$

$$p_i^k, q_i^k \quad \geq \quad 0 \qquad \forall i \leq n, k \leq m.$$

Statement 4: The following linear programming problem is bounded:

$$(D) \qquad Min_{(\lambda^k)_{k \leq m}, (\alpha^{kl})_{\substack{k,l \leq m \\ k \neq l}}, (\beta^{kl})_{k,l \leq m}, (\gamma^{kl})_{k,l \leq m}, (\delta^{kl})_{\substack{k,l \leq m \\ k \neq l}}, (\mu^k)_{k \leq m}, (\nu^k)_{k \leq m}}$$

$$\sum_{k=1}^{m} 0 \cdot \lambda^k + \sum_{k=1}^{m} \sum_{\substack{l=1 \\ l \neq k}}^{m} 0 \cdot \alpha^{kl} + \sum_{k=1}^{m} \sum_{l=1}^{m} 0 \cdot \beta^{kl} +$$

$$\sum_{k=1}^{m} \sum_{l=1}^{m} 0 \cdot \gamma^{kl} + \sum_{k=1}^{m} \sum_{\substack{l=1 \\ l \neq k}}^{m} 0 \cdot \delta^{kl} + \sum_{k=1}^{m} \mu^k + \sum_{k=1}^{m} \nu^k$$

subject to

$$x_i^k \lambda^k + x_i^k \left[\sum_{\substack{l=1 \\ l \neq k}}^m \alpha^{kl} + \sum_{l=1}^m \beta^{kl} \right] - \sum_{\substack{l=1 \\ l \neq k}}^m x_i^l \alpha^{lk} - \sum_{l=1}^m y_i^l \gamma^{lk} + \mu^k \geq 0 \qquad \forall i \leq n, k \leq m \qquad \text{(D1)}$$

$$-y_i^k \lambda^k + y_i^k \left[\sum_{l=1}^m \gamma^{kl} + \sum_{\substack{l=1 \\ l \neq k}}^m \delta^{kl} \right] - \sum_{l=1}^m x_i^l \beta^{lk} - \sum_{\substack{l=1 \\ l \neq k}}^m y_i^l \delta^{lk} + \nu^k \geq 0 \qquad \forall i \leq n, k \leq m \qquad \text{(D2)}$$

$$\left(\lambda^k\right)_{k \leq m}, \left(\alpha^{kl}\right)_{\substack{k,l \leq m \\ k \neq l}}, \left(\beta^{kl}\right)_{k,l \leq m}, \left(\gamma^{kl}\right)_{k,l \leq m}, \left(\delta^{kl}\right)_{\substack{k,l \leq m \\ k \neq l}} \geq 0.$$

Statement 5: For every $\left(\lambda^k\right)_{k \leq m}, \left(\alpha^{kl}\right)_{\substack{k,l \leq m \\ k \neq l}}, \left(\beta^{kl}\right)_{k,l \leq m}, \left(\gamma^{kl}\right)_{k,l \leq m}, \left(\delta^{kl}\right)_{\substack{k,l \leq m \\ k \neq l}} \geq 0$ and every real numbers $\left(\mu^k\right)_{k \leq m}, \left(\nu^k\right)_{k \leq m}$,
IF

$$x_i^k \lambda^k + x_i^k \left[\sum_{\substack{l=1 \\ l \neq k}}^m \alpha^{kl} + \sum_{l=1}^m \beta^{kl} \right] - \sum_{\substack{l=1 \\ l \neq k}}^m x_i^l \alpha^{lk} - \sum_{l=1}^m y_i^l \gamma^{lk} + \mu^k \geq 0 \qquad \forall i \leq n, k \leq m$$

$$-y_i^k \lambda^k + y_i^k \left[\sum_{l=1}^m \gamma^{kl} + \sum_{\substack{l=1 \\ l \neq k}}^m \delta^{kl} \right] - \sum_{l=1}^m x_i^l \beta^{lk} - \sum_{\substack{l=1 \\ l \neq k}}^m y_i^l \delta^{lk} + \nu^k \geq 0 \qquad \forall i \leq n, k \leq m$$

THEN

$$\sum_{k=1}^m \mu^k + \sum_{k=1}^m \nu^k \geq 0.$$

Statement 6: For every $\left(\lambda^k\right)_{k \leq m}, \left(\alpha^{kl}\right)_{\substack{k,l \leq m \\ k \neq l}}, \left(\beta^{kl}\right)_{k,l \leq m}, \left(\gamma^{kl}\right)_{k,l \leq m}, \left(\delta^{kl}\right)_{\substack{k,l \leq m \\ k \neq l}} \geq 0$, we have

$$\sum_{k=1}^m \mu^k + \sum_{k=1}^m \nu^k \geq 0$$

where

$$\mu^k \equiv \max_{1 \leq i \leq n} \left[-x_i^k \lambda^k - x_i^k \left(\sum_{\substack{l=1 \\ l \neq k}}^m \alpha^{kl} + \sum_{l=1}^m \beta^{kl} \right) + \sum_{\substack{l=1 \\ l \neq k}}^m x_i^l \alpha^{lk} + \sum_{l=1}^m y_i^l \gamma^{lk} \right] \text{(4)}$$

$$\nu^k \equiv \max_{1 \leq i \leq n} \left[y_i^k \lambda^k - y_i^k \left(\sum_{l=1}^m \gamma^{kl} + \sum_{\substack{l=1 \\ l \neq k}}^m \delta^{kl} \right) + \sum_{l=1}^m x_i^l \beta^{lk} + \sum_{\substack{l=1 \\ l \neq k}}^m y_i^l \delta^{lk} \right]$$

for every $i \le n, k \le m$.

Statement 7: For every $\left(\lambda^k\right)_{k\le m}, \left(\alpha^{kl}\right)_{\substack{k,l\le m\\ k\neq l}}, \left(\beta^{kl}\right)_{k,l\le m}, \left(\gamma^{kl}\right)_{k,l\le m}, \left(\delta^{kl}\right)_{\substack{k,l\le m\\ k\neq l}} \ge 0$, for every $k \le m$ there exist $i_k, j_k \le n$ such that

$$\sum_{k=1}^{m} \left[\begin{array}{c} -x_{i_k}^k \lambda^k - x_{i_k}^k \left(\sum_{\substack{l=1\\ l\neq k}}^m \alpha^{kl} + \sum_{l=1}^m \beta^{kl} \right) + \sum_{\substack{l=1\\ l\neq k}}^m x_{i_k}^l \alpha^{lk} + \sum_{l=1}^m y_{i_k}^l \gamma^{lk} + \\ y_{j_k}^k \lambda^k - y_{j_k}^k \left(\sum_{l=1}^m \gamma^{kl} + \sum_{\substack{l=1\\ l\neq k}}^m \delta^{kl} \right) + \sum_{l=1}^m x_{j_k}^l \beta^{lk} + \sum_{\substack{l=1\\ l\neq k}}^m y_{j_k}^l \delta^{lk} \end{array} \right] \ge 0.$$

Statement 8: For every $\left(\lambda^k\right)_{k\le m}, \left(\alpha^{kl}\right)_{k,l\le m}, \left(\beta^{kl}\right)_{k,l\le m}, \left(\gamma^{kl}\right)_{k,l\le m}, \left(\delta^{kl}\right)_{k,l\le m} \ge 0$, for every $k \le m$ there exist $i_k, j_k \le n$ such that

$$\sum_{k=1}^{m} \lambda^k \left(y_{j_k}^k - x_{i_k}^k \right) \ \ge \sum_{k=1}^{m} \left[\begin{array}{c} \sum_{l=1}^m \alpha^{kl} \left(x_{i_k}^k - x_{i_l}^k \right) \\ +\sum_{l=1}^m \beta^{kl} \left(x_{i_k}^k - x_{j_l}^k \right) \\ +\sum_{l=1}^m \gamma^{kl} \left(y_{j_k}^k - y_{i_l}^k \right) \\ +\sum_{l=1}^m \delta^{kl} \left(y_{j_k}^k - y_{j_l}^k \right) \end{array} \right].$$

Proof that Statements (1) and (2) are equivalent:

Assume first that Statement (1) holds for some set C. For each k, choose $p^k, q^k \in C$ such that

$$\begin{aligned} \min_{p\in C} p \cdot x^k &= p^k \cdot x^k \\ \min_{p\in C} p \cdot y^k &= q^k \cdot y^k. \end{aligned}$$

Since $p^k \in \arg\min_{p\in C} p \cdot x^k$, and $q^k \in \arg\min_{p\in C} p \cdot y^k$, (3) follows. And Statement (1) guarantees that the minimal payoff for x is bounded above by the minimal payoff for y, that is, that (2) holds as well.

Conversely, assume that Statement (2) holds. Define

$$C = conv\left[\left\{p^k\right\}_{k\le m} \cup \left\{q^k\right\}_{k\le m} \right]$$

and observe that it satisfies the requirement of Statement (1).

Proof that Statements (2) and (3) are equivalent:

The feasible set of (P) in Statement (3) is a restatement of Statement (2), with (i) constraints (P1) corresponding to (2), (ii) constraints (P2) and (P3) guaranteeing the first set of inequalities in (3), having to do with $p^k \in \arg\min_{p\in C} p \cdot x^k$, and (iii) $q^k \in \arg\min_{p\in C} p \cdot y^k$ being equivalent to constraints (P4) and (P5). Finally, constraints (P6) and (P7), combined with the nonnegativity constraints, imply that p^k and q^k are indeed probability vectors for each k.

Proof that Statements (3) and (4) are equivalent:

(D) is the dual problem of (P), with $\left(\lambda^k\right)_{k\leq m}$ corresponding to the inequalities (P1), and
$\left(\alpha^{kl}\right)_{\substack{k,l\leq m\\ k\neq l}}, \left(\beta^{kl}\right)_{k,l\leq m}, \left(\gamma^{kl}\right)_{k,l\leq m}, \left(\delta^{kl}\right)_{\substack{k,l\leq m\\ k\neq l}}$ corresponding to the sets of inequalities (P2),(P3),(P4), and (P5), respectively, and $\left(\mu^k\right)_{k\leq m}, \left(\nu^k\right)_{k\leq m}$ corresponding to the sets of equalities (P6) and (P7), respectively. (P) is a Max problem, and so (D) is a Min problem. The sets of inequalities (P1)-(P5) have the natural direction, and thus the corresponding variables $\left(\lambda^k\right)_{k\leq m}, \left(\alpha^{kl}\right)_{\substack{k,l\leq m\\ k\neq l}}, \left(\beta^{kl}\right)_{k,l\leq m}, \left(\gamma^{kl}\right)_{k,l\leq m}, \left(\delta^{kl}\right)_{\substack{k,l\leq m\\ k\neq l}}$ are constrained to be nonnegative, while the equalities (P6) and (P7) leave $\left(\mu^k\right)_{k\leq m}, \left(\nu^k\right)_{k\leq m}$ unconstrained. The constraints of (D) denoted (D1) correspond to $\left(p_i^k\right)$, while those denoted (D2) correspond to $\left(q_i^k\right)$. As both $\left(p_i^k\right)$ and $\left(q_i^k\right)$ are sets of nonnegative variables in (P), the direction of the inequalities in (D) is the natural one, namely $\geq$.

Note that the objective function in (P) was set of be identically zero, and thus all right sides in (D) are zero and the problem is homogenous. (Note that its feasible set includes the origin and it is a convex cone.) As the right side in (P1)-(P5) (in (P)) are zero, the variables $\left(\lambda^k\right)_{k\leq m}, \left(\alpha^{kl}\right)_{\substack{k,l\leq m\\ k\neq l}}, \left(\beta^{kl}\right)_{k,l\leq m}, \left(\gamma^{kl}\right)_{k,l\leq m}, \left(\delta^{kl}\right)_{\substack{k,l\leq m\\ k\neq l}}$ have zero coefficients in the objective function of (D), so that the latter is

$$Min_{(\lambda^k)_{k\leq m},(\alpha^{kl})_{\substack{k,l\leq m\\ k\neq l}},(\beta^{kl})_{k,l\leq m},(\gamma^{kl})_{k,l\leq m},(\delta^{kl})_{\substack{k,l\leq m\\ k\neq l}},(\mu^k)_{k\leq m},(\nu^k)_{k\leq m}} \sum_{k=1}^{m}\mu^k+\sum_{k=1}^{m}\nu^k.$$

Proof that Statements (4) and (5) are equivalent:

As Problem (D) is homogeneous, it is bounded if and only if it is bounded below by 0. This means that for every set of variables in the feasible set (i.e., that satisfy the constraints (D1) and (D2) and the relevant nonnegativity constraints), the objective function is nonnegative.

Proof that Statements (5) and (6) are equivalent:

For every $k \leq m$, μ^k is bounded below by n constraints in (D1), while ν^k is bounded below by n constraints in (D2). Apart from these constraints, $\left(\mu^k\right)_{k\leq m}, \left(\nu^k\right)_{k\leq m}$ are unconstrained, and therefore the inequality $\sum_{k=1}^m \mu^k + \sum_{k=1}^m \nu^k \geq 0$ would hold if and only if it holds for their minimal values. These minimal values are the maxima of the lower bounds, as given in (4).

Proof that Statements (6) and (7) are equivalent:

If Statement (6) holds, one can choose i_k and j_k to be in the argmax of the expressions defining μ^k and ν^k, respectively, to obtain the conclusion of Statement (7). Conversely, if Statement (7) holds and such i_k and j_k exist, the summation of the maxima is obviously nonnegative as well.

Proof that Statements (7) and (8) are equivalent:

Statement (7) can also be written as: For every $\left(\lambda^k\right)_{k\leq m}, \left(\alpha^{kl}\right)_{\substack{k,l\leq m\\ k\neq l}}, \left(\beta^{kl}\right)_{k,l\leq m}, \left(\gamma^{kl}\right)_{k,l\leq m}, \left(\delta^{kl}\right)_{\substack{k,l\leq m\\ k\neq l}} \geq 0$, for every $k\leq m$ there exist $i_k, j_k \leq n$ such that

$$\sum_{k=1}^{m}\left[\begin{array}{c}\lambda^k\left(y_{j_k}^k-x_{i_k}^k\right)\\ -\sum_{\substack{l=1\\ l\neq k}}^{m}\alpha^{kl}x_{i_k}^k+\sum_{\substack{l=1\\ l\neq k}}^{m}\alpha^{lk}x_{i_k}^l\\ -\sum_{l=1}^{m}\beta^{kl}x_{i_k}^k+\sum_{l=1}^{m}\beta^{lk}x_{j_k}^l\\ -\sum_{l=1}^{m}\gamma^{kl}y_{j_k}^k+\sum_{l=1}^{m}\gamma^{lk}y_{i_k}^l\\ -\sum_{\substack{l=1\\ l\neq k}}^{m}\delta^{kl}y_{j_k}^k+\sum_{\substack{l=1\\ l\neq k}}^{m}\delta^{lk}y_{j_k}^l\end{array}\right]\geq 0$$

or – allowing for arbitrary $\left(\alpha^{kk},\delta^{kk}\right)_k$ –

$$\sum_{k=1}^{m}\lambda^k\left(y_{j_k}^k-x_{i_k}^k\right)\geq\sum_{k=1}^{m}\left[\begin{array}{c}\sum_{l=1}^{m}\alpha^{kl}x_{i_k}^k-\sum_{l=1}^{m}\alpha^{lk}x_{i_k}^l\\ +\sum_{l=1}^{m}\beta^{kl}x_{i_k}^k-\sum_{l=1}^{m}\beta^{lk}x_{j_k}^l\\ +\sum_{l=1}^{m}\gamma^{kl}y_{j_k}^k-\sum_{l=1}^{m}\gamma^{lk}y_{i_k}^l\\ +\sum_{l=1}^{m}\delta^{kl}y_{j_k}^k-\sum_{l=1}^{m}\delta^{lk}y_{j_k}^l\end{array}\right].$$

By rearranging the terms (so that α^{lk} appears in the l'th row), we obtain the equivalent inequality

$$\sum_{k=1}^{m}\lambda^k\left(y_{j_k}^k-x_{i_k}^k\right)\geq\sum_{k=1}^{m}\left[\begin{array}{c}\sum_{l=1}^{m}\alpha^{kl}x_{i_k}^k-\sum_{l=1}^{m}\alpha^{kl}x_{i_l}^k\\ +\sum_{l=1}^{m}\beta^{kl}x_{i_k}^k-\sum_{l=1}^{m}\beta^{kl}x_{j_l}^k\\ +\sum_{l=1}^{m}\gamma^{kl}y_{j_k}^k-\sum_{l=1}^{m}\gamma^{kl}y_{i_l}^k\\ +\sum_{l=1}^{m}\delta^{kl}y_{j_k}^k-\sum_{l=1}^{m}\delta^{kl}y_{j_l}^k\end{array}\right]$$

or

$$\sum_{k=1}^{m}\lambda^k\left(y_{j_k}^k-x_{i_k}^k\right)\geq\sum_{k=1}^{m}\left[\begin{array}{c}\sum_{l=1}^{m}\alpha^{kl}\left(x_{i_k}^k-x_{i_l}^k\right)\\ +\sum_{l=1}^{m}\beta^{kl}\left(x_{i_k}^k-x_{j_l}^k\right)\\ +\sum_{l=1}^{m}\gamma^{kl}\left(y_{j_k}^k-y_{i_l}^k\right)\\ +\sum_{l=1}^{m}\delta^{kl}\left(y_{j_k}^k-y_{j_l}^k\right)\end{array}\right].$$

6.2 Proof of Remark 1

Assume indeed that there is but a single agent, $m=1$. We first observe that Condition (B) boils down to

- Condition B(1): For all $\lambda,\beta,\gamma\geq 0$, there exists states i and j with

$$\lambda(y_j-x_i)\geq\beta(x_i-x_j)+\gamma(y_j-y_i).$$

To see this, consider the right side of Condition (B). As there is only one individual, $k=l=1$ and it follows that i_k must equal i_l and, similarly, $j_k=j_l$. Hence the expressions involving α and δ vanish. Condition (B) then becomes Condition B(1).

We now wish to show that (when $m=1$), Condition B(1) holds iff there exists a state ℓ with $y_\ell\geq x_\ell$. First assume that such an ℓ exists. To see that Condition B(1) holds, let there be given $\lambda,\beta,\gamma\geq 0$ and choose $i=j=\ell$ so that the inequality above becomes $\lambda(y_\ell-x_\ell)\geq 0$ which indeed holds.

Conversely, if Condition B(1) holds, apply it to $\lambda=\gamma=1$ and $\beta=0$ to obtain i and j such that $y_j-x_i\geq y_j-y_i$ or $y_i\geq x_i$ (so that the condition holds for $\ell=i$).

6.3 Proof of Remark 3

Assume that Condition (GSS') holds. Given $\left(\lambda^k\right)_{k\leq m}, \left(\alpha^{kl}\right)_{k,l\leq m}, \left(\beta^{kl}\right)_{k,l\leq m}, \left(\gamma^{kl}\right)_{k,l\leq m}, \left(\delta^{kl}\right)_{k,l\leq m} \geq 0$, use Condition (GSS') for $\left(\lambda^k\right)_{k\leq m} \geq 0$ to find a state ℓ such that $\sum_{k=1}^m \lambda^k(y_\ell^k - x_\ell^k) \geq 0$. Setting, for each k, $i_k, j_k = \ell$, the inequality (1) boils down to $\sum_{k=1}^m \lambda^k(y_\ell^k - x_\ell^k) \geq 0$, so that Condition (B) holds.

6.4 Proof of Remark 4

We already established that Condition (GSS') implies Condition (B), whether $\left(x^k\right)_k$ are constant (across states for each agent). To complete the proof we need to show that, if $\left(x^k\right)_k$ are indeed constant (that is, $x_i^k = x_j^k \ \forall i, j, k$) the converse also holds. Let x be such an allocation, and let there be given $\left(\lambda^k\right)_{k\leq m} \geq 0$ as in Condition (GSS'). Set all of $\left(\alpha^{kl}\right)_{k,l\leq m}, \left(\beta^{kl}\right)_{k,l\leq m}, \left(\gamma^{kl}\right)_{k,l\leq m}, \left(\delta^{kl}\right)_{k,l\leq m}$, apart from $\left(\delta^{k1}\right)_{k\leq m}$, to zero, and $\delta^{k1} = \lambda^k$ for $k \leq m$. Apply Condition (B) to $\left(\lambda^k\right)_{k\leq m}, \left(\alpha^{kl}\right)_{k,l\leq m}, \left(\beta^{kl}\right)_{k,l\leq m}, \left(\gamma^{kl}\right)_{k,l\leq m}, \left(\delta^{kl}\right)_{k,l\leq m}$ (with the original $\left(\lambda^k\right)_{k\leq m}$ given in the antecedent of Condition (GSS')). Inequality (1) then implies

$$\sum_{k=1}^m \lambda^k \left(y_{j_k}^k - x_{i_k}^k\right) \geq \sum_{k=1}^m \lambda^k \left(y_{j_k}^k - y_{j_1}^k\right)$$

or $\sum_{k=1}^m \lambda^k \left(y_{j_1}^k - x_{i_k}^k\right) \geq 0$. However, as x^k is constant across states, we have, in particular, $x_{i_k}^k = x_{j_1}^k$ and we can set $\ell = j_1$ to obtain the desired result.

References

[1] Shiri Alon and Gabrille Gayer. Utilitarian preferences with multiple priors. *Econometrica*, 84(3):1181–1201, 2016.

[2] Markus Brunnermeier, Alp Simsek, and Wei Xiong. A welfare criterion for models with distorted beliefs. *Quarterly Journal of Economics*, 129:1711–1752, 2014.

[3] Eric Danon, Thibault Gajdos, Brian Hill, and Jean-Marc Tallon. Robust social decisions. *American Economic Review*, 106:24–7–2425, 2016.

[4] Itzhak Gilboa and Larry Samuelson. What were you thinking? Revealed preference theory as coherence test. Technical report, HEC Paris/Tel Aviv Unversity and Yale University, 2020.

[5] Itzhak Gilboa, Larry Samuelson, and David Schmeidler. No-betting Pareto dominance. *Econometrica*, 82(4):1405–1442, 2014.

[6] Antony Millner, Simon Dietz, and Geoffrey Heal. Scientific ambiguity and climate policy. *Environmental and Resource Economics*, 55:21–46, 2013.