No-Betting Pareto Dominance*

Itzhak Gilboa[†], Larry Samuelson[‡], and David Schmeidler[§]

December 9, 2012

Abstract

We argue that, in the presence of uncertainty, the notion of Pareto dominance is not as compelling as under certainty. In particular, voluntary trade that is based on differences in tastes is commonly accepted as favorable, because no agent involved in it can be wrong about her tastes. By contrast, voluntary trade that is based on incompatible beliefs may indicate that at least one agent is wrong about her beliefs. We propose a weaker, *No-Betting*, notion of Pareto domination, which requires, on top of unanimity of preference, the existence of shared beliefs that can rationalize such preference for each agent.

*We thank Markus Brunnermeier, José Heleno Faro, Johannes Hörner, Ben Polak, Andy Postlewaite, Anna Rubinchik, Alp Simsek, Glen Weyl, and Wei Xiong for discussions and comments on the 2008 version of this paper. Specifically, two results in this paper were obtained in response to specific issues raised by Anna Rubinchik and Glen Weyl. ISF Grant 396/10 (Gilboa and Schmeidler), ERC Grant 269754 (Gilboa), and NSF Grant SES-1153893 (Samuelson) are gratefully acknowledged.

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

[‡]Department of Economics, Yale University. larry.samuelson@yale.edu

[§]The Interdisciplinary Center, Herzliya, Tel-Aviv University, and The Ohio State University. davidschmeidler@gmail.com

No-Betting Pareto Dominance

Contents

No-Betting Pareto Dominance

1 Introduction

1.1 Motivation

Standard economic lore suggests that free trade is a good thing. Voluntary trade improves everyone's lot, at least if the latter is evaluated by the revealed preference paradigm. Moreover, barriers to trade lead to Pareto-inefficient outcomes in realistic and robust examples. Thus, the absence of certain markets is often criticized as a hindrance to optimality; should these markets come into being, Pareto-improving allocations would become feasible.

This paper addresses a difficulty in this line of reasoning, when applied to trade under uncertainty. Specifically, in the absence of objective probabilities, we find this argument in favor of free trade much weaker than it is when all alternatives are certain, or at least have known distributions. Consider the following examples.

Example 1. Mary and John have one banana and one mango each. Their utility functions are linear in quantities of the two goods, but Mary prefers bananas, and John prefers mangos. To be concrete, Mary is indifferent between 1 unit of banana and 2 units of mango, and John is indifferent between 2 units of banana and 1 unit of mango. In the absence of a market, they can only consume their initial endowments. In the competitive equilibrium of this simple economy, they obtain a Pareto optimal allocation in which Mary consumes only bananas and John consumes only mangos. □

Example 2. Ann and Bob have one dollar each. There are two states of the world: in state 1 the price of oil a year from now is above \$100 a barrel, and in state 2 it is below \$100 a barrel. Ann and Bob are risk neutral. Ann thinks that state 1 has probability 2/3 and Bob thinks state 1 has probability

1/3. If there are no financial markets, Ann and Bob will have to consume their initial endowments only, namely \$1 whatever is the price of oil in a year. In the competitive equilibrium of this simple economy, they obtain a Pareto optimal allocation in which Ann has no money in state 2 and Bob has no money in state 1. □

Example 3. Ruth is a young computer scientist with an idea for a start-up company. If successful, the company would net well over \$10 million. Ruth assigns to this event a probability of 90% and she seeks investors. She approaches Tom, who runs a venture capital fund, asking for seed money of \$100,000. Tom is somewhat amused by Ruth's optimism. Still, her idea is promising and he believes that she will succeed with probability of 10%. He decides to take the risk and make the investment. □

Let us first analyze the first two examples. Clearly, they map to the same Arrow-Debreu (1954) general equilibrium model: there are two goods $\{1, 2\}$, and two agents $\{a, b\}$; the utility functions are given by

$$\begin{aligned} u_a(x_1, x_2) &= \frac{2}{3}x_1 + \frac{1}{3}x_2 \\ u_b(x_1, x_2) &= \frac{1}{3}x_1 + \frac{2}{3}x_2 \end{aligned}$$

and the initial endowments are

$$e_a = e_b = (1, 1).$$

In equilibrium, goods one and two trade one-for-one, and person a consumes both units of good 1 while person b consumes both units of good 2. This equilibrium is Pareto optimal, and Pareto dominates the initial allocation.

However, it is not obvious that Pareto domination has the same meaning in both examples. In Example 1, there is no uncertainty and the differences between the two consumers are only in tastes. If Mary prefers bananas and

John prefers mangos, they are better off when they switch one banana for one mango. By contrast, in Example 2 Ann and Bob are both better off, but only because they have different subjective beliefs about the price of oil in the future.

Evidently, in Example 2 Ann and Bob cannot both be right: if the probability of state 1 is 2/3, it cannot be 1/3 as well. One may wonder what is meant by this probability in the first place. Perhaps Ann and Bob should not have probabilistic beliefs over the future price of oil. But *if* they do, and if these beliefs have any concrete meaning, *then* these beliefs are incompatible: at least one of them is wrong. The unanimous preference for trade in this example follows from the fact that the difference in beliefs "cancels out" the difference in tastes.

Next, contrast Example 2 with Example 3. In both cases uncertainty is involved. Moreover, in both cases the agents who trade entertain different subjective beliefs, and thus they cannot both be right. However, there are important differences between the two examples: in Example 2 the agents are not exposed to risk a priori. Each has an endowment that is risk-free, that is, she is "fully-insured" against the source of uncertainty. By contrast, in Example 3 one agent bears risk a priori: Ruth has an asset that will be worth a lot in one state and little in another, and trade allows her to share this risk with Tom. Therefore, Example 2 has the flavor of a pure bet, whereas Example 3 does not. Another distinction between the two examples is that in Example 2 the difference in beliefs is crucial for trade to take place: one cannot come up with a joint prior belief that would make both Ann and Bob better off by trading. This is not the case in Example 3: even though the agents disagree on beliefs, one could assume, for the sake of the argument, that Tom shares Ruth's optimism: if he is willing to invest under his moderate assessment of the probability of success, he would definitely be willing to invest were his beliefs more optimistic. Thus, Example 3 can be justified as voluntary trade between agents who are not necessarily wrong about their beliefs.

1.2 No-Betting Pareto

This paper proposes a refined notion of Pareto domination for uncertain allocations. Specifically, we wish to distinguish between Pareto domination that hinges on incompatible beliefs and Pareto domination that can be justified by shared beliefs. We do not take issue with Pareto domination under certainty (as in Example 1). Also, we find Pareto domination compelling under uncertainty, *if* agents' preferences can be justified not only according to their actual, potentially different beliefs, but also according to hypothetical *shared* beliefs (as in Example 3). However, we argue that Pareto domination is less compelling when, in the face of uncertainty, unanimous preference for one alternative over another can only be justified by variability in beliefs. As we show, these situations are closely related to pure bets (as in Example 2).

The difficulty with Pareto domination that results from different beliefs has been discussed by various authors over the years (see a brief survey in the next sub-section). Recently, with the growing sophistication of financial assets, and especially following the financial crisis that started in 2008, there is a growing literature on the topic.

Some of this literature focuses on agents' inability to comprehend the risks they are facing, or on psychological phenomena such as over-confidence. We do not think that the problem with Pareto domination is restricted to agents who are irrational in one way or another. There are many situations in which rationality does not single out particular beliefs, and in those circumstances there will be agents who may wish to trade based on differences in their beliefs. In other words, agents need not be confused or over-optimistic in order to engage in such trade; it suffices that there be some dispersion in beliefs for the market to "find" the agents who are willing to trade.[1]

[1]For similar reasons, it is not obvious that rationality necessitates Bayesian beliefs. See Gilboa, Postlewaite, and Schmeidler (2008, 2009, 2010) and Gilboa, Lieberman, and

We conclude this section with a survey of the literature. We then present our model and the refined notion of Pareto domination in Section 2. Section 3 provides two characterizations of No-Betting-Pareto domination: the first identifies those trades that are allowed by our refined definition, whereas the second shows that the trades that are excluded by our definition have a flavor of pure betting. Section 4 comments on properties of the No-Betting-Pareto domination, showing that it is, in general, not transitive, but that its transitive closure also cannot favor bets, and commenting on its computability. Section 5 comments on the relationship between the concept presented here and utilitarian aggregation. Finally, Section 6 concludes.

1.3 Related Literature

Many people have been bothered by the interpretation of Pareto domination when beliefs differ. The difference between trade as in Examples 1 and 2 above was already pointed out by Stiglitz (1989), minimizing the importance of Pareto inefficiency that might result from taxation of financial trade. Mongin (1997) referred to Pareto domination as in Example 2 as *spurious unanimity.*

Indications that it is more difficult to aggregate preferences under subjective uncertainty than under either certainty or risk have also appeared in the social choice literature. Harsanyi's (1955) celebrated result showed that, in the context of risk (that is, known, objective probabilities), if all individuals as well as society are von-Neumann-Morgenstern expected utility maximizers (von Neumann and Morgenstern, 1944), a mild Pareto condition implies that society's utility function is a linear combination of those of the individuals. When probabilities are not given, the literature typically resorts to Savage (1954), who provided an axiomatic justification of subjective expected utility maximization, namely, the maximization of a utility function

Schmeidler (2009). However, we use here "rationality" in the common sense, namely, satisfying Savage's axioms.

according to a probability measure, where both the utility and the probability are derived from preferences. However, Hylland and Zeckhauser (1979) and Mongin (1995) found that an extension of Harsanyi's theorem to the case of uncertainty cannot be obtained. An impossibility theorem shows that one cannot simultaneously aggregate utilities and probabilities in such a way that society will satisfy the same decision theoretic axioms as the individuals.

Gilboa, Samet, and Schmeidler (2004) (hereafter GSS) argue that, due to the spurious unanimity problem, Pareto domination is not compelling in the context of subjective beliefs. They offer the example of a duel between two gentlemen, each of which entertains subjective beliefs that he is going to win (and kill the other) with probability 90% (and die with probability 10%). Each would flee town if he thought that his probability of dying exceeded 20%. But, optimistic as they are, they both prefer a duel to a non-duel. Should society adopt these preferences, as the Pareto condition suggests? GSS argue that society should not agree with these preferences *only because* of the Pareto argument. While the two individuals agree, their agreement results from a "cancelling out" of differences in tastes and differences in beliefs. There is no way to get them to agree on the preferred choice as well as on the reasoning that leads to it. If they were to agree on the reasoning (and the probabilities), the differences in tastes would imply that at least one of them would prefer to cancel the duel.[2]

Recently, this difficulty with the notion of Pareto domination has been noted by several authors, partly in the context of trade in financial markets. Weyl (2007) points out that arbitrage might be harmful in case agents are "confused". Posner and Weyl (2012) call for a regulatory authority, akin to the FDA, that would need to approve trade in new financial assets, guaranteeing that it does not cause harm. This problem is also discussed in Kreps

[2]GSS weaken the Pareto condition, so that it only applies to choices over which there is no disagreement over probabilities. GSS show that this weak condition is sufficient, in the presence of certain conditions, to derive social utility and social probability that are averages of the individual ones.

(2012). Brunnermeier, Simsek, and Xiong (2012) (hereafter BSX) develop a "belief-free" welfare criterion for markets in the presence of individuals who might entertain wrong beliefs. In particular, applying their criterion to financial markets with over-optimistic traders can result in speculative trade becoming normatively inferior.[3]

Our approach differs from that of BSX in that they propose a new concept of domination, by which one can say, in cases of pure betting (as in Example 2 above), that no-betting dominates betting. Thus, their definition of belief-free domination may override Pareto domination: while all agents may wish to bet, the BSX belief-free criterion may have opposite preferences. By contrast, we only weaken Pareto domination so that it will no longer be true that betting dominates no-betting. BSX's approach is closer to the social choice literature, in attempting to come up with a reasonable social preference relation that, while not necessarily complete, will be able to rank alternatives that are not ranked by standard Pareto domination. Our approach is closer to general equilibrium analysis, in that we do not attempt to rank alternatives that are incomparable according to Pareto domination.

Another difference between the two approaches is that BSX consider only beliefs in the convex hull of the agents' beliefs. This reflects an implicit supposition that the "true" probability measure is in this set, that is, that it cannot be the case that all agents are wrong. By contrast, we make no implicit or explicit reference to any probability being "true". We allow any conceivable probability to justify trade, as long as it can do so simultaneously for all agents.

2 The Model

There is a set of agents $N = \{1, ..., n\}$, a measurable state space (S, Σ), and a set of outcomes X. An outcome specifies all the aspects relevant to

[3] See also Simsek (2012), who discusses financial innovation where trade is motivated both by risk sharing and by speculation.

all agents. It is often convenient to assume that X consists of real-valued vectors, denoting each individual's consumption bundle, but at this point we do not impose any conditions on the structure of X. Thus, as in the social choice literature, X can be viewed of consisting of general outcomes that the agents might experience.

The alternatives compared are *simple acts*: functions from states to outcomes whose images are finite and measurable with respect to the discrete topology on X. We denote

$$F = \left\{ f : S \to X \;\middle|\; \begin{array}{c} f \text{ is simple and} \\ \Sigma\text{-measurable} \end{array} \right\}.$$

The restriction to simple acts guarantees that acts will be bounded in utility for each agent, and for any utility function.

Each agent i has a preference order $\succsim_i$ over F. We assume throughout this paper that the agents are expected utility maximizers a la Savage. Agent i is characterized by a utility function $u_i : X \to \mathbb{R}$ and a probability measure p_i on (S, Σ), and $\succsim_i$ is represented by maximization of $\int_S u_i(f(s))dp_i$. We assume that the agents can be represented as expected utility maximizers to emphasize that our arguments do not hinge on any type of so-called bounded rationality of the agents.

The standard notion of Pareto domination, denoted by $\succ_P$, is defined as follows:

Definition 1 *$f \succ_P g$ iff for all $i \in N$, $f \succsim_i g$, and for some $k \in N$, $f \succ_k g$.*

Throughout the paper we consider pairs of acts, $(f, g) \in F^2$. A pair (f, g) is interpreted as a suggested swap in which the agents give up act g in return for f. Such a swap would involve some individuals but not others. Given a pair (f, g), agent $i \in N$ is said to be *involved* in (f, g) if $u_i(f(\cdot)) \neq u_i(g(\cdot))$, that is, if there exists at least one state s at which the agent is not indifferent between $f(s)$ and $g(s)$. Let $N(f, g) \subset N$ denote the agents who are involved in the pair (f, g). Observe that, for given $f, g \in F$, the definition of $N(f, g)$ depends on the agents' utilities, $(u_i)_i$, but not on their beliefs, $(p_i)_i$.

Definition 2 *A pair* (f, g) *is an* improvement *if* $N(f, g) \neq \varnothing$ *and, for all* $i \in N(f, g)$, $f \succ_i g$.

We use the term *improvement* to emphasize the fact that the agents in the economy would swap g for f voluntarily. We will also use the terminology f *improves upon* g, denoted by $f \succ_* g$. Our main interest lies in improvements for which $|N(f, g)| \geq 2$, though the cases in which $|N(f, g)| = 1$ are not ruled out.

Notice that we require *strict* preference for the agents involved in the improvement. The relation $f \succ_* g$ is thus more restrictive than standard Pareto domination, which allows some agents, for whom $u_i(f(\cdot)) \neq u_i(g(\cdot))$, to be indifferent between f and g. We find that, once one makes an explicit distinction between the agents who are involved in the improvement and those who aren't, strict preference for the former appears to be a natural condition: clearly, if there are agents who are not affected by the proposed swap, they would be indifferent to it; but we require that those agents who are involved, that is, whose cooperation is needed for the swap (f, g) to take place, strictly prefer f to g. In particular, we are reluctant to assume that indifferent agents are willing to actively participate in the trade.

Our weaker notion of domination is defined as follows:

Definition 3 *For two alternatives* $f, g \in F$, *we say that* f No-Betting Pareto dominates g, *denoted* $f \succ_{NBP} g$, *if:*

(i) f *improves upon* g*;*

(ii) There exists a probability measure p_0 *such that, for all* $i \in N(f, g)$,

$$\int_S u_i(f(s))dp_0 > \int_S u_i(g(s))dp_0.$$

Observe that our definition does not assume that the agents agree on the distributions of the alternatives f and g. The actual beliefs of the agents, determining their actual preferences, may be quite different. Condition (i) of the definition requires that the agents involved prefer f to g according to their actual beliefs. Condition (ii), by contrast, requires that one be able to find a single probability measure, according to which all involved agents prefer swapping g for f. That is, one can find hypothetical beliefs, which, when ascribed to all relevant agents, can rationalize the preference for f over g. As in Example 3 above, two partners may invest in a business opportunity about which one is much more optimistic than the other. Their actual beliefs, therefore, differ. However, as long as there are some beliefs (say, of the more optimistic one) that justify the investment for both, the alternative of investment would No-Betting-Pareto dominate that of no-investment.

Clearly, Condition (i) implies that f Pareto dominates g (recall that Condition (i) also implies that $N(f,g) \neq \varnothing$). Thus, if one uses our stronger notion of Pareto dominance, $\succ_{NBP}$, rather than the standard one, one gets a larger set of Pareto optimal outcomes. In particular, the first welfare theorem still holds, though the second does not.

3 Characterizations

3.1 Combining Agents

The following result characterizes pairs (f,g) that satisfy condition (ii) of the definition of No-Betting-Pareto domination.

Theorem 1 *Consider acts f and g with $N(f,g) \neq \varnothing$. There exists a probability vector p_0 such that, for all $i \in N(f,g)$,*

$$\int_S u_i(f(s))dp_0 > \int_S u_i(g(s))dp_0$$

if and only if, for every distribution over the set of agents involved, $\lambda \in \Delta(N(f,g))$, there exists a state $s \in S$, such that

$$\sum_{i \in N(f,g)} \lambda(i)\, u_i(f(s)) > \sum_{i \in N(f,g)} \lambda(i)\, u_i(g(s)).$$

To interpret this result, assume that a set of agents $N(f,g)$ wish to swap g for f. Presumably, each one of them has a higher expected utility under f than under g (according to the agent's subjective beliefs). In particular, it is necessary that each agent be able to point to a state s at which she is better off with f than with g. The proposition says that for f also to No-Betting-Pareto dominate g, this condition should be satisfied for all "convex combinations" of the agents involved, where a combination is defined by a distribution λ over the agents' utility functions.

A convex combination of agents, λ, can be interpreted in two famously related ways. First, we may take a utilitarian interpretation, according to which $\sum_{i\in N}\lambda(i)u_i(\cdot)$ is a social welfare function defined by some averaging of the agents' utilities. Second, we may think of an individual behind the "veil of ignorance", believing that she may be agent i with probability $\lambda(i)$, and calculating her expected utility ex-ante. In both interpretations the condition states that not only the actual agents, but also all convex combinations thereof can justify the improvement by pointing to a state of the world that would make them at least as well off with the proposed improvement.

3.2 Bets

We would like to argue that agents cannot make themselves better off, under the No-Betting-Pareto criterion, by betting with one another. Intuitively, a bet is a transfer of resources between agents that is not driven by production, different tastes or risk sharing. To capture the fact that a bet does not involve production, we need to endow the set of outcomes with additional structure. Assume then that $X = L^n$ where L is a partially ordered linear space, where $x = (x_1, ..., x_n) \in X$ specifies an allocation, x_i, of each agent i. In such a set-up, one can express the fact than an improvement (f,g) is a mere allocation of existing resources by requiring that

$$\sum_{i\in N(f,g)} f(s)_i \leq \sum_{i\in N(f,g)} g(s)_i \qquad \forall s\in S. \tag{1}$$

In this case, we say that the pair (f,g) is *feasible*.

For simplicity, we focus on the case $L=\mathbb{R}$, where x_i denotes agent i's wealth. Further, assume that each agent's utility function depends only on her own wealth. We abuse notation and denote this function by u_i as well, so that, for each $i\in N$ and $x\in X$, $u_i\left((x_1,...,x_n)\right)=u_i\left(x_i\right)$. Finally, we assume that each u_i is differentiable, strictly monotone and (weakly) concave in its real-valued argument.

In this unidimensional set-up, trade cannot be driven by differences in tastes, as all agents are assumed to want more of the only good, namely, money. Hence we can define betting as follows.

Definition 4 *A feasible improvement* (f,g) *is a* bet *if* $g(s)_i$ *does not depend on* s *for* $i\in N(f,g)$.

Section 3.4 explains how, if L were multidimensional, this definition would need to be modified so as to exclude Pareto-improving barter that may take place at each state of the world independently of the others. However, when we consider financial markets, and $L=\mathbb{R}$, feasible Pareto improvements among risk averse agents can only be driven by differences in beliefs or by risk-sharing. The requirement that g be independent of s (for all i) precludes the risk-sharing motivation, thereby justifying the definition of (f,g) as a bet. We can now state:

Proposition 1 *If* (f,g) *is a bet, then it cannot be the case that* $f\succ_{NBP} g$.

Proposition 1 partly justifies the term "No-Betting-Pareto", as it shows that Condition (ii) of the definition of $\succ_{NBP}$ rules out Pareto improvements that are bets. However, this condition also rules out many Pareto improvements that are not bets. One may wonder what other Pareto improvements do not qualify as improvements according to $\succ_{NBP}$.

In the following subsection we characterize these *excluded improvements* and show that, in a certain sense, they can be thought of as bets as well.

3.3 Characterization of Excluded Improvements

Consider an improvement (f, g) and assume that $f \succ_{NBP} g$ does not hold. That is, $f \succ_* g$, but there does not exist a probability p_0 such that

$$\int_S u_i(f(s))dp_0 > \int_S u_i(g(s))dp_0$$

for all $i \in N(f, g)$. Thus, the swap of f for g, which is a Pareto improvement, is ruled out by our more restrictive definition. It turns out that the agents who are interested in such a swap would also have been interested in betting:

Theorem 2 *Fix utilities $(u_i)_i$. Assume that the beliefs $(p_i)_i$ are such that $f \succ_* g$, but that $f \succ_{NBP} g$ does not hold. Let $g' \in F$ be such that $g'(s)_i$ is independent of s for each $i \in N(f, g)$. Then there exists an act $f' = f'(f, g, g') \in F$ such that (f', g') is a bet for the utilities $(u_i)_i$ and any beliefs $(p_i)_i$ such that $f \succ_* g$.*

If (f', g') is a bet, then by definition we have $|N(f', g')| \geq 1$, and the feasibility constraint and the assumption that all utilities are monotone then ensures that $|N(f', g')| \geq 2$. However, $N(f', g')$ may be a proper subset of $N(f, g)$.

The theorem states that, if our definition of No-Betting-Pareto improvement rules out a Pareto-improving swap, then the agents involved in it would have been willing to engage also in pure betting, had their allocations been independent of the state of the world. That is, had the agents already held the full-insurance allocation g', one could have found an act f' such that (f', g') would be a bet. This bet need not involve all agents in $N(f, g)$, but it involves at least two. Importantly, this bet depends only on the utilities

$(u_i)_i$ and the acts f, g, but not on the beliefs $(p_i)_i$. Indeed, it is easy to see that, given the actual beliefs p_i, one can find a bet that would have been accepted by the agents with these beliefs. In fact, to this end it suffices to take two agents whose beliefs differ. However, the statement of the theorem is stronger: one can find such a bet that would be accepted by at least two of the agents, independently of their actual beliefs, as long as these beliefs satisfy $f \succ_* g$. Thus, a hypothetical bookie who would have tried to offer such a bet and make a sure profit based on the agents' differences in beliefs could do so without knowing the agents' actual beliefs: it is sufficient to know that these beliefs make them prefer f over g.[4]

One can interpret the theorem using an imaginary scenario, according to which agents who wish to trade have to seek the approval of a market maker, whose job is to verify that Condition (ii) holds, that is, that trade is not a result of spurious unanimity. Assume that the market maker rejects a proposed trade, because there is no joint belief that can justify it. The agents appeal and argue that they wish to trade not in order to bet, but in order to share risks. Indeed, they point out to the market maker that they hold risky positions: g is not assumed to be constant, and thus it doesn't offer them full insurance, while f presumably does better in this respect. However, the market maker would then reply, "According to your interest in trade, I know something about your beliefs, and, in particular, I know that, even if you were to share risks and be fully insured, you would still be interested in betting. In other words, even once you finish all the risk-sharing trades one can imagine, a smart bookie will be able to make a sure profit by offering you bets you'd accept. Hence, I suspect that the proposed trade already contains a non-negligible aspect of betting and I do not approve it."

Clearly, this imaginary dialog need not take place in reality, nor is it suggested here that financial markets be regulated by market makers who verify

[4]Note, however, that the price that the bookie would require of each agent for the bet will depend on the belief p_i.

that each trade No-Betting-Pareto dominates the status quo. The scenario above is merely a rhetorical device, intended to support the reasonability of the concept we propose: relying on Theorem 2, such a scenario indicates that the trades that our definition does not allow have a flavor of betting.

3.4 A Comment on Exchange Economies

As suggested above, if the set L determining each agent's utility is multi-dimensional, the definition of a bet needs to be modified. For example, if $L = \mathbb{R}^K$, denoting bundles composed of K goods, one may have a constant act g and a constant act f such that, at each state s, $f(s)$ Pareto dominates $g(s)$ due to an exchange of goods under certainty as in Example 1 (dealing with mangos and bananas). One way to rule out this possibility is to define (f, g) as a bet if it is an improvement, $g(s)$ is independent of s, and g constitutes a Pareto optimal allocation. In such a set-up a counterpart of Theorem 2 can be proved, under the assumption that for at least one good k, $g(s)_{i.k} > 0$ for all i.

4 Properties

4.1 Transitivity

Condition (ii) of the definition of No-Betting-Pareto domination involves an existential quantifier, and this raises the question, is the relation transitive? The negative answer is given by:

Proposition 2 *The relation $\succ_{NBP}$ is acyclic but it need not be transitive.*

This result means that two consecutive Pareto improvements, each of which is not a matter of spurious unanimity, may result in a Pareto improvement that is spurious in the sense that no shared beliefs can justify it.

Denoting the transitive closure of $\succ_{NBP}$ by $\succ_{NBP}^t$, we observed that $\succ_{NBP} \subsetneq \succ_{NBP}^t$. It is natural to ask, how large can the relation $\succ_{NBP}^t$ be?

Is it the case that every improvement can be obtained by a sequence of No-Betting-Pareto improving exchanges? It turns out that the answer is in the affirmative under the following condition.

The range of u

$$range(u) = \{\, u(f) \mid f \in F \} \subset \mathbb{R}^n$$

is *rectangular* if the following condition holds: for every $(f_i)_i \in F^n$, there exists $f^* \in F$ such that, for all $i \in N$, $u_i(f^*(s)) = u_i(f_i(s))$ for all $s \in S$. Rectangularity means that, if certain utilities can be obtained for each agent separately, then the profile of these utilities can also be obtained for all of them simultaneously. We can now state:

Proposition 3 *Assume that* $range(u)$ *is rectangular and convex. Then* $\succ^t_{NBP} = \succ_*$.

The Proposition states that for every improvement (f, g) there exists a finite sequence $h_0 = g$, $h_1, ..., h_L = f$ such that $h_l \succ_{NBP} h_{l-1}$ for $1 \leq l \leq L$. This might suggest that, while our definition attempts to rule out certain swaps, it does not do so very effectively: any swap that the agents eventually wish to perform ($f \succ_* g$) can be carried out by a sequence of swaps, each of which qualifies as a No-Betting-Pareto improvement.

However, rectangularity is a very strong condition. In particular, it would be in conflict with any reasonable feasibility constraints. Assuming, as in Subsection 3.2 that $X = \mathbb{R}^n$, one may consider only feasible improvements (f, g). Clearly, limiting attention to $\{\, u(f) \mid (f, g)$ is a feasible improvement$\}$, rectangularity does not hold and neither does the conclusion of Proposition 3.

More generally, when $X = \mathbb{R}^n$, we may refine the definition of $\succ_{NBP}$ to consider only feasible improvements. Say that f *feasibly No-Betting-Pareto dominates* g, $f \succ_{fNBP} g$, if (f, g) is a feasible improvement, and $f \succ_{NBP} g$. Let $\succ^t_{fNBP}$ be the transitive closure of this relation. Then we mention:

Proposition 4 *If (f,g) is a bet, then it cannot be the case that $f \succ^t_{fNBP} g$.*

Thus, restricting attention to feasible improvements allows us to strengthen Proposition 1: with these improvements, even a sequence of No-Betting-Pareto dominations cannot "simulate" a bet.

Finally, we mention that, even if rectangularity were satisfied, Proposition 3 does not suggest a realistic way of implementing bets by sequences of No-Betting-Pareto improvements. To perform a sequence of such swaps, the agents involved need to plan the sequence and follow it. This might not be practical for various problems of coordination. Further, in the absence of a commitment device (which would make a sequence of swaps equivalent to a single one), agents may not trust other agents to continue trading along the pre-specified sequence.

4.2 Computation

It is worthy of mention that Condition (ii) is not difficult to verify from a computational viewpoint. To state this result one has to determine how acts are represented by finite strings of bits. Since the acts discussed are simple, it is natural to think of them as finite vectors. Specifically, given f and g, there is a finite measurable partition of S, $(A_j)_{j \le J}$, such that both f and g are constant over each A_j. Thus, we use the notation $f(A_j), g(A_j)$ to denote the elements of X that f and g, respectively, assume over A_j, for each $j \le J$. Next, assume that the utility values $(u_i(f(A_j)))_j, (u_i(g(A_j)))_j$ are rational numbers for every i. Under these assumptions, the following result states that verifying whether Condition (ii) is verified for f and g is an "easy" task.

Proposition 5 *Given rational numbers, $(u_i(f(A_j)))_{i,j}, (u_i(g(A_j)))_{i,j}$ it can be checked in polynomial time complexity whether Condition (ii) holds.*

Thus, the imaginary scenario in which a market maker needs to approve swaps may be implausible for a variety of reasons, but complexity is not one of them: if a set of agents propose an exchange (f,g), their incentive compatibility constraint guarantees that $f \succ_* g$. To check whether it is also the case that $f \succ_{NBP} g$, one needs to verify Condition (ii). As Proposition 5 states, this is a simple computational task, given the utility profiles of the agents under f and under g.

5 Relation to Utilitarian Aggregation

As mentioned in the Introduction, the current paper shares much of its motivation with Gilboa, Samet, and Schmeidler (2004) (GSS). That paper employed a restricted unanimity condition, stating that society should necessarily agree with all individuals' preferences (where the latter agree) only when these preferences concern alternatives over which there are no disagreements in beliefs. The current paper also restricts a unanimity-style condition, namely, Pareto dominance, to agreements in beliefs. It may therefore be useful to clarify the relationship between the two papers.

In GSS it is implicitly assumed that the entire preference relation of each individual i, $\succsim_i$, is observable, and the question is, what conditions should the preference relation of society, $\succsim_0$, satisfy. Since for each i the entire preference relation is observable, and it is assumed to satisfy Savage's axioms, the social planner can also figure out each individual's probability measure, p_i, and tell, for each act f, whether it is an act on whose distribution all individuals agree. The restricted Pareto condition suggested in GSS states that society should find f as desirable as g when all individuals do so, *if* both f and g are such acts (but not necessarily if one of f or g induces different distributions over outcomes according to different individuals' beliefs). The result of that paper is that, when one restricts the Pareto condition in this way, the simultaneous aggregation of utilities and of probabilities becomes possible, and, moreover, linear aggregation of both is, under some conditions, necessitated by the restricted Pareto condition.

By contrast, the present paper does not assume that individuals' entire preference relations, or probabilities, are observable. Nor does it ascribe to society a complete preference over alternatives. It merely discusses a particular instance of unanimous preferences, $f \succsim_i g$ for all i (with strict preference for at least one), and asks whether society should agree with that particular ranking, that is, whether we should have $f \succ_0 g$ for that pair f, g. Importantly, shared beliefs appear in this paper in a very different way than in GSS: whereas in the latter shared beliefs over the outcomes of the acts f, g is assumed (for the Pareto condition to apply), here we consider acts over whose distributions individuals may well disagree. However, it is required that one could come up with shared *hypothetical* beliefs for the individuals that would still rationalize trade for each of them.

To compare the two approaches more sharply, one may ignore the observability of beliefs and the completeness of society's preference and pose a more concrete problem. Suppose that for each individual i we have a relation $\succsim_i$ defined by the maximization of the expectation of a function u_i relative to a probability p_i. Assume that we strengthen GSS's assumptions (adding conditions of monotonicity and symmetry) such that, given utilities $(u_i)_{i=1}^n$ and probabilities $(p_i)_{i=1}^n$ of the agents, society maximizes expected utility with respect to the utility function and the probability measure given by

$$u_0 = \frac{1}{n}\sum_{i=1}^{n} u_i$$

$$p_0 = \frac{1}{n}\sum_{i=1}^{n} p_i.$$

Let $\succsim_0$ be the resulting ordering, which is obviously complete, and, moreover, satisfies the rest of Savage's axioms with the possible exception of P6 (which requires that p_0 be a non-atomic measure). Let $\succ_0$ denote the asymmetric part of $\succsim_0$.

One may now ask whether there is any relationship between the utilitarian strict preference ($\succ_0$) and No-Betting Pareto domination ($\succ_{NBP}$). The negative answer is given by the following.

Proposition 6 *The relation $\succ_0$ need not imply $\succ_{NBP}$ nor vice versa.*

The two relations are quite different also from a conceptual point of view: GSS deals with an attempt to simultaneously aggregate tastes and beliefs, and its main motivation are group decisions. When a country has to choose an economic policy, decide whether to use nuclear power plants, or whether to wage a war, the decision cannot be decentralized; it has to be made for all individuals as a group. In this context, GSS show that the natural idea, of simultaneous averaging of utilities and of probabilities, is necessitated by a reasonable version of the unanimity (Pareto) axiom. However, these "averaged" preferences are not necessarily very relevant for decentralized decisions. When economic agents interact in markets, each can make her own decisions according to her tastes and beliefs, and there is no need to define an "averaged" individual or a representative agent. Hence, economists would tend to eschew the task of defining a social welfare function or a complete preference order for society as a whole. Rather, a weaker notion such as Pareto dominance can be defined, restricting normative claims to those that can be made in the language of this partial relation. The current paper belongs in this tradition. It differs from the classical literature in its definition of "dominance". In an attempt to avoid trade that is basically a bet, our new definition further restricts the notion of dominance, making the social preference relation even further from completeness than is the standard notion of Pareto dominance.

6 Discussion

6.1 Pareto Rankings

When one discusses pure consumption goods, as in Example 1, it seems compelling that one does not wish to settle for given allocations if Pareto superior ones are feasible. The first welfare theorem then provides an argument in support of complete competitive markets: they offer at least one way in which Pareto-dominated allocations can be avoided. As is well known, Pareto optimality is a weak normative concept, which remains silent on important issues such as equality, poverty, and well-being. Further, the conditions of the welfare theorems are hardly met in reality, with known classes of robust examples that give rise to sub-optimal equilibria. Yet, these qualifications notwithstanding, the welfare theorems do provide a powerful argument in favor of complete competitive markets.

When uncertainty is considered, it is very tempting to model the state of the world as one of the features of a good and reduce the problem to a known one. It is an elegant exercise that suggests that the argument in favor of complete competitive markets for consumption goods should also extend to any markets involving uncertainty, including financial markets. But in these markets, where a strong speculative component exists, beliefs tend to vary across agents. We argue that the welfare analysis should be revisited in this context. In particular, the standard argument against incomplete markets again suggests that the absence of certain assets may lead to Pareto dominated allocations, and hence only if we have a set of assets that spans the space of functions over the state space (such as Arrow securities) can we trust free trade to guarantee Pareto optimality. However, when highly complex financial derivatives are discussed, higher-level beliefs will typically be implicitly involved in the definition of a state. One suspects that an underlying state space that is rich enough to describe such beliefs ("states of the world" as opposed to "states of nature") will allow for a non-negligible amount of speculation, alongside risk sharing. Our definition of No-Betting-Pareto domination attempts to draw the line between the two. It suggests that, in the context of higher order beliefs, incomplete markets are not necessarily inferior to complete markets.

It is evident that our argument relies on the assumption that beliefs do differ. If all agents shared the same beliefs, Pareto domination under uncertainty would be as convincing as it is (or isn't) under certainty. However, the claim that all agents share the same prior beliefs, championed by Harsanyi (1967-68), does not appear realistic. The agreeing-to-disagree and no-trade results (Aumann, 1976, Milgrom and Stokey, 1982) show that the common prior assumption implies that rational agents (in a model that is common knowledge among them) cannot agree to disagree and should not trade in financial markets, even as a result of the arrival of new information. The prevalence of different beliefs and the large volumes of trade in financial markets suggest either that the common prior assumption does not hold, or that rationality is not common knowledge (or both). In any event, Pareto dominance becomes a problematic concept.

Should financial markets be regulated, as suggested by Posner and Weyl (2012)? We do not find that theoretical arguments provide a compelling answer to this question. There are weighty arguments for regulation, especially if agents might be prone to psychological biases, and there are also weighty arguments against regulation. The present contribution does not attempt to resolve this issue. Rather, we only wish to fine-tune a certain theoretical argument that might be brought forth in the context of this debate. Thus, without taking a stance on desired policy, we argue that one standard argument for free markets does not apply in this context without an appropriate qualification.

6.2 Extensions

Our definition of a bet (f, g) assumes that the given allocation, g, is constant across the state space. This is obviously restrictive. For example, assume that two agents are considering a bet on the outcome of a soccer match. It so happens that their current wealth does not depend on this match in any way. Yet, their current allocations are far from constant, as the two are exposed to various risks, ranging from their health to stock market crashes.

To capture this type of exchange in the definition of a bet, one has to allow the existing allocations g to depend on s, but to be independent of the exchange. That is, the variable $f - g$ should be stochastically independent of g according to all the probability measures considered. In other words, one may assume that the state space is a product of two spaces, $S = S_1 \times S_2$ such that g is measurable with respect to S_1, and consider only probabilities obtained as a product of a measure p_1 on S_1 and a measure p_2 on S_2. Relative to such a model, ours can be viewed as a reduced form model, where our entire discussion is conditioned on a state $s_1 \in S_1$.

7 Appendix: Proofs and Related Analysis

7.1 Proof of Theorem 1

This is a standard application of a duality/separation argument. Let there be given two acts f, g. As each of them is simple and measurable, there is a finite measurable partition of S, $(A_j)_{j\leq J}$, such that both f and g are constant over each A_j. Thus, we use the notation $f(A_j), g(A_j)$ to denote the elements of X that f and g, respectively, assume over A_j, for each $j \leq J$.

The theorem characterizes condition (ii) of the definition of No-Betting-Pareto domination, namely that there be a probability vector p_0 such that, for all i,

$$\int_S u_i(f(s))dp_0 > \int_S u_i(g(s))dp_0. \tag{2}$$

We first note that (2) holds if and only if there exists a probability vector $(p_0(j))_{j\leq J}$, such that, for all i,

$$\sum_{j\leq J} p_0(j)\, u_i(g(A_j)) > \sum_{j\leq J} p_0(j)\, u_i(g(A_j)). \tag{3}$$

In particular, if a measure p_0 that satisfies (2) exists, it induces a probability vector $(p_0(j))_{j\leq J}$ (over $(A_j)_{j\leq J}$) that satisfies (3). Conversely, if a vector $(p_0(j))_{j\leq J}$ satisfying (3) exists, it can be extended to a measure p_0 on (S, Σ) such that (2) holds. (Since f and g are constant over each A_j, the choice of the extension does not matter.)

When is there a probability vector $(p_0(j))_{j\leq J}$ satisfying (3)? Consider a two-person zero-sum game in which player I chooses an event in $(A_j)_{j\leq J}$ and player II chooses an agent in $N(f,g)$. The payoff to player I, should she choose A_j and player II choose $i \in N(f,g)$, is $u_i(f(A_j)) - u_i(g(A_j))$. Then (3) is equivalent to the existence of a mixed strategy of player I, $p_0 \in \Delta\left((A_j)_{j\leq J}\right)$ such that, for every pure strategy of player II, $i \in N$,

$$\sum_{j\leq J} p_0(j)\left[u_i(f(A_j)) - u_i(g(A_j))\right] > 0$$

or: there exists $p_0 \in \Delta\left((A_j)_{j \leq J}\right)$ such that for all mixed strategy of player II, $\lambda \in \Delta(N)$,

$$\sum_{j \leq J} p_0(j) \sum_{i \in N} \lambda(i) \left[u_i(f(A_j)) - u_i(g(A_j))\right] > 0.$$

In other words, $\exists p_0 \in \Delta\left((A_j)_{j \leq J}\right)$ such that $\forall \lambda \in \Delta(N)$

$$E_{p_0} E_\lambda \left[u_i(f(A)) - u_i(g(A))\right] > 0,$$

where A denotes a generic member of $(A_j)_{j \leq J}$. The above is equivalent to

$$\max_{p \in \Delta\left((A_j)_{j \leq J}\right)} \min_{\lambda \in \Delta(N)} E_p E_\lambda \left[u_i(f(A)) - u_i(g(A))\right] > 0,$$

which, by the minmax theorem for zero-sum games, is equivalent to

$$\min_{\lambda \in \Delta(N)} \max_{p \in \Delta\left((A_j)_{j \leq J}\right)} E_p E_\lambda \left[u_i(f(A)) - u_i(g(A))\right] > 0,$$

that is, to the claim that $\forall \lambda \in \Delta(N)$ there exists $p \in \Delta\left((A_j)_{j \leq J}\right)$ such that

$$E_\lambda E_p \left[u_i(f(A)) - u_i(g(A))\right] > 0.$$

It follows that (3) holds if and only if for every $\lambda \in \Delta(N)$ there exists $p \in \Delta\left((A_j)_{j \leq J}\right)$ such that

$$\sum_{i \in N} \lambda(i) \sum_{j \leq J} p(j) \left[u_i(f(A_j)) - u_i(g(A_j))\right] > 0.$$

However, for each $\lambda \in \Delta(N)$, such a $p \in \Delta\left((A_j)_{j \leq J}\right)$ exists if and only if there exists such a p that is a unit vector, namely, if and only if there exists $j \leq J$ such that

$$\sum_{i \in N} \lambda(i) \left[u_i(f(A_j)) - u_i(g(A_j))\right] > 0,$$

and this is the case if and only if there exists a state $s \in S$ such that

$$\sum_{i \in N} \lambda(i) \left[u_i(f(s)) - u_i(g(s))\right] > 0.$$

$\blacksquare$

Observe that, should one use the weak inequality version of Condition (ii), a similar characterization holds: there exists a probability vector p_0 such that, for all i,

$$\int_S u_i(f(s))dp_0 \geq \int_S u_i(g(s))dp_0$$

if and only if for every $\lambda \in \Delta(N)$ there exists a state s such that

$$\sum_{i \in N} \lambda(i) \left[u_i(f(s)) - u_i(g(s))\right] \geq 0.$$

7.2 Proof of Proposition 1

We first show that $f \succ_{NBP} g$ cannot hold if (f, g) is a bet. Let there be given a bet (f, g). That is, $f \succ_i g$ for all $i \in N(f, g)$ and

(i) $g(s)_i$ is independent of s for each i;

(ii) $\sum_i f(s)_i \leq \sum_i g(s)_i$ for all s.

We provide two short proofs. First, observe that, if it were the case that $f \succ_{NBP} g$, there would be a belief p_0 such that

$$\int_S u_i(f(s)_i)dp_0 > \int_S u_i(g(s)_i)dp_0$$

for all $i \in N(f, g)$. For each $i \in N(f, g)$, let $\bar{g}_i = g(s)_i$ and $\bar{u}_i = u_i(g(s)_i)$ for all s. Then we have

$$E_{p_0}\left(u_i\left(f_i\right)\right) > E_{p_0}\left(u_i\left(g_i\right)\right) = \bar{u}_i$$

and, since u is concave,

$$u_i\left(E_{p_0}\left(f_i\right)\right) \geq E_{p_0}\left(u_i\left(f_i\right)\right)$$

thus

$$u_i\left(E_{p_0}\left(f_i\right)\right) > \bar{u}_i$$

and, because u is strictly monotone,

$$E_{p_0}\left(f_i\right) > \bar{g}_i.$$

Summation over $i \in N\left(f,g\right)$ yields

$$\sum_i E_{p_0}\left(f_i\right) = E_{p_0}\left(\sum_i f_i\right) > \sum_i \bar{g}_i$$

which is a contradiction because $\left(\sum_i f_i\right)(s) \leq \sum_i \bar{g}_i$ for all s.

The second proof makes use of Theorem 1. To this end, consider the vector of weights $\lambda = \left(\lambda_i\right)_i$ defined by

$$\lambda_i = \begin{cases} \frac{1}{|N(f,g)|} & i \in N\left(f,g\right) \\ 0 & \text{otherwise} \end{cases}.$$

Because $\sum_i f(s)_i \leq \sum_i g(s)_i$ for all s, we also have $\sum_{i \in N(f,g)} f(s)_i \leq \sum_{i \in N(f,g)} g(s)_i$ and it follows that the λ-weighted utility under f cannot exceed that corresponding to g. Thus, the λ-weighted "average" agent cannot point to a state where she is strictly better off under f than under g. ■

7.3 Proof of Theorem 2

Let utilities $\left(u_i\right)_i$ be given. Consider an improvement (f,g) and assume that $f \succ_{NBP} g$ does not hold. That is, there does not exist a probability p_0 such that

$$\int_S u_i(f(s))dp_0 > \int_S u_i(g(s))dp_0$$

for all $i \in N(f,g)$.

Suppose that $\left(A_j\right)_{j=1}^{J}$ is a finite, measurable partition of S, which is a refinement of the two partitions of S defined by f^{-1} and g^{-1}. In other words, f and g are constant on each A_j. Let $f\left(A_j\right)$, $g\left(A_j\right) \in X$ denote their values, correspondingly, on the elements of the partition, for $j \leq J$. Consider a probability over (S, Σ), restricted to the elements of the partition (and their unions). With a minor abuse of notation this probability is still denoted by p, and we write $p(j)$ instead of $p(A_j)$. Let Δ^{J-1} denote the simplex of all such probabilities.

Each $i \in N(f, g)$ would strictly prefer f to g whenever her belief p is in

$$C_i = \left\{ p \in \Delta^{J-1} \left| \sum_{j \leq J} p(j) \left(u_i \left(f \left(A_j \right) \right) - u_i \left(g \left(A_j \right) \right) \right) > 0 \right. \right\}.$$

Observe that, since $f \succ_* g$ (i.e., $f \succ_i g$ $\forall i \in N(f, g)$), it has to be the case that $p_i \in C_i$ $\forall i \in N(f, g)$. Clearly, for such f, g, $f \succ_{NBP} g$ does not hold if and only if $\cap_{i \in N(f,g)} C_i = \varnothing$.

For simplicity of notation, assume $N = N(f, g)$. Without loss of generality, assume that the state space is $\{1, ..., J\}$, that is, that $A_j = \{j\}$. Also without loss of generality, assume that $g'(j)_i = 0$ for all $i \in N$, $j \leq J$.

We mention:

Claim 0: For each $i \in N$, C_i has a non-empty interior relative to the simplex Δ^{J-1}.

Proof: Since $f \succ_i g$, we know that C_i is non-empty, as agent i's actual beliefs p_i lie in C_i. Then C_i has a non-empty interior relative to the simplex, as it is the non-empty intersection of an open half-space and the simplex. □

We need to construct an act f' such that (f', g') is a bet, that is, such that $f' \succ_i g'$ for all $i \in N(f', g')$ and all $p_i \in C_i$, and $\sum_i f'(j)_i = 0$ for all j. To this end, we start by constructing an act f'' such that $\sum_i f''(j)_i = 0$, and, for every $i \in N(f'', g')$, $\sum_i p_i(j) f''(j)_i > 0$. (The last step of the proof would consist of defining f' as a multiple of f'' by a small positive constant.)

Step 1: First, we fix beliefs $p_i \in C_i$ and construct a bet (f'', g') for these beliefs. This would also prove a weaker version of the theorem, in which a bookie can find a bet, if the bookie knows the actual beliefs (p_i) (and not only that they lie in the respective C_i).

Define, for $k \geq 1$ and $i \in N$,

$$C_i^k = \left\{ p \in \Delta^{J-1} \left| \sum_{j \leq J} p(j) \left(u_i \left(f\left(A_j\right)\right) - u_i \left(g\left(A_j\right)\right)\right) \geq \frac{1}{k} \right.\right\}$$

so that, for all k, $C_i^k \subset C_i^{k+1} \subset C_i$ and $C_i = \cap_k C_i^k$. Since we have $\cap_{i \in N} C_i = \varnothing$, it is certainly true that $\cap_{i \in N} C_i^k = \varnothing$ for all k. However, C_i^k is a non-empty, convex, and (as opposed to C_i) also compact subset of Δ^{J-1}. When such compact and convex sets of priors have an empty intersection, it is known that one can find a bet that they would all accept, as long as their beliefs are in the specified sets of priors. Specifically, Theorem 2 in Billot et al. (2000, p. 688) states that there are linear functionals h_i, such that h_i is strictly positive on C_i^k, and $\sum h_i = 0$.[5] Thus, for each k there exists a $n \times J$ matrix, $\left(h_{i,j}^k\right)_{i,j}$ of real numbers, such that

$$\sum_i h_{i,j}^k = 0 \qquad \forall j$$

and

$$\sum_j p(j) h_{i,j}^k > 0 \qquad \forall i, \quad \forall p \in C_i^k.$$

Since, for each i, $p_i \in C_i$, for each i there exists $K = K(i)$ such that $p_i \in C_i^k$ for $k \geq K$. Let $K_0 = \max_i K(i)$, and note that $f''(j)_i = h_{i,j}^{K_0}$ satisfies $\sum_i p_i(j) f''(j)_i > 0$ and $\sum_i f''(j)_i = 0$ as required.

Step 2: We now wish to show that the construction of f'' above can be done in a uniform way: there exists an f'' such that $\sum_i f''(j)_i = 0$ and $\sum_i p_i(j) f''(j)_i > 0$ for all $p_i \in C_i$ and all $i \in N(f'', g')$. (Observe, however,

[5] Similar theorems have been proved by Bewley (1989) and Samet (1998). Billot et al. provide a stronger result, also saying that the hyperplanes corresponding to the functionals h_i can be chosen so that they intersect at a point in the convex hull of the sets of priors, but this geometric fact is not used here.

that while in Step 1 we obtained a bet that involved all agents, here we may find that $N(f'', g') \subsetneq N$.)

Since we intend to consider a converging sub-sequence of matrices $\left(h_{i,j}^k\right)_{i,j}$, it will be convenient to consider matrices that satisfy weak inequalities. However, to veer away from the origin, we will restrict attention to matrices of norm 1. Let H denote the set of all such matrices h that satisfy

$$\sum_i h_{i,j} = 0 \qquad \forall j \tag{4}$$

$$\sum_{i,j} (h_{i,j})^2 = 1 \tag{5}$$

and,

$$\sum_j p(j) h_{i,j} \geq 0 \qquad \forall i, \quad \forall p \in C_i. \tag{6}$$

Claim 2.1: $H \neq \varnothing$.

Proof: Defining C_i^k and $\left(h_{i,j}^k\right)_{i,j}$ as above, one may assume without loss of generality that $\left(h_{i,j}^k\right)_{i,j}$ is on the unit disc, that is, that

$$\sum_{i,j} \left(h_{i,j}^k\right)^2 = 1$$

so that $h^k = \left(h_{i,j}^k\right)_{i,j} \in H$.

Because the unit disc is compact, there exists a subsequence of $\left(h^k\right)_k$ that converges to a matrix h^*. This point satisfies conditions (4, 5) because it is the limit of points that satisfy these conditions. The matrix h^* also satisfies (6) because it is the limit of matrices that satisfy this inequality (strictly) on a subset that converges to C_i. Explicitly, for any $p \in C_i$ there exists K such that for $k \geq K$, $p \in C_i^k$ and $\sum_j p(j) h_{i,j}^k > 0$, which implies $\sum_j p(j) h_{i,j}^* \geq 0$. It follows that $h^* \in H$ and $H \neq \varnothing$. □

For $h \in H$ let the set of agents who would be involved in h, were it offered as a bet, be denoted by

$$D(h) = \{ i \in N \mid \exists j, \quad h_{i,j} \neq 0 \}.$$

Clearly, $D(h) \neq \varnothing$ for $h \in H$, as h is on the unit disc and therefore cannot be 0. Also, $D(h)$ cannot be a singleton because of (4).

Claim 2.2: For $h \in H$ there is no $i \in D(h)$ such that $h_{i,j} \leq 0 \; \forall j$.

Proof: Suppose, to the contrary, that h and i satisfy $h_{i,j} \leq 0$. As $i \in D(h)$, $h_{i\cdot}$ isn't identically zero. Hence there is a j such that $h_{i,j} < 0$. In view of Claim 0, there is a $p \in C_i$ that is strictly positive. For such a p, $\sum_j p(j) h_{i,j} < 0$, contradicting (6). □

Claim 2.3: Let $h \in H$ be such that $D(h)$ is minimal (with respect to set inclusion). Then there is no $i \in D(h)$ such that $h_{i,j} \geq 0 \; \forall j$.

Proof: Assume, to the contrary, that h and i satisfy $h_{i,j} \geq 0$. As $i \in D(h)$, $h_{i\cdot}$ isn't identically zero. Hence there are j's such that $h_{i,j} > 0$. We wish to construct another matrix $h' \in H$ such that $D(h') = D(h) \setminus \{i\}$, contradicting the minimality of $D(h)$.

By (4) we know that there exists $k \in D(h) \setminus \{i\}$. Define

$$h''_{r,j} = \begin{cases} 0 & r = i \\ h_{k,j} + h_{i,j} & r = k \\ h_{r,j} & \text{otherwise} \end{cases} .$$

It is easy to verify that h'' satisfies (4). To see that (6) also holds, observe that, for i (6) is satisfied as an equality, for k the left side of (6) could have only increased, as compared to the left side of h, while it is unchanged for $r \notin \{i, k\}$.

Next we wish to show that h'' is not identically zero. If it were, we would have $h_{k,j} = -h_{i,j}$ for all j. But, since $h_{i,j} \geq 0$ (for all j), this would imply $h_{k,j} \leq 0$ (for all j), in contradiction to Claim 2.2.

It follows that h'' can be re-normalized to guarantee (5) without violating (4,6), obtaining $h' \in H$ with $D(h') \subsetneq D(h)$. □

Claim 2.4: Let $h \in H$ have a minimal $D(h)$ (with respect to set inclusion) over H. Let $i \in D(h)$. Then $(h_{i,j})$ contains both positive and negative entries.

Proof: Combine Claims 2.2 and 2.3. □

Claim 2.5: Let $h \in H$ have a minimal $D(h)$ (with respect to set inclusion) over H. Let $i \in D(h)$ and $p \in C_i$. Then $\sum_j p(j)h_{i,j} > 0$.

Proof: Because $h \in H$, we know that $\sum_j p(j)h_{i,j} \geq 0$ holds for all $p \in C_i$. Assume that it were satisfied as an equality. Distinguish between two cases (in fact, the argument for Case 2 applies also in Case 1, but the argument for the latter is simple enough to be worth mentioning):

Case 1: p is in the relative interior of Δ^{J-1} (hence also in the interior of C_i relative to Δ^{J-1}). In this case, by Claim 2.4, there exist j, j' such that $h_{i,j} < 0 < h_{i,j'}$. One can find a small enough $\varepsilon > 0$ such that $p_\varepsilon = p + \varepsilon e_j - \varepsilon e_{j'} \in C_i$ where e_j is the j-unit vector. For such a p_ε, $\sum_j p_\varepsilon(j)h_{i,j} < 0$, a contradiction to (6).

Case 2: p is on the boundary of Δ^{J-1}. Consider the problem

$$
\begin{aligned}
& Min_p \sum_j p(j)h_{i,j} \\
& \quad s.t. \\
& \qquad \sum_{j \leq J} p(j)\left(u_i\left(f\left(A_j\right)\right) - u_i\left(g\left(A_j\right)\right)\right) \ \geq \ 0 \qquad (1) \\
& \qquad p \ \in \ \Delta^{J-1}.
\end{aligned}
$$

Since $h \in H$, the optimal value of this problem is non-negative. Since $\sum_j p(j)h_{i,j} = 0$, p is a solution to the problem. However, because $p \in C_i$, constraint (1) is inactive at p. Given that this is a linear programming problem, removing an inactive constraint cannot render p sub-optimal. Hence p is also an optimal solution to $Min_p \sum_j p(j)h_{i,j}$ subject to $p \in \Delta^{J-1}$. But this implies that $\sum_j p(j)h_{i,j} \geq 0$ for all $p \in \Delta^{J-1}$. This, in turn, implies that $h_{i,j} \geq 0$ for all j, contradicting Claim 2.3. □

To complete the proof of Step 2, all we need to do is define $f'' = h$ for some $h \in H$ for which $D(h)$ is minimal with respect to set inclusion, and observe that $N(f'', g') = D(h)$. □

Step 3: Finally, we turn to construct the act f' such that (f', g') is a bet, that is, such that $\sum_i f'(j)_i \equiv 0$ and, for every $i \in N(f', g')$, $\sum_i p_i(j) u_i(f'(j)_i) > 0$. Consider an act $f'_\alpha = \alpha f''$ for $\alpha > 0$. Clearly, $\sum_i f'_\alpha(j)_i = 0$ for all j and all α. As u_i are differentiable, for a small enough α the conclusion $\sum_i p_i(j) u_i(f'_\alpha(j)_i) > 0$ follows. ■

7.4 Proof of Proposition 2

It is obvious that $\succ_{NBP}$ does not admit cycles, because strict preference for each agent i, $\succ_i$, is acyclic.

To see that transitivity may fail, consider the following example.

Let there be two agents $N = \{1, 2\}$ and two states $S = \{s, t\}$. Let the agents' beliefs be $p_1 = (1, 0)$, $p_2 = (0, 1)$ and let g, h and f be acts with the following utility profiles:

		State s	State t
$g:$	Agent 1	0	0
	Agent 2	0	0

		State s	State t
$h:$	Agent 1	2	-1
	Agent 2	-3	2

		State s	State t
$f:$	Agent 1	4	-4
	Agent 2	-4	4

First observe that $f \succ_i h \succ_i g$ according to the agents' actual beliefs. Moreover, agent 1 will find h better than g for any belief $(p, 1-p)$ such that $p > \frac{1}{3}$ and agent 1 will find f better than h for any belief with $p > \frac{3}{5}$. Agent 2, by contrast, will prefer h to g whenever $p < \frac{2}{5}$ and f to h for $p < \frac{2}{3}$. Thus, both agents prefer h to g for $p \in \left(\frac{1}{3}, \frac{2}{5}\right)$ and f to h for $p \in \left(\frac{3}{5}, \frac{2}{3}\right)$. However, f cannot No-Betting-Pareto dominate g as there is no belief for which both agents prefer f to g. (One can also use Proposition 1 to observe that, for

$\lambda = 0.5$, the λ-average of utilities is identical under f as under g at each state.) ■

7.5 Proof of Proposition 3

Assume that $range(u)$ is rectangular and convex. We need to show that $\succ^t_{NBP} = \succ_*$. The inclusion $\succ^t_{NBP} \subset \succ_*$ is immediate: for f, g, $f \succ_{NBP} g$ implies that $f \succ_i g$ for all i, that is, $f \succ_* g$. By transitivity of $\succ_i$ and $\succ_*$, $\succ^t_{NBP} \subset \succ_*$.

To see the converse, assume that $f \succ_* g$. We need to construct a finite sequence $h_0 = g$, $h_1, ..., h_L = f$ such that $h_l \succ_{NBP} h_{l-1}$ for $1 \leq l \leq L$. The basic idea is quite simple: setting $L = |N(f, g)|$, we improve the outcome vector of the agents in $N(f, g)$ one at a time, so that only agent i gets a different utility vector under h_i as compared to h_{i-1}, for $i = 1, ..., L$. In other words, agent i gets $u_i(g(\cdot))$ under h_l for $l < i$ and $u_i(f(\cdot))$ under h_l for $l \geq i$. This will be possible thanks to the fact that $range(u)$ is rectangular. To show that there exists one probability, p_0, according to which h_i is at least as desirable as h_{i-1} for all agents, one may take p_0 to be p_i. Since $f \succ_i g$, we know that agent i is strictly better off under h_i than under h_{i-1} according to $p_0 = p_i$. The other agents obtain the same utility vector, and are thus indifferent between h_i and h_{i-1} according to all probabilities, and, in particular, according to $p_0 = p_i$. However, according to this construction only agent i *strictly* prefers h_i to h_{i-1}. Therefore, we modify the definition of $h_1, ..., h_L$, making use of convexity of $range(u)$, to guarantee strict preferences for all agents throughout the sequence.

Let there be given an improvement (f, g). Let $(A_j)_{j \leq J}$ be a measurable partition of S so that both f and g are constant over each A_j. Without loss of generality assume that $A_j = \{j\}$ and that $N = N(f, g)$. For $i, k \in N$, let $u_k(h'_i)$ as explained above:

$$u_k(h'_i(j)) = \begin{cases} u_k(g(j)) & i < k \\ u_k(f(j)) & i \geq k \end{cases} \qquad \forall j \in J.$$

Again, such h_i' exist because of the rectangularity condition.

Next we construct $(u_k(h_i))_{k,i}$ from $(u_k(h_i'))_{k,i}$ as follows. Given that $f \succ_i g$ for all i, we have

$$\sum_{j \leq J} p_i(j) \left[u_i(f(j)) - u_i(g(j))\right] > 0$$

for all $i \in N$. Let $\varepsilon > 0$ be small enough such that, for every $i \in N$,

$$\sum_{j \leq J} p_i(j) \left[u_i(f(j)) - (n-1)\varepsilon - u_i(g(j))\right] > 0 \tag{7}$$

i.e., $0 < \varepsilon < \frac{1}{n-1} \sum_{j \leq J} p_i(j) \left[u_i(f(j)) - u_i(g(j))\right]$.

Choose $(h_i)_{1 \leq i < L}$ such that

$$u_k(h_i(j)) = \begin{cases} u_k(g(j)) + i\varepsilon & i < k \\ u_k(f(j)) - (L-i)\varepsilon & i \geq k \end{cases}$$

with $h_L = f$. Observe that such $(h_i)_{1 \leq i < L}$ exist because their utility vectors are in the convex hull of those of $(h_i')_{1 \leq i < L}$.

It follows that, for all $i \leq L$, all $k \in N \backslash \{i\}$, and all $j \leq J$,

$$u_k(h_i(j)) - u_k(h_{i-1}(j)) = \varepsilon > 0$$

so that, for agent k, h_i strictly dominates h_{i-1}. In particular, whatever are agent k's beliefs, she strictly prefers h_i to h_{i-1}. In particular, this is true both for agent k's actual beliefs p_k and for agent i's beliefs, p_i. As for $k = i$, (7) guarantees that agent i also prefers h_i to h_{i-1}. Thus, all agents prefer h_i to h_{i-1} both given their actual beliefs and given $p_0 = p_i$, and thus $h_i \succ_{NBP} h_{i-1}$. ■

7.6 Proof of Proposition 4

Let (f, g) be a bet. If $f \succ^t_{fNBP} g$, we would have

$$f = h_L \succ_{fNBP} h_{L-1} \succ_{fNBP} \dots \succ_{fNBP} h_1 \succ_{fNBP} g.$$

However, no h_1 can feasibly-No-Betting-Pareto dominate g as in the proof of Proposition 1. ■

7.7 Proof of Proposition 5

Given the rational numbers $\left(u_i\left(f(A_j)\right), u_i\left(g(A_j)\right)\right)_{i,j}$ we need to check if there exists a probability vector $p_0 \in \Delta\left(\left(A_j\right)_{j\leq J}\right)$ such that, for all $i \in N\left(f, g\right)$,

$$\sum_{j\leq J} p_0(j)\left[u_i(f(A_j)) - u_i(g(A_j))\right] > 0.$$

Observe first that one can easily identify the set $N\left(f, g\right)$. Consider the maximization problem

$$\begin{aligned} &Max_{p_0(1),\ldots,p_0(J)} \quad y \\ &\quad s.t. \\ &\qquad \sum_{j\leq J} p_0(j)\left[u_i(f(A_j)) - u_i(g(A_j))\right] - y \ \geq \ 0 \qquad \forall i \in N \\ &\qquad \sum_{j\leq J} p_0(j) \ = \ 1 \\ &\qquad p_0(j) \ \geq \ 0 \qquad \forall j \leq J. \end{aligned}$$

The optimal value of this problem is positive if and only if there exists a probability vector $p_0 \in \Delta\left(\left(A_j\right)_{j\leq J}\right)$ such that $\sum_{j\leq J} p_0(j)\left[u_i(f(A_j)) - u_i(g(A_j))\right] > 0$ for every $i \in N\left(f, g\right)$, which is easy to verify because linear programming can be solved in polynomial time complexity. ■

7.8 Proof of Proposition 6

Let there be two agents, $1, 2$, and two states s, t. Agent 1 has beliefs $p_1 = (1, 0)$ and agent 2 has beliefs $p_2 = (0, 1)$. The analysis does not hinge on these extreme beliefs, as all preferences will be strict. Throughout the proof, let us suppose that g induces the matrix of utilities

		State s	State t
g :	Agent 1	0	0
	Agent 2	0	0

The first two examples show that we might have $f \succ_0 g$ but not $f \succ_{NBP} g$. The third one shows the converse.

7.8.1 Example

We first show that $f \succ_0 g$ might occur when there is no Pareto domination of f over g even according to the standard definition. Let the utilities induced by f be

		State s	State t
$f:$	Agent 1	-9	11
	Agent 2	11	-9

Clearly, if we average utilities we get $+1$ for sure, and society's (u_0, p_0)-expected-utility maximization would favor the f over g, that is $f \succ_0 g$. However, none of the agents prefers f to g, and therefore f does not Pareto dominate g, let alone No-Betting-Pareto dominate it.

7.8.2 Example

In this example f Pareto dominates g according to the standard definition. That is, we will have $f \succ_0 g$, $f \succ^i g$ for $i = 1, 2$, but not $f \succ_{NBP} g$. For this example, define

		State s	State t
$f:$	Agent 1	600	-64
	Agent 2	-240	20

Clearly, agent 1, who believes that the state is s, prefers f to g, and this is also true of agent 2, who is sure that the state is t ($f \succ^i g$ for $i = 1, 2$). Also, if we average the utilities we get

	State s	State t
Average Agent	180	-22

so that f is better than g according to the average utility and the average belief $p_0 = (.5, .5)$.

Is it the case that $f \succ_{NBP} g$? We claim that the answer is negative. Indeed, by Theorem 1, $f \succ_{NBP} g$ would hold if and only if for every $\lambda \in [0, 1]$, the hypothetical agent with utility

$$u_\lambda = \lambda u_1 + (1 - \lambda)\, u_2$$

should be able to point to a state of the world where, for her, f is strictly better than g. But for $\lambda = 0.25$ we get

	State s	State t
$u_{0.25}$	-30	-1

and this agent cannot point to a state where she's better off with f than with g. Hence $f \succ_{NBP} g$ does not hold.

7.8.3 Example

Conversely, consider now an example where $f \succ_{NBP} g$ (and, in fact, $f \succ_i g$ for $i = 1, 2$), but where $f \succ_0 g$ does not hold. For this example, define

		State s	State t
f :	Agent 1	10	-100
	Agent 2	0	10

Clearly, agent 1, who thinks that the state is s, prefers f to g. For agent 2, f weakly dominates g, and she prefers f to g whenever she assigns a positive probability to state t. Because she assigns probability 1 to this state, she surely prefers f to g. Hence $f \succ_i g$ for $i = 1, 2$. Moreover, there exists a probability vector, say $(0.95, 0.05)$, for which both agents prefer f to g. Hence $f \succ_{NBP} g$ is established. However, when we consider the average utility we get

	State s	State t
Average Agent	5	-55

and for the average probability $(0.5, 0.5)$ act f results in a lower expected utility than does g. Hence $f \succ_0 g$ does not hold. ■

8 References

Arrow, K. J. and G. Debreu (1954), "Existence of an Equilibrium for a Competitive Economy", *Econometrica*, **22**: 265-290.

Aumann, R. J. (1976), "Agreeing to Disagree", *The Annals of Statistics*, **4**: 1236-1239.

Bewley, T. (1989), "Market Innovation and Entrepreneurship: A Knightian View", Discussion Paper 905, Cowles Foundation for Economic Research.

Billot, A., A. Chateauneuf, I. Gilboa and J.-M. Tallon (2000), "Sharing Beliefs: Between Agreeing and Disagreeing", *Econometrica*, **68**: 685-694.

Brunnermeier, M. K., A. Simsek, and W. Xiong (2012), "A Welfare Criterion for Models with Distorted Beliefs", working paper.

Gilboa, I., O. Lieberman, and D. Schmeidler (2009), "On the Definition of Objective Probabilities by Empirical Similarity", *Synthese,* **172**: 79-95.

Gilboa, I., A. Postlewaite, and D. Schmeidler (2008), "Probabilities in Economic Modeling", *Journal of Economic Perspectives*, **22**: 173-188.

——— (2009), "Is it Always Rational to Satisfy Savage's Axioms?", *Economics and Philosophy*, **25**: 285-296.

——— (2010), "Rationality of Belief", *Synthese*, forthcoming.

Gilboa, I., D. Samet, and D. Schmeidler (2004), "Utilitarian Aggregation of Beliefs and Tastes", *Journal of Political Economy*, **112**, 932-938.

Harsanyi, J. C. (1955), "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility", *Journal of Political Economy*, **63**, 309-321.

——— (1967-68), "Games of Incomplete Information Played by Bayesian Players", *Management Science*, **14**: 159-182, 320-334, 486-502.

Hylland A. and R. Zeckhauser (1979), "The Impossibility of Bayesian Group Decision Making with Separate Aggregation of Beliefs and Values", *Econometrica*, **47**, 1321-1336.

Kreps, D. M. (2012), *Microeconomic Foundations I: Choice and Competitive Markets*. Princeton: Princeton University Press.

Milgrom, P. and N. Stokey (1982), "Information, Trade, and Common Knowledge", *Journal of Economic Theory*, **26**: 17-27.

Mongin, P. (1995), "Consistent Bayesian Aggregation", *Journal of Economic Theory*, **66**, 313-351.

——— (1997), "Spurious Unanimity and the Pareto Principle", working paper, Cergy-Pontoise, France: Universite Cergy-Pontoise, Theorie Economique, Modelisation et Applications and Centre National de Recherche Scientifique, presented at the Conference on Utilitarianism, New Orleans.

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

Samet, D. (1998), "Common Priors and Separation of Convex Sets", *Games and Economic Behavior*, **24**, 172-174.

Savage, L. J. (1954), *The Foundations of Statistics*. New York: John Wiley and Sons. (Second addition in 1972, Dover)

Simsek, A. (2012), "Speculation and Risk Sharing with New Financial Assets", working paper.

Stiglitz, J. (1989), "Using Tax Policy to Curb Speculative Trading," *Journal of Financial Services Research*, **3**: 101-115.

von Neumann, J. and O. Morgenstern (1944), *Theory of Games and Economic Behavior*. Princeton, N.J.: Princeton University Press.

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