Pareto Efficiency with Different Beliefs*

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

December 7, 2013

Abstract

Pareto efficiency is not as compelling when people hold different beliefs as it is under common beliefs or certainty. Gilboa, Samuelson, and Schmeidler (2013) have suggested that the standard Pareto relation be weakened by imposing the additional constraint that, in order for one allocation to dominate another, there should exist a single hypothetical belief under which all agents prefer the former to the latter. In the present work we propose an alternative definition whereby Pareto efficiency is supplemented by the requirement that according to each agent's belief the former alternative is preferred to the latter for all other agents. This paper analyzes and compares these and other definitions.

*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. ISF Grant 396/10 (Gilboa and Schmeidler), ERC Grant 269754 (Gilboa), and NSF Grant SES-1153893 (Samuelson) are gratefully acknowledged.

[†]Department of Economics, Bar-Ilan University. gabi.gayer@gmail.com

[‡]Tel-Aviv University and HEC, Paris. 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

Pareto Efficiency with Different Beliefs

Contents

Pareto Efficiency with Different Beliefs:

1 Introduction

One of the most basic insights of economic theory is that constraints on trade tend to generate inefficiency. At an individual level, revealed preference theory holds that if Ann chooses one alternative over another, we can infer that Ann prefers the alternative she chooses. The attendant normative implication is that leaving Ann free to make her choices can only make her better off. Similarly, if Ann and Bob choose to trade with one another, we can infer that both are better off with the trade than without, and hence that we have no reason to interfere with such trades. At the aggregate level, this intuition is captured by the welfare theorems, stating that equilibrium allocations of competitive markets are Pareto efficient, coupled with examples in which constraints on market outcomes lead to inefficient allocations.

The welfare theorems rest on stringent assumptions. In the literature on "market failure" there is a long list of circumstances under which various of these assumptions fail, creating a role for market intervention. Markets may not be competitive, and we may therefore want an antitrust policy. There may be externalities, giving rise to a need for regulatory agencies. Public goods may be underprovided, motivating government subsidies. Information may be incomplete, giving rise to licensing or disclosure requirements. Even when Pareto efficiency is ensured, distributional considerations may prompt a desire to redistribute the fruits of the economy, even at an efficiency cost. At a more fundamental level, people may be irrational or confused or plagued by complexity constraints, opening a role for benevolent paternalism. There may be also circumstances, studied by a large literature in psychology, under which people do not know their tastes, questioning the very foundation of revealed preference and Pareto efficiency arguments.

We agree that the market imperfections mentioned in the previous paragraph are important, and we agree that perfect rationality is an elusive ideal. However, this paper argues that, even apart from market failures and even in the presence of perfectly rational agents, the logic of Pareto efficiency may not be compelling. In particular, we argue that Pareto efficiency is less compelling when considering trade that is motivated by differences in beliefs, as is often the case in financial markets.

There are several possibilities to modify the Pareto relation to address this concern. Gilboa, Samuelson, and Schmeidler (2013) have introduced the No-Betting Pareto relation. In this paper this notion of Gilboa, Samuelson, and Schmeidler (2013) is examined in addition to another notion of unanimity-based dominance. We close with a brief comparison to similar notions that have appeared in the literature, most notably those offered by Brunnermeier, Simsek, and Xiong (2012).

2 Tastes vs. Beliefs

2.1 Examples

Consider the following examples (taken from Gilboa, Samuelson, and Schmeidler (2013)).

Example 1. Alice and Bob have one apple (good 1) and one banana (good 2) each. Their utility functions are linear in quantities of the two goods, but their tastes differ. Alice is indifferent between one apple and two bananas, and hence has utility function

$$u_A(x_1, x_2) = 2x_1 + x_2,$$

while Bob is willing to trade two apples for a banana, and hence has utility function

$$u_B(x_1, x_2) = x_1 + 2x_2.$$

In the absence of a market, they can only consume their initial endowments. Trade allows them to reach a Pareto efficient allocation in which Alice and Bob are both better off than when consuming their endowment, with Alice consuming only apples and Bob consuming only bananas. There are other Pareto efficient allocations, but this one is the unique competitive equilibrium of the economy. □

Example 2. Ann and Bill each have a financial asset that will pay one dollar in one year's time. There are two states of the world. In state 1 the stock market will rise by more than ten percent over the next year, in state 2 it will not do so. Ann and Bill are both risk neutral, but they differ in their beliefs. Ann thinks that state 1 has probability 2/3 and Bill thinks state 1 has probability 1/3. Thus, their expected utilities are given by

$$\begin{aligned} u_A(x_1, x_2) &= \frac{2}{3}x_1 + \frac{1}{3}x_2 \\ u_B(x_1, x_2) &= \frac{1}{3}x_1 + \frac{2}{3}x_2. \end{aligned}$$

If there are no financial markets, Ann and Bill will consume their initial endowments, namely \$1 whatever happens to the market in the coming year. Trade allows them to reach a Pareto efficient allocation in which Ann and Bill are both better off than when consuming their endowment, in which Ann has no money in state 2 and Bill has no money in state 1. This is the unique competitive equilibrium of the economy. □

These two examples map to the same Arrow-Debreu (1954) general equilibrium model. There are two goods $\{1, 2\}$ and two agents $\{A, B\}$, with preferences represented by the utility functions

$$\begin{aligned} u_A(x_1, x_2) &= 2x_1 + x_2 \\ u_B(x_1, x_2) &= x_1 + 2x_2 \end{aligned}$$

and initial endowments

$$e_A = e_B = (1, 1).$$

In equilibrium, goods 1 and 2 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.

It appears as if the exchange is a win-win situation in each example, giving both agents an allocation that they prefer to their endowment. This type of example is the standard argument for the benefits of voluntary exchange in a variety of circumstances, from international trade to financial markets. However, we suggest that the formal similarity between the two examples is misleading. In Example 1 trade is driven by difference in tastes. An observer asked to explain the trade in this circumstance could do no better than to state *de gustibus non est disputandum*. We do not know why Alice's taste buds prefer apples whereas Bob's are more pleased by bananas, but we have no reason to question these tastes and every reason to think that trade should reasonably reflect such tastes. Importantly, we do not expect a discussion, however erudite, to change the tastes of either person, and we believe these differences in preferences can reasonably be taken as given.[1]

By contrast, trade in Example 2 is driven by differences in beliefs. There is surely room for a debate to change others' beliefs. Indeed, this may be the main goal of debate in general. In the example at hand, rather than let Ann and Bill trade, one may stop them and say, "You are basically betting against each other. If you like betting, maybe you should also bet on a toss of a fair coin. But if you do not enjoy betting for its own sake, and instead are risk neutral or even risk averse, it is impossible that both of you would strictly gain from this trade. One of you needs the probability of state 1 to be above 0.5 to gain from trade, and the other need that probability to be below 0.5. Obviously, it cannot be both!"

By way of analogy, suppose that agent A has a ring that she is willing

[1] If Ann preferred apples to bananas because she thought that the former were healthier than the latter, and Bob were of the opposite conviction, they would differ in beliefs, and then there would indeed be room for debate between them. However, we assume here that the differences are in primitive tastes only.

to sell for \$100. Agent B finds the ring extremely beautiful and is willing to pay \$100 for it. In this case, as in Example 1, trade is Pareto improving. Moreover, as trade is driven by differences in tastes alone, it poses no theoretical difficulties. Suppose instead that agent B doesn't like the ring per se, but that he believes that it is endowed with magical powers that would allow him to foresee the future, and so is willing to pay \$100 for it. Agent A, by contrast, thinks that the ring has no magical power whatsoever, and maybe also that foreseeing the future is theoretically impossible. A is willing to sell the ring, which she considers useless, for \$100, but she wouldn't engage in this trade had she believed that the ring actually could foresee the future. In this case, as in Example 2, trade is not driven by differences in tastes, but in differences in beliefs. And, assuming that these beliefs are well-defined, trade is only possible if one of the agents is wrong. In this case the trade of the ring would feel rather different than in the first case.

The ring story has a flavor of a hoax. As we doubt that there are rings that foresee the future, selling the ring to B appears deceitful. The stock market trade in Example 2 is a little different, as we cannot tell that one of the agents is obviously in the wrong. However, the two stories have a common feature: for trade to take place, the agents have to entertain contradicting beliefs. Our basic claim is that this type of Pareto-improving trade is more problematic than one based on differences in tastes alone.

2.2 Common Reasoning and Good Faith

There are at least two ways of formulating the difference between Examples 1 and 2, capturing different aspects of Example 2 and leading to different notions of Pareto domination. Suppose first that one assumes the paternalistic role of a social planner, or an outside observer who is asked whether she is willing to endorse trade. For the sake of the argument, assume that A and B are siblings and that the outside observer is their mother. In Example 1, Mom has no reason to raise any objections to their trading. In Example 2, by contrast, she may suspect that at least one of her children is in the wrong, and may not want one of her children to take advantage of the other's mistake. She may be unable to ascertain which child is wrong, but she can find out whether one of them *must* be wrong by asking the children to come up with a plausible, consistent story that would explain why they wish to trade. That is, she can ask the children not only to reach unanimity over the *conclusion* that trade is desirable, but also on the *reasoning* that leads to this conclusion. In the classical economic tradition of subjective expected utility maximization, such reasoning is represented by a probability vector. Mom may thus require that there be at least one probability vector that can simultaneously explain why both children wish to trade. Thus, the agents, irrespective of their actual beliefs, should be able to point to hypothetical beliefs that, if shared, would result in both of them having a higher expected utility after trade than before. This is the essence of the No-Betting Pareto domination suggested in Gilboa, Samuelson, and Schmeidler (2013).

Suppose instead that Ann and Bill are unrelated, and that Ann consults her father about the proposed trade. Her father asks her about her payoff and about Bill's, and then about their beliefs, and finally concludes, "Ann, I'm not sure that this trade is done in good faith. You think that you know what you're doing, and I tend to trust you, but that means that your partner is wrong. Or, at the very least, you believe that he's wrong, and you take advantage of his mistake. After all, if your beliefs are the right ones, you should have told Bill to stay away from the deal."

This line of reasoning is different from the previous one. Here Ann's father, talking only to her, expects her to rise to a rather high moral standard, and only engage in trades that are, to the best of her judgment, beneficial to all involved. This standard suggests the following refinement of Pareto domination: when all agents wish to trade, one should also verify that all agents would be better off with trade than without according to each agent's belief.

2.3 Sharing Risks

This section presents two examples that clarify these concepts. The first is an example of trade with different beliefs, which is endorsed by both concepts. This example (also taken from Gilboa, Samuelson, and Schmeidler (2013)) shows that restricting trade to these refined Pareto improvements still allows agents to share risks. The subsequent example highlights the differences between the two notions.

Example 3. Agnes is an entrepreneur with an idea for a start-up company, involving an initial investment of \$1 million. If successful, the company will be worth \$10 million. If unsuccessful, the company will net \$0. Agnes assigns probability 90% to success and has the initial investment (but no more), but she is still unwilling to fund the project herself and hence become destitute with probability of 10%. Agnes approaches Barry, who runs a venture capital fund, asking him to provide the initial \$1 million in return for half of the resulting profit. Barry believes the probability of success to be only .3. However, he is risk neutral, and even at this more pessimistic probability, Barry finds the venture worthwhile. Agnes and Barry enter into a funding arrangement that allows them to share the risk. □

As in Example 2, Example 3 involves uncertainty and a disagreement between the agents' beliefs. However, in this case the disagreement in beliefs is not crucial for trade to take place. As far as Agnes is concerned, signing the deal is a weakly dominant strategy. Relative to her current wealth, the deal offers her additional \$0 in one state of the world and additional \$4.5 million in the other. She is strictly better off with the deal (than without) for any belief that attaches a positive probability to success. Barry, by contrast, will lose \$1 million in one state of the world and will gain \$4.5 million in another. Being risk neutral, he will have a higher expected utility for the deal (relative to no-deal) as long as the probability of success is higher than $\frac{1}{5.5}$ (roughly 18%).

The deal in Example 3 is endorsed by both criteria: it is No-Betting Pareto improvement (over no-deal) because there exists a belief—say, that success has probability 50%—that makes the deal's expected utility higher (than no-deal) for both agents. Moreover, the deal also Unanimity Pareto dominates no-deal. Agnes is better off with the deal for any probability of success $p > 0$, which includes both her beliefs ($p = 0.9$) and Barry's ($p = 0.3$). And Barry will be better off as long as $p > \frac{1}{5.5}$, which again includes both beliefs.

The point of this example is that risk sharing is not precluded by either of our refinements of Pareto domination. In this example, Agnes may be viewed as holding an asset that is worth -1 (millions of dollars) in one state, and $+9$ in another, giving her a risky position (should she embark on her project). Barry holds a riskless asset. For some range of beliefs, it will be in Barry's interest to take the risk off Agnes's shoulders in return for some of the potential profits. For another set of beliefs, it will be in Agnes's interest to give up these profits for the insurance she gets from Barry. Since both their beliefs are in these ranges, the risk-sharing deal is endorsed by both refinements of the Pareto criterion.

It is clearly the case that Unanimity Pareto domination implies No-Betting Pareto domination. The latter criterion requires that *there be a single* probability that justifies trade for all agents, whereas the former requires that *all probabilities* of the agents (hence, all probabilities in their convex hull) satisfy this criterion. The following example shows that No-Betting Pareto domination does not imply Unanimity Pareto domination.

Example 4. Ada is also an entrepreneur with an idea for a start-up company, involving an initial investment of \$1 million and a very considerable time investment. For the company to succeed, Ada will have to work some 16 hours a day for three years. If successful, the company will be worth \$10 million. If unsuccessful, the company will net \$0. Ada assigns probability

90% to success. She has no funds to invest, but she is willing to work hard for half of the profits. She approaches Babbage, who—like Barry above—runs a venture capital fund, asking him to provide the initial $1 million in return for half of the resulting profit. Babbage is risk neutral. As Barry in Example 3, Babbage believes the probability of success to be only 30%. However, even at this more pessimistic probability, Babbage finds the venture worthwhile. Ada and Babbage enter into a funding arrangement that allows them to share the risk. □

In this example trade Pareto dominates no-trade according to the standard definition, as well as according to the No-Betting Pareto definition: there exist beliefs—say, Ada's optimistic beliefs that success has probability of .9—according to which both agents would be willing to trade. Indeed, if Babbage is willing to trade at probability of success of .3, he would certainly be willing to do so at probability .9 (whereas Ada is known to be willing to trade given her own beliefs). But trade need not dominate no-trade according to the more stringent, unanimity principle. Instead, one may well imagine Babbage saying, "I'm certainly glad that there are people like Ada, willing to work night and day to make their dreams come true, and I surely benefit from their existence in the market. But, were Ada my daughter, I'd tell her she's making a mistake. Should Ada embark on this project, she won't be seeing her children for three years. I believe the probability of success is 30%, and that this does not provide a sufficiently high expected return to justify the cost. I wouldn't do it if I were in Ada's place."

Example 4 thus highlights the difference between the two criteria. The No-Betting criterion only requires that an outside observer not be able to rule out the trade as sheer betting, while the Unanimity criterion is more demanding, requiring that each agent involved think that each other agent is not mistaken in participating in the deal.

3 No-Betting Pareto and Unanimity Pareto

3.1 Definitions

This section presents the two Pareto conditions.[2] There is a set of agents $N = \{1, ..., n\}$ and a state space $S = \{1, ..., s\}$. A *social outcome*, $x = (x_1, ..., x_n) \in \mathbb{R}^n$, specifies a wealth level, x_i, for each agent i.

The alternatives compared are functions from states to social outcomes, which can be viewed as matrices, specifying a level of wealth for each individual at each state. Formally, the set of alternatives is

$$F \subset \{ f : S \to \mathbb{R}^n \}$$

where $f_i(j)$ is i's wealth at state j.

Each agent i has a preference order $\succsim_i$ over F, depending (only) on her own wealth. The agents are expected utility maximizers. Specifically, each agent i has (1) a utility function $u_i : \mathbb{R} \to \mathbb{R}$, which is differentiable, strictly monotone and (weakly) concave, and (2) a probability vector p_i on j such that she maximizes $\sum_{j \in S} p_i(j)\, u_i(f_i(j))$.

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

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

A "trade" is pair of acts $(f, g) \in F^2$ in which the agents give up act g in return for f. Such a trade will in general involve some individuals but not others. Agent $i \in N$ is said to be *involved* in the trade (f, g) if $f_i(\cdot) \neq g_i(\cdot)$—that is, if there exists at least one state j at which agent i's wealth is different for f than for g. Let $N(f, g) \subset N$ denote the agents who are involved in the trade (f, g). Observe that the definition of $N(f, g)$ does not depend on the agents' beliefs, $(p_i)_i$.

[2] Gilboa, Samuelson and Schmeidler (2013) present the No-Betting-Pareto criterion in a more general set-up allowing for abstract outcomes and an infinite state space.

Definition 2 *A trade* (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 trade g for f voluntarily—no confusion or coercion is required. 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 $f_i(\cdot) \neq g_i(\cdot)$, to be indifferent between f and g. However, having made an explicit distinction between the agents who are involved in the improvement and those who aren't, we find that strict preference for the former appears to be a natural condition. Moreover, given any trade (f, g) with the property that $f \succeq_i g$ for all $i \in N(f, g)$ but without all such preferences being strict, we can find an act f' arbitrarily close to f with $N(f', g) = N(f, g)$ and with $f' \succ_i g$ for all $i \in N(f', g)$. Defining an improvement to involve strict preference for every involved agent thus imposes no essential constraints.

The two weaker notions of domination are 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* $i \in N(f, g)$*,*

$$\sum_{j \in S} p_0(j)\, u_i(f_i(j)) > \sum_{j \in S} p_0(j)\, u_i(g_i(j)).$$

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 trading g for f.

Definition 4 *For two alternatives $f, g \in F$, we say that f* Unanimity Pareto dominates *g, denoted $f \succ_U g$, if:*

(i) f improves upon g;

(ii) For every probability measure p_k, $k = 1, ..., n$, for all $i \in N(f, g)$,

$$\sum_{j \in S} p_k(j)\, u_i(f_i(k)) > \sum_{j \in S} p_k(j)\, u_i(g_i(j)).$$

Condition (ii) requires each agent to believe that f is an improvement over g for all other agents. This is equivalent to the requirement that for each agent, f be an improvement over g according to *all* other agents' beliefs.

3.2 Illustration

In this section we adapt a simple example from Gilboa, Samuelson, and Schmeidler (2013), that allows us to visualize the No-Betting Pareto condition in a 2x2 Edgeworth box. We then proceed to illustrate the new, Unanimity condition in this context.

Suppose that there two agents, A and B, and two states, 1 and 2. The aggregate endowment is the same in the two states, so there is no aggregate uncertainty. The agents have identical, strictly concave utility functions. The alternatives are given as points in an Edgeworth box. The diagonal in Figure 1 is the set of full-insurance allocations. Along this diagonal the slopes of the indifference curves of a given agent are identical, and are determined solely by the probabilities the agent attaches to the two states.

If the agents share identical beliefs, then the diagonal is the set of Pareto optimal allocations. In this case efficiency calls for the agents to fully insure one another, so that neither agent bears any risk. Given this coincidence of beliefs, this is also the set of allocations that are No-Betting Pareto efficient and Unanimity Pareto efficient.

Suppose A thinks state 2 is more likely than does B. Then the set of Pareto efficient allocations will no longer be the diagonal, but instead will be given by a contract curve that lies above the diagonal. Because agent A

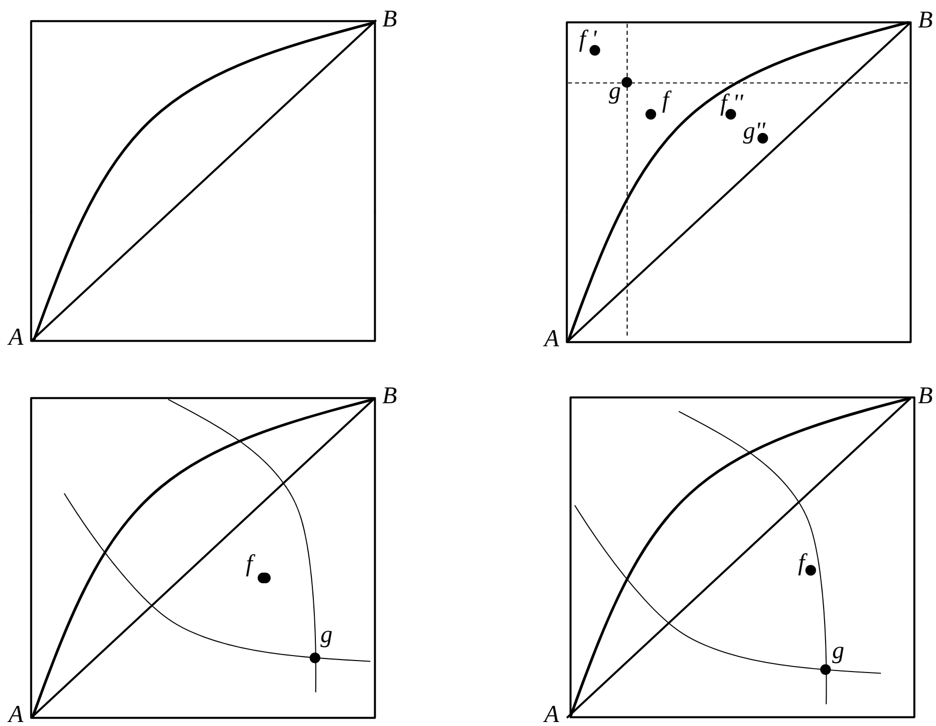

Figure 1: The horizontal axis in each panel identifies the allocation of money between agents A and B in state 1, with the vertical axis doing the same for state 2. The total endowment is constant across states. The diagonal is the set of full insurance allocations. Agent A views state 2 as being more likely that does agent B, and hence the set of Pareto efficient allocations is given by the contract curve that lies above the diagonal. In the top right panel, f Pareto and No-Betting Pareto dominates g, while f' dominates g in neither sense. Allocation f'' Pareto dominates g'' but does not No-Betting Pareto dominate g''. In each of the bottom panels, f Pareto dominates g and No-Betting-Pareto dominates g. Allocation f Unanimity Pareto dominates g in the left panel, but not the right panel.

thinks state 2 is more likely than does B, both can gain by trading away from a fully-insured allocation on the diagonal to an allocation in which A consumes more in state 2 than in state 1, while B consumes more in state 1 than in state 2. This causes both agents to bear some risk, and the set of Pareto efficient allocations balances this risk against the ability of trade to exploit the agents' differing beliefs. The top left panel of Figure 1 illustrates the contract curve.

First, consider No-Betting-Pareto. Consider the allocation g in the top right panel of Figure 1. The allocation f' lies further from the contract curve than does g, and hence f' cannot improve on g, ensuring that f' neither Pareto nor No-Betting Pareto dominates g. The allocation f lies closer to the contract curve than g, and so it is possible (but not necessary) that f is an improvement. Suppose that it is an improvement. Gilboa, Samuelson, and Schmeidler (2013) show that for an allocation g above the diagonal, another allocation f above the diagonal satisfies the second condition for No-Betting-Pareto dominance if and only if f lies to the southeast of g. Hence, in the top right panel, f No-Betting-Pareto dominates g. Indeed, for allocations above the contract curve, Pareto and No-Betting Pareto domination agree. In contrast, consider the points g'' and f'', assuming that the latter is a Pareto improvement over the former. The allocation f'' does not lie closer to the diagonal and hence does not satisfy the second No-Betting-Pareto criterion, and here the Pareto and No-Betting-Pareto criteria disagree. The set of allocations that are undominated under the No-Betting Pareto criterion consists of those contained in the flattened lens bounded by the diagonal and the contract curve.

Consider now the Unanimity Pareto condition. First consider allocations f and g that both lie on the same side of the diagonal as does the contract curve, as in the top right panel of Figure 1. In this case f No-Betting Pareto dominates g if and only if f Unanimity dominates g. Clearly, the "if" part follows from the general definitions. We should therefore convince ourselves that No-Betting Pareto domination implies Unanimity Pareto domination. By the previous analysis, if f No-Betting Pareto dominates g, then f is closer to the diagonal than is g, so that, by switching from g to f, agent A gives up income in state 2 (which she finds relatively more likely) for income in state 1 (which she finds less likely relative to agent B). If agent A is willing to make this swap given her own beliefs, she would definitely be willing to make it given agent B's beliefs (assigning a higher probability to state 1, where income increases by the swap, and lower to state 2, where the swap decreases income). Similarly, agent B, who's willing to swap g for f according to her own beliefs, will be surely willing to do so given agent A's beliefs.

The No-Betting Pareto and Unanimity Pareto criteria do not coincide if f and g lie on the side of the diagonal opposite the contract curve. In this case one can easily see that the set of allocations f that Unanimity Pareto dominate a given allocation g (and are, therefore, closer to the diagonal than is g) may be a strict subset of the allocations that No-Betting Pareto dominate g. The two bottom panels of Figure 1 both show cases in which allocation f Pareto dominates g and No-Betting-Pareto dominates g. The right panel shows a case in which agent B realizes virtually no gains, under her own beliefs, when switching from g to f. This switch is even less attractive when evaluated according to agent A's beliefs, and for specifications of f sufficiently close to B's indifference curve, will surely be disadvantageous under A's beliefs. In this case f will not Unanimity Pareto dominate g. In contrast, there are specifications of f, such as that illustrated in the lower left panel, that Pareto, No-Betting Pareto, and Unanimity Pareto dominate g.

These comparisons reflect the following principles. When f and g are on the same side of the diagonal as the contract curve, the trade can be justified by either No-Betting or Unanimity Pareto only if it moves the agents closer to the diagonal, so that each agent gives up income in a state that she finds relatively more likely. Thus, their only motive for trade is to insure each other against the risk they are exposed to. By contrast, when the agents are on the other side of the certainty line (than is the contract curve) the swap towards the diagonal makes each of them richer in the state that she finds relatively more likely. While this type of trade still has an insurance aspect—attested to by the move toward the diagonal—it may also hinge on disagreement, as each agent thinks that she gains by trade more than the other agent thinks she does. It is under these circumstances that the Unanimity criterion proves to be more restrictive than the No-Betting criterion. The latter only requires that both parties reduce the risk that they are exposed to, whereas the former also demands that they do it in a way each finds beneficial when holding the other's beliefs.

4 Implementation

What could the notions of No-Betting-Pareto or Unanimity Pareto be used for, and what issues arise in their use? We first note a key implication of both criteria, then consider some issues that might arise in their implementation.

4.1 Betting

We say that a trade from g to f is a *bet* if g is a full-insurance allocation, i.e., if $g_i(j) = g_i(j')$ for every agent $i \in N$ and states $j, j' \in S$. Gilboa, Samuelson and Schmeidler (2013) prove the following:

Proposition 1 *If (f, g) is a bet, then it cannot be that f No-Betting Pareto dominates g.*

It is obviously then also the case that Unanimity Pareto precludes betting. The No-Betting Pareto criterion also eliminates some Pareto rankings that are not bets. Gilboa, Samuelson and Schmeidler (2013) characterize the Pareto rankings that are precluded by the No-Betting Pareto criterion, showing that these rankings can be viewed as "generalized bets."

Bets push agents away from risk sharing, and can be mutually advantageous for risk-averse agents only if they hold incompatible beliefs. The potential usefulness of our refinements of the Pareto criterion thus arises out of the desire and ability to separate bets, which rely crucially on differing beliefs, from other trades.

4.2 Observability

Suppose we have a collection of agents who are willing trade from g to f, so that we know that f is a Pareto improvement over g. What else do we require to conclude that f No-Betting Pareto or Unanimity Pareto dominates g?

To verify that an alternative f Unanimity Pareto dominates an alternative g, one needs to know the agents' utilities and probabilities. As for No-Betting Pareto, only the utility functions need be known. The No-Betting criterion may thus be easier to verify, from an informational perspective, than the Unanimity one.

4.3 Transitivity

A binary relation that captures some notion of "dominance" is naturally expected to be transitive. Indeed, if the trade from g to f satisfies some version of the Pareto dominance relation, as does the trade from f to h, then one would like to know that h is "better" than g according to the same criterion.

This is clearly the case with Unanimity Pareto. If f Unanimity Pareto dominates g, then all individuals must prefer f to g according to each and every probability (entertained by any of them). The same is true for h and f and hence for h and g, establishing that the relation is transitive.

By contrast, this is not the case when No-Betting Pareto domination is concerned:

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

Hence, two consecutive Pareto improvements, each of which can be justified by shared beliefs, 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, given by

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

is *rectangular* if the following condition holds: for every collection of acts $\left\{f^k\right\}_{k=1}^{n} \in F^n$, there exists $f^* \in F$ such that, for all $i \in N$, $u_i\,(f_i^*\,(j)) = u_i\,(f_i^i\,(j))$ for every state j. 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* $\succ_{NBP}^{t} = \succ_{*}$.

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 trades, it does not do so very effectively: any trade that the agents eventually wish to perform ($f \succ_{*} g$) can be carried out by a sequence of trades, each of which qualifies as a No-Betting-Pareto improvement.

However, rectangularity is a very strong condition. Let us say that the trade (f, g) is *feasible* if for every state j, $\sum_i f_i(j) = \sum_i g_i(j)$, so that the aggregate wealth in each state is preserved. If we limit attention to $\{\, u\,(f) \mid (f, g)$ is a feasible improvement$\}$, then rectangularity does not hold and neither does the conclusion of Proposition 3. More generally, let us 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:

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.

4.4 Computation

An obvious advantage of the standard Pareto condition is that it requires no computation on the part of a central planner. As long as voluntary trade is concerned, each agent may be viewed as if she verified that trade is beneficial to her and hence that the trade is a Pareto improvement. The task of verifying that trade satisfies the Pareto criterion has effectively been decentralized among the agents. But when we start tinkering with the standard definition, one might ask, is it at all possible to verify whether our new definitions hold?

Let us assume here that utilities and probabilities of all agents are known, and focus on the computational task involved. We observe that under our assumptions, Unanimity Pareto domination is easy to verify: given alternatives f and g, one has only to verify n^2 inequalities (matching each of the n utility functions with each of the n probability vectors).

No-Betting Pareto domination requires a bit more calculation. Condition (i) in the definition can again be effectively decentralized. Condition (ii) requires that there be a probability vector that justifies trade for all n utility functions, but it does not give us a clue as to which probability vectors should be tested. However, the following result establishes that verifying Condition (ii) is an "easy" task.

Proposition 5 *Given rational numbers,* $(u_i(f_i(j)))_{i,j}, (u_i(g_i(j)))_{i,j}$ *it can be checked in polynomial time complexity whether Condition (ii) for the No-Betting Pareto criterion holds.*

Thus, an imaginary scenario in which a market maker needs to approve trades 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 $f \succ_* g$, while Proposition 5 ensures that verifying condition (ii) is a simple computational task.

4.5 Meddling in Markets?

How could these notions be applied? Our inclination is to be conservative. If we are to consider an alternative to the Pareto criterion, we would prefer it to have two characteristics. First, the alternative criterion is a subset of the Pareto criterion, in the sense that it may decline to rank some trades that the Pareto criterion ranks, but ranks no previously unranked trades. Second, the alternative criterion excludes a ranking only if we are *certain* there is no way to justify the comparison. We accordingly consider the more permissive No-Betting-Pareto criterion. No-Betting Pareto and Unanimity Pareto both satisfy the first criterion. No-Betting Pareto eliminates fewer rankings than does Unanimity Pareto, and hence No-Betting Pareto is accordingly less aggressive in terms of the second criterion. Coupling this observation with the fact that less information is required to evaluate the No-Betting Pareto relation, we concentrate here on No-Betting Pareto.

We imagine a scenario in which a monitoring authority must approve proposed mergers, acquisitions, or financial deals, either in advance or as part of occasional audits. We can further imagine that in order to approve a proposal, the monitoring authority requires not only that each party indicate they are willing to participate in the proposal, ensuring the proposal is a Pareto improvement, but also that the parties can present a model identifying the states of the world, their endowments, the net trade, and a single belief under which no party loses from the trade.

The monitoring authority need not know the agents' beliefs, but ascertaining that "no party loses" from the trade requires knowing something about their utility functions. We assume that the utility functions used in this calculation would be drawn from a fixed set of standard utility functions commonly used to study financial markets. For example, the agents may be required to evaluate their payoffs according to CARA or CRRA utility functions, with the risk-aversion parameters drawn from a fixed range of "appropriate" parameters. One of the functions of the regulator would be to formulate the guidelines specifying the relevant functions and parameter ranges. In this sense, the regulator would be acting much as does an antitrust authority when identifying the guidelines under which mergers are to be evaluated, or a financial regulator identifying the guidelines under which a bank's compliance with reserve requirements are to be evaluated. The range of acceptable risk aversion parameters may be tailored to the type of agent engaged in the proposal. For example, a public employees pension fund may exhibit (or be required to exhibit) a greater degree of risk aversion than an investment bank, which may in turn exhibit greater risk aversion than a hedge fund. Of course one cannot hope to precisely capture preferences with selections from a simple menu, just as one cannot hope that a single set of merger guidelines or capital requirements will fit all circumstances. However, it is not obvious what it means for a pension fund or hedge fund to have preferences, even though one can talk about the appropriate degree of risk for such a fund, and hence the approach described here may be as close as one can come to making meaningful statements. Similarly, the state space will typically fall short of resolving all uncertainty, which may again be an unobtainable ideal. Instead, one might expect the parties to work with rather course state spaces, perhaps in some cases including only two states, perhaps of the form "the stock market increases by at least ten percent over the next year" and "the market does not do so." The requirement would then be that the parties to the transaction can identify a single belief that rationalizes the transaction for all of them. There would be no need for debate about what the parties actually believe, with the discussion concerning only whether such a belief exists.

We believe the case for restricting trade, and especially for restricting trade that is equivalent to betting, is particularly acute when dealing with institutions that manage other people's money. Retirement funds, for example, should be very careful in making financial decisions that may result in their members losing their future incomes. We would find it troublesome if two such funds were to engage in trade that is not No-Betting-Pareto improving, or if a single fund were to invest in assets with the property that the pension fund and the manager of the assets can both hope to gain only because they hold different beliefs. We view the concept of No-Betting-Pareto as a first step in thinking about how to make trade in financial markets more responsible than allowed by the standard general equilibrium model. A second step may be the Unanimity Pareto criterion requiring that no party believes that it is exploiting the other parties.

5 Discussion

5.1 Foundations

Economists have constructed an elegant and elaborate theory of welfare economics on the foundation of Pareto efficiency. Given our indication that we sometimes find the notion of Pareto efficiency unconvincing, it is important to be clear about the foundations of our view of welfare economics.

First, we interpret beliefs and utilities differently. This difference is reflected in the fact that we are willing to view unanimously acceptable trades driven by differences in utilities as obviously improving welfare, while being suspicious of trades based on differences in belief. It is not obvious that we should make such a distinction, and indeed not obvious that we should be interpreting utilities and beliefs at all. A possible interpretation of Savage's (1954) representation theorem is that the utilities and probabilities appearing in decision theory have no meaning. They are simply mathematical constructs that allow us to parsimoniously describe behavior (as long as certain axioms are satisfied), with no connection to intentional concepts or mental phenomena. By contrast, we believe that, especially for normative purposes, beliefs and utilities are meaningful concepts, going beyond mere mathematical objects used in a representation of preferences. In particular, normative analysis seems to rely on attaching to the utility function some meaning having to do with welfare or desirability, a meaning that cannot be attached to beliefs. Moreover, while our analysis focuses on ex-ante Pareto domination, embedding our discussion in a dynamic set-up would further clarify the distinction between the two concepts: the utility function is used also ex-post, and it is still a function that we would like to increase for some agents, if this can be done at no utility cost to others. By contrast, the probability vector is either updated or becomes meaningless after information has been revealed. Thus, both the nature of normative exercises and the dynamic unfolding of information suggest that utilities and probabilities should be treated differently. Moreover, people routinely talk about beliefs, and attempt to convince others as to what their beliefs should be, but rarely do the same for utilities. People will readily argue about the likelihood that eating a certain food will have various health effects, but will not argue that someone else "should" like carrots better than broccoli.

One obvious way for probabilities to have meaning is for probabilities to be objective, with differing beliefs then reflecting the fact that people have not yet discovered the underlying "true" probability. While we believe that probabilities have meaning, we do not believe that probabilities can always be usefully viewed as objective. If they were objective, the appropriate Pareto criterion under uncertainty would be straightforward: find the true probability (and perhaps reveal it to the agents) and use it to evaluate any allocation that comes along. Unfortunately, we do not know how to calculate the probability of, say, the stock market rising by more than ten percent over the course of the next year in a way that would qualify as objective.

The absence of objective probabilities poses no conceptual problems if agents agree on the relevant probabilities. In particular, if all agents shared the same beliefs, the No-Betting-Pareto, the Unanimity Pareto, and conventional Pareto criteria would all coincide. However, the "Harsanyi doctrine" (Harsanyi, 1967-68) that all agents can be taken to have the same prior beliefs strikes us as unrealistic. Aumann's (1976) agreeing-to-disagree result shows that rational agents sharing a common prior (in a model that is common knowledge among them) cannot hold diverging posterior beliefs, and the no-trade theorem (Milgrom and Stokey, 1982) identifies circumstances under which such agents cannot trade in financial markets, even as a result of the arrival of new information. We take the prevalence of different beliefs and the large volumes of trade in financial markets as convincing evidence that the common prior assumption does not hold.[3]

To conclude, while we recognize that beliefs differ, we do not view probabilities subjective in the sense that tastes are. It is meaningful for people to debate probabilities and attempt to convince one another about reasonable levels of such probabilities. One cannot attach an objective probability to the event that North Korea will launch a nuclear attack on a neighbor in the next two years, but one could bring evidence and arguments to bear on this question that would generally prompt similar revisions in peoples' probabilities, indicating that the latter are not purely subjective. In a similar vein, we believe that a financial intermediary can reasonably make a case that a particular asset is valuable based on the likelihood that various events will occur, beginning with something like "let me show you why this stock

[3]An alternative explanation is that people do have common priors, but rationality is not common knowledge. Pareto efficiency is immediately problematic in the absence of rationality; our interest lies in studying Pareto efficiency among rational but disagreeing agents.

is sure to go up". We would be troubled by a financial intermediary who could balance his books only by having the agents on the opposite ends of the trades he brokers hold quite different beliefs, making the "let me show you why this stock is sure to go up" case on one end of the trade and "let me explain why this stock is sure to fall" on the other end.

5.2 Related Literature

Others have argued that the concept of Pareto domination is troubling when beliefs differ. The difference between gains from trade based on tastes (as in Example 1) and differences in beliefs (as in Example 2) was discussed by Stiglitz (1989), in the context of an argument that inefficiencies arising from the taxation of financial trades might not be too troubling. Mongin (1997) referred to the type Pareto domination appearing in Example 2 as *spurious unanimity.*

More recently, concerns about the notion of Pareto domination have been raised in the context of trade in financial markets. Weyl (2007) points out that arbitrage might be harmful when agents are "confused". Posner and Weyl (2012) call for a regulatory authority, analogous 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 (2012).

The paper most closely related to ours, with a quite similar motivation and somewhat different details, is Brunnermeier, Simsek, and Xiong (2012).[4] Brunnermeier, Simsek, and Xiong offer two approaches to making Pareto comparisons under uncertainty. The first, the *expected social welfare criterion*, applies a unanimity criterion on beliefs to a weighted sum of agents' utilities. As is typically the case with criteria based on social welfare functions, this criterion will rank some alternatives that are not ranked by the standard Pareto criterion. We recognize that such value judgments must of-

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

ten be made, but cling to our belief that the point of departure for a criterion motivating interventions in financial markets should be a refinement of the Pareto criterion.

Brunnermeier, Simsek, and Xiong offer another approach that is independent of a specification of the welfare weights and that is a more obvious comparison to our notions. They define f to be *belief-neutral inefficient* if for every belief p in the convex hull of the agent's beliefs, there is an alternative g with the property that every agent prefers g to f given belief p (with weak preference assumed for all agents and strict for at least one of them). Notice that the alternative g is allowed to depend on the belief p. An alternative f is *belief-neutral efficient* if there is no alternative g and belief p in the convex hull of the agents' beliefs with the property that every agent prefers g to f, given belief p. Clearly, if f is belief-neutral inefficient, it cannot be belief-neutral efficient. However, there may well be allocations f that are neither belief-neutral inefficient nor belief-neutral efficient.

Brunnermeier, Simsek, and Xiong's emphasis on comparisons that hold for all beliefs in the convex hull of the agents' beliefs makes it most directly comparable to our Unanimity Pareto notion.

- - If f is Unanimity Pareto inefficient, then it is belief-neutral inefficient.
  - (The contrapositive of the above:) If f is belief-neutral efficient, or unclassified (that is, not belief-neutral inefficient), then it is Unanimity Pareto efficient.
  - if f is Unanimity Pareto efficient, nothing is implied about its belief-neutral efficiency: it may be belief neutral efficient, belief neutral inefficient, or may be unclassified by the belief-neutral Pareto criterion.

The belief-neutral Pareto criterion thus differs from Unanimity Pareto by transferring some efficient outcomes into the inefficient category, or not classifying them as either efficient or inefficient.

6 Appendix: Proofs

6.1 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 $S = \{j, t\}$. 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 j	State t
g :	Agent 1	0	0
	Agent 2	0	0

		State j	State t
h :	Agent 1	2	-1
	Agent 2	-3	2

		State j	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. ■

6.2 Proof of Proposition 3

Assume that $range(u)$ is rectangular and convex. We need to show that $\succ_{NBP}^{t}=\succ_{*}$. The inclusion $\succ_{NBP}^{t}\subset\succ_{*}$ 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 $\succ_{*}$, $\succ_{NBP}^{t}\subset\succ_{*}$.

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 $u_i\left(g\left(\cdot\right)\right)$ under h_l for $l<i$ and $u_i\left(f\left(\cdot\right)\right)$ under h_l for $l\geq i$. This will be possible thanks to the fact that $range\left(u\right)$ 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 $range\left(u\right)$, to guarantee strict preferences for all agents throughout the sequence.

Let there be given an improvement (f,g). Without loss of generality assume that $N=N\left(f,g\right)$. For $i,k\in N$, let $u_k\left(h_i'\right)$ as explained above:

$$u_k\left(h_i'\left(j\right)\right)=\begin{cases}u_k\left(g\left(j\right)\right) & i<k\\ u_k\left(f\left(j\right)\right) & i\geq k\end{cases}\qquad\forall j\in S.$$

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

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

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

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

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

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

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

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

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 $j \leq s$,

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

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 $k = i$, (1) 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 $h_i \succ_{NBP} h_{i-1}$. ■

6.3 Proof of Proposition 4

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

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

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

6.4 Proof of Proposition 5

Given the rational numbers $\left(u_i\left(f\left(j\right)\right), u_i\left(g(j)\right)\right)_{i,j}$ we need to check if there exists a probability vector $p_0 \in \Delta^{s-1}$ such that, for all $i \in N\left(f, g\right)$,

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

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

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

The optimal value of this problem is positive if and only if there exists a probability vector $p_0 \in \Delta^{s-1}$ such that $\sum_{j \le s} p_0(j)\left[u_i(f(j)) - u_i(g(j))\right] > 0$ for every $i \in N(f,g)$, which is easy to verify because linear programming can be solved in polynomial time complexity. ■

7 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.

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

Gilboa, I., L. Samuelson, and D. Schmeidler (2013), "No-Betting Pareto Domination", working paper.

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.

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.

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