Pareto Efficiency with Different Beliefs

Gabrielle Gayer, Itzhak Gilboa, Larry Samuelson, and David Schmeidler

ABSTRACT

Pareto efficiency is not as compelling when people hold different beliefs as it is under common beliefs or certainty. In the present paper we propose to restrict the standard Pareto relation by imposing the following constraint: in order for one allocation to dominate another, all agents must prefer the former to the latter according to each agent's belief. In contrast to this unanimity Pareto criterion, the no-betting Pareto concept suggested elsewhere by Gilboa, Samuelson, and Schmeidler supplements the standard Pareto relation with the requirement that there should exist a single hypothetical belief under which all agents prefer the former to the latter. This paper analyzes and compares these and other definitions.

1. INTRODUCTION

One of the fundamental insights of economic theory is that voluntary trade is good for all. At an individual level, revealed-preference theory holds that if Ann chooses one alternative over another, then we can infer that Ann prefers the alternative she chooses. Similarly, if Ann and Bob choose to trade with one another, we can infer that both are better off with the trade than without. At the aggregate level, this intuition is captured by the welfare theorems, establishing a precise connection be-

GABRIELLE GAYER is Lecturer of Economics at Bar-Ilan University. ITZHAK GILBOA is Chair for Economic Theory and Decision Theory at Tel-Aviv University. LARRY SAMUELSON is A. Douglas Melamed Professor of Economics at Yale University. DAVID SCHMEIDLER is Professor of Economics at the Interdisciplinary Center, Herzliya, and Professor Emeritus at Tel-Aviv University. 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. Israel Science Foundation grant 396/10 (Gilboa and Schmeidler), European Research Council grant 269754 (Gilboa), and National Science Foundation grant SES-1153893 (Samuelson) are gratefully acknowledged.

[*Journal of Legal Studies*, vol. 43 (June 2014)]
 0047-2530/2014/4302-0019$10.00

tween the equilibrium allocations of competitive markets and Pareto-efficient allocations. The normative implications are that we should prefer a Pareto-dominating allocation to a Pareto-inferior one and toward that end should leave people free to make their choices and leave markets free to work their magic.

The welfare theorems rest on stringent assumptions. The literature on market failure contains a long list of circumstances under which various of these assumptions fail, which creates a role for market intervention. Markets may not be competitive, and we may therefore want an antitrust policy. There may be externalities, which gives rise to a need for regulatory agencies. Public goods may be underprovided, which motivates government subsidies or provision. Information may be incomplete, which gives rise to licensing or disclosure requirements. At a more basic level, people may be irrational or confused or plagued by complexity constraints, which opens 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, which brings into question the very foundation of revealed-preference and Pareto-efficiency arguments.

We agree that market imperfections 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 for modifying the Pareto relation to address this concern. Gilboa, Samuelson, and Schmeidler (2014) introduce the no-betting Pareto relation. This paper reexamines the no-betting Pareto relation as well as 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 VERSUS BELIEFS

2.1. Examples

Example 1. We begin with an illustration of the classic link between voluntary trade and Pareto improvements. There are two agents, Alice and Bob, each endowed with $\frac{1}{2}$ unit of each of two goods. They have Cobb-Douglas utility functions, given by

$$u_A(x_1, x_2) = \frac{2}{3}\ln x_1 + \frac{1}{3}\ln x_2$$

and

$$u_B(x_1, x_2) = \frac{1}{3}\ln x_1 + \frac{2}{3}\ln x_2,$$

where x_i is the amount of good i consumed. Alice and Bob thus have identical endowments but different tastes, with Alice being more favorably disposed toward good 1 and Bob more favorably disposed toward good 2. There are gains from trade. In the unique competitive equilibrium of this economy, Alice consumes $(\frac{2}{3}, \frac{1}{3})$, while Bob consumes $(\frac{1}{3}, \frac{2}{3})$.

Example 2. Now suppose that Ann and Bill consume a single good, "money," with identical utility functions $u(x) = \ln x$. Ann and Bill are each endowed with $\frac{1}{2}$ unit of money. However, each is aware that war in Asia may be imminent. Ann thinks the possibility of immediate war is $\frac{2}{3}$, while Bill thinks the probability is only $\frac{1}{3}$. Ann proposes that Bill agree to transfer some money to Ann if war breaks out, while Ann transfers some money to Bill if peace prevails. Indeed, letting x_1 be the amount of money consumed in the event of war and x_2 be the amount consumed in the event of peace, their expected utility functions are given by

$$u_A(x_1, x_2) = \frac{2}{3}\ln x_1 + \frac{1}{3}\ln x_2$$

and

$$u_B(x_1, x_2) = \frac{1}{3}\ln x_1 + \frac{2}{3}x_2.$$

Once again, there are gains from trade. In the unique competitive equilibrium of this economy, Alice consumes $(\frac{2}{3}, \frac{1}{3})$, while Bob consumes $(\frac{1}{3}, \frac{2}{3})$.

These two examples obviously map to the same Arrow-Debreu (Arrow and Debreu 1954) general equilibrium model. 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 differences 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 is relatively partial to good 1 and Bob is relatively partial to good 2, 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 that these differences in preferences can reasonably be taken as given.

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 simply go to a casino. But if you do not enjoy betting for its own sake, then there is no way that both of you can gain from this trade. In order to benefit from this deal, one of you needs the probability of war to be strictly greater than .5, and one needs it to be strictly less than .5. Obviously, it cannot be both." A win-win outcome is impossible here.

To construct a more extreme example, suppose that Ann has a ring that she is willing to sell Bob for $100. Bob does not like the ring per se but believes the ring is endowed with magical powers that would allow him to foresee the future and so is willing to pay $100 for it. Ann thinks that the ring has no magical power whatsoever, and maybe also that foreseeing the future is theoretically impossible, but would not part with the ring if she thought it actually could foresee the future. In this case, as in example 2, trade is not driven by differences in tastes but by differences in beliefs. And, assuming that these beliefs are well-defined, trade is possible only if one of the agents is wrong. The ring story differs from the financial market trade in example 2 because the beliefs are more extreme in the case of the ring, with each agent certain that he or she is correct, and because the analyst is likely to know which agent is correct. Indeed, the ring story has a flavor of a hoax. As we doubt that there are rings that foresee the future, selling the ring to Bob appears deceitful. However, the two stories share a 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, suppose that the observer is on the board of directors of a pension fund and that Ann and Bill are managers working for the fund, each in charge of a portion of the fund's assets, who propose to trade with one another. One possibility is that Ann and Bill have discovered that each of their funds is exposed to some risk that they can diversify by trading. In this case, the board has no reason to object.

Suppose, however, that Ann and Bill propose the type of trade encountered in example 2. In this case, the board member may suspect that at least one of the managers is in the wrong and may not want one manager to take advantage of the other's mistake. The board members may be unable to ascertain which manager is wrong, but they can find out whether one of them must be wrong by asking the managers to come up with a plausible, consistent story that would explain why they wish to trade. That is, the board members can ask the managers to reach unanimity not only 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. The board may thus require that there be at least one probability vector that can simultaneously explain why both managers 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 (2014).

Suppose instead that Ann and Bill manage money on behalf of different government agencies and that Ann consults her supervisor about the proposed trade. Her supervisor 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 respect your expertise, but that means that your counterpart 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 supervisor, talking only to her, expects her to rise to a rather high moral standard and engage only 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. This is the basis of unanimity Pareto domination.

2.3. Sharing Risks

This section presents an example that clarifies these concepts and their relationship to risk-sharing trade.

Example 3. Agnes and Barry are both expected-utility maximizers, with a Bernoulli utility function over monetary outcomes given by $u(x) = \ln x$. Agnes and Barry are thus both risk averse and in particular have utilities that become arbitrarily negative as consumption approaches 0, which perhaps reflects an inability to survive with no consumption. Each is endowed with 1 unit of money.

Agnes is an entrepreneur with an opportunity to invest her endowment in a start-up company. If successful, it will return a net payoff of $e^5 \approx 148$. If unsuccessful, Agnes will lose the unit with which she is currently endowed. Agnes thinks the project will succeed with a probability of .50, giving a tempting expected payoff, but there is no probability of success short of 1 that will tempt her to undertake the project and risk the downside of 0.

Now Agnes and Barry propose a trade. Barry will share half the cost of the project in the event of a failure, which gives each a consumption of $\frac{1}{2}$, and Agnes will share half the proceeds in the event of a success. Barry is less optimistic about the project than is Agnes, thinking that the probability of success is only .20. However, it is a quick calculation that the expected utility of this common lottery faced by Agnes and Barry is positive as long as the probability of success exceeds .14, as is the case for both Agnes and Barry. In this case, both agents are willing to trade, and there is a range of common beliefs, including both agents' beliefs, at which both would be willing to trade. The no-betting Pareto criterion and the unanimity Pareto criterion would endorse the trade.

This result shows that risk sharing is not precluded by either of our refinements of Pareto domination. Agnes may be viewed as holding an asset that exposes her to risk. Barry holds a riskless asset. For some range of beliefs, it will be in Barry's interest to take some of 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 of their beliefs are in both ranges, the risk-sharing deal is endorsed by both refinements of the Pareto criterion.

Now suppose that Agnes, perhaps uncertain of Barry's willingness to trade but anxious to escape the tyranny of a possible 0 payoff, offers to bear $\frac{3}{4}$ of the loss if the project is a failure and to transfer $\frac{3}{4}$ of the gains to Barry if it is a success. Barry obviously finds this enhanced deal profitable, and Agnes is sufficiently optimistic that the deal is also still advantageous for her. Moreover, the higher the probability of state 2, the better the deal for both Agnes and Barry, so there is clearly a common belief, such as Agnes's, at which both are willing to trade. The trade thus no-betting Pareto dominates not trading. However, it does not unanimity Pareto dominate no trading. Instead, the lower bound on the probability of success for which the trade is profitable for Agnes is about .28, which exceeds Barry's belief. Hence, the unanimity Pareto condition rejects the trade.

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. As we have just seen, no-betting Pareto domination does not imply unanimity Pareto domination. The no-betting criterion requires only 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 DOMINATION

3.1. Definitions

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

The allocations compared are functions from states to social out-

1. Gilboa, Samuelson, and Schmeidler (2014) present the no-betting Pareto criterion in a more general setup allowing for abstract outcomes and an infinite state space.

comes, which can be viewed as matrices, specifying a level of wealth for each individual at each state. Formally, the set of allocations is

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

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. Each agent i has a utility function $u_i{:}\mathbb{R} \rightarrow \mathbb{R}$ that is differentiable, strictly monotone, and (weakly) concave and a probability vector p_i on j such that she maximizes

$$\sum_{j \in S} p_i(j) u_i[f_i(j)].$$

A trade is a pair of allocations $(f, g) \in F^2$ in which the agents give up allocation 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, $(p_i)_{i \in N}$.

Definition 1. For two allocations $f, g \in F$, we say that

1) allocation f Pareto dominates g, denoted $f \succ g$, if for all $i \in N(f, g)$, $f \succ_i g$.

2) Allocation f no-betting Pareto dominates g, denoted $f \succ_{\text{NBP}} g$, if

i) $f \succ g$ and

ii) there exists a probability measure p_0 such that $\forall i \in N(f, g)$ and

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

3) Allocation f unanimity Pareto dominates g, denoted $f \succ_{\text{U}} g$, if for every probability measure p_k, $k = 1, \ldots, n$, $\forall i \in N(f, g)$,

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

Notice that we require strict preference for the agents involved in a trade when defining Pareto domination. 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 trade and those who are not, 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 allocation 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 a Pareto-improving trade to involve strict preference for every involved agent thus imposes no essential constraints.

Both condition i of no-betting Pareto domination and the unanimity Pareto condition (for $k = i$) require that the agents involved prefer f to g according to their actual beliefs. Condition ii of no-betting Pareto domination, by contrast, requires that one be able to find a single probability measure according to which all involved agents prefer trading g for f. The unanimity Pareto condition 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 constitutes a Pareto improvement over g according to all other agents' beliefs. Unanimity Pareto domination thus implies no-betting Pareto domination, which in turns implies the ordinary Pareto ranking.

3.2. Illustration

In this section we adapt a simple example from Gilboa, Samuelson, and Schmeidler (2014) that allows us to visualize the no-betting Pareto condition in a 2×2 Edgeworth box. We then proceed to illustrate the new unanimity condition in this context.

Suppose that there are 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 allocations are given as points in an Edgeworth box. The diagonal in Figure 1 is the set of full-insurance allocations.[2] Along this diagonal the slopes of the indifference curves of a given agent are identical and are determined solely by the probabilities that the agent attaches to the two states.

The analysis is straightforward if the agents have identical beliefs. The sets of Pareto-efficient, no-betting-Pareto-efficient, and unanimity-Pareto-efficient allocations coincide in this case and are given by the

2. The horizontal axis in each panel of Figure 1 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 than does agent B, and hence the set of Pareto-efficient allocations is given by a 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.

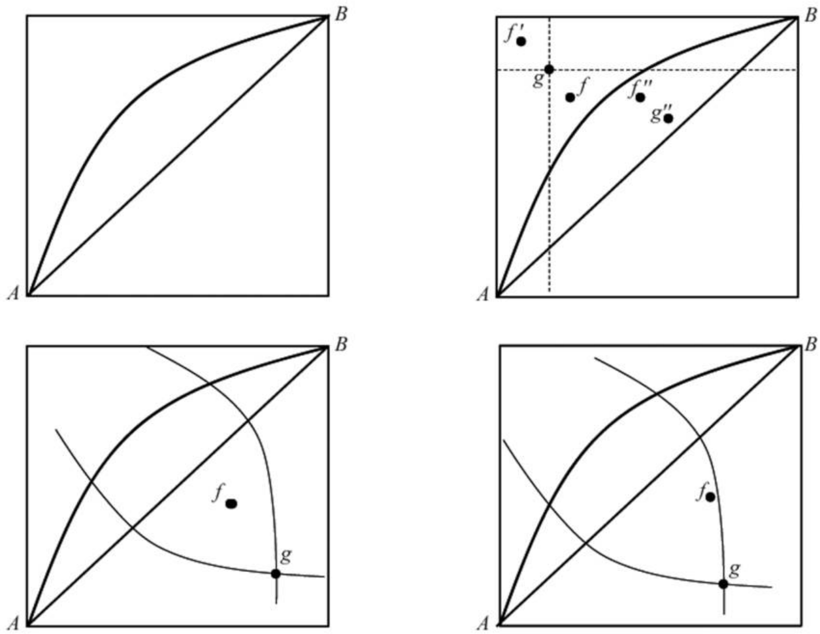

Figure 1. Allocation of money between agents A and B

diagonal. Efficiency calls for the agents to fully insure one another so that neither agent bears any risk.

When the agents hold different beliefs, the sets of Pareto-efficient, no-betting-Pareto-efficient, and unanimity-Pareto-efficient allocations are no longer identical. To illustrate, suppose that A thinks that state 2 is more likely than does B, and begin with the set of Pareto-efficient allocations. This set will be given by a curve that lies above the diagonal, illustrated in the top left panel of Figure 1. Because agent A thinks that 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.

Now consider no-betting Pareto domination. Consider the allocation g in the top right panel of Figure 1. The allocation f' lies farther from the Pareto curve than does g, and hence f' can neither Pareto dominate nor no-betting Pareto dominate g. In particular, no pair of indifference curves through g can both pass below f'. The allocation f lies closer to the Pareto curve than g, and so it is possible (but not necessary) that f Pareto dominates g. If f Pareto dominates g, then we turn to the second condition for no-betting Pareto dominance. Gilboa, Samuelson, and Schmeidler (2014) 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 Pareto 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 thus consists of those contained in the flattened lens bounded by the diagonal and the Pareto curve.

Consider now the unanimity Pareto condition. First consider allocations f and g that both lie on the same side of the diagonal as the Pareto 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, since unanimity Pareto domination in general imposes stronger requirements than does no-betting Pareto domination. 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 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 than does 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 the swap increases income, and a lower probability to state 2, where the swap decreases income). Similarly, agent B, who is willing to swap g for f according to her own beliefs, will surely be 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 Pareto 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 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 results. When f and g are on the same side of the diagonal as the Pareto curve, the trade can be justified by either no-betting or unanimity Pareto domination 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 to which they are exposed. By contrast, when the agents are on the other side of the certainty line (than the Pareto curve), the movement toward the diagonal makes each agent 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 requires only that both parties reduce the risk to which they are exposed, whereas the former also demands that they do it in a way that each finds beneficial when holding the other's beliefs.

4. IMPLEMENTATION

What could the notions of no-betting Pareto or unanimity Pareto domination be used for, and what issues arise in their use? We first note a key implication of both criteria and 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, that is, if $g_i(j) = g_i(j')$ for every agent $i \in N$ and states $j, j' \in S$. Notice that the designation of a trade as a bet depends only on the characteristics of the trade rather than the preferences of the agents. Gilboa, Samuelson, and Schmeidler (2014) justify the definition of no-betting Pareto domination by the fact that bets cannot be Pareto improving according to this definition:

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

Bets push agents away from risk sharing and can be strictly mutually advantageous only for agents who 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. The no-betting Pareto criterion also eliminates some Pareto rankings that are not bets. Gilboa, Samuelson, and Schmeidler (2014) characterize the Pareto rankings that are precluded by the no-betting Pareto criterion, showing that these rankings can be viewed as generalized bets. It is obviously also the case that unanimity Pareto domination precludes betting, since unanimity Pareto domination requires that all agents be better off according to each of their beliefs, and proposition 1 implies that this cannot be the case for any probability.

4.2. Observability

Suppose that we have a collection of agents who are willing to trade from g to f, so 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 allocation f unanimity Pareto dominates an allocation g, one needs to know the agents' utilities and probabilities. As for no-betting Pareto domination, 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. Meddling in Markets?

How could these notions be applied? Should one wish to use the concepts discussed for regulation of financial markets, it seems wise to be conservative and start with the minimal possible deviation from standard lore and common practice. If we are to consider an alternative to the Pareto criterion that is a relatively minor modification thereof, 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 that there is no way to justify the comparison. We therefore discuss here only the no-betting Pareto criterion: whereas both no-betting Pareto and unanimity Pareto domination satisfy the first condition (being a subset of standard Pareto domination), no-betting Pareto domination eliminates fewer rankings than does unanimity Pareto domination, and hence no-betting Pareto domination suggests a less radical departure from common conceptualizations. Coupling this observation with the fact that less information is required to evaluate the no-betting Pareto relation, we concentrate here on this relation.

As do Posner and Weyl (2013), we imagine a scenario of a financial monitoring authority, in our case an authority that 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 that it is willing to participate in the proposal, ensuring that 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 are 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 constant absolute risk aversion or constant relative risk aversion 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 agents 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 of financial institutions by specifying states of the world, utilities, and probabilities. Yet we would find it problematic if, post hoc, such an institution cannot justify its decisions by describing them in such a model or can do so only by listing unreasonable states or choosing outlandish beliefs or degrees of risk aversion. Thus, we imagine a scenario in which institutions reach their decisions in whatever way they may, but they are held accountable to be able to justify these decisions by economic models. In such a scenario we add the requirement that the parties to the transaction also be able to 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 that the question of whether trade should be restricted, and especially when trade is equivalent to betting, is particularly important when dealing with institutions that manage other people's money. Pension funds and mutual funds, especially index funds, typically identify enhanced risk management as one of their key advantages, and we view the funds as entering implicit contracts with their investors to provide appropriate risk management. These funds have a legitimate interest in trades that allow them to effectively insure their investors. However, 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. In particular, we do not view a trade that can turn out well for one party only if the other party's beliefs are mistaken as consistent with the provision of insurance. In a frictionless world, investors would face no difficulties in identifying and preventing problematic transactions. In practice, we view the concept of no-betting Pareto domination 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 that we some-times 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 the representation theorem of Savage (1954) 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, a normative analysis relies on attaching to the utility function some meaning having to do with welfare or desirability, a meaning that cannot be attached to beliefs.[3] Moreover, while our analysis focuses on ex ante Pareto domination, embedding our discussion in a dynamic setup 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 more than broccoli.

3. The development of the theory of revealed preference, interpreting utilities simply as descriptions of consistent choices, was accompanied by warnings that doing so eviscerated the foundations of welfare economics. For example, Sen (1973, p. 253) writes, "In economic analysis individual preferences seem to enter in two different roles: preferences come in as determinants of behaviour and they also come in as the basis of welfare judgments. . . . This dual link between choice and preference on the one hand and preference and welfare on the other is crucial to the normative aspects of general equilibrium theory." No such link can exist if utilities have no meaning.

One obvious way for probabilities to have meaning is for them 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 they 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 10 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 beliefs, the no-betting Pareto, the unanimity Pareto, and conventional Pareto criteria would all coincide. However, the Harsanyi doctrine (Harsanyi 1967, 1968a, 1968b) 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 prior belief (in a model that is common knowledge among them) cannot hold diverging posterior beliefs (that are commonly known among them), 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.[4]

To conclude, while we recognize that beliefs differ, we do not view probabilities as being 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 2 years, but one could bring evidence and arguments to bear on this question that would generally prompt similar revisions in peoples' probabilities, which indicates 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 on the basis of the likelihood that various events will occur,

4. An alternative explanation is that people do have common prior beliefs, 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.

beginning with something like, “Let me show you why this stock 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 the “let me explain why this stock is sure to fall” case on the other end.

5.2. Related Literature

Others argue that the concept of Pareto domination is troubling when beliefs differ. The distinction between gains from trade based on tastes (as in example 1) and based on differences in beliefs (as in example 2) is 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 (forthcoming) refers to the type of 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 (2013) call for a regulatory authority, analogous to the Food and Drug Administration, 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 building blocks of our concepts are familiar. A hallmark of the legal definition of fraud is an element of deception (compare Bigelow 1887). If there exists a common belief under which all involved parties would gain from a trade, then one can imagine the trade being conducted without deception. On the other hand, if it were meaningful to talk about true probabilities and an agent knew such probabilities, then the agent advocating a Pareto improvement excluded by the no-betting Pareto criterion would have to act fraudulently, and we can imagine the financial monitoring authority screening for such fraudulent trades. Related to this, al-Suwailem (2006) argues that the difference between risk arising from value-adding and wealth-creating activities and gambling lies at the heart of Islamic finance but offers quite different criteria for the two categories. The unanimity Pareto criterion, in contrast, is reminiscent of Kant’s categorical imperative (Kant [1797] 2005) and the golden rule (Wattles 1996) in that it endorses trade only if each participant believes that all participants will gain from the trade. This is in keeping with the higher moral standard discussed in Section 2.2 in that there must be an interpretation of a trade not only that avoids fraud but under which everyone believes that everyone can gain from the trade.

The paper most closely related to ours, with a quite similar motivation and somewhat different details, is Brunnermeier, Simsek, and Xiong (2012).[5] 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 ranks some allocations that are not ranked by the standard Pareto criterion and hence are not ranked by either the no-betting or unanimity Pareto criterion. We recognize that such value judgments must often be made[6] but believe that the point of departure for such value judgments should be a refinement of the Pareto criterion.

Brunnermeier, Simsek, and Xiong (2012) offer another approach that is independent of a specification of the welfare weights and 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 allocation g with the property that every agent prefers g to f given belief p (with weak preference assumed for all agents and strict preference assumed for at least one of them). Notice that the allocation g is allowed to depend on the belief p. An allocation f is belief-neutral efficient if there is no allocation g and belief p in the convex hull of the agent's 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 agent's beliefs makes their belief-neutral notion most directly comparable to our unanimity Pareto notion:

1. If f is unanimity Pareto inefficient, then it is belief-neutral inefficient.

2. If (contrapositive of 1) f is belief-neutral efficient or unclassified (that is, not belief-neutral inefficient), then it is unanimity Pareto efficient.

5. See also Simsek (2012), who discusses financial innovation when trade is motivated both by risk sharing and by speculation.

6. See Blume et al. (2013) for a welfare criterion for heterogeneous beliefs that has a Rawlsian flavor.

3. If f is unanimity Pareto efficient, nothing is implied about its belief-neutral efficiency: it may be belief-neutral efficient, belief-neutral inefficient, or unclassified by the belief-neutral Pareto criterion.

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

We can compare Brunnermeier, Simsek, and Xiong's (2012) belief-neutral concept to the no-betting Pareto concept as follows:

1. An allocation g may be belief-neutral inefficient but may not be Pareto or no-betting Pareto dominated because, for an allocation f to satisfy $f \succ_{\mathrm{NBP}} g$, it has to be the case that $f \succ g$ (according to the agents' actual beliefs), which is not needed for belief-neutral inefficiency of g.

2. An allocation g may be no-betting Pareto dominated by an allocation f but may not be belief-neutral inefficient, as the latter condition requires that for every probability p in the convex hull of the agents' beliefs there be an allocation f_p that dominates g (for all agents under p), whereas $f \succ_{\mathrm{NBP}} g$ implies this domination for only one probability, which may not even be in this convex hull.

3. An allocation g may be belief-neutral efficient, but f may no-betting Pareto dominate g because the second condition of $f \succ_{\mathrm{NBP}} g$ may hold only for beliefs outside of the convex hull of the agents' beliefs (as above).

4. There may be no allocation f that no-betting Pareto dominates the allocation g, but g may nonetheless fail to be belief-neutral efficient, because it is possible that an allocation f dominates g for some belief p in the convex hull of the agents' beliefs (hence, g is not belief-neutral inefficient) but f does not Pareto dominate g according to the agents' original (and different) beliefs.

To conclude, the alternative concepts so far suggested in the literature do not seem to be tightly related. Each obviously captures certain considerations and misses others. While it is probably too early to attempt to choose one of them as a guiding principle on the basis of theoretical reasoning alone, we believe that the type of considerations they raise may be useful in thinking about the regulation of financial markets.

REFERENCES

al-Suwailem, Sami. 2006. Hedging in Islamic Finance. Occasional Paper No. 10. Islamic Development Bank, Jeddah.

Arrow, Kenneth J., and Gerard Debreu. 1954. Existence of an Equilibrium for a Competitive Economy. *Econometrica* 22:265–90.

Aumann, Robert J. 1976. Agreeing to Disagree. *Annals of Statistics* 4:1236–39.

Bigelow, Melville M. 1887. Definition of Fraud. *Law Quarterly Review* 3:419–28.

Blume, Lawrence E., Timothy Cogley, David A. Easley, Thomas J. Sargent, and Viktor Tsyrennikov. 2013. Welfare, Paternalism, and Market Incompleteness. Working paper. Cornell University, Department of Economics, Ithaca, N.Y.

Brunnermeier, Markus K., Alp Simsek, and Wei Xiong. 2012. A Welfare Criterion for Models with Distorted Beliefs. Working paper. Princeton University, Department of Economics, Princeton, N.J.

Gilboa, Itzhak, Larry Samuelson, and David Schmeidler. 2014. No-Betting-Pareto Domination. *Econometrica* 82:1405–42.

Harsanyi, John C. 1967. Games with Incomplete Information Played by "Bayesian" Players, Part 1: The Basic Model. *Management Science* 14:159–82.

———. 1968a. Games with Incomplete Information Played by "Bayesian" Players, Part 2: Bayesian Equilibrium Points. *Management Science* 14:320–33.

———. 1968b. Games with Incomplete Information Played by "Bayesian" Players, Part 3: The Basic Probability Distribution of the Game. *Management Science* 14:486–502.

Kant, Immanuel. [1797] 2005. *Fundamental Principles of the Metaphysics of Morals*. Mineola, N.Y.: Dover.

Kreps, David M. 2012. *Microeconomic Foundations I: Choice and Competitive Markets*. Princeton, N.J.: Princeton University Press.

Milgrom, Paul, and Nancy Stokey. 1982. Information, Trade, and Common Knowledge. *Journal of Economic Theory* 26:17–27.

Mongin, Philippe. Forthcoming. Spurious Unanimity and the Pareto Principle. *Economics and Philosophy.*

Posner, Eric A., and E. Glen Weyl. 2013. An FDA for Financial Innovation: Applying the Insurable Interest Doctrine to Twenty-First-Century Financial Markets. *Northwestern University Law Review* 107:1307–57.

Savage, Leonard J. 1954. *The Foundations of Statistics*. New York: John Wiley & Sons.

Sen, Amartya. 1973. Behavior and the Concept of Preference. *Economica* 40: 241–59.

Simsek, Alp. 2012. Speculation and Risk Sharing with New Financial Assets. Working paper. Massachusetts Institute of Technology, Department of Economics, Cambridge, Mass.

Stiglitz, Joseph E. 1989. Using Tax Policy to Curb Speculative Short-Term Trading. *Journal of Financial Services Research* 3:101–15.

Wattles, Jeffrey. 1996. *The Golden Rule*. New York: Oxford University Press.

Weyl, E. Glen. 2007. Is Arbitrage Socially Beneficial? Working paper. University of Chicago, Department of Economics, Chicago.