ECONOMETRICA

JOURNAL OF THE ECONOMETRIC SOCIETY

An International Society for the Advancement of Economic Theory in its Relation to Statistics and Mathematics

http://www.econometricsociety.org/

Econometrica, Vol. 82, No. 4 (July, 2014), 1405–1442

NO-BETTING-PARETO DOMINANCE

ITZHAK GILBOA
HEC Paris, 78351 Jouy-en-Josas Cedex, France and Berglas School of Economics, Tel Aviv University, Tel Aviv 6997801, Israel

LARRY SAMUELSON
Yale University, New Haven, CT 06520, U.S.A.

DAVID SCHMEIDLER
School of Economics, IDC Herzliya, Israel and School of Mathematical Sciences, Tel Aviv University, Tel Aviv 69203, Israel

Econometrica, Vol. 82, No. 4 (July, 2014), 1405–1442

NO-BETTING-PARETO DOMINANCE

BY ITZHAK GILBOA, LARRY SAMUELSON, DAVID SCHMEIDLER[1]

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

KEYWORDS: Pareto efficiency, betting, speculation, Pareto dominance, beliefs.

1. INTRODUCTION

1.1. *Motivation*

STANDARD ECONOMIC LORE SUGGESTS that Pareto-improving trades are obviously a good thing—how can one argue with making everyone (at least weakly) better off? This paper argues that when outcomes are uncertain and agents have different beliefs, the Pareto-improving criterion is less compelling than when agents have common beliefs. Consider the following examples.

EXAMPLE 1: Alice and Bob have one apple and one banana each. Their utility functions are linear. Alice is indifferent between one apple and two bananas, and Bob is indifferent between two apples and one banana. In the competitive equilibrium of this simple economy, Alice and Bob obtain a Pareto optimal allocation in which Alice consumes only apples and Bob consumes only bananas. Each prefers this outcome to the initial endowment.

EXAMPLE 2: Ann and Bill have one dollar each. There are two states of the world: in state 1, the price of oil a year from now is above \$100 a barrel, and in state 2, it is no more than \$100 a barrel. Ann and Bill are risk neutral. Ann thinks that state 1 has probability $2/3$ and Bill thinks state 1 has probability $1/3$. In the competitive equilibrium of this simple economy, Ann and Bill obtain a Pareto optimal allocation in which Ann has no money in state 2 and Bill has no money in state 1. Each prefers this outcome to the initial endowment, that is, to consuming \$1 whatever is the price of oil in a year.

[1]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. We thank the editor and three referees for helpful comments and suggestions. ISF Grant 396/10 (Gilboa and Schmeidler), ERC Grant 269754 (Gilboa), and NSF Grant SES-1153893 (Samuelson) are gratefully acknowledged.

 DOI: 10.3982/ECTA11281

These examples map to the same Arrow–Debreu (1954) general equilibrium model: there are two goods $\{1, 2\}$, and two agents $\{A, B\}$. The utility functions are given by

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

and the initial endowments are

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

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 is not obvious that Pareto domination has the same meaning in both examples. In Example 1, there is no uncertainty and the two consumers differ only in tastes. If Alice prefers apples and Bob prefers bananas, they are better off when they trade one apple for one banana. Indeed, Alice and Bob would both be better off even if transactions costs or trading frictions consumed some of their goods. Importantly, an outside observer who witnesses the trade between them has no reason to try to convince Alice to keep her banana or Bob to keep his apple. *De gustibus non est disputandum*: the agents are the only judges of their tastes over final consumption goods.

By contrast, in Example 2, Ann and Bill are both better off after trade, but only because they have different beliefs about the future price of oil. Ann and Bill cannot both be right: if the probability of state 1 is $2/3$, it cannot be $1/3$ as well. Moreover, there is no common belief under which both can gain from trade, and in the presence of transactions costs no common belief under which neither loses from trade.

It is not clear how one should interpret the probability of states in this setting, and perhaps Ann and Bill should not have probabilistic beliefs over events such as the future price of oil at all. But if they do, and if these beliefs have any concrete meaning, then these beliefs are incompatible. To see this more clearly, imagine that the weight vectors $((2/3, 1/3), (1/3, 2/3))$ are replaced by $((1, 0), (0, 1))$ in both Example 1 and Example 2. Thus, Alice has no value for bananas and Bob has none for apples. The trade between them makes perfect sense, and again an outside observer has no reason to dissuade them from trading the fruits. True, both apples and bananas have nutritional value, but the latter is not the agents' concerns: the agents do not hold mistaken beliefs about nutrition; rather, each of the agents simply does not like the taste of one of the fruits. In Example 2, by contrast, each agent is convinced of a proposition that is incompatible with the other agent's conviction. In fact, they behave as if they *knew* certain facts for sure, but these facts are contradictory. If we were to ask each agent, why do you behave as if you knew a proposition p, "*de gustibus non est disputandum*" would hardly qualify as a rational reply, when we know that as soon as the state of the world is revealed, one of the agents will be proven wrong.

This paper proposes a *No-Betting-Pareto* refinement of Pareto domination for uncertain allocations that distinguishes these two situations. One way to do so would be to restrict Pareto domination to cases in which all parties involved have the same beliefs. This, however, may needlessly restrict agents in their attempts to share risk. We do not find trade among agents whose beliefs differ problematic, as long as the trade does not hinge on these differences in beliefs. The following example illustrates.

EXAMPLE 3: Agnes is a computer scientist with an idea for a start-up company. If successful, the company will net \$10 million. If unsuccessful, she will lose her initial investment of \$1 million. Agnes assigns probability .9 to a success, but is not willing to take the risk and invest her own funds. She 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, but he is risk neutral and therefore willing to take the risk and make the investment.

Once again, the agents have different beliefs. However, in this case, there is a range of beliefs, including Agnes's belief, at which both agents would be willing to trade. In contrast to Example 2, where the difference in beliefs is crucial for trade to take place, the exchange in Example 3 can be justified as voluntary trade between agents who do not necessarily have mistaken beliefs. Their beliefs do differ, but this difference is not essential to the trade. Agnes has an asset that will be worth a lot in one state and little in another, and trade allows her to share this risk with Barry. This is the sort of risk sharing that gives rise to productive trade in financial markets, and our refinement of Pareto domination will allow such trade.

1.2. *No-Betting-Pareto*

We do not take issue with Pareto domination under certainty, as in Example 1. We also find Pareto domination under uncertainty compelling if the agents share common beliefs. Our task is then to address cases in which beliefs differ in such a way that the type of betting illustrated by Example 2 is not classified as an improvement, while the risk sharing of Example 3 is.

In order to distinguish between betting and risk sharing, we offer the following definition: allocation f *No-Betting-Pareto dominates* allocation g if every agent (who is not indifferent between the two at all states of the world) prefers f to g, *and* if there exists a single belief with the property that, if all agents held this belief, then they would still prefer f to g.

There are two parts to this definition. The first imposes a standard Pareto criterion, and ensures that No-Betting-Pareto refines Pareto efficiency. The second component requires the existence of a single belief under which all agents prefer f to g. Theorems 1 and 2 below show that this second requirement captures a sense in which f allows more risk sharing than does g. As both theorems characterize only the second requirement, they may be useful for thinking about risk sharing when beliefs differ, irrespective of Pareto rankings.

The competitive allocation in Example 1 trivially No-Betting-Pareto dominates the initial endowment, since beliefs are degenerate in this case. More generally, when all the agents have the same beliefs, No-Betting-Pareto domination is basically identical to the usual notion of ex ante Pareto domination.[2] However, the competitive allocation in Example 2 Pareto dominates the initial endowment but does not No-Betting-Pareto dominate it, because no belief can make both agents strictly prefer the former over the latter. In Example 3, the risk-sharing deal No-Betting-Pareto dominates the endowment, even though the agents have different beliefs, because there is a common belief under which both would prefer the deal.

The distinction between Pareto domination and No-Betting-Pareto domination can be viewed as a manifestation of a more general principle, under which unanimity about a given claim—say, that trade is desirable—becomes more compelling when unanimity about the reasoning that leads to it is also possible. A unanimous conclusion loses much of its appeal if, by accepting it, one has to concede that at least some of the agents involved must be wrong in the reasoning process that led them to that conclusion.

1.3. *Underlying Assumptions*

We examine agents who satisfy Savage's (1954) axioms and can therefore be viewed as maximizing expected utility with respect to a probability measure.[3] Thus, our agents satisfy a rather demanding standard of rationality. We agree that Pareto comparisons may be problematic when agents are confused or inconsistent or otherwise irrational, but our focus here is on differences in beliefs.

Our definitions and interpretation implicitly assume that the utility functions of the agents have to do with desirability, so that higher utility values are (other things equal) a good thing. Similarly, we interpret the probability measures as related to the notion of beliefs, so that a higher subjective probability of an event suggests that the agent under consideration finds this event more likely.

[2]For some subtleties, see the discussion in Section 2 below.

[3]The formal model we use differs from Savage's in some technical details, such as allowing for finite state spaces and general sigma-algebras.

This is the heart of the distinction between Examples 1 and 2 above. The technical formalities of these two examples are identical, but they capture different notions—saying that one likes an apple better than a banana is different from saying that one state is more likely than another.[4]

Of the various interpretations of probabilities, there are two views under which our analysis is useless. The first suggests that probabilities are scientifically objective quantities, which can be measured or estimated in such a way that all rational individuals would agree on them, up to negligible measurement errors. According to this view, our No-Betting-Pareto criterion is unobjectionable but pointless: if all agents will always agree on the probabilities of the states of the world, there will be no substantial distinction between our concept and the standard one. If the agents ever do disagree on the probabilities, then one need only to point out the measurement error to them and they will converge to a unanimous estimate.

We argue that for many phenomena of interest, probabilities cannot be claimed to be scientifically measurable quantities.[5] People express their views about market conditions, economic outlooks, or inflation predictions and provide probabilistic assessments that often vary greatly from one person to another. Moreover, these differences are not taken to reflect simply imprecise measurements. Rather, there does not seem to be a meaningful way of determining or even defining the probability of a war erupting in the Middle East over the coming year in the same way that one can determine the probability of, say, a fair coin coming up head in at least one of three tosses. Even in the case of climate change, where estimates rely on a remarkable body of solid scientific knowledge, predictions vary (see, for instance, Heal and Millner (2013)). We also take the large volume of trade in financial markets, despite agreeing to disagree (Aumann (1976)) and no-trade (Milgrom and Stokey (1982)) results, as testifying to differences in beliefs. Moreover, it seems that differences in beliefs are partly the reason for the existence of betting markets, ranging from bets on horse races to election outcomes. In short, we readily admit that when probabilities of states of the world can be scientifically computed or estimated, our analysis is irrelevant. However, we find that many economic problems of interest do not belong to this category.

At the other extreme is the view that probabilities are meaningless constructs that are part of a mathematical representation of preferences but have no meaning outside the context of this representation. This view arises most

[4]For example, see Aumann (1987, p. 13), who wrote "... utilities directly express tastes, which are inherently personal. It would be silly to talk about 'impersonal tastes', tastes that are 'objective' or 'unbiased'. But it is not at all silly to talk about unbiased probability estimates, and even to strive to achieve them. On the contrary, people are often criticized for wishful thinking—for letting their preferences color their judgments. One cannot sensibly ask for expert advice on what one's tastes should be; but one may well ask for expert advice on probabilities."

[5]This is not an ontic claim, but rather a realistic working assumption adopted by (for example) most observers of financial markets.

clearly out of Savage's (1954) simultaneous derivation of a utility function and probabilities that together rationalize behavior satisfying his axioms. According to this interpretation, it is a categorical mistake to ask what these constructs represent in the real world, or to suggest that they capture mental phenomena such as beliefs. And it follows that differences in subjective probabilities need not be treated differently than differences in utilities, and that there is no difference between trade in Example 1 and in Example 2 above.

This view is logically coherent and often theoretically appealing. But it is not the only possible interpretation of probabilities, or even of subjective probabilities. Savage's (1954) derivation of subjective expected utility maximization takes only preferences between acts as its point of departure, but his development is compatible with a view of probabilities as ranking events by their likelihood.[6] Similarly, his development is consistent with a view of probabilities and utilities as conceptually different.[7]

The view that probabilities are not empty, theoretical constructs runs throughout economics. The main goal of the separation between utility and probability in the representation of preferences is the presumption that these are different entities, with the first (utility) representing characteristics of the model that do not change, whereas the second (probabilities) may change with new information. We see this in the fact that comparative statics exercises in economics overwhelmingly hold utility functions fixed, while varying factors such as prices and incomes but also probabilities. We also note that there are many situations in which the preferences assumed in Savage's model may not be observable, but in which one might still extract information about probabilities from other types of data. For example, suppose one is making plans for a retirement that is two decades off, including the possibility of a retirement home in Greenland, and hence is interested in an expert's best estimate of the effect of global warming on the average temperature on Greenland's ice cover. One might find it simpler to ask the expert what she thinks rather than observe her choices between acts that would determine consequences 20 years later, and one would view this information as valuable precisely because it is not

[6]Savage used his axiom P4 to define a binary relation over events, denoted by $\geq$, that is interpreted as "at least as likely as." Under P1–P3 and P5, this relation satisfies de Finetti's axioms of "qualitative probability," which is a binary relation over events that is taken as a primitive (though it is implicitly assumed to be related to behavior via preferences over bets). Thus, Savage's theorem can be viewed as first completing de Finetti's project, showing that a qualitative probability relation that is both *fine* and *tight* can be represented by a probability measure, and then proceeding to relate this probability to decisions via the expected utility paradigm. Kreps (1988) presented Savage's theorem in this order.

[7]See, for example, footnote 13 in Aumann (1987, p. 13), starting with, "That Bayesian decision theory a la Savage derives both utilities and probabilities from preferences does not imply that it does not discriminate conceptually between these two concepts."

purely subjective.[8] Finally, people more readily talk about what probabilities "should be," suggest that someone should alter her belief, or argue that a belief is irrational, than they do about utility. A bet that pays out if a Republican wins the next presidential election and otherwise collects might well be criticized on the grounds that the implicit odds are unreasonable. The probability is thus not treated as a meaningless mathematical construct, though we see no way to view it as objective.

We thus find neither of the two preceding interpretations of probabilities—as directly observable and scientifically measurable quantities, or as mere mathematical constructs that have no meaning outside the expected utility model—satisfactory or practically helpful in the discussion of economic phenomena such as financial markets. For an example of the interpretation we have in mind, consider the standard mindset of statistical inference according to which there exists a "true" underlying data generating process, which may in principle be knowable, but which is not known by the agents. Instead, the agents may have different views about the data generating process. They may be Bayesian statisticians, holding different prior beliefs about certain parameters of the model. However, some or even all of them, Bayesian or not, may have misspecified the model, and thus make predictions based on a mistaken statistical model. As a result, beliefs differ. At the same time, the agents are willing to learn, either by Bayesian updating, by hypotheses testing, or by other means, and they can sometimes be convinced that their initial beliefs were wrong. In situations such as this, it makes sense to ask whether trade, as in Example 2, hinges upon differences in beliefs, or whether, as in Example 3, it can be justified by a shared belief for all agents.

We have a special interest in concepts that do not rely on perfect observability of probabilities. That is, we assume that agents have beliefs that can be captured by probability measures, but these need not be directly observable to other agents.[9] By contrast, we believe that utility profiles will effectively be observable in many applications, in which case a willingness to trade will provide some information about an agent's belief, but will typically not pin down a single probability measure. For concreteness, consider trade of financial assets by institutions that manage funds for clients, such as retirement funds or hedge funds. Suppose that these institutions have to make their holdings and their trades publicly available. Moreover, imagine that they must stand ready, if challenged, to justify their trading decisions by showing that these decisions maximized expected utility given a certain belief and a predetermined utility function. Thus, utilities are effectively observable, and trade maximizes expected utility relative to some probability, but the latter is typically not directly observable.

[8]Clearly, if one observes the expert behaving inconsistently with her stated beliefs, then one should be suspicious of the latter, but this does not imply that stated probabilities should be completely ignored.

[9]However, our analysis is equally applicable when probabilities are observable.

To sum, we consider economic agents who are rational in the sense that they maximize subjective expected utility. They have utility functions that reflect the desirability of various outcomes. However, they are aware that their estimates of probabilities are not reliable and they are not surprised that other agents have different estimates. Our argument is that this should make us think more carefully about trades based on differences in beliefs. Should it not also make the agents think more carefully about whether they trust their own probabilities, and whether they wish to be subjective expected utility maximizers in the first place? Perhaps so. Economic agents may therefore use other models, which allow them to express their doubts about the appropriate probabilities to be used in their trading. In this paper, however, we retain the standard, Bayesian, expected utility maximization in order to isolate considerations arising out of differences in beliefs.

1.4. *Alternative Definitions*

There are several reasonable alternatives for our definition of No-Betting-Pareto. For example, one may consider a variation on our criterion in which the single belief under which all agents prefer f to g must lie in the convex hull of the agents' beliefs. Brunnermeier, Simsek, and Xiong (2012), discussed in Section 4.2, imposed such a convex-hull condition. Another alternative (suggested by Gayer (2013)) is to restrict Pareto domination to the cases in which an alternative guarantees a higher expected utility than another, for each agent, given *all* agents' beliefs. Blume, Cogley, Easley, Sargent, and Tsyrennikov (2013) ascribed sets of possible beliefs to a planner and to each agent, and then assessed acts according to the minimum (over the planner's beliefs and the agent's beliefs) utility of the agent with the smallest expected utility.

Our No-Betting-Pareto criterion does not require the common belief under which all agents prefer f to g to be one of the agents' beliefs or to lie in the convex hull of the agents' beliefs. In contrast, the alternatives mentioned in the preceding paragraph respect unanimity among the agents: if all agents agree that the probability lies in a certain half-space, only probabilities in this half-space are considered for further discussion. While this property is appealing, it is also subject to question: once we admit that each agent may be wrong in her beliefs, it is difficult to dismiss the possibility that they are all wrong. In particular, if we were to view each agent's belief as a point estimate of a "true" probability, a confidence set for this belief would typically allow beliefs outside the convex hull of the point estimates.

Our criterion satisfies two additional desiderata. First, it is compatible with a cautious approach, imposing a relatively minimal restriction on the standard notion of Pareto dominance. Our criterion does not rank any pair of alternatives that are not already ranked by the standard Pareto criterion, and rejects the standard ranking only when there exists *no* common belief that could rationalize the ranking, whether in the convex hull of the agents' beliefs or not.

In this sense, our definition is conservative: it rules out relatively few Pareto-dominance instances. Given the general acceptance of the Pareto dominance criterion, it seems fruitful to start by a minimal modification, restricting attention to Pareto rankings and rejecting such rankings only when the argument against them is compelling.

Second, to verify whether our definition holds, one does not need to know the agents' beliefs. Given the agents' utility profiles before and after trade, one may check whether a single probability can justify this trade for all of them without wondering what the actual beliefs were.[10] This may prove to have practical advantages. Beliefs may not be directly observable, and definitions that use them may be more difficult to test, may raise manipulability problems, and so forth.[11]

1.5. *Related Literature*

It is well recognized 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 of Pareto domination appearing in Example 2 as *spurious unanimity*.

The growing sophistication of financial assets and the recent financial crisis have attracted increased attention to the question of the efficiency of financial markets. Weyl (2007) pointed out that arbitrage might be harmful when agents are "confused." Posner and Weyl (2012) called for a regulatory authority, akin to the FDA, that would need to approve trade in new financial assets, guaranteeing that such trade would not cause more harm than good. Kreps (2012) and Brunnermeier, Simsek, and Xiong (2012) also discussed distorted beliefs.[12]

We agree that cognitive and affective phenomena such as confusion or overconfidence are important, and that they should be taken into consideration when discussing regulation of financial markets. But we do not think that the conceptual difficulty with Pareto domination is restricted to agents who are irrational in one way or another. To highlight this point, we discuss agents who are subjective expected utility maximizers, each satisfying all of Savage's axioms.[13] We believe that the conceptual underpinnings of voluntary trade

[10]We mention in passing that this task is computationally easy.

[11]Utilities may also be difficult to observe. As we mentioned at the end of the preceding subsection and discuss in Section 4.4, one application we envision is that in which the No-Betting-Pareto criterion is applied to pension funds or other investment funds, with utility functions drawn from a class set by regulators to identify appropriate risk levels.

[12]The latter paper is particularly relevant to ours, and the two are compared in Section 4.2.

[13]Again, we refer here to Savage's conceptual axioms P1–P5, rather than to P6 or to the structural assumption about the algebra of events consisting of all subsets of the state space.

among such agents needs to be examined even if the agents are not confused or overoptimistic.

Indications that it is more difficult to aggregate preferences under subjective uncertainty than under either certainty or risk have also appeared in the social choice literature. Harsanyi's (1955) celebrated result showed that, in the context of risk (i.e., known, objective probabilities), if all individuals as well as society are von Neumann–Morgenstern expected utility maximizers (von Neumann and Morgenstern (1944)), a mild Pareto condition implies that society's utility function is a linear combination of those of the individuals. However, Hylland and Zeckhauser (1979) and Mongin (1995) found that an extension of Harsanyi's theorem to the case of uncertainty cannot be obtained. An impossibility theorem shows that under a mild Pareto condition, one cannot simultaneously aggregate utilities and probabilities in such a way that society will satisfy the same decision theoretic axioms as do the individuals.

Gilboa, Samet, and Schmeidler (GSS, 2004) introduced a restricted Pareto condition under which society finds f as desirable as g when all individuals do so, but only when these preferences concern alternatives over which there are no disagreements in beliefs. The result of their paper is that, when one restricts the Pareto condition in this way, preferences can be aggregated into a complete societal preference respecting the restricted Pareto criterion. By contrast, the present paper does not ascribe to society a complete preference over alternatives. It discusses a particular instance of unanimous preferences, $f \succsim_i g$ for all i (with strict preference for all agents i who are affected by the swap of g for f), and asks whether society should agree with that particular ranking. Importantly, GSS applied the Pareto criterion only when agents have identical beliefs over the relevant alternatives, whereas we rank acts over whose distributions individuals may well disagree, as long as there is a shared *hypothetical* belief that would still rationalize trade for each of them.

The complete social preference relation in GSS (2004) and No-Betting-Pareto domination are quite different from a conceptual point of view. When a country has to choose an economic policy, decide whether to use nuclear power plants, or decide whether to wage a war, the decision cannot be decentralized; it has to be made for all individuals as a group. In this context, GSS showed that the natural idea of simultaneous averaging of utilities and of probabilities is necessitated by a reasonable version of the unanimity (Pareto) axiom. However, these "averaged" preferences are not very relevant for decentralized decisions. Hence, economists would tend to eschew the task of defining a social welfare function or a complete preference order for society as a whole in deference to a weaker notion such as Pareto dominance. The current paper belongs in this tradition, while differing from the classical literature in its definition of "dominance."

2. THE MODEL

There is a set of agents $N = \{1, \dots, n\}$, a measurable state space (S, Σ), and a set of outcomes X. An outcome specifies all the aspects relevant to all agents. It is often convenient to assume that X consists of real-valued vectors, denoting each individual's consumption bundle, but at this point we do not impose any conditions on the structure of X.

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

$$F = \{f : S \to X \mid f \text{ is simple and } \Sigma\text{-measurable}\}.$$

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

Each agent i has a preference order $\succsim_i$ over F. Agent i is characterized by a utility function $u_i : X \to \mathbb{R}$ and a probability measure[14] p_i on (S, Σ), and $\succsim_i$ is represented by the maximization of $\int_S u_i(f(s))\, dp_i$. We assume that the agents can be represented as expected utility maximizers to emphasize that our arguments do not hinge on any type of so-called bounded rationality of the agents. We assume that the agents agree on the state space S, returning to the significance of this assumption in Section 4.3.

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

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

2.1. *No-Betting-Pareto*

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

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

[14] For simplicity, we adopt the standard definition of a probability measure, which implies sigma-additivity. As we consider only simple acts, no modifications are needed to consider measures that are only finitely additive.

We use the term *improvement* to emphasize the fact that the agents in the economy would swap g for f voluntarily. We will also use the terminology f *improves upon* g, denoted by $f \succ_* g$.

Our main interest lies in improvements for which $|N(f, g)| \geq 2$, though the cases in which $|N(f, g)| = 1$ are not ruled out.

Notice that we require *strict* preference for the agents involved in the improvement. The relation $f \succ_* g$ is thus more restrictive than standard Pareto domination, which allows some agents, for whom $u_i(f(s)) \neq u_i(g(s))$ for some s, to be indifferent between f and g. This definition eliminates some technical complications, but we also find it intuitive. Notice that we allow the existence of agents who are unaffected by the swap (f, g), and hence necessarily indifferent; our requirement is that those actually affected by the swap have a strict preference for it. In particular, we are reluctant to assume that indifferent agents are willing to actively participate in the trade. In addition, Section 3 imposes additional structure on the outcome space under which this requirement is innocuous, in the sense that for any (f, g) with $f \succsim_i g$ for all $i \in N(f, g)$, there exists f' arbitrarily close to f with $f' \succ_i g$ for all $i \in N(f, g)$.

Our weaker notion of domination is defined as follows:

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

(i) f improves upon g;

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

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

Observe that our definition does not assume that the agents agree on the distributions of the alternatives f and g. The actual beliefs of the agents, determining their actual preferences, may be quite different.

Condition (i) of the definition requires that the agents involved prefer f to g according to their actual beliefs. Clearly, Condition (i) implies that f Pareto dominates g (recall that Condition (i) also implies that $N(f, g) \neq \emptyset$), ensuring that our notion is a refinement of Pareto domination. Thus, if one uses our stronger notion of Pareto dominance, $\succ_{\text{NBP}}$, rather than the standard one, one gets a larger set of Pareto optimal outcomes. In particular, the first welfare theorem still holds, though the second does not.

Condition (ii) requires that one be able to find a single probability measure, according to which all involved agents prefer swapping g for f. That is, one can find hypothetical beliefs, which, when ascribed to all relevant agents, can rationalize the preference for f over g. This second requirement fails in Example 2 above, where the proposed trade is simply a bet on different beliefs, but is satisfied in Example 3, where the proposed trade allows risk sharing. Two partners may thus invest in a business opportunity about which one is much more optimistic than the other, as long as there are some beliefs (say, of the more optimistic one) that justify the investment for both. Theorem 2 below shows that these are examples of a more general result, namely that this second criterion allows us to classify allocations in terms of the risk-sharing opportunities they provide.

2.2. *Illustration*

Much of our analysis can be illustrated by the following *Trade Example*. Suppose there two agents, A and B, and two states, 1 and 2. An alternative f is defined by the amount of money that each agent has in each state. The aggregate endowment is the same in the two states, that is, there is no aggregate uncertainty. The agents have identical utility functions u defined on their own wealth, that are strictly monotone and strictly concave. We can then depict alternatives in an Edgeworth box, as shown in Figure 1. The diagonal in Figure 1

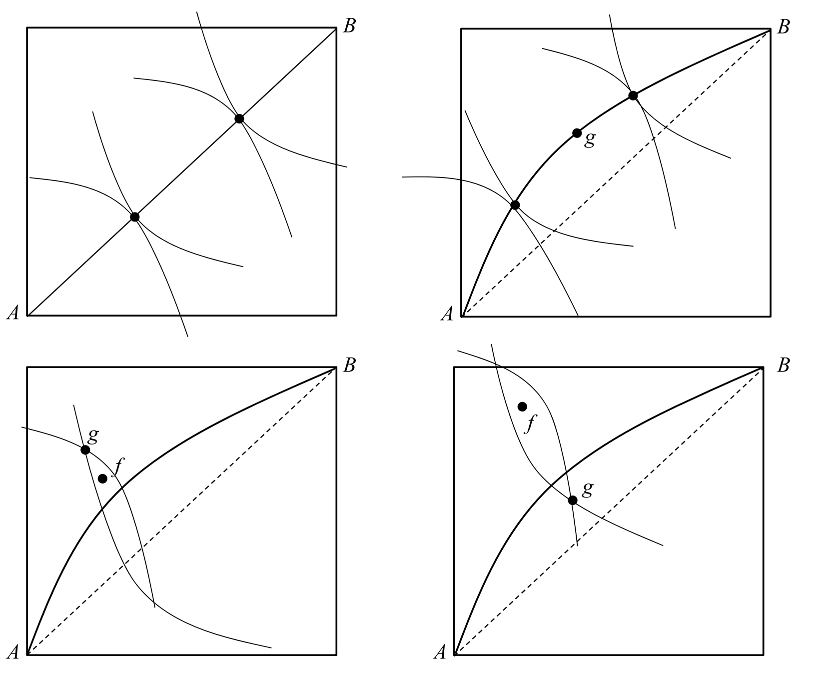

FIGURE 1.—The horizontal axis in each panel identifies the allocation of money between agents A and B in state 1, with the vertical axis doing the same for state 2. The total endowment is constant across states. If the agents have identical beliefs, then the set of Pareto efficient allocations coincides with the set of full-insurance allocations, and hence is given by the diagonal, as in the top left panel. If the agents' beliefs differ, then the Pareto efficient allocations lie on a curve that differs from the diagonal, as in the top right and bottom two panels. These three panels have identical beliefs and (sets of) Pareto efficient allocations. The two bottom panels show cases in which f Pareto dominates g.

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 by the agent's probability for the two states.

Condition (i) in the definition of No-Betting-Pareto dominance is the standard Pareto efficiency condition. Pareto efficiency is described by a curve, along which the slopes of the two agents' indifference curves are identical. The set of Pareto efficient allocations coincides with the full-insurance diagonal if the agents' beliefs are identical, as in the top left panel of Figure 1. If the agents entertain different beliefs, then no interior point on the diagonal can be Pareto efficient, as in the top right panel.

In the top right panel in Figure 1, alternative g is Pareto efficient and no f dominates it (in the standard Pareto sense, let alone in the No-Betting-Pareto refinement). The two bottom panels of Figure 1 show alternatives g that are dominated by other alternatives f. The distinction between these two panels is that, in the bottom left panel, f is closer to the diagonal than is g, whereas in the bottom right panel, f is farther from the diagonal than is g. The standard notion of Pareto domination does not distinguish between the two, but we now explain how our refined notion does.

Consider alternative g in the interior of the box, strictly above the diagonal, as in Figure 2. Let $P(g)$ be the interior of the triangle that is to the right and below g but above the diagonal, and let $Q(g)$ be the interior of the rectangle that is to the left and above g. We note the following:

PROPOSITION 1: *Every f in $P(g)$ satisfies Condition* (ii) *of Definition 3*, *while every f in $Q(g)$ does not*.

PROOF: Consider first $f \in P(g)$ as shown in the left panel. Choose p_0 to be the probability according to which f has the same expected value as does g for agent A, and hence also for agent B. Both agents would then be indifferent between f and g, were they risk neutral. Given that they are risk averse, and

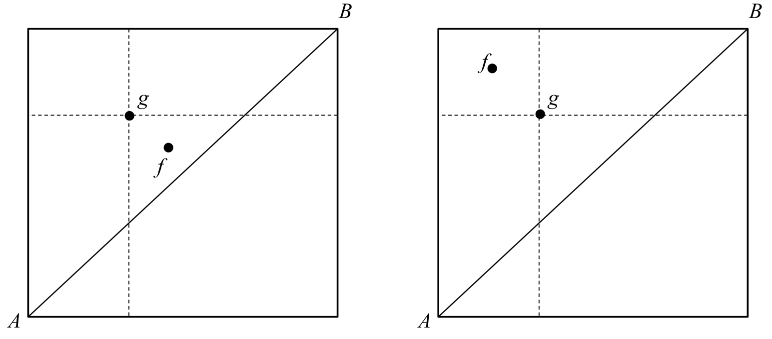

FIGURE 2.—Illustration of Condition (ii) of the definition of No-Betting-Pareto dominance. Condition (ii) is satisfied in the left panel, but fails in the right panel.

that g is a (p_0)-mean-preserving spread of f, both agents strictly prefer f to g under p_0, and hence Condition (ii) holds.

By contrast, consider $f \in Q(g)$ as shown in the right panel. Assume that Condition (ii) were to hold for the two agents with beliefs p_0. Given that these risk-averse agents strictly prefer f to g, despite the fact that f is a spread of g, such strict preferences would certainly hold were the agents expected value maximizers. However, for no p_0 can the difference $f - g$ simultaneously increase the p_0-expected value for both agents. *Q.E.D.*

Condition (ii) of Definition 3 thus divides the feasible trades (on the same side of the diagonal) into those that are "closer" to the full-insurance diagonal than is g, versus those that are "farther away" from the diagonal. Confronted with agents who prefer f to g, and thus satisfy Condition (i) of Definition 3, we can conclude that f No-Betting-Pareto dominates g if and only if f is closer to the full-insurance diagonal than is g. Moreover, we can make this assessment without knowing the agents' beliefs or indifference curves. In the bottom two panels of Figure 1, allocations above the curve of efficient allocations (see the left panel) are No-Betting-Pareto dominated (because allocations on the Pareto efficient curve, for example, are preferred by both agents and satisfy Condition (ii)). Allocations between the Pareto efficient curve and the diagonal (right panel) are not No-Betting-Pareto dominated (because the allocations satisfying Condition (ii) lie closer to the diagonal and are not preferred by both agents).

3. CHARACTERIZATIONS

In this section, we generalize the characterization of No-Betting-Pareto domination, developed for the Trade Example in Section 2.2, to arbitrary numbers of states and (heterogeneous) agents. The appropriate generalization is not obvious, as there is no longer an obvious definition of a "move toward full insurance." We offer two characterizations. The first is general, while the second is specific to an environment that contains the Trade Example of Section 2.2 as a special case.

3.1. *Combining Agents*

We use the Trade Example of Section 2.2 to introduce our first characterization. Consider Figure 3. The left panel shows an allocation g and an allocation $f \in P(g)$, for which there exists a p_0 with $\int_S u_i(f(s))\,dp_0 > \int_S u_i(g(s))\,dp_0$. Continue the line segment gf until it intersects the diagonal, say at a point f^*.

For $\lambda \in [0, 1]$, consider the "λ-averaged agent" whose utility in each state is the weighted average that places weight $\lambda \in [0, 1]$ on the utility of agent A in

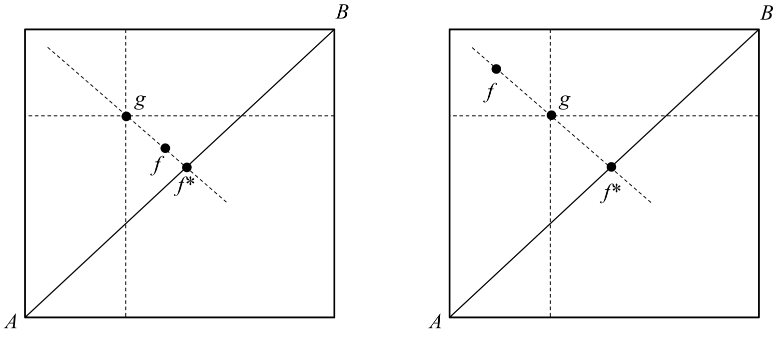

FIGURE 3.—Illustration of Theorem 1, in the setting of the Trade Example of Section 2.2. In the left panel, $f \succ_{\text{NBP}} g$, and the text explains why there is an agent who is better off in both states under f than under g. This is not possible in the right panel, where we do not have $f \succ_{\text{NBP}} g$.

that state and weight $1-\lambda$ on the utility of agent B in that state. Thus, if total endowment in each state is E and agent A has x in a particular state, the utility of the λ-averaged agent in that state is $\lambda u(x)+(1-\lambda)u(E-x)$, where u is the common utility function.

For every $x \in [0, E]$, there exists a λ such that x is the (unique) maximizer of $\lambda u(x)+(1-\lambda)u(E-x)$. (This λ is defined by $\lambda/(1-\lambda) = u'(E-x)/u'(x)$.) Let λ^* be the coefficient for which $f^*(1)$ (i.e., the amount consumed by A in state 1) in the left panel maximizes $\lambda u(x)+(1-\lambda)u(E-x)$. Observe that $f^*(2) = f^*(1)$ also maximizes this expression. It follows that, for this λ^*, the λ^*-averaged agent is better off at f^* than at g in both states $1, 2$.

Next we observe that, for every other λ, the λ-averaged agent is better off at f^* than at g in at least one of the states $1, 2$. Indeed, if $\lambda > \lambda^*$, this obviously will be the case in state 1, and for $\lambda < \lambda^*$, this will be the case in state 2. Finally, it then follows from the convexity of preference, or the concavity of u, that for any λ, the λ-averaged agent is also better off at f than at g in at least one state.

Consider now the right panel, where $f \in Q(f)$ and does not satisfy Condition (ii) of Definition 3. Continue the line segment fg until it hits the diagonal at a point f^*, and define λ^* as above. The λ^*-averaged agent is better off at f^* than at g in both states $1, 2$. Similarly, as $\lambda u(x)+(1-\lambda)u(E-x)$ is single-peaked, this agent is also better off at g than at f in both states. Hence, for this $\lambda = \lambda^*$, there does not exist a state in which the averaged agent is better off at f than at g.

The upshot of this discussion is that, in the Trade Example, f satisfies Condition (ii) of Definition 3 (relative to g) if and only if, for *every* λ, there exists a state s in which the λ-averaged agent is better off at f than at g. This turns out to be a general result, independently of the number of states and of agents and independently of the specification of X.

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

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

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

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

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

A convex combination of agents, λ, can be interpreted in two famously related ways. First, we may take a utilitarian interpretation, according to which $\sum_{i \in N} \lambda(i) u_i(\cdot)$ is a social welfare function defined by some averaging of the agents' utilities (cf. Harsanyi (1955)). Second, we may think of an individual behind the "veil of ignorance," believing that she may be agent i with probability $\lambda(i)$, and calculating her expected utility ex ante (cf. Harsanyi (1953)). In both interpretations, the condition states that not only the actual agents, but also all convex combinations thereof can justify the improvement by pointing to a state of the world that would make them at least as well off with the proposed improvement.

3.2. *Bets*

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

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

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

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

In this unidimensional setup, trade cannot be driven by differences in tastes, as all agents are assumed to want more of the only good. It then remains only to exclude risk sharing, and hence we can define betting as follows.

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

The requirement that g be independent of s (for all i) precludes the risk-sharing motivation, thereby justifying the definition of (f, g) as a bet. In the Trade Example of Section 2.2, an improvement (f, g) is a bet if and only if g lies on the diagonal. We can now state the following.

PROPOSITION 2: *If (f, g) is a bet, then it cannot be the case that $f \succ_{\text{NBP}} g$.*

Proposition 2 partly justifies the term "No-Betting-Pareto," as it shows that Condition (ii) of the definition of $\succ_{\text{NBP}}$ rules out Pareto improvements that are bets.

3.3. *Characterization of Excluded Improvements*

The No-Betting-Pareto criterion excludes Pareto improvements that are bets. What else does it exclude? This section develops another characterization of the second requirement in the No-Betting-Pareto notion. We can interpret this requirement as ranking allocations according to the extent of their risk sharing. Consider again the Trade Example of Section 2.2, illustrated in Figure 4. Recall that the set of Pareto efficient alternatives under common beliefs is the full-insurance diagonal, and that under any beliefs, any move from the diagonal to a point off the diagonal is a bet.

Consider the left panel of Figure 4. Suppose that an outside observer were to observe the trade of f for g. In this case, the outside observer has no evidence that the agents' beliefs differ. Indeed, the left panel is compatible with the agents' beliefs being identical. But this is not the case in the right panel. In the right panel, f can be an improvement over g (with convex preferences) only if agent A's indifference curves along the diagonal are flatter than B's indifference curves (because A thinks state 2 is more likely than does B). Indeed, agent A's indifference curve must be flatter at any point on the diagonal than at g, and agent B's indifference curve must be steeper, and hence the trade $f - g$, acceptable by both agents at g, indicates a direction of change that

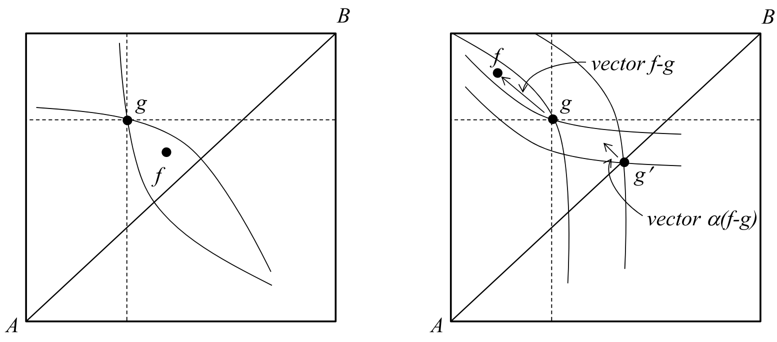

FIGURE 4.—Motivation for Theorem 2, in the setting of the Trade Example of Figure 1. Movement from g to f pushes the agents toward full insurance in the left panel, but in the right panel exacerbates the bet that brings the agents from the full-insurance allocation g' to g.

would certainly be acceptable by both agents at any full-insurance allocation g' along the diagonal. More precisely, at *any* point g' in the interior of the diagonal, the trade $\alpha(f-g)$ would be a strict improvement for both agents for some positive α. Thus, a distinguishing feature of the trade $f-g$ in the right panel is that, armed only with the information that the agents are willing to make the trade, a designer can construct a book against them, that is, can formulate a trade that both agents will strictly prefer to accept and that yields the designer a certain profit, *starting at any full-insurance allocation* g'.

This comparison reflects a key distinction between Pareto improvements that are NBP-improvements and Pareto improvements that are not: the former do not reveal differences in beliefs. In particular, if the agents trade, and keep trading until they are both fully insured, no information would have been revealed that an observer could use to offer the agents to bet against each other. By contrast, the Pareto improvements that do not satisfy our criterion are precisely those that do reveal such information. When the agents trade as in the right panel, even if by some process they manage to be fully insured, a bookie who observed the trade can then come up with a bet that they would both accept. Specifically, if the bookie offers financial assets corresponding to the functions $(f_i(s)-g_i(s))_i$, there will be a positive amount that will be demanded by each agent respectively. Since their preferences are strict, this also means that the bookie can make a sure profit by offering the agents their part of the bet, deducting a fee for herself.

In the Trade Example, No-Betting-Pareto domination can be precisely characterized as prohibiting Pareto improvements that move the agents further away from full insurance. As we have noted, the concept of a movement away from full insurance does not readily generalize to settings with many agents, many states, and heterogeneous utility functions. However, we have just established that in the Trade Example, there is an equivalence between the existence of an improvement that moves the agents away from full insurance and the agents being willing to abandon any full-insurance allocation in favor of a bet. The existence of a single bet acceptable to the agents at any full-insurance allocation is a well-defined condition in the general setup. We now show that the Pareto improvements satisfying this condition are precisely those excluded by the No-Betting-Pareto criterion.

THEOREM 2: *Let there be given utilities* $(u_i)_i$ *and two alternatives* $f, g \in F$. *The following are equivalent*:

(i) *There does* not *exist a probability vector* p_0 *such that*, *for all* $i \in N(f, g)$,

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

(ii) *There exists an alternative* $d \in F$ *satisfying*

$$\sum_{i\in N} d(s)_i = 0 \quad \forall s \in S,$$

that also has the following property: *for every* $g' \in F$ *such that* $g'(s)_i$ *is independent of* s *for each* $i \in N(f, g)$ *and lies in the interior of* R, *and for every beliefs* $(p_i)_i$ *such that* $f \succ_* g$, *there exists* $\alpha > 0$ *such that* $(g' + \alpha d, g')$ *is a bet* (*for the utilities* $(u_i)_i$ *and the beliefs* $(p_i)_i$).

This theorem thus provides another characterization of the pairs (f, g) that satisfy Condition (ii) of Definition 3. This time we state the result in terms of the *failure* to satisfy Condition (ii): this condition does not hold (clause (i) of the theorem) if and only if the following is true: given the fact that the relevant agents prefer f to g, and given their utility functions, one may start at any state-independent allocation g', and construct $f' \equiv g' + \alpha d$ such that (f', g') is a bet. Note that in this case, since bets are defined by strict preferences, we have $|N(f', g')| \geq 1$, and the feasibility constraint and the assumption that all utilities are monotone then ensure that $|N(f', g')| \geq 2$. However, $N(f', g')$ may be a proper subset of $N(f, g)$.

The constructed bet d need not, in general, be a multiple of the trade $(f - g)$ that is preferred by both agents at g: with many states, it is possible that some agents prefer f to g despite the risk that the former entails, because this risk occurs at states where their wealth is high, but they may not be willing to take the same risk starting from another allocation g'.

At first glance, there is no surprise that, if we have $f \succ_* g$ but not $f \succ_{\text{NBP}} g$, one can construct a bet that fully insured agents would take. Indeed, if we have $f \succ_* g$ but not $f \succ_{\text{NBP}} g$, we know that some agents have different beliefs, and it suffices to have two agents with different beliefs to be able to find a bet that they would be willing to take. However, finding such a bet would require one to know who are the two agents whose probabilities differ, and what these probabilities are, and the bet constructed would typically depend on these agents and on their beliefs. By contrast, Theorem 2 guarantees the existence of a bet independently of the identity of the agents and of their probabilities: a bookie may offer the agents to buy assets corresponding to their share in d, and these assets are independent of the beliefs (and of g'). Whatever are the probabilities p_i's, as long as they are known to lie in the respective half-spaces (so that p_i makes agent i prefer f to g), the agents would strictly prefer to bet by swapping g' for $f' = g' + \alpha d$ for some $\alpha > 0$. While the volume of trade α may depend on the beliefs, the financial assets d are uniformly constructed for all p_i's such that $f \succ_* g$.[15]

4. DISCUSSION

4.1. *Pareto Rankings*

When uncertainty is considered, it is tempting to model the state of the world as one of the features of a good and analyze an economy with uncertainty as one analyzes an economy with more goods but no uncertainty (see Debreu (1959)). This expands the scope of the welfare theorems to economies with uncertainty. One may then argue that under uncertainty as in the case of certainty, complete and competitive markets have the advantage of guaranteeing Pareto efficient allocations, whereas incomplete or regulated markets run the risk of resulting in Pareto inefficient equilibria.

However, we argue that trade in financial markets has a strong speculative component, fueled by differences in beliefs across agents, and that welfare analysis should be revisited in such contexts. For example, suppose that we are discussing the introduction of a new financial asset, currently not traded in the market. While it is tempting to assume that such a new market can only make agents better off, as it brings the economy closer to the complete markets ideal, it has already been pointed out by Hart (1975) that adding assets to an incomplete market need not lead to Pareto improvements. Our point is, however, different: once we question the notion of Pareto domination, it is no longer self-evident that we should seek Pareto improvements, even if we can ensure them. In particular, if the inferior allocation is only dominated by allocations that are generated by pure bets, one may well prefer the inferior allocation to the supposedly dominating ones.

The growing complexity of financial assets makes the market completeness question quite relevant. Using complex financial derivatives, one may trade on not only the fundamentals of the economy, but also arbitrarily complex combinations of fundamentals, assets, beliefs about fundamentals, beliefs about

[15] The proof of Theorem 2 is constructive. One can therefore imagine a scenario in which automated bookies, observing swaps such as (f, g), are quick to offer bets such as (f', g'). This observation fortunately falls short of offering a practical method for making profit by exploiting differences in belief. Yet, it may help us see why some Pareto improvements are not as compellingly desirable as others.

assets, and so on. Those who trade futures contracts are expressing their beliefs about the fundamentals of the economy. Those who trade options on future contracts are expressing their beliefs about others' beliefs. One could, in principle, combine these options to construct assets that depend upon beliefs about beliefs about others' beliefs, and so on. The chain of such possibilities has no end, ensuring that a given (finite) set of financial markets must be incomplete. Nonetheless, it is not obvious that adding more sophisticated assets to the market will be helpful, especially those that simply allow opportunities to bet.

Should financial markets be regulated, as suggested by Posner and Weyl (2012)? Without taking a stance on desired policy, our goal here is only to refine a theoretical argument. We claim that one standard argument for free markets, namely, that only complete and free markets are guaranteed to lead to Pareto efficient allocations, does not apply in this context without an appropriate qualification.

4.2. *A Comparison With Brunnermeier, Simsek, and Xiong*

Brunnermeier, Simsek, and Xiong (2012) (hereafter BSX) also developed a welfare criterion for markets in the presence of individuals who might entertain different beliefs.[16] This section offers a comparison between their model and ours.

BSX offered two criteria: one is utilitarian in spirit, and the other has the flavor of Pareto efficiency. The first, the *Expected Social Welfare* criterion, is based on the expected utility of an "average" agent, namely, a hypothetical agent whose utility function is $\sum_{i=1}^{n} \lambda_i u_i$ for a set of nonnegative weights $\lambda_1, \dots, \lambda_n$. Given differences in beliefs, it is not obvious which probability measure should be used for the calculation of expected average utility. BSX's criterion solves this problem by a unanimity approach, requiring that one alternative be superior to another given the belief of each agent, or, equivalently, for each probability in the convex hull of the agents' beliefs. Thus, the expected social welfare criterion of BSX uses an aggregation over agents' utility functions, and a unanimity rule for the agents' beliefs.

When beliefs differ, the expected social welfare criterion is logically independent of No-Betting-Pareto domination. One direction of this comparison is familiar, and has nothing to do with beliefs, instead reflecting the basic difference between social welfare functions and Pareto comparisons. Once one has committed to welfare weights $\lambda_1, \dots, \lambda_n$, presumably in order to obtain a complete ranking of alternatives, one may have attached a higher welfare to f than to g without there being Pareto domination, let alone No-Betting-Pareto domination between them. The other direction illustrates the differing role

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

of beliefs. The statement that f No-Betting-Pareto dominates g requires that there is at least one belief at which all agents would prefer f to g, but this need not be true for all probabilities in the convex hull of the agents' beliefs. (In fact, it may not hold for any probability in this convex hull, as discussed in the following subsection.)

BSX offered another approach that does not require a specification of the welfare weights and hence is closer in spirit and offers a more obvious comparison to our notion. They defined f to be *belief-neutral inefficient* if, for every belief p in the convex hull of the agents' 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, so that alternatives are sorted into three sets (efficient, inefficient, and unclassified).[17]

These possibilities reflect some conceptual differences, concerning (i) whether a concept of domination should insist on retaining Pareto efficiency, (ii) whether the relevant ranking is required to hold for one or for all beliefs, and (iii) the set for which the relevant beliefs are drawn. First, an alternative g is belief-neutral inefficient only if it is dominated according to each of the agents' beliefs, but belief-neutral inefficiency allows different alternatives to dominate g for different beliefs, opening the possibility that an alternative can be belief-neutral inefficient without being Pareto dominated. In contrast, No-Betting-Pareto dominance requires (and refines) Pareto domination. Second, to be belief-neutral inefficient, an allocation f must be dominated for every belief in the set of "reasonable beliefs" (though not by the same alternative g in each case), while No-Betting-Pareto dominance requires dominance only for a single (shared) belief. Finally, belief-neutral inefficiency identifies the set of reasonable beliefs as the convex hull of the agents' beliefs, while No-Betting-Pareto dominance takes all possible beliefs to be fair game.

These conceptual differences reflect the differing motivations of the two concepts. BSX motivated their notion with the idea that the agents' beliefs may be incorrect, perhaps because they have made mistakes in processing information, but the planner does not know the correct belief. Having recognized that the agents' beliefs may be incorrect, the planner attaches no particular prominence to those beliefs, and hence is willing to conclude that f is belief-neutral

[17] Brunnermeier, Simsek, and Xiong (2012, Proposition 1) made precise the sense in which their belief-neutral efficiency notion frees the expected social welfare criterion from dependence on a single set of welfare weights.

inefficient without being able to identify an alternative g that dominates f under the agents' actual beliefs. Because the planner does not know the agents' beliefs, the planner insists on rankings that are robust, to the extent that they hold for all reasonable beliefs.

No-Betting-Pareto dominance takes a more agnostic approach to agents' beliefs and a more cautious approach to dominance. We are unwilling to identify an alternative f as dominated without being able to identify a single alternative g that dominates it, and are unwilling to insist that an agent should find g better than f if the agent's own beliefs do not support such a comparison. Instead, we are interested in identifying dominance relations that reflect spurious unanimity, in the sense that they can arise only because agents hold different beliefs. We would also like to build robustness into our test for spurious unanimity, meaning that we would like to be independent of the agents' actual (perhaps unobservable) beliefs. We thus identify a ranking as reflecting spurious unanimity, and disqualify it from our dominance relation, if there exists no single belief that could motivate the ranking.

The differences between the notions of belief-neutral efficiency (and inefficiency) and No-Betting-Pareto are illustrated by the following possibilities:

1. An allocation g may be belief-neutral inefficient, but not NBP-dominated, as there may be no single f that Pareto dominates g according to the agents' original beliefs.

2. Turning this around, an allocation g may be NBP-dominated by an allocation f, but not belief-neutral inefficient, as g may not be dominated by any f' according to some beliefs in the convex hull of the agents' beliefs.

3. An allocation g may be belief-neutral efficient, but f may No-Betting-Pareto dominate g because there is a belief outside the convex hull of the agents' beliefs under which all prefer f to g. (We return to this point below.)

4. An allocation g may be NBP-efficient, that is, undominated by any f, but not belief-neutral efficient, because g may be dominated by an allocation f for some (hypothetically shared) belief p, but that f does not Pareto dominate f according to the agents' original (and different) beliefs.

4.3. *The Range of Beliefs*

The third difference between the belief-neutral and No-Betting-Pareto notions, identified but not discussed in Section 4.2, is the range of the beliefs involved. To argue that $f \succ_{\text{NBP}} g$, one may use any belief whatsoever when showing that there exists a belief under which all agents prefer f to g. In contrast, BSX restricted attention to beliefs in the convex hull of the agents' beliefs. Our motivation for considering all beliefs is that, once differences in beliefs have been detected, it is clear that some agents are wrong, and then we cannot dismiss the possibility that they are all wrong. Consider the following example.

EXAMPLE 4: Suppose that f and g are given by the following:

			f		g	
State	p_A	p_B	U_A	U_B	U_A	U_B
1	.4	.3	2	1	1	2
2	.3	.325	2	2	2	1
3	.3	.375	1	2	2	2

where p_A and p_B are the probabilities Amy and Bruce attach to the various states, and U_A and U_B are their payoffs in these states. Amy and Bruce both prefer f to g. Indeed, Amy prefers f to g iff $p_A(1) > p_A(3)$ and Bruce prefers f to g iff $p_B(2) > p_B(1)$. Clearly, there are beliefs that satisfy both inequalities, for example, $(.325, .375, .3)$. However, no such belief is in the convex hull of Amy and Bruce's beliefs, as both p_A and p_B agree that state 3 is at least as likely as state 2.

The possibility illustrated by Example 4 motivates our reluctance to restrict attention to some particular half-space of beliefs, simply because it includes all of the agents' beliefs. One agent has the belief $(.4, .3, .3)$ and the belief $(.3, .325, .375)$. Does this mean that we can exclude the possibility that the appropriate belief is $(.325, .375, .3)$?

On the other hand, restricting attention to the convex hull of beliefs (as did BSX) makes the model more robust to irrelevant states of the world. The following example illustrates.

EXAMPLE 5: Suppose there are three states. Abigail and Bart each have an endowment of two units of income in each state. Abigail and Bart are both expected income maximizers. They consider investing their endowment in a joint product giving net returns of $(3, 0)$ in state 1, $(0, 3)$ in state 2, and $(4, 4)$ in state 3. Abigail believes state 1 will surely occur, while Bart believes state 2 will surely occur, and hence each is willing to undertake the investment—the outcome of the investment Pareto dominates the endowment. It also No-Betting-Pareto dominates the endowment. There are many beliefs, one of which is that the three states are equally likely, under which both agents would be willing to undertake the investment.

In Example 5, while there are many beliefs that would induce both agents to engage in the trade, none of them lies in the convex hull of the agents' beliefs. The Pareto criterion of BSX thus does not rank the two alternatives, while No-Betting-Pareto does. To defend the latter ranking, one may note that Abigail is absolutely sure that state 1 will occur, whereas Bart is willing to stake his life on state 2 being the case. An outside observer has to raise an eyebrow and say that at least one of them is wrong, and, to be on the safe side, it seems reasonable to doubt their beliefs—including what they happen to agree upon (namely, that state 3 is impossible). Hence, it seems cautious to allow NBP dominations even if they use probabilities that assign a positive weight to state 3.

However, this raises a potential modeling difficulty. Suppose that we consider a model in which f Pareto dominates g but not NBP-dominates it. Let us add to the model another state (such as state 3 in Example 5), according to which alternative f guarantees all agents eternal bliss. We can view the agents in the original model as assigning this state probability zero. The BSX notions, as well as Gayer's notion, which only consider the convex hull of the agents' beliefs, will be robust to the inclusion of this new state in the model, but NBP domination will not: given the new state, f now NBP dominates g, whereas it did not before this state was introduced. It follows that our definition of NBP domination requires some care in defining the state space, and that we must take seriously the seemingly innocuous assumption that the agents agree on the state space.

Notice the implicit tension here—we assume that we can identify the state space, while assuming either that we cannot identify the agents' beliefs or that these beliefs provide no clues as to what beliefs with which we should work. We must have some common structure if we are to attempt welfare comparisons at all, but we think it unrealistic that agents will have identical and observable views of all aspects of their environment. We expect the analyst to more readily find a meaningful consensus on states than on probabilities, a situation we model by assuming a common state space while placing no restriction on applicable beliefs.

4.4. *Implications*

The main goal of this paper is to make a theoretical contribution to the debate about free markets, and, in particular, to highlight a theoretical difference between two different interpretations of the same mathematical model. Our analysis relies on highly idealized assumptions about the agents' rationality, their common conception of a state space, and so forth, and we therefore prefer to be cautious in suggesting policy recommendations.

The above notwithstanding, it may prove useful to ask if and how our analysis might affect policy decisions. Suppose, perhaps partly as a result of the financial crisis of 2007/8, lawmakers wish to change or tighten the regulation of financial markets. They need to deal with many issues that are not addressed by this paper, including incomplete information and incentive problems, complexity of financial assets, and so forth. However, they might also wonder whether, even in the simplest of cases, free trade in financial assets is necessarily to be endorsed. We suggest that they might reasonably consider limiting trade in financial assets if this trade comes too close to pure betting, and that the

No-Betting-Pareto criterion is a reasonable starting point for identifying such trades.

How can this notion be applied? We can imagine a scenario in which proposed mergers, acquisitions, or financial swaps either must be approved before executed, or must be justified in response to audits. We can further imagine that in order to approve a proposal, the monitoring authority ascertains not only that each party views the proposal as beneficial, in the sense that they enter it willingly, but also that the parties can present (at least post hoc) 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. In this exercise, one may assume that the utility functions, defined on monetary outcomes, are predetermined to be in a given class, such as CARA or CRRA, with parameters calibrated to reflect the risk aversion appropriate for the relevant party.[18] We imagine the regulatory guidelines specifying the relevant functions and parameter ranges. The basic point here is that the parties to trade should be able to point to a shared probability according to which they all benefit from trade.

A scenario in which CFOs appear before a judge or regulator and compute integrals might seem a bit outlandish, but is not too far removed from the type of analysis performed when firms make cases for merger approval or respond to antitrust allegations. When the deals proposed are large enough to threaten the solvency of major financial institutions, such a scenario might be more realistic than it first appears.

To get an idea of how this might work, suppose that a pension fund and an investment bank currently have positions described by an alternative g, and consider a trade that would bring them to the allocation f. Both are willing to trade, signaling to a regulator that both prefer f to g, but the regulator will, in general, know nothing more about their preferences. The regulator can then ask whether, given the appropriate utility functions, there exists a shared belief under which both parties gain from the exchange. If the situation were as simple as the 2×2 case illustrated in Figure 2, this would be trivial, and would require confirming only whether f lies closer to the diagonal than does g. In practice, we imagine the transactions will be quite complicated and the outcome will accordingly be less obvious. The regulator must then frame the assessment in terms of shared beliefs, with approval hinging upon whether the No-Betting-Pareto criterion is satisfied, that is, on whether there exists a shared belief under which all parties would be willing to trade.

What will be the outcome of this process? Much will depend upon the status quo allocation and the process by which alternative allocations come into

[18]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.

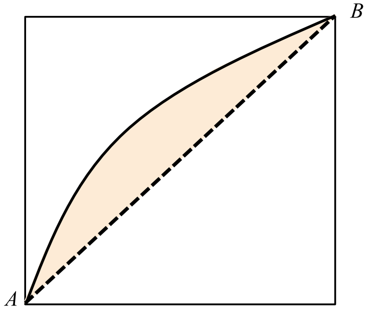

FIGURE 5.—The shaded lens, bounded by the full-insurance allocations (the diagonal) and the Pareto efficient allocations (the upper curve) is the set of points that are not No-Betting Pareto dominated.

consideration. Consider the 2×2 case, now illustrated in Figure 5. Everything outside the shaded lens is NBP-dominated, because there is room for risk sharing, and the NBP-Pareto criterion will endorse movements toward the lens. It is then reasonable to suppose that a process of proposing and evaluating allocations will bring us into this lens. Once there, there are no additional No-Betting-Pareto improvements. If we find ourselves at an allocation near the diagonal, there will be (nearly) full risk sharing, and the No-Betting-Pareto criterion will prevent Pareto-improving movements toward the curve of Pareto-efficient points that exhibit more speculation. However, there will be no pressure to move away from the latter curve, precisely because its constituents are Pareto efficient.

We find the case for regulating trade, with restrictions on trade that is equivalent to betting, to be particularly compelling when it comes to institutions that manage other people's money. Though we have explained in Section 1.3 why we do not hold such a view, we can imagine someone arguing that consumer sovereignty should apply to people's beliefs just as it does to their tastes, and hence, people who knowingly enter trades based on different beliefs should be allowed to do so. We cannot imagine the same argument being made on behalf of a financial institution to whom people have delegated their financial decisions. For example, we would find it troubling that a pension fund would offer its members a return that is only half as high as the market's on the strength of the argument that the fund, on behalf of those investors, has adopted a belief that puts considerably more weight on favorable outcomes than does the market. While our paper is not intended to offer practical solutions to such problems, we do believe that the concept of No-Betting-Pareto may be a first step in thinking about how to make trade in financial markets more responsible than allowed by the standard general equilibrium model.

4.5. *An Extension*

Our definition of a bet (f, g) assumes that the given allocation, g, is constant across the state space. This is obviously restrictive. For example, assume that two agents are considering a bet on the outcome of a soccer match. It so happens that their current wealth does not depend on this match in any way. Yet, their current allocations are far from constant, as the two are exposed to various risks, ranging from their health to stock market crashes. Thus, Proposition 2 does not apply to such a swap (f, g). Yet, it would be nice to know that, for such swaps, f does not No-Betting-Pareto dominate g.

To capture this type of exchange in the definition of a bet, and correspondingly generalize Proposition 2, one has to allow the existing allocations g to depend on s, but to be independent of the exchange. That is, the variable $f - g$ should be stochastically independent of g according to all the probability measures considered. In other words, one may assume that the state space is a product of two spaces, $S = S_1 \times S_2$ such that g is measurable with respect to S_1, and consider only probabilities obtained as a product of a measure p_1 on S_1 and a measure p_2 on S_2. Relative to such a model, ours can be viewed as a reduced form model, where our entire discussion is conditioned on a state $s_1 \in S_1$.

APPENDIX: PROOFS

A.1. *Proof of Theorem 1*

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

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

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

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

$$\sum_{j \le J} p_0(j) u_i\bigl(g(A_j)\bigr) > \sum_{j \le J} p_0(j) u_i\bigl(g(A_j)\bigr). \tag{3}$$

In particular, if a measure p_0 that satisfies (2) exists, it induces a probability vector $(p_0(j))_{j \le J}$ (over $(A_j)_{j \le J}$) that satisfies (3). Conversely, if a vector $(p_0(j))_{j \le J}$ satisfying (3) exists, it can be extended to a measure p_0 on (S, Σ)

such that (2) holds. (Since f and g are constant over each A_j, the choice of the extension does not matter.)

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Q.E.D.

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

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

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

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

A.2. *Proof of Proposition 2*

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

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

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

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

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

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

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

and, since u is concave,

$$u_i\big(E_{p_0}(f_i)\big) \geq E_{p_0}\big(u_i(f_i)\big);$$

thus

$$u_i\big(E_{p_0}(f_i)\big) > \bar{u}_i,$$

and, because u is strictly monotone,

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

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

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

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

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

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

Because $\sum_i f(s)_i \le \sum_i g(s)_i$ for all s, we also have

$$\sum_{i \in N(f,g)} f(s)_i \le \sum_{i \in N(f,g)} g(s)_i$$

and it follows that the λ-weighted utility under f cannot exceed that corresponding to g. Thus, the λ-weighted "average" agent cannot point to a state where she is strictly better off under f than under g. *Q.E.D.*

A.3. *Proof of Theorem 2*

Let utilities $(u_i)_i$ and alternatives f, g be given. We first observe that (ii) implies (i). Indeed, if (i) does not hold, then there exists p_0 such that

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

for all $i \in N(f, g)$. Setting $p_i = p_0$, we obtain beliefs for which $f \succ_* g$. However, for any $g' \in F$ such that $g'(s)_i$ is independent of s (for each $i \in N(f, g)$), and any d with $\sum_{i \in N} d(s)_i = 0$ (for all s), defining $f' = g' + d$, it cannot be the case that all agents in $N(f', g')$ strictly prefer f' to g' because the agents are (weakly) risk averse. Thus, there is no bet (f', g') that can be constructed for the beliefs $(p_i)_i$, contrary to (ii).

We now turn to the main part of the theorem, namely, that (i) implies (ii). Assume, then, that (i) holds, that is, that there does not exist a probability p_0 such that

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

for all $i \in N(f, g)$. We need to construct an alternative d such that $\sum_{i \in N} d(s)_i = 0$ (for all s) and, whenever the beliefs $(p_i)_i$ imply $f \succ_* g$, then, for every g'

(in the interior of the diagonal) there exists $\alpha > 0$ such that $f' \succ_* g'$, where $f' = g' + \alpha d$.

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

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

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

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

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

We mention the following:

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

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

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

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

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

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

so that, for all k, $C_i^k \subset C_i^{k+1} \subset C_i$ and $C_i = \bigcap_k C_i^k$. Since we have $\bigcap_{i\in N} C_i = \emptyset$, it is certainly true that $\bigcap_{i\in N} C_i^k = \emptyset$ for all k. However, C_i^k is a nonempty, convex, and (as opposed to C_i) also compact subset of Δ^{J-1}. When such compact and convex sets of priors have an empty intersection, it is known that one can find a bet that they would all accept, as long as their beliefs are in the specified sets of priors. Specifically, Theorem 2 in Billot, Chateauneuf, Gilboa, and Tallon (2000, p. 688) states that there are linear functionals h_i, such that h_i is strictly positive on C_i^k, and $\sum h_i = 0$.[19] Thus, for each k, there exists an $n \times J$ matrix, $(h_{i,j}^k)_{i,j}$, of real numbers, such that

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

and

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

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

STEP 2: We now wish to show that the construction of f'' above can be done in a uniform way: there exists an f'' such that $\sum_i f''(j)_i = 0$ and $\sum_i p_i(j) f''(j)_i > 0$ for all $p_i \in C_i$ and all $i \in N(f'', g')$. (Observe, however, that while in Step 1 we obtained a bet that involved all agents, here we may find that $N(f'', g') \subsetneq N$.)

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

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

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

and

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

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

CLAIM 2.1: $H \neq \emptyset$.

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

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

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

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

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

$$D(h) = \left\{i \in N \middle| \exists j, h_{i,j} \neq 0\right\}.$$

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

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

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

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

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

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

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

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

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

It follows that h'' can be renormalized to guarantee (5) without violating (4), (6), obtaining $h' \in H$ with $D(h') \subsetneq D(h)$. *Q.E.D.*

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

PROOF: Combine Claims 2.2 and 2.3. *Q.E.D.*

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

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

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

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

$$\begin{aligned} &\operatorname*{Min}_p \sum_j p(j)h_{i,j} \\ &s.t. \\ &\quad \sum_{j\leq J} p(j)\big(u_i\big(f(A_j)\big) - u_i\big(g(A_j)\big)\big) \geq 0, \qquad (7) \\ &\quad p \in \Delta^{J-1}. \end{aligned}$$

Since $h \in H$, the optimal value of this problem is nonnegative. Since $\sum_j p(j)h_{i,j} = 0$, p is a solution to the problem. However, because $p \in C_i$, constraint (7) is inactive at p. Given that this is a linear programming problem, removing an inactive constraint cannot render p sub-optimal. Hence p is also an optimal solution to $\operatorname{Min}_p \sum_j p(j)h_{i,j}$ subject to $p \in \Delta^{J-1}$. But this implies that $\sum_j p(j)h_{i,j} \geq 0$ for all $p \in \Delta^{J-1}$. This, in turn, implies that $h_{i,j} \geq 0$ for all j, contradicting Claim 2.3. *Q.E.D.*

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

STEP 3: Finally, consider an act $f'_\alpha = \alpha d$ for $\alpha > 0$. Clearly, $\sum_i f'_\alpha(j)_i = 0$ for all j and all α. As u_i are differentiable, for a small enough α the conclusion $\sum_i p_i(j)u_i(f'_\alpha(j)_i) > 0$ follows. *Q.E.D.*

REFERENCES

ARROW, K. J., AND G. DEBREU (1954): "Existence of an Equilibrium for a Competitive Economy," *Econometrica*, 22, 265–290. [1406]

AUMANN, R. J. (1976): "Agreeing to Disagree," *The Annals of Statistics*, 4, 1236–1239. [1409]

——— (1987): "Correlated Equilibrium as an Expression of Bayesian Rationality," *Econometrica*, 55, 1–18. [1409,1410]

BEWLEY, T. (1989): "Market Innovation and Entrepreneurship: A Knightian View," Discussion Paper 905, Cowles Foundation for Economic Research. [1438]

BILLOT, A., A. CHATEAUNEUF, I. GILBOA, AND J.-M. TALLON (2000): "Sharing Beliefs: Between Agreeing and Disagreeing," *Econometrica*, 68, 685–694. [1438]

BLUME, L. E., T. COGLEY, D. A. EASLEY, T. J. SARGENT, AND V. TSYRENNIKOV (2013): "Welfare, Paternalism, and Market Incompleteness," Working Paper. [1412]

BRUNNERMEIER, M. K., A. SIMSEK, AND W. XIONG (2012): "A Welfare Criterion for Models With Distorted Beliefs," Working Paper. [1412,1413,1426,1427]

DEBREU, G. (1959): *Theory of Value*. New Haven: Yale University Press. [1425]

GAYER, G. (2013): Personal Communication. [1412]

GILBOA, I., D. SAMET, AND D. SCHMEIDLER (2004): "Utilitarian Aggregation of Beliefs and Tastes," *Journal of Political Economy*, 112, 932–938. [1414]

HARSANYI, J. C. (1953): "Cardinal Welfare in Welfare Economics and in the Theory of Risk Taking," *Journal of Political Economy*, 61, 434–435. [1421]

——— (1955): "Cardinal Welfare, Individualistic Ethics, and Interpersonal Comparison of Utility," *Journal of Political Economy*, 63, 309–321. [1414,1421]

HART, O. (1975): "On the Optimality of Equilibrium When the Market Structure Is Incomplete," *Journal of Economic Theory*, 11, 418–443. [1425]

HYLLAND, A., AND R. ZECKHAUSER (1979): "The Impossibility of Bayesian Group Decision Making With Separate Aggregation of Beliefs and Values," *Econometrica*, 47, 1321–1336. [1414]

HEAL, G., AND A. MILLNER (2013): "Uncertainty and Decision in Climate Change Economics," Working Paper. [1409]

KREPS, D. M. (1988): *Notes on the Theory of Choice*. Underground Classics in Economics. Boulder: Westview Press. [1410]

——— (2012): *Microeconomic Foundations I: Choice and Competitive Markets*. Princeton: Princeton University Press. [1413]

MILGROM, P., AND N. STOKEY (1982): "Information, Trade, and Common Knowledge," *Journal of Economic Theory*, 26, 17–27. [1409]

MONGIN, P. (1995): "Consistent Bayesian Aggregation," *Journal of Economic Theory*, 66, 313–351. [1414]

——— (1997): "Spurious Unanimity and the Pareto Principle," Working Paper, Theorie Economique, Modelisation et Applications, Université de Cergy-Pontoise and Centre National de Recherche Scientifique. Presented at the Conference on Utilitarianism, New Orleans. [1413]

POSNER, E. A., AND G. WEYL (2012): "An FDA for Financial Innovation: Applying the Insurable Interest Doctrine to 21st-Century Financial Markets," Working Paper. [1413,1426]

SAMET, D. (1998): “Common Priors and Separation of Convex Sets,” *Games and Economic Behavior*, 24, 172–174. [1438]
SAVAGE, L. J. (1954): *The Foundations of Statistics*. New York: Wiley. (Second edition in 1972, New York: Dover.) [1408,1410]
SIMSEK, A. (2012): “Speculation and Risk Sharing With New Financial Assets,” Working Paper. [1426]
STIGLITZ, J. (1989): “Using Tax Policy to Curb Speculative Trading,” *Journal of Financial Services Research*, 3, 101–115. [1413]
VON NEUMANN, J., AND O. MORGENSTERN (1944): *Theory of Games and Economic Behavior*. Princeton: Princeton University Press. [1414]
WEYL, G. (2007): “Is Arbitrage Socially Beneficial?” Working Paper. [1413]

Dept. of Economics and Decision Sciences, HEC Paris, 78351 Jouy-en-Josas Cedex, France and Berglas School of Economics, Tel Aviv University, Tel Aviv 6997801, Israel; tzachigilboa@gmail.com,

Dept. of Economics, Yale University, 30 Hillhouse Avenue, New Haven, CT 06520, U.S.A.; larry.samuelson@yale.edu,

and

School of Economics, IDC Herzliya, Israel and School of Mathematical Sciences, Tel Aviv University, Tagore 40, Tel Aviv 69203, Israel; davidschmeidler@gmail.com.

Manuscript received December, 2012; final revision received March, 2014.