A Case for Incomplete Markets[*]

Lawrence E. Blume[†]   Timothy Cogley[‡]   David A. Easley[§]

Thomas J. Sargent[¶]   Viktor Tsyrennikov[‖]

June 10, 2015

Abstract

We propose a new welfare criterion that allows us to rank alternative financial market structures in the presence of belief heterogeneity. We analyze economies with complete and incomplete financial markets and/or restricted trading possibilities in the form of borrowing limits or transaction costs. We describe circumstances under which various restrictions on financial markets are desirable according to our welfare criterion.

Keywords: social welfare, heterogeneous beliefs, spurious unanimity, speculation, pessimism, incomplete markets, financial regulation

1 Introduction

A conventional wisdom in the economics profession is that complete markets are good. The welfare theorems state that complete markets outcomes are

Pareto optimal and that any optimal allocation can be realized by trade in complete markets with an appropriate lump-sum transfer scheme. So putting limits on trade, by foreclosing trading opportunities, leaves potential mutual gains unrealized. This "wisdom" has practical consequences. Arguments for the privatization of social security and against the regulation of financial markets rely in part on the assertion that barriers to trade are bad things.

Complete markets have their critics. Some say that traders have market power and that their exploitation can be limited only by constraining trade. Others argue that market outcomes, though optimal, are bad in other ways; since lump-sum transfers are impossible, the sacrifice of dead-weight loss is necessary to achieve other goals. These critiques are empirical. The degree of market power could be large or small. Lump-sum transfers are not so much impossible as they are difficult to execute. Consequently, these concerns are typically considered to be second-order.

We offer here a different and perhaps more fundamental critique. When markets allocate contingent claims among expected-utility-maximizing agents, Pareto optimality with *ex ante* beliefs is an inappropriate welfare criterion except in the instance where all traders have identical beliefs over states of the world. This critique is detailed in section 3, after an infinite-horizon model of trade in a single consumption good with complete markets is developed in section 2. If the "true distribution of states" were known to an omniscient social planner, Pareto calculations with correct beliefs is an obvious fix. Omniscient social planners are rare, however, and without them there is no alternative welfare requirement that obviously ameliorates the issues raised in section 3. We investigate the magnitude of the problem through simulations of competitive equilibria. The simulations of sections 5, 6 and 7 examine several policy alternatives to complete markets in Markovian instances of the model of financial restrictions developed in section 4, and explore the size and location of the set of potentially true distributions for which the policies would lead to a true welfare improvement with respect to several distinct welfare criteria. We conclude in section 8 with a discussion of the theoretical and the policy implications of our findings.

2   The model

We assume that time is discrete and begins at date 0. At each date a state is drawn from the set $\mathcal{S} = \{1, \ldots, S\}$. The set of all sequences of states is Σ with

representative sequence $\sigma = (s_0, s_1, ...)$, called a *path*. Let $\sigma^t = (s_0, ..., s_t)$ denote the partial history through date t. We use $\tilde{\sigma}|\sigma^t$ to indicate that a path $\tilde{\sigma}$ coincides with a path σ up through period t.

The set Σ together with its product sigma-field is the measurable space on which everything is built. Let P^0 denote the "true" probability measure on Σ. For any probability measure P on Σ, $P_t(\sigma)$ is the (marginal) probability of the partial history σ^t: $P_t(\sigma) = P(\{\sigma^t\} \times S \times S \times \cdots)$.

In the next few paragraphs we introduce a number of random variables of the form $x_t(\sigma)$. All such random variables are assumed to be date-t measurable; that is, their value depends only on the realization of states through date t. Formally, $\mathcal{F}_t$ is the σ-field of events measurable at date t, and each $x_t(\sigma)$ is assumed to be $\mathcal{F}_t$-measurable.

An economy contains I consumers, each with consumption set $\mathbb{R}_+$. A *consumption plan* $c : \Sigma \to \prod_{t=0}^{\infty} \mathbb{R}_+$ is a sequence of $\mathbb{R}_+$-valued functions $\{c_t(\sigma)\}_{t=0}^{\infty}$ in which each c_t is $\mathcal{F}_t$-measurable. Each consumer is endowed with a particular consumption plan, called the *endowment stream*. Consumer i's endowment stream is denoted e^i. The aggregate endowment stream is denoted by $\bar{e}$:

$$\bar{e}_t(\sigma) = \sum_{i=1}^{I} e_t^i(\sigma).$$

An *allocation* is a profile of consumption plans, one for each individual. The allocation $(c^1, \ldots, c^I)$ is *feasible* if for all σ and t, $\sum_i c_t^i(\sigma) - e_t^i(\sigma) = 0$.

Consumer i's preferences on consumption plans are described by a *belief* or *forecast distribution* P^i, a probability distribution on Σ, a discount factor $0 < \beta_i < 1$, and a payoff function $u_i : \mathbb{R}_{++} \to \mathbb{R}$. The utility consumer i assigns to consumption plan c is the expectation of the average discounted value of the sequence of payoff realizations:

$$U_{P^i}^i(c) = (1 - \beta_i) E_{P^i} \left\{ \sum_{t=0}^{\infty} \beta_i^t u_i(c_t(\sigma)) \right\}. \tag{1}$$

Notice that beliefs are indexed by individual names. Different individuals may believe different things about the future, and these beliefs need not coincide with what will actually happen. The *true state process* is a stochastic process on $\mathcal{S}$, characterized by a probability distribution P^0 on Σ, and it may be the case that for no distinct $i, j \geq 0$ does $P^i = P^j$. We will impose some constraints on how different beliefs can be.

We assume the following properties of the payoff function:

A1. *Each $u_i : \mathbb{R}_{++} \to (-\infty, \infty)$ is C^1, strictly increasing and strictly concave.*

A2. *Each u_i satisfies an Inada condition at 0: $\lim_{c \downarrow 0} u_i'(c) = \infty$.*

We assume the following properties of the aggregate endowment:

A3. *The aggregate endowment is uniformly bounded from above and away from 0:*

$$\infty > F = \sup_{t,\sigma} \bar{e}_t(\sigma) \geqslant \inf_{t,\sigma} \bar{e}_t(\sigma) = f > 0.$$

Finally, we assume that anything is possible at any date, and that individuals believe this to be true:

A 4. *For all individuals i, all dates t and all paths σ, the distributions $P_t^i(s_t|\sigma^{t-1})$ for $i \geq 0$ have full support.*

We will often refer to agents as being optimistic or pessimistic. We say that a type-i agent is *optimistic* after history σ^t if $E^i[e_{t+1}^i|\sigma^t] > E^0[e_{t+1}^i|\sigma^t]$. *Pessimism* is defined analogously.

3   Welfare economics of heterogeneous beliefs

The welfare analysis of market outcomes begins with the Pareto order, taking preferences as given. "Tastes," say Stigler and Becker [1977, p. 76], "are the unchallengeable axioms of a man's behavior: he may properly (usefully) be criticized for inefficiency in satisfying his desires, but the desires themselves are data." Tastes, they say, "are not capable of being changed by persuasion."

In contingent-claims markets, "Pareto optimality" is taken to be with respect to *ex ante* preferences (tastes); that is, *ex ante*, or time-0, expected utility. While we do certainly agree that tastes for apples and oranges, work and leisure, etc., are to be taken as given, we dispute the claim that *ex ante* preferences on contingent claims are above dispute. Time-separable expected utility representations of these preferences have three components: attitudes towards risk, the rate of time preference, and beliefs about the realization of states. While risk attitudes and discount factors may be unarguable, beliefs are not. When market participants have different beliefs, not all can be right, and those who are wrong are making decisions that they would regard as incorrect if only they had correct beliefs. Finally, if beliefs were

4

an indisputable attribute of tastes there would be no role for expectation-managing policies, and the rational expectations hypothesis would never be conceived.[1]

3.1 Spurious unanimity

Ithaca NY, the home of three of us, has a pedestrian mall. It is still serviceable, but would benefit from renovation. The work, however, will be costly. Suppose that half the town believes that revitalization will enhance Ithaca's attraction as a summer tourist destination. This group believes that crowds of tourists will bring more business opportunities and badly needed tax revenues. The other half of the town believes that revitalization will make downtown more pleasant without materially perturbing downtown's summer population density, thereby enhancing the quality of life. The town is unanimous in its support for the project. Is unanimity of preference a good argument for undertaking the project? Not according to Mongin [2005], who calls this problem "spurious unanimity". He argues that not only preferences themselves, but the reasons why people hold the preferences they have, need to be considered in making welfare claims. This is clear in the mall-renovation case. Suppose that many editorials have appeared in the local newspaper, many public meetings have been held, and the issue has been thoroughly aired. It is common knowledge, then, that individuals believe different things. It is common knowledge, then, that if the mall is renovated, half the town will be unhappy with the result. It is common knowledge that the renovation cannot be an *ex post* Pareto improvement. There is disagreement only over who is in which half. Suppose there are N different possible states of the world rather than 2, and that the population is divided equally into N groups. Individuals in any group will benefit from a proposal only if "their" state of the world occurs and will be harmed otherwise, and each individual is sure that the state beneficial to him will occur. It is then common knowledge that only fraction $1/N$ of the population will be made better off, that fraction $N - 1/N$ will be made worse off. Imagine that N is large. The justification of the proposal by *ex ante* Pareto optimality is not at all

---

[1]The assumption of "common knowledge of prior beliefs" is often used, following Aumann [1976], to justify common beliefs. Common prior arguments are critically discussed in Morris [1995]. To his analysis we add that the entire apparatus of belief about beliefs about beliefs is simply misplaced in models of trade in large anonymous markets, wherein one individual may have no idea of who or what is on the other side of his transaction.

compelling.

The problem of spurious unanimity is even more compelling when expected utility decision makers choose over alternatives with random payoffs. Imagine now that two decision-makers are choosing between two policies, $\mathcal{A}$ and $\mathcal{B}$. Policy $\mathcal{A}$ gives outcome a on event E and b on E^c. Policy $\mathcal{B}$ is the mirror-image; it gives outcome b on E and a on E^c. Individuals 1 and 2 each have a payoff function and a prior belief, which are as follows:

$$
\begin{array}{llll}
\text{Individual 1} & u_1(a) = 1, & u_1(b) = 0, & \rho_1(E) = 0.99, \\
\text{Individual 2} & u_2(a) = 0, & u_2(b) = 1, & \rho_1(E) = 0.01.
\end{array}
$$

Each individual prefers policy $\mathcal{A}$ to policy $\mathcal{B}$. Unanimity is a consequence of their divergent beliefs. Given their payoff functions, if they shared a common belief they could never agree on a policy except in the case where they both believe each state is equally likely.

Of course, if one believes that all individuals have common, correct beliefs then spurious unanimity is not an issue and ex ante Pareto optimality is an appropriate welfare criterion. We do not find this restriction on beliefs compelling. It certainly does not follow from Savage [1954 (1972] subjective expected utility theorem. It is instead a restriction on preferences which goes far beyond the notion of rationality embedded in Savage. In fact, even the ability to reason about the probabilities in Savage's representation as beliefs is not as simple as is normally supposed. We discuss this issue in Appendix A.1.

3.2 The ex ante welfare economics of contingent claims

Because beliefs are not above dispute, we are concerned with two Pareto orders. The usual welfare analysis is concerned with the *ex ante Pareto order*, and because individuals would choose to adopt the true distribution if only they knew it, we are also concerned with the *true Pareto order* which is the order that obtains when each individual computes expected utility with the true distribution $\mathcal{P}^0$.

If individuals disagree, then in economies of the type described in Section 2, these two orders differ. That is, *ex ante* optimal contingent claims for given beliefs $P^1, \ldots, P^I$, with $P^i \neq P^j$, for some i and j, cannot be true Pareto optimal for any P^0.

Proposition 1. *If the economy contains two individuals i and j such that for some t and some path σ, $P_t^i(\sigma) \neq P_t^j(\sigma)$, then no ex ante Pareto optimal allocation in which $c^i, c^j \neq 0$ can be optimal for any true distribution P^0.*

Proof. If $P_t^i(\sigma) \neq P_t^j(\sigma)$, then there must exist some other path σ' such that $P_t^i(\sigma')/P_t^j(\sigma') \neq P_t^i(\sigma)/P_t^j(\sigma)$, else probabilities cannot sum to one. The first-order conditions for optimality on path σ imply that

$$\frac{u_i'\big(c_t^i(\sigma)\big)}{u_j'\big(c_t^j(\sigma)\big)} = \frac{\lambda_i \beta_j^t}{\lambda_j \beta_i^t} \frac{P_t^j(\sigma)}{P_t^i(\sigma)},$$

where the λ's, multipliers for the Pareto optimization problem, are both positive as $c^i, c^j \neq 0$. Suppose now that the allocation is true Pareto optimal for some P^0. Then first-order conditions imply that there will be positive multipliers γ_i and γ_j such that

$$\frac{u_i'\big(c_t^i(\sigma)\big)}{u_j'\big(c_t^j(\sigma)\big)} = \frac{\gamma_i}{\gamma_j} \frac{\beta_j^t}{\beta_i^t}.$$

Consequently the vectors $(\gamma_i \beta_j^t, \gamma_j \beta_i^t)$ and $\big(\lambda_i \beta_j^t P_t^j(\sigma), \lambda_j \beta_i^t P_t^i(\sigma)\big)$ are proportional.

Now consider path σ'. Since the allocation is truly optimal, it must be the case that:

$$\frac{u_i'\big(c_t^i(\sigma')\big)}{u_j'\big(c_t^j(\sigma')\big)} = \frac{\gamma_i}{\gamma_j} \frac{\beta_j^t}{\beta_i^t}.$$

Since the allocation is also *ex ante* optimal:

$$\frac{u_i'\big(c_t^i(\sigma')\big)}{u_j'\big(c_t^j(\sigma')\big)} = \frac{\lambda_i \beta_j^t}{\lambda_j \beta_i^t} \frac{P_t^j(\sigma')}{P_t^i(\sigma')}.$$

Thus $P_t^i(\sigma')/P_t^j(\sigma') = P_t^i(\sigma)/P_t^j(\sigma)$, which is a contradiction. $\qquad\square$

When discount factors are identical, there is in fact a simple necessary condition for true Pareto optimality: Everyone's consumption is bounded away from 0.

Corollary 1. *If individuals have identical discount factors, if the allocation c is true-Pareto optimal, and if for all i, $c^i \neq 0$, then for each individual i and all σ, $\liminf_t c_t^i(\sigma) > 0$.*

Proof. This follows immediately from the fact that the first order conditions are independent of P^0, that the welfare weights are positive, and that aggregate endowments are uniformly bounded above and below across paths. □

Another necessary condition for true optimality is that there is no speculation on irrelevant states (frivolous uncertainty).

Corollary 2. *Suppose that c is true-Pareto optimal, that $c^i \neq 0$ for all i, and that the endowment allocation at date t is constant on some event E, that is, for $\sigma, \sigma' \in E$, $e_t(\sigma) = e_t(\sigma')$. Then for all individuals $c_t^i(\sigma) = c_t^i(\sigma')$.*

Proof. Since the allocation is true-Pareto optimal and $c^i \neq 0$ for all i, it must be the case that:

$$\frac{u_i'\big(c_t^i(\sigma)\big)}{u_j\big(c_t^j(\sigma)\big)} = \frac{\gamma_i}{\gamma_j}\frac{\beta_j^t}{\beta_i^t} = \frac{u_i'\big(c_t^i(\sigma')\big)}{u_j\big(c_t^j(\sigma')\big)}, \quad \forall i, j.$$

Then $e_t(\sigma) = e_t(\sigma')$ on E and the fact that the allocation c is feasible imply the desired result. □

Proposition 1 and the first welfare theorem suggest that the introduction of some kind of market incompleteness could be welfare-improving, that is, incomplete markets could yield allocations that true-Pareto dominate the complete-markets allocation. Interestingly, someone whose beliefs are correct cannot be *ex ante* hurt by any true-Pareto improvement. So, as long as majority of the population have correct beliefs, proposals that are true-Pareto improvements should gain political support.

Unfortunately, the mechanism design problem depends critically on the true distribution P^0. It is easy to construct examples where there is no allocation that true-Pareto dominates a given *ex ante* optimal allocation for every possible P^0.[2] Since individuals in the market do not have privileged

---

[2]Consider the following example. Two agents with logarithmic preferences believe that the distribution over two possible states is (0.6,0.4) and (0.4,0.6), respectively. Agent i is endowed with $1 - e$ units of consumption good in state i and e units otherwise. In the competitive equilibrium (CE), consumption of agents 1 and 2 are (0.6,0.4) and (0.4,0.6). The even split is the allocation in which the agents consume 0.5 in each state. It true Pareto dominates the CE allocation only if the true distribution is sufficiently close to $(0.5, 0.5)$. If the probability of state 1 under the true distribution exceeds $\bar{p}_1 \equiv ln(1.25)/ln(1.5) \approx 0.55$ or is below $1 - \bar{p}_1$, then the even split no longer true-Pareto dominates the CE allocation. In fact, in this case there is no other allocation that true-Pareto dominates the CE allocation for all belief assignments.

knowledge of the true distribution, it would be unreasonable to assume that market designers would have any better knowledge. That is, we want to do distribution-independent market design.

Our solution to this problem is to explore the parameter space. We show that there are market institutions that outperform complete markets over much of the parameter space. "Outperform" here has three meanings. For the market interventions we consider, through simulation we delineate regions of the model's parameter space where the intervention is true Pareto superior, where it is better according to a Rawlsian welfare aggregator, and where it is better according to a Bergson-Samuelson social welfare function where the welfare weights are those that solve the *ex ante* Pareto optimality problem.

3.3   Spurious unanimity: other approaches

Others have addressed the problem of spurious unanimity in contingent claims allocations. Brunnermeier et al. [forthcoming] introduce *belief-neutral* Pareto optimality. They identify a set of "reasonable beliefs", potential true distributions, which is the convex hull of the set of individuals' beliefs. Allocation x is then belief-neutral Pareto superior to allocation y if x is true Pareto superior to y for every true distribution in the set of reasonable beliefs. The intersection of a collection of Pareto orders is, generally speaking, incredibly incomplete. Brunnermeier et al. [forthcoming] reduce incompleteness by examining partial orders induced by Bergson-Samuelson social welfare functions, taking weighted averages of each profile of true expected utilities.

Gilboa et al. [2014] offer a somewhat complicated alternative. Allocation x *no-bet* Pareto improves upon y if x *ex ante* Pareto improves upon y and if there exists a potentially true probability distribution such that each individual whose position is *ex ante* improved in the move from y to x also truly prefers x to y. This is a direct attempt to remove from Paretian calculations the speculative component to trade that is introduced when beliefs disagree. The no-bet Pareto relation, while acyclic, can be intransitive.

These two proposals delineate the trade offs that arise when considering potential true distributions. Requiring Pareto improvement with respect to a large class of potential true distributions for all welfare comparisons thickens the contract curve; few welfare comparisons can be made. Relaxing this ordinal uniformity condition, however, and allowing different distributions

for different comparisons, will, generally speaking, introduce intransitivities.

Duffie [2014] also addresses the issue of trading generated by heterogeneous beliefs and appropriate policy responses to it. He raises issues of how to evaluate welfare in this context and considers the tradeoffs between reducing speculation and trading to hedge risk, provide liquidity or use information.

We do not see any particularly compelling way to undertake welfare analysis when beliefs are heterogeneous. This includes *ex ante* Pareto optimality. So in this paper we carry out the more limited task of identifying sets of beliefs and potentially true distributions for which given market restrictions are in some sense welfare-improving in some simple examples. We believe that if, in a carefully calibrated model of economic activity, for some market restriction the set of potentially true distributions for which it is a welfare improvement is large, then there is a strong *prima facie* case for introducing it.

4 Financial markets, competitive equilibria

In this section, we describe optimization problems of an agent under different financial market designs.

4.1 The complete markets economy

The first and the key market design is (dynamically) complete financial markets. Let $Q_t(\sigma)$ be the date-t price of an Arrow security that pays along path σ. The number of Arrow securities purchased by a type-i agent in period t along history σ is denoted by $a_t^i(\sigma)$. Then a type-i agent faces the following budget constraint at each date t

$$c_t^i(\sigma) + \sum_{\tilde{\sigma}|\sigma^t} Q_t(\tilde{\sigma})a_{t+1}^i(\tilde{\sigma}) = a_t^i(\sigma) + e_t^i(\sigma). \tag{2a}$$

Purchases of Arrow securities are subject to *natural borrowing limits* at each date t

$$a_{t+1}^i(\sigma) \geqslant -N_{t+1}^i(\sigma), \tag{2b}$$

constructed as follows. Define the j-period ahead price $Q_t^j(\sigma) = \Pi_{k=0}^{j-1}Q_{t+k}(\sigma)$. Then a natural borrowing limit equals the date-t value of the continuation

of an agent's endowment plan:

$$N_t^i(\sigma) = \sum_{j=0}^{\infty} \sum_{\tilde{\sigma}|\sigma^t} Q_t^j(\tilde{\sigma}) e_{t+j}^i(\tilde{\sigma}). \tag{3}$$

Natural borrowing limits never bind in a competitive equilibrium if a period utility function satisfies our Inada condition (A2). A type-i agent chooses consumption and asset trading plans to maximize life-time utility (1) subject to constraints (2a) and (2b).

We define prices of two assets to which we refer later. The price of a risk-free bond at date t is

$$q_t^b(\sigma) = \sum_{\tilde{\sigma}|\sigma^t} Q_t(\tilde{\sigma}). \tag{4a}$$

The price at date t on path σ of a claim to the aggregate endowment is:

$$q_t^e(\sigma) = \sum_{j=0}^{\infty} \sum_{\tilde{\sigma}|\sigma^t} Q_t^j(\tilde{\sigma}) e_{t+j}(\tilde{\sigma}). \tag{4b}$$

Definition. The complete financial markets (CM) design is a set of S financial markets where market j trades an Arrow security that pays one unit of consumption good next period if state j realizes. Trading is subject to natural borrowing limits defined above.

Definition. An agent is said to survive if his consumption remains strictly positive forever almost surely-P^0.

In addition to standard complete markets, we analyze several other designs: complete markets with borrowing limits (CMB), complete markets in which transactions are taxed (CMT), and markets trading only a risk-free bond subject to a borrowing limit (B). We think of these intermediate designs as partially regulated financial markets and aim to shed light on the relative desirability of different restrictions.

4.2   A bond economy

Definition. A bond-only financial market design (B) consists of a single market that trades a risk-free bond subject to an exogenous borrowing limit.

In the bond economy, a type-i agent faces the following constraints:

$$c_t^i(\sigma) + q_t^b(\sigma)b_{t+1}^i(\sigma) = b_t^i(\sigma) + e_t^i(\sigma), \tag{5a}$$

$$b_{t+1}^i(\sigma) \geqslant -B_{t+1}^i(\sigma), \tag{5b}$$

where $q_t^b(\sigma)$ denotes the date-t price of a risk free bond, $b_t^i(\sigma)$ represents the date-t bond purchases of agent i, and $B_{t+1}^i(\sigma)$ is an *exogenous* borrowing limit. These borrowing limits have to be sufficiently tight to make sure that all loans are repaid with certainty. Borrowing limits must be tighter than the worst-case date-t value of the continuation of an agent-i's endowment plan:

$$\inf_{\tilde{\sigma}|\sigma^t} \left[e_t^i(\tilde{\sigma}) + \sum_{j=0}^{\infty} \Pi_{k=0}^{j-1} q_{t+k}^b(\tilde{\sigma}) e_{t+1+j}^i(\tilde{\sigma}) \right].$$

The above borrowing limit is the largest limit that can (potentially) be imposed after history σ^t on a type-i agent in the bond-only economy. However, unlike in the complete markets economy, an endogenous borrowing limit cannot be determined before solving for a competitive equilibrium. Hence, an exogenous borrowing limit must be imposed instead.

4.3   Borrowing limits

Definition. The complete financial markets with a borrowing limit (CMB) design is a set of S financial markets where market j trades an Arrow security that pays one unit of consumption good next period if state j realizes and zero otherwise. Trading is subject to an exogenous borrowing limit.

Under the complete markets with a borrowing limit design trading is subject to an exogenous borrowing limit B that is tighter than the natural borrowing limits in (2b)

$$a_{t+1}^i(\sigma) \geqslant -B. \tag{6}$$

The financial markets are complete in the sense that a full set of Arrow securities is traded. Yet, when a tight borrowing limit is imposed, insurance possibilities are restricted. Speculation opportunities are also limited, which tames the survival forces analyzed by Blume and Easley [2006] that otherwise would drive the consumption of agents with less accurate beliefs to zero asymptotically.

4.4 Transaction costs

Definition. The complete financial markets with a transaction cost (CMT) design is a set of S financial markets where market j trades an Arrow security that pays one unit of consumption good next period if state j realizes and zero otherwise. Trading is subject to a transaction cost.

Under this design the budget constraint (2a) is replaced with the following

$$c_t^i(\sigma) + \sum_{\tilde{\sigma}|\sigma^t} Q_t(\tilde{\sigma}) a_{t+1}^i(\tilde{\sigma}) + \tau \cdot \sum_{\tilde{\sigma}|\sigma^t} [a_{t+1}^i(\tilde{\sigma}) - a_t^i(\sigma)]^2$$
$$= a_t^i(\sigma) + e_t^i(\sigma) + T_t(\sigma)/2, \qquad (7)$$

where $T_t(\sigma)$ is the total transaction tax revenue. Our transaction tax design embeds two important assumptions. First, the transaction tax is assumed to be a quadratic function of security purchases to insure continuity of demands for securities. Second, we rebate the transaction tax back to investors as equal lump sums.

A transaction cost limits speculation opportunities, as does a borrowing limit, but agents are not guaranteed to survive. The two alternatives differ in how they control potential welfare losses: a transaction cost slows the rate at which agents can lose wealth, and a borrowing limit imposes a bound on how much wealth can be lost.

5   Welfare criteria

We require some welfare criterion in order to discuss the welfare consequences of market restrictions. Although we do not want to commit to a particular criterion, any acceptable criterion should satisfy some obvious desiderata. First, as we have argued, it cannot be based on individual welfare computed ex ante. We consider individuals who chose optimally, given their preferences, but we evaluate their welfare using the true probability on states. Second, we do not want to evaluate social welfare using any particular truth as we see no justification for assuming that the social planner, who we view as choosing market restrictions, knows the truth when individuals do not know it. So, we evaluate welfare over a set of possible truths. Third, we do not want to design market restrictions that work only for particular configurations of individual beliefs. Once we drop the usual restriction that beliefs are correct, we see

no justification for placing joint restrictions on, possibly incorrect, beliefs of individuals. Therefore we evaluate welfare over a set of individual beliefs. Finally, for any given individual beliefs and truth, we need to aggregate individual payoffs. Here too we see no compelling argument for any particular aggregator and we instead consider several possibilities: one based on a Rawlsian criterion as well as one based on a Bergson-Samuelson criterion. We admit at the outset that this approach requires that we take utility to be interpersonally comparable. The burgeoning literature on *ex ante* optimality with heterogeneous beliefs has yet to produce a satisfactory ordinal way of proceeding; so at the present time social welfare functions are the only way to go forward. *Definition. A utility aggregator $\mathcal{W} : R^I \to R$ is a non-decreasing continuous function such that $\mathcal{W}(U) \in [\min_i U_i, \max_i U_i], \forall U \in R^I$.*

The Pareto welfare criterion uses:

$$\mathcal{W}(U_1, .., U_I) = \sum_{i=1}^{I} \theta^i U_i \tag{8}$$

for some exogenously given vector of Pareto weights $\theta \in \Delta^I$. Another possibility is the Rawlsian utility aggregator:

$$\mathcal{W}(U) = \min_i U_i. \tag{9}$$

Definition. Let $\mathcal{B}$ be a set of admissible beliefs and let $\mathcal{P} = (P^1, ..., P^I) \in \mathcal{B}^I$ denote a belief assignment. Let $P^0 \in \mathcal{B}^0$ be a data generating process, where $\mathcal{B}^0$ is a set of admissible data generating processes. Let $c(\mathcal{P}|\mathcal{M})$ be a competitive equilibrium allocation under a financial market structure $\mathcal{M}$ and a belief assignment $\mathcal{P}$. Then the social welfare function using a utility aggregator $\mathcal{W}$ is:

$$\min_{P^0 \in \mathcal{B}^0} \min_{\mathcal{P} \in \mathcal{B}^I} \mathcal{W}\left(\left(U_{i,P^0}(c^i(\mathcal{P}|\mathcal{M})) \right)_{i=1}^{I} \right). \tag{10}$$

This welfare criterion emerges from three analytic choices. The first choice is that of $\mathcal{W}$. Fix the beliefs assignment $\mathcal{P} = (P^1, ..., P^I)$ and the true data generating process P^0. By choosing a market structure $\mathcal{M}$ the designer effectively selects an allocation $(c^1(\mathcal{P}|\mathcal{M}), ..., c^I(\mathcal{P}|\mathcal{M}))$ and the associated distribution of utilities $(U_{1,P^0}, ..., U_{I,P^0})$. A utility aggregator $\mathcal{W}$ transforms

the distribution of utilities into a social welfare measure. A designer using (9) would choose a financial market structure that benefits the least-advantaged members of society. That is the designer would adhere to one of the principles of justice proposed in Rawls [1971].[3] A designer using (8) would act similarly to a Pareto planner. But the utility aggregator (8) presents a new degree of arbitrariness: What weights should a designer use? One could choose $\theta_i = 1/I, \forall i$, the choice that is attractive in *ex-ante* symmetric environments. One could also choose θ to be a vector of "market weights".[4] These two choices are special cases of Bergson-Samuelson social welfare functions. However, any set of weights is arbitrary. The paternalistic designer using the Rawlsian aggregator (9) is spared the obligation of deciding a fair set of weights.

The second choice is the minimization over belief assignments. Our motivation is that in reality many configurations of beliefs are possible. Each such configuration may support a different financial market design. For example, some agents may be optimistic and undertake excessively risky investments that could drive them quickly out of financial markets. Imposing borrowing limits may be desirable in this case. On the other hand, pessimistic agents might over-invest in safe assets. They would still be driven out of financial markets, but perhaps at a slower rate. Financial regulation in this case would have to strike a balance between saving agents from financial ruin and encouraging diversification. A rational planner would need to take into account any available information on the distribution of beliefs in the population. We minimize over beliefs for two reasons. First, our goal in this paper is to provide examples of how complete markets can go astray, and computing worst cases satisfies that need. Second, one can view minimization over

---

[3]Rawls [1971] argues that a fair social choice can only be made in a hypothetical "original position":

> No one knows his place in society, his class position or social status, nor does anyone know his fortune in the distribution of natural assets and abilities, his intelligence, strength, and the like. I shall even assume that the parties do not know their conceptions of the good or their special psychological propensities. The principles of justice are chosen behind a veil of ignorance.

For our purposes, replace "principles of justice" with "design of financial markets." The veil of ignorance advocated by Rawls allows devising a set of rules that are independent of the current economic fundamentals – beliefs assignment, true data generating process, and wealth distribution.

[4]This is a vector of weights for which the Pareto and the competitive allocations coincide under $P^i = P^0, \forall i$. See section 7 for more details.

belief assignments as a Rawlsian approach to social welfare in which each belief assignment corresponds to a different feasible configuration of society and the planner has no *ex ante* information about the beliefs held by the traders.[5]

6 Examples

We now present the leading example that we use to illustrate economic forces that operate in economies with heterogeneous beliefs. We apply our welfare criterion to compare complete markets with various incomplete market settings. In this section, we investigate social welfare using the Rawlsian utility aggregator (9). In section 7, we consider the Pareto criterion using market weights. The two criteria lead to remarkably similar results.

In our economy agents share a common utility function

$$u(c) = c^{1-\gamma}/(1-\gamma),$$

where $\gamma = 2$.[6] There are two types of agents and three states: $\sigma_t \in \{0, 1, 2\}$. The economy begins in state 0 and then exits to states 1 and 2.[7] Endowments are

$$(e_t^1(\sigma), e_t^2(\sigma)) = \begin{cases} (0.5, 0.5) & if \ \sigma_t = 0 \\ (e_h, e_l) & if \ \sigma_t = 1 \\ (e_l, e_h) & if \ \sigma_t = 2 \end{cases}, \quad \forall t, \sigma. \tag{11}$$

We assume that $e_h > e_l$. Although there is no aggregate uncertainty, individuals face idiosyncratic risk.

Beliefs are specified as follows:

$$\Pi^i = \begin{bmatrix} 0 & 0.5 & 0.5 \\ 0 & p^i & 1-p^i \\ 0 & p^i & 1-p^i \end{bmatrix}, \tag{12}$$

---

[5]Other approaches are possible. Phelan and Rustichini [2015], in a different context, take an alternative approach in which each individual (for us each individual defined by some beliefs) is treated a separate person and Pareto optimality takes each of these "people" directly into account.

[6]Robustness of our results to the specification of preferences is considered in Appendix A.3.

[7]The only purpose of the transitory state 0 is symmetry. It insures that agents begin with identical endowments and can trade prior to the first realization of states 1 and 2.

where Π^0 denotes the true probability transition matrix. Subjective probabilities over histories $P_t^i(\sigma)$ are computed using individual transition matrices.[8]

6.1   Complete markets economy

First, we describe a competitive equilibrium in the complete markets economy when beliefs are homogeneous, but not necessarily correct. Because there is no aggregate uncertainty and preferences are homothetic, both agents consume a constant amount. The competitive equilibrium allocation is:

$$(c_t^1(\sigma), c_t^2(\sigma)) = (0.5 + \beta^2(\mu_e - 0.5), 0.5 - \beta^2(\mu_e - 0.5)), \quad \forall t, \sigma, \qquad (13)$$

where $\mu_e \equiv pe_l + (1 - p)e_h$ is the expected endowment evaluated using the common beliefs, p. An agent achieves a constant consumption plan by buying an amount $A_j \equiv 0.5 - e_j + \beta(\mu_e - 0.5)$ of Arrow securities paying in the state where its income is e_j. The quantity of Arrow securities traded in equilibrium, $|A_j|$, is small relative to the natural borrowing limit: $N_t^i(\sigma) = e_t^i(\sigma) + \beta\mu_e/(1 - \beta)$.

Second, we describe a competitive equilibrium with complete markets and heterogeneous beliefs. Suppose that $p^1 = p^0$ and $p^2 \neq p^0$. In this case, not only do agents not consume constant amounts, but as shown by Blume and Easley [2006], consumption of a type-2 agent converges to zero:

$$\limsup_{t \to \infty} c_t^2(\sigma) = 0 \quad P^0 a.s. \qquad (14)$$

Following Blume and Easley [2006], we say that type-2 agents do not survive. The eventual immiseration of agents with incorrect beliefs when market are complete is the source of an instinct that market restrictions could be useful.

Agents invest in Arrow securities for two reasons: income hedging and disagreement. Suppose $p^2 > p^0 = p^1$. To hedge income fluctuations, a type-2

---

[8]The beliefs on sample paths induced by this structure do not involve learning. Our individuals believe that the exogenous states follow an iid process and each person i is certain about p^i. The analysis can be extended to include learning. Our general results carry over to this extension as can be seen from the analysis in Blume and Easley [2006]. Computing equilibria is, however, much more difficult. Learning in which individuals' beliefs converge to the true data generating process makes complete markets more attractive, but it does not eliminate the desirability of restrictions on financial markets. Individuals care about the entire process, not just the limit, and so what matters is the speed of learning versus the rate at which the future is discounted.
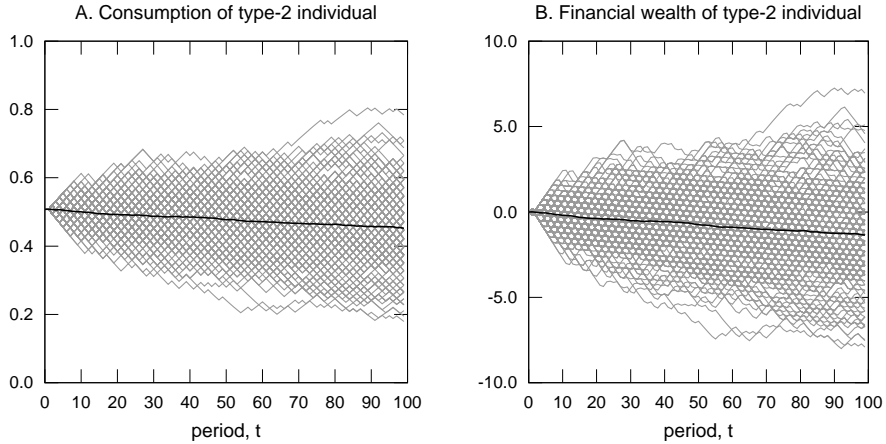
Figure 1: Sample paths of consumption and financial wealth of a type-2 agent in the complete markets economy.

Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, p^2 = 0.55$.

agent buys Arrow securities that pay in state 1 (when his income is low) and sells Arrow securities that pay in state 2 (when his income high). Because a type-2 agent overestimates the probability of state 1, he buys extra securities that pay in this state. So he over-invests in securities that pay in state 1 and under-invests in securities that pay in state 2. These additional trades are "speculative."[9] As a result of these trades, a type-2 agent's consumption increases every time state 1 realizes. The opposite happens if state 2 realizes. State 1 is less likely than a type-2 agent anticipates. So his investments pay off less than he expects, he loses wealth on average, and his consumption converges to zero.

Figure 1 plots 200 sample paths of consumption (panel A) and financial wealth (panel B) of a type-2 agent for a simple example of the complete markets economy. The solid line in each panel denotes the average across sample paths. Both consumption and wealth drift towards their respective lower bounds. The speed of convergence is slow: for example, after 100 periods a type-2 agent's consumption decreases from 0.493 to 0.432 along the average path. The decline in financial wealth is more substantial, falling from 0 to -1.524 (or roughly three average individual annual incomes) along

---

[9]Speculation is trading activity that is motivated by differences in beliefs and would be absent had all agents had the same beliefs.
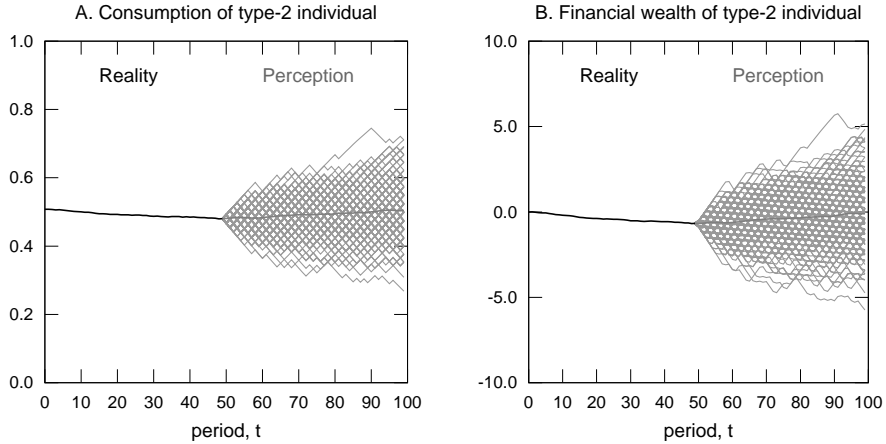
18

Figure 2: Actual vs perceived sample paths of consumption and financial wealth of a type-2 agent in the complete markets economy.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, p^2 = 0.55.$

the average path.

Despite a decline in his consumption and financial wealth, a type-2 agent believes that what happens to him is simply bad luck. Figure 2 demonstrates the difference between actual and perceived outcomes. This figure plots expected, from a point of view of type-2 agent, evolution of his consumption and financial wealth in periods 51-100 assuming that during periods 0-50 he followed the "average path." Not surprisingly, he expects to prosper. This is a manifestation of another result in Blume and Easley [2006] which applied to our example shows that agent 2 believes that his consumption will converge almost surely to the entire aggregate endowment:

$$\limsup_{t \to \infty} c_t^i(\sigma) = 1 \quad P^i a.s.$$

Finally, we present welfare levels for the two types of agents in our example. As a benchmark, we compute welfare in the complete markets economy when beliefs are homogeneous and coincide with the truth. Assuming $p^0 = 0.50$, this benchmark level of welfare, denoted by W^*, is -2 for each type. Subjective welfare levels in the heterogeneous beliefs economy are -1.943 and -2.124, respectively, for type-1 and type-2 agents. A type-1 agent, whose beliefs coincide with the truth, expects higher welfare than

19

W^*. He is better off in the economy with diverse beliefs as his "speculative" financial trades allow him to accumulate wealth. A type-2 agent expects welfare level that is lower than W^*. This happens because the type-2 agent believes that his endowment stream has a relatively low value. Objective welfare levels (expected utility of equilibrium consumptions computed using the truth) are -1.947 and -2.129, respectively for a type-1 and type-2 agent.

In this example, belief diversity has a substantial impact on welfare: relative to the common beliefs benchmark, a reduction in a type-2 agent's welfare is equivalent to a permanent 6.45% decline in his consumption.[10] So welfare of a type-2 agent is low, and hence according to the Rawlsian aggregator, social welfare is low. Two sources contribute to this outcome: consumption volatility and a downward trend in a type-2 agent's consumption. To quantify the contribution of each source, we note that the welfare of a type-2 agent computed along the "average path" is -2.091. Thus, low welfare of a type-2 agent is caused largely by a diminishing trend in his consumption rather than by increased consumption volatility.[11]

6.2   Bond economy

In the bond-only economy, agents can save or borrow by buying or selling bonds, but they cannot transfer income across states. To insure that an equilibrium exists, we impose a borrowing limit as explained in section 4.2. Since it is impossible to devise *a priori* a borrowing limit that would never bind, we impose an exogenous, yet generous, limit of 16 average individual annual incomes: $B_t^i(\sigma) = 8, \forall t, \sigma$.

Continuing with the example from the previous section, we simulate equilibrium consumption and wealth dynamics in the bond economy. As shown in figure 3, consumption and financial wealth for the type-2 agent now grow on average. Consumption increases from an average of 0.492 to 0.526 (panel

---

[10]Costs of aggregate fluctuations in a standard RBC model are typically found to be below 0.1%.

[11]It is natural to ask what would happen in this economy if a type-2 agent were optimistic. To answer this we studied the case with $p^0 = p^1 = 0.50, p^2 = 0.45$. Welfare levels in this case are: $U_{P0}^1 = U_{P1}^1 = -2.002, U_{P0}^2 = -2.063$ and $U_{P2}^2 = -2.058$. Here a type-2 agent still has the lower welfare in the economy, but it is not as low. This happens largely because optimism increases the value of his endowment plan. So his consumption while decreasing on average starts from a value above 0.5. If we replaced his consumption plan with an average plan his welfare would be -2.024. So here the welfare loss is attributed mainly to increased consumption volatility. See also section A.2.1.
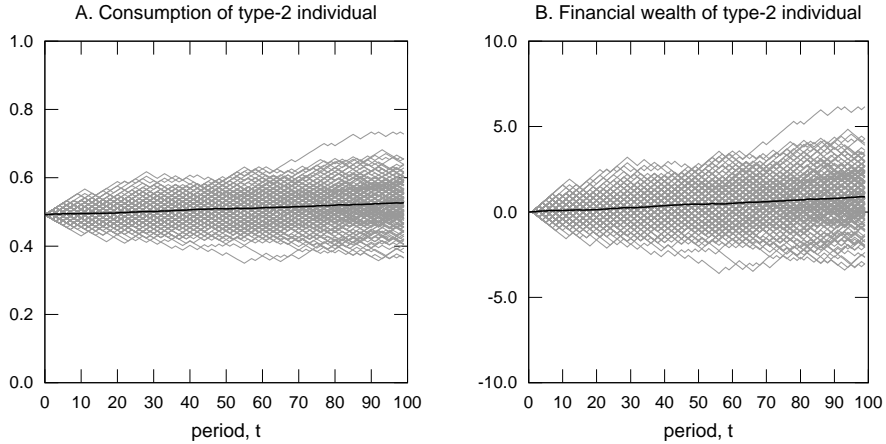
Figure 3: Sample paths of consumption and financial wealth of a type-2 agent in the bond economy.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, p^2 = 0.55, B = 8$.

A), and financial wealth rises from an average of 0 to 0.878, or 1.76 average individual annual incomes (panel B). As explained in Cogley et al. [2014], this occurs because the type-2 agent is pessimistic and buys bonds as a precautionary store of value.

Subjective welfare levels are -2.004 and -2.011, respectively, for the type-1 and type-2 agents. So both agents expect to be worse off than in the complete markets economy in which agents have common, correct beliefs. Objective welfare levels show that despite accumulating financial wealth, a type-2 agent has lower welfare. This occurs because pessimism motivates a type-2 agent to postpone consumption far into the future, which lowers expected utility.

6.3 Bond-only vs complete markets

If ($p^1 = p^0 = 0.5, p^2 = 0.55$) were the only admissible beliefs, our welfare criterion (with the Rawlsian aggregator) would select the bond-only design over the complete markets design. The former awards a substantial welfare level to both types because it limits speculation while still allowing resources to be transferred across periods. Under complete markets, type-1 agents take advantage of the poor forecasting abilities of type-2 agents, eventually

21

driving them to destitution.

Matters are more complicated when we consider a larger set of admissible beliefs. For instance, suppose $(p^1, p^2) \in [0.45, 0.55]^2$, and $p^0 = 0.5$.[12] Figure 4 plots the welfare surface $\min_i [U_{P^0}(c^i(p^1, p^2 | \mathcal{M}))]$ for this belief set.[13] The lowest welfare level under the bond-only design is -2.011, and it is achieved at $(p^1, p^2) = (0.45, 0.45)$ and $(0.55, 0.55)$. At these "critical points" (depicted by black points in the figure), beliefs are homogeneous but wrong.
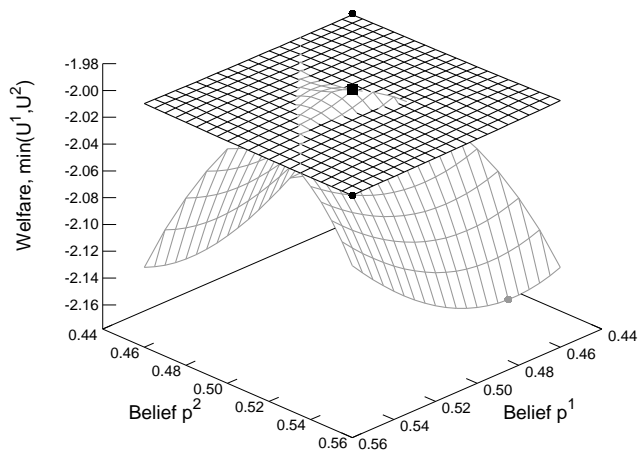
Figure 4: Welfare in example 1: the bond-only (black) vs the complete markets (gray) design. Square point denotes the unconstrained maximum: $(p^1, p^2, W^*) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare under the corresponding design.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 8$.

The lowest welfare in the complete markets economy is -2.139, and it is achieved at $(p^1, p^2) = (0.45, 0.525)$ and $(0.475, 0.55)$ (portrayed by gray points in the figure). At the critical points, beliefs are nearly maximally different. Consider the belief assignment $(p^1, p^2) = (0.45, 0.525)$. With these beliefs the

---

[12]Note that for now, we consider only one possible true data generating process. In section 8, we relax this restriction.

[13]The shape of this welfare surface is explained in Appendix A.2.

type-1 agent has lower welfare. Two forces act against him. First, his beliefs are less accurate, so his consumption is eventually driven to zero. Second, he is more pessimistic than a type-2 agent, and his endowment stream is valued less – he is subject to a negative wealth effect. But a type-2 agent is also pessimistic, and this activates a wealth effect that reduces a type-2 agent's welfare.

In this example, our welfare criterion (using the Rawlsian aggregator) selects the bond-only design over the complete markets design because:

$$-2.011 = \min_{P^1,P^2} \min_i U_{P^0}^i(c^i(\mathcal{P}|B)) > \min_{P^1,P^2} \min_i U_{P^0}^i(c^i(\mathcal{P}|CM)) = -2.139.$$

The complete markets design would be preferred if the set of admissible beliefs were concentrated tightly enough about the truth, for example, if it were reduced to $[0.49, 0.51]^2$.

This is not surprising as the complete markets design is, of course, preferred to the bond-only design with common, correct beliefs. It is surprising, though, that the bond-only design performs so robustly, at least when there is no aggregate risk.

Before we continue with our next market design we would like to describe the choice made by a designer who is guided by the true-Pareto welfare criterion. Figure 5 plots the set of beliefs for which the bond-only market design true-Pareto dominates complete markets. The dark gray area denotes belief assignments for which the bond-only CE allocation true-Pareto dominates the complete markets CE allocation. Naturally, this occurs where disagreement is strongest. As explained above, restricting financial trade to risk-free bonds effectively shuts down speculation, thereby increasing everyone's utility. The portion of the region at the bottom right corner is larger than at the top left corner. This is because these beliefs make agents optimistic and more willing to speculate. Hence, in our example, regulation is desirable over a larger set of parameters when agents are optimistic. The light gray area denotes belief assignments under which the two market designs cannot be ranked because one agent gains while another loses. It is useful to emphasize that this area is more than 50% of the whole set of beliefs. It shrinks dramatically when the truth is varied. The white area denotes belief assignments under which the complete markets dominate the bond-only design. This area includes beliefs that coincide with or are close to the truth. It also includes a narrow area parallel to the "agreement diagonal." In this portion of the parameter space, the effect of disagreement is offset by the bias in beliefs towards one of the

states. To illustrate, consider point $(p^1, p^2) = (0.475, 0.450)$. At this point agent 1 is closer to the truth and he is rewarded in financial markets that are unregulated. However, beliefs are stacked against him as both agents believe that he is relatively unlikely to receive high endowment. These two effects happen to offset each other leaving both agents relatively well off.
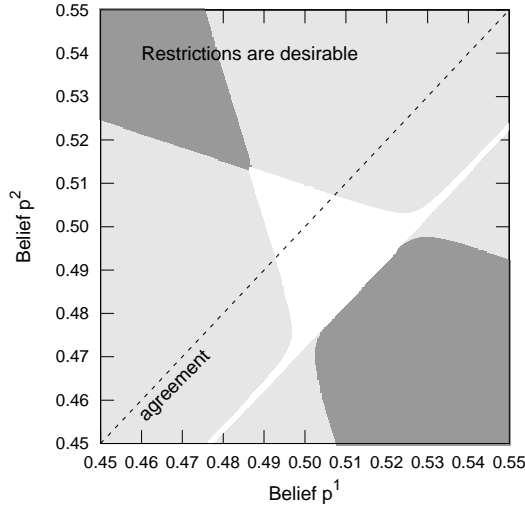

Figure 5: True-Pareto ranking: $\mathcal{B} \succ CM$ (dark gray), $CM \succ \mathcal{B}$ (white), allocations cannot be ranked (light gray).
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 1$.

6.4 Borrowing limits

We continue to assume that $p^0 = 0.50$ and that the admissible set of belief assignments is $(p^1, p^2) \in [0.45, 0.55]^2$. We impose a borrowing limit $B = 1$, equivalent to two average individual annual incomes. Figure 6 shows the social welfare surface for this environment (black) and contrasts it with the benchmark complete markets design (gray).

The square depicts the maximum achievable welfare in the two economies. It is reached at $(p^1, p^2) = (0.5, 0.5)$ in both cases and is equal to $W^* = -2$. When agents agree, there is little trading and borrowing limits are slack.

The two circles portray the minimum welfare achieved under the respective market designs. Under the design with borrowing limits, the low-
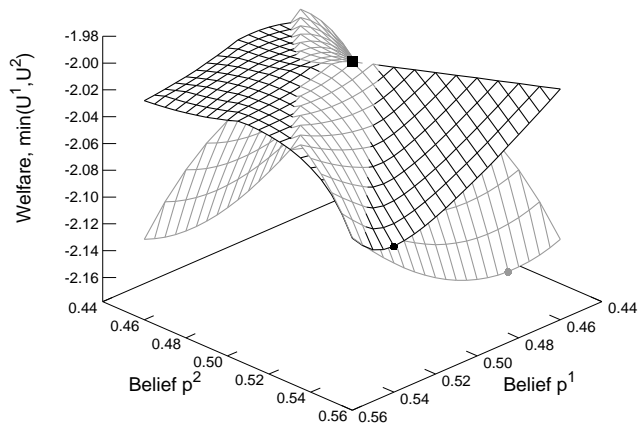
Figure 6: Welfare in example 2: the complete markets with borrowing limits (black) vs the complete markets design (gray). Square point denotes the unconstrained maximum: $(p^1, p^2, W^*) = (0.5, 0.5, -2)$. Circles denote belief assignments that attain the lowest welfare under the corresponding design. Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 1$.

est welfare levels are achieved at either $(p^1, p^2) = (0.45, 0.48)$ or $(p^1, p^2 = (0.53, 0.55)$. As in the bond economy, belief heterogeneity ceases to be the critical force defining the lowest welfare in the economy. Instead, at the critical belief assignments, agents nearly agree on one of the types being poor. For example, at the point $(p^1, p^2) = (0.45, 0.48)$, everyone agrees that a type-1 agent is less likely to receive high endowments. Moreover, a type-1 agent's beliefs are less accurate. For both reasons, his and the society's welfare is lower. At $(p^1, p^2) = (0.53, 0.55)$ it is a type-2 agent who suffers. Tightening the borrowing limit significantly lessens speculation and attenuates survival forces. For this example, society's welfare increases from -2.139 to -2.083, a difference equivalent to a 2.7% permanent increase in consumption.

Next we turn to an economy in which the type-1 agent knows the truth and the type-2 agent is pessimistic, $(p^1, p^2) = (0.50, 0.55)$. We compute means and standard deviations (in parentheses) of the following variables:

25

type-2 agent's financial wealth $a_t^2(\sigma)$, his consumption $c_t^2(\sigma)$ and prices of a risk-free bond $q_t^b(\sigma)$ and a claim to the aggregate endowment $q_t^e(\sigma)$ (see (4) for a definition).[14] We contrast two designs: complete markets with (restrictive) $B = 1$ and (relaxed) $B = 8$ borrowing limits. Table 1 summarizes our findings. First, financial wealth of the type-2 agent is 3.79 times less volatile under $B = 1$ than under $B = 8$. Second, consumption of the type-2 agent stays closer to 0.5 and it is also 2.43 times less volatile than under $B = 8$. A more nearly equal and less volatile distribution of consumption is the source of welfare gains in the design with the tight borrowing limit. Third, prices of the two financial assets are increased and also more volatile. That is, by tightening the borrowing limit the designer drives volatility out of consumption and into prices. This outcome illustrates that a goal of financial price stability may conflict with social welfare maximization.

	a^2	c^2	$ln(q^b)$	$ln(q^e)$
$B = 1$	-0.135	0.483	-0.037	3.275
	(0.612)	(0.038)	(0.008)	(0.030)
$B = 8$	-5.248	0.284	-0.041	3.182
	(2.317)	(0.092)	(0.002)	(0.022)

Table 1: Mean and standard deviation (in parentheses) for the complete markets with borrowing limit design

6.5   Transaction tax

Figure 7 shows welfare for our example under three market designs: complete markets with a natural borrowing limit, complete markets with an exogenous borrowing limit $B = 8$, and complete markets with $B = 8$ plus a transaction tax $\tau = 0.05$.[15] Welfare levels for the first two designs are very close, suggesting that competitive equilibrium allocations under $B = 8$ are close to allocations under the natural borrowing limits. Imposing a transaction tax on top of the borrowing limit increases society's welfare from -2.134 to -2.079, an amount equivalent to a permanent 2.6% increase in consumption.

---

[14]We simulated 11,000 periods starting from a random state and $(a_0^1, a_0^2) = (0,0)$. We discarded the first 1,000 observations.

[15]The natural borrowing limits are difficult to compute in the presence of a transaction cost. So, a generous exogenous borrowing limit is imposed instead.
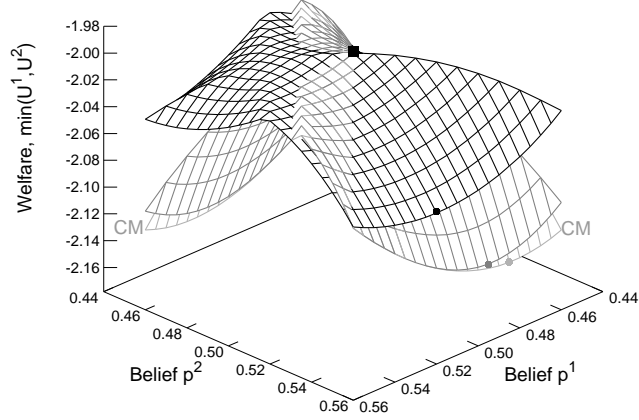
Figure 7: Welfare in example 3: complete markets with borrowing limits and transaction tax (black) vs complete markets with borrowing limit (dark gray) vs complete markets. Square point denotes the unconstrained maximum: $(p^1, p^2, W) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare under the corresponding market design.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 8, \tau = 0.05$.

7 Bergson-Samuelson criterion

Next, we examine an alternative utility aggregator that makes use of individuals' "market weights." This means that for any belief assignment $\mathcal{P}$ we first solve for the competitive equilibrium. Then we compute the vector of Pareto weights for which the competitive and the Pareto allocations coincide. Let $\theta^i(\mathcal{P})$ denote the implied Pareto weight, which we also call the market

weight, of type-i agent.[16] The corresponding social welfare function is:

$$\min_{P^0 \in \mathcal{B}^0} \min_{\mathcal{P} \in \mathcal{B}^I} \sum_{i=1}^{I} \theta^i U_{i,P^0}(c^i(\mathcal{P}|\mathcal{M})).$$

The social welfare criterion with this utility aggregator replaces the lowest welfare in the society with a particular weighted average of individual welfare levels. Under this criterion, the social welfare of an allocation cannot be driven by a small but disadvantaged group because its Pareto weight would, in general, be small. Nevertheless, using this aggregator we obtain qualitative results that are similar to those derived using the Rawlsian criterion.

Figure 8 plots the welfare weight of agent 2 for $(p^1, p^2) \in [0.45, 0.55]^2 \equiv \mathcal{B}$. The weights of the two types sum to one. The weight equals 0.5 on the diagonal where agents are symmetric opposites: $p^1 = 1 - p^2$. Consider now moving away from the diagonal towards $(p^1, p^2) = (0.45, 0.45)$. Agent 2 becomes more optimistic and agent 1 more pessimistic; so a type-2 agent's weight increases and a type-1 agent's weight decreases. This occurs because prices reflect the common belief that type 2 is more likely to receive high endowment. So the type 2 agent is wealthier and his market weight is higher.[17]

We now compare welfares of the bond-only economy and the complete markets economy. We continue to fix $p^0 = 0.5$. Figure 9 plots social welfare $\sum_i [\theta^i(\mathcal{P}) U_{P^0}(c^i(\mathcal{P}|\mathcal{M}))]$ for the two financial markets designs. When beliefs coincide with the truth, $p^1 = p^2 = p^0$, the welfare of each type agent is -2 – the maximum achievable under any market design (depicted by the gray square point). Social welfare is close to this benchmark when agents have common beliefs, even if those beliefs are wrong, i.e., on the diagonal with $p^1 = p^2$. Close to the diagonal, welfare under the bond-only design and under the complete markets design are similar. But the latter is higher because disagreement is small and survival forces are weak. So while agents are driven out of financial markets, this occurs slowly.

As we move away from the common beliefs diagonal, social welfare stays robustly high under the bond-only design, but declines under the complete

---

[16]With logarithmic preferences Pareto weights are date-0 wealth shares so the weight of agent i is the proportion of the aggregate wealth owned by him. This suggests yet another possibility, namely, to use wealth shares from the complete markets competitive equilibrium.

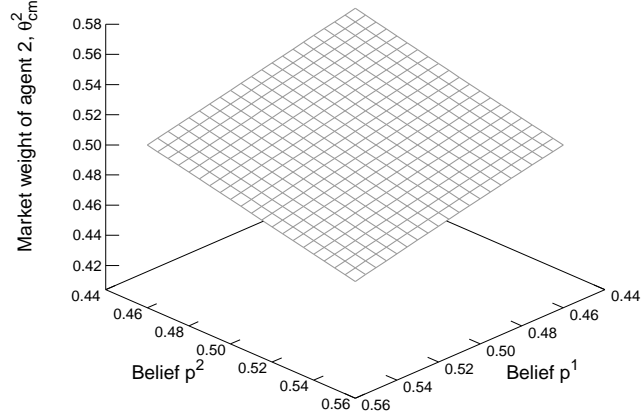[17]This is a manifestation of the wealth effect that may sometimes dominate survival/speculative forces.

Figure 8: Implied Pareto weight of the type-2 agent θ^2 for the complete markets economy.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5$.

markets design. The reason for the robust performance of the bond-only design is that it limits survival forces. That is, differences in beliefs have only a limited effect on the equilibrium outcome when only a risk-free bond is traded. The lowest welfare under the bond-only design is achieved at $(p^1, p^2) = (0.55, 0.55)$ and $(p^1, p^2) = (0.45, 0.45)$. At these belief assignments, both types incorrectly believe that one of them is more likely to receive a high endowment. As the common belief is reflected in the bond price, the believed-to-be-poor type turns out to be poor in fact. So social welfare is low because the discrepancy between agents' individual welfare levels is large. The lowest welfare under the complete markets design is achieved at $(p^1, p^2) = (0.45, 0.55)$ and $(p^1, p^2) = (0.55, 0.45)$ (depicted by gray circles in the figure). At these points beliefs are maximally heterogeneous. So speculative motives are strong and survival forces occasionally drive each agent arbitrarily close to losing all of his wealth. As a result, consumption is volatile and social welfare is low. We conclude that for this example the

bond-only design dominates the complete markets design:

$$-2.132 = \min \mathcal{P} \sum_{i=1}^{I} \theta^i U_{P^0}^i(\mathcal{P}|CM) < \min \mathcal{P} \sum_{i=1}^{I} \theta^i U_{P^0}^i(\mathcal{P}|\mathcal{B}) = -2.011.$$
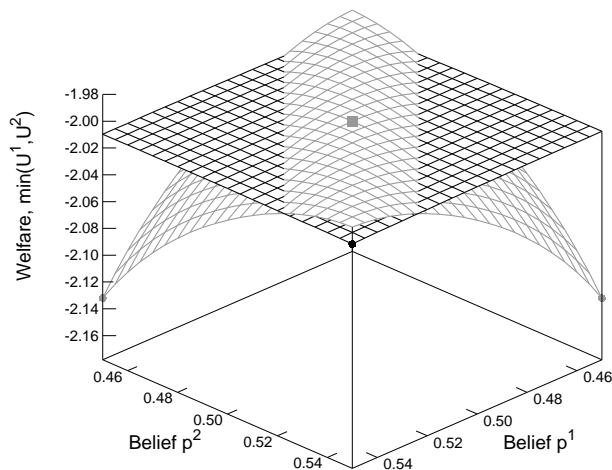

Figure 9: Welfare in example 1: the bond-only (black) vs the complete markets (gray) design. Square point denotes the unconstrained maximum: $(p^1, p^2, W^*) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare under the corresponding design.
Parameters: $\beta = 0.96$, $e_l = 1/3$, $e_h = 2/3$, $p^0 = 0.5$, $B = 8$.

We next compare the complete markets economy with borrowing limits and the unrestricted complete markets design. Figure 10 plots welfare surfaces corresponding to the two designs. We impose a borrowing limit of $B = 1$ or two average annual incomes. This limit is restrictive compared with the natural borrowing limits, but it would not bind in a competitive equilibrium with complete markets if both agents had correct beliefs. With diverse beliefs, on the other hand, any tight exogenous borrowing limit must be binding. Moreover, any borrowing limit is more restrictive when agents
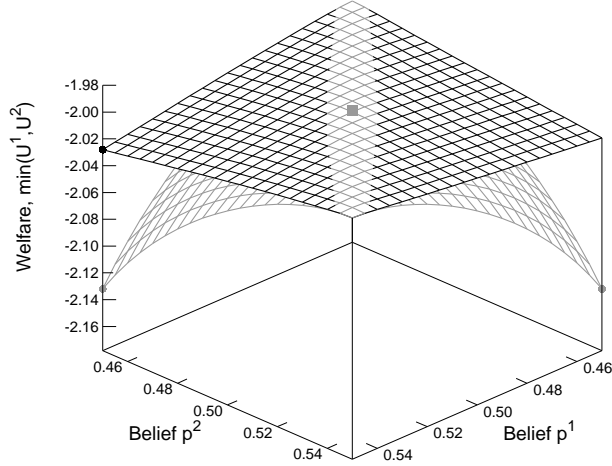
30

Figure 10: Welfare in example 2: the complete markets with borrowing limits (black) vs the complete markets design (gray). Square point denotes the unconstrained maximum: $(p^1, p^2, W^*) = (0.5, 0.5, -2)$. Circles denote belief assignments that attain the lowest welfare under the corresponding design. Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 1$.

are pessimistic than when they are optimistic.[18] For this reason, the worst belief assignment is $(0.45, 0.55)$ – when disagreement is maximal and both types of agent are pessimistic. A tight borrowing limit restricts the maximal amount of wealth that can be lost by any agent and bounds consumption away from zero. Yet borrowing limits are attained only when wealth becomes unevenly distributed. When disagreement is small, wealth is always close to being equally distributed so borrowing limits rarely bind. For this reason, the two market designs deliver similar welfare levels when disagreement is small in which case imposing even the tight borrowing limit $B = 1$ is nearly inconsequential. We conclude that the design with the borrowing

---

[18]This occurs because insurance and speculative motives align and prompt both types to purchase more Arrow securities for the low income state as in Tsyrennikov [2012].

limit $B = 1$ dominates the design with natural borrowing limits:

$$-2.132 = \min \mathcal{P} \sum_{i=1}^{I} \theta^i U_{P^0}^i(\mathcal{P}|CMT) < \min \mathcal{P} \sum_{i=1}^{I} \theta^i U_{P^0}^i(\mathcal{P}|CMB) = -2.028.$$

Finally, we analyze the effect of a transaction tax on each financial transaction as specified in (7). This tax limits speculation, but it also restricts hedging possibilities. As above, the new friction matters more when agents are pessimistic. So the worst belief assignment under both designs is $(p^1, p^2) = (0.45, 0.55)$; see figure 11. Welfare is lowest in this case because 1) survival forces are strongest, and 2) agents' pessimism prompts agents to trade/hedge more actively, making the transaction tax harmful. Nonetheless, imposing a transaction tax of $\kappa = 0.05$ increases social welfare relative to complete markets:

$$-2.118 = \min \mathcal{P} \sum_{i=1}^{I} \theta^i U_{P^0}^i(\mathcal{P}|CM) < \min \mathcal{P} \sum_{i=1}^{I} \theta^i U_{P^0}^i(\mathcal{P}|CMB) = -2.049.$$

8   Dependence on P^0

Our welfare criterion uses three min operators. But so far we have demonstrated the use of the criterion only for a singleton $\mathcal{B}^0$. In this section, we confront our designer with multiple data-generating processes: $|\mathcal{B}^0| > 1$. Recall that our theoretical results hold for any P^0 and, hence, for any $\mathcal{B}^0$. Our numerical examples also show that welfare varies more under the complete markets design than under the designs with financial restrictions. By introducing ambiguity about P^0 via expansion of $\mathcal{B}^0$, we expect welfare gains from financial restrictions to increase. These expectations are confirmed by the results reported in table 2. Here our welfare criterion uses the Rawlsian aggregator.[19]

In constructing Table 2, we have assumed that $(p^1, p^2) \in [0.45, 0.55]^2$. That is, for each choice of $p^0 \in [0.45, 0.50]$ we report $\min_{P^1, P^2} \min_i W_{P^0}^i(c^i(\mathcal{P}|\cdot))$.[20]

---

[19]We use the following notation: $W_{P^0}(\mathcal{M}) = \min_{P^1, P^2} \min_i U_{P^0}^i(c^i|\mathcal{M})$.

[20]The results for $p^0 \in [0.50, 0.55]$ are symmetric. So both at $p^0 = 0.55$ and at $p^0 = 0.45$ we get $W(CMB) = -2.171, W(CM) = -2.545$. Only the identity of the less well-off agent changes.
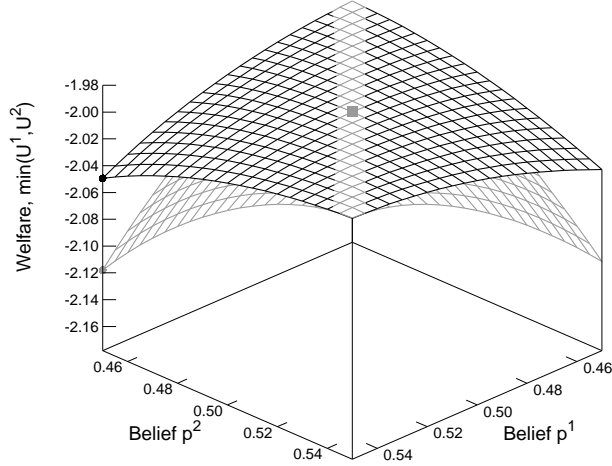
Figure 11: Welfare in example 3: complete markets with borrowing limits and transaction tax (black) vs complete markets with borrowing limit (gray). Square point denotes the unconstrained maximum: $(p^1, p^2, W) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare under the corresponding market design.

Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 8, \tau = 0.05$.

Columns 2 through 5 present welfare under the unrestricted complete markets design (section 4.1), the bond economy (section 4.2), complete markets with borrowing limits (section 6.4), and complete markets with transaction tax (section 6.5) respectively. All of these financial designs achieve the lowest welfare at $p^0 = 0.45$. Welfare under the complete markets is $W(CM) = -2.545$, the lowest among our financial designs.

The best performing design is the bond-only economy that achieves welfare level $W(B) = -2.084$. It offers an improvement over the complete markets design equivalent to a permanent 22.1% increase in consumption. The design with borrowing limit $B = 1$ dominates the design with transaction tax $\tau = 0.05$: $W(CMB) = -2.121 > -2.234 = W(CMT)$. The former offers substantial improvement over the complete markets design, equivalent to a permanent 20.0% increase in consumption. But it under performs relative to

P^0	$W_{P^0}(CM)$	$W_{P^0}(B)$ $B = 8$	$W_{P^0}(CMB)$ $B = 1$	$W_{P^0}(CMT)$ $B = 8, \tau = 5\%$
1	2	3	4	5
0.45	**-2.545**	**-2.084**	**-2.121**	**-2.234**
0.46	-2.439	-2.068	-2.113	-2.193
0.47	-2.347	-2.053	-2.106	-2.154
0.48	-2.267	-2.039	-2.098	-2.117
0.49	-2.195	-2.025	-2.090	-2.094
0.50	-2.139	-2.011	-2.083	-2.079

Table 2: Welfare level under different P^0: the designs with financial restrictions (B,CMB,CMT) vs the complete markets design (CM).
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5$.

the bond-only design. The design with a transaction tax does not perform well when $p^0 \neq 0.5$. This happens because as p^0 diverges from 0.5 agents must take larger financial positions to hedge income fluctuations. These are costly due to the transaction tax.[21]

The worst-case beliefs assignment for the complete markets design is $(p^0, p^1, p^2) = (0.45, 0.45, 0.55)$. This point assigns correct beliefs to type-1 agents and maximally wrong beliefs to type-2 agents. This worst-case choice of beliefs maximizes the strength of survival forces. Type-2 agents have the lowest welfare. For the bond-only design. the worst-case assignment of beliefs is $(p^0, p^1, p^2) = (0.45, 0.55, 0.535)$. Type-1 agents have the lowest welfare. Under this belief assignment, type-1 agents wrongly believe that they are more likely to receive a high endowment. So they dis-save and end up consuming less than type-2 agents. In addition, type-1 agents have less accurate beliefs that guide them to worse financial decisions. But because the bond return adjusts and because there are limited speculation opportunities, type-1 agents lose wealth very slowly. This makes the bond-only economy a substantially more robust design than the complete markets. Under com-

---

[21]To build intuition, consider the case with correct and homogeneous beliefs. Recall that in the initial state, $z = 0$, both agents receive the same income 0.5. When $p^0 = 0.5$ agents trade to reallocate income across states. When $p^0 \neq 0.5$ agents get an additional motive to trade: to reallocate income across time. This motive appears because expected individual income is no longer 0.5 and agents want to borrow or lend against the future income. Because trading is costly, agents end up with 'suboptimal' positions. See also derivations in section 6.

plete markets with a borrowing limit, the worst-case assignment of beliefs is $(p^0, p^1, p^2) = (0.45, 0.515, 0.55)$. Type-2 agents have the lowest welfare, first, because their beliefs are less accurate and, second, because both types agree that type-2 agents are less likely to receive a high endowment. This forces type-2 agents to stay close to a restrictive borrowing limit. However, unlike outcomes under the complete markets design, the strict borrowing limit $B = 1$ allows type-2 agents to rebuild their financial wealth quickly. Under complete markets with a transaction tax, the worst-case assignment of beliefs is $(p^0, p^1, p^2) = (0.45, 0.45, 0.55)$. Type-2 agents have the lowest welfare. This design is better than complete markets because a transaction tax limits speculation. But a transaction tax also limits the speed of type-2 agent's recovery once he runs into financial trouble. This makes the design with a transaction tax worse than the other restrictions.

The larger $\mathcal{B}^0$ and $\mathcal{B}$, the starker are the welfare differences. Reasonable choices of $\mathcal{B}^0$ and $\mathcal{B}$ can be constructed using error detection probabilities as in Hansen and Sargent [2007].[22]

8.1 Putting our welfare criterion to work

A benefit of our welfare criterion is that it can be immediately applied to determine optimal financial market restrictions. As an example, we demonstrate how to compute an optimal borrowing limit.[23] Consider the complete markets financial design with borrowing limits. Previously we imposed an exogenous borrowing limit $B = 1$. We now compute the optimal borrowing limit:

$$B^* = \arg\max_{B} \min_{P^1, P^2, P^0} \min_i W_{P^0}^i(CMB). \qquad (15)$$

In our example, the optimal borrowing limit B^* is 36% of an average annual income. In the economy with homogeneous beliefs, agents would borrow 33% of an average annual income.[24] So the optimal borrowing limit is just over what is needed to hedge income fluctuations.

---

[22]This approach allows forming a set of models that are reasonably hard to distinguish using a log-likelihood ratio test and a finite data sample.

[23]The magnitudes computed in this example are meant only as an illustration of how to apply our welfare criterion. Serious policy proposals would need a more realistic model.

[24]This is not the natural borrowing limit but an equilibrium borrowing amount.

9   Concluding remarks

We propose a framework and a welfare criterion for evaluation of different financial market designs. Our setting is an endowment economy in which agents may hold heterogeneous beliefs. We imagine a social planner who chooses a financial market design to maximize social welfare before beliefs and the true data generating process are assigned.

We use our criterion to study a simple economy. Our analysis illustrates the trade-offs between welfare-reducing speculation and welfare-improving insurance possibilities. Complete financial markets allow maximal insurance possibilities, but for economies with heterogeneous beliefs they also allow social welfare reducing speculation. We find that in the economies that we study, financial market designs with simple restrictions like limits on the set of traded assets, borrowing limits, and transaction taxes offer substantial welfare gains relative to a complete financial markets benchmark. In our examples, gains can be as large as those stemming from a 6% permanent increase in consumption.

Our simulations do not allow for the possibility of traders who learn the truth, but the analysis of learning in Blume and Easley [2006] suggests that the story would not be much changed. They find that among Bayesian learners, those selecting across a countable set of models containing the true model do not vanish. In the more interesting case of individuals selecting across a space of models of dimension at least 2, those selecting across a given model space vanish in the presence of individuals selecting across a lower-dimensional set. Some learners learn faster than others, and the slow learners do not learn quickly enough to keep their consumption from becoming negligible. And finally, in the presence of any learners, individuals who do not learn disappear. Here too, limiting the bets that traders can take slows down or prevents financial ruin of slow learners.

We believe that the most important limitation of our analysis is the absence of incentive effects. That is, in our analysis restrictions imposed on the financial markets have no effects on the set of feasible allocations. Relaxing this feature is arguably the most profitable direction for future research in this line of work.

References

Robert J. Aumann. Agreeing to disagree. *Annals of Statistics*, 4(6):1236–1239, 1976.

Lawrence E. Blume and David Easley. If you're so smart, why aren't you rich? belief selection in complete and incomplete markets. *Econometrica*, 74(4):929–966, 2006.

Jaroslav Borovicka. Survival and long-run dynamics with heterogeneous beliefs under recursive preferences. NYU manuscript, 2012.

Markus Brunnermeier, Alp Simsek, and Wei Xiong. A welfare criterion for models with distorted beliefs. *Quarterly Journal of Economics*, forthcoming.

Tim Cogley, Thomas Sargent, and Viktor Tsyrennikov. Wealth dynamics in a bond economy with heterogeneous beliefs. *Economic Journal*, 124(575):1–30, 2014.

D. Duffie. Challenges to a policy treatment of speculative trading motivated by differences in beliefs. *Journal of Legal Studies*, 43(2):173–182, 2014.

Larry G. Epstein and Stanley E. Zin. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework. *Econometrica*, 57(4):937–969, 1989.

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

Lars Hansen and Thomas Sargent. *Robustness*. Princeton University Press, 2007.

Edi Karni and Philippe Mongin. On the determination of subjective probabilities by choices. *Management Science*, 46(2):233–248, 2000.

Edi Karni and David Schmeidler. On the uniqueness of subjective probabilities. *Economic Theory*, 3(2):267–77, 1993.

Philippe Mongin. Spurious unanimity and the pareto principle. LSE manuscript, 2005.

Stephen Morris. The common prior assumption in economic theory. *Economics and Philosophy*, 1(2):227–253, 1995.

C. Phelan and A. Rustichini. Pareto efficiency and identity. NBER Working paper 20883, 2015.

John Rawls. *A Theory of Justice*. Belknap Press of Harvard University Press, 1971.

Leonard J. Savage. *The Foundation of Statistics*. New York: John Wiley, and New York: Dover Publications, 1954 (1972).

George J. Stigler and Gary S. Becker. De gustibus est disputandum. *American Economic Review*, 67(2):76–90, 1977.

Viktor Tsyrennikov. Heterogeneous beliefs, wealth distribution, and asset markets with risk of default. *American Economic Review, Papers and Proceeedings*, 102(3):156–160, 2012.

A   Appendix

A.1   Beliefs

As much of our analysis deals with beliefs and their correctness it is important to be aware of the foundations for our ability to theorize about individuals' beliefs. This foundation is Savage [1954 (1972] subjective expected utility representation theorem that delivers for each preference order satisfying his axioms a payoff function and a probability vector that together generate an additively separable representation over state-contingent payoffs. Although Savage's theorem does not compel any particular interpretation, economists and game theorists typically take the payoff function as representing tastes, such as attitudes towards risk, and the probability distribution as representing beliefs.

The argument that one can extract beliefs from preferences depends critically on the supposed uniqueness of the probability distribution in Savage's representation theorem. Unfortunately, uniqueness requires an assumption about the representation that cannot be verified by observed choice behavior alone. Suppose that a preference order for acts mapping states $s \in S$ to

outcomes $y \in Y$ has an expected utility representation: a payoff function $u : Y \to \mathbf{R}$ and a probability distribution p on S. The uniqueness theorem states that if v and q combine to give another expected utility representation of the same preference, then v is a positive affine transformation of u and q equals p. This result, however, is limited to state-independent payoff representations (not state-independent preferences, but a restriction to representations that are state-independent).[25] Savage's axioms include a "state-independence" assumption, that preferences conditional on two distinct non-null states are identical. This allows the possibility of a state-independent payoff function, but it does not rule out state-dependent payoffs. Together with the other axioms, the only requirement it imposes on $v : S \times Y \to \mathbf{R}$ is that the function $v(s', \cdot) : Y \to \mathbf{R}$ is a positive affine transformation of $v(s, \cdot) : Y \to \mathbf{R}$ (whenever s and s' are both non-null).

Interpreting probabilities in expected utility representations as likelihood assessments requires uniqueness of the probability distribution in the larger class of state-dependent expected utility representations. Pinning the positive affine transformations down to translations is the necessary condition for deriving uniqueness, but this requires an extension to the structure of preferences that is not revealed in choice behavior.[26] If one insists that individual preferences have expected utility representations, then the commitment that individuals have identical beliefs can only be justified by non-choice considerations even when individual preferences can be represented by (perhaps different) state-independent payoff functions and a common probability distribution.

A.2 Complete markets design

In this section we explain the shape of the welfare surface under the complete markets design. Two forces are key to understanding this surface. The first is the survival force: the type of agent with the least accurate beliefs has his wealth drift downward and likely to have the lowest welfare. The second is the wealth effect: an equilibrium price system is affected by the configuration

---

[25]If we allow that tastes can depend upon states, so that payoff functions can map $S \times Y$ into $\mathbf{R}$, then the only thing unique about the probability distribution is its support. For any q with support identical to p, there is a state-dependent payoff function v such that v and q combine to represent the preferences.

[26]See Karni and Schmeidler [1993] and Karni and Mongin [2000].

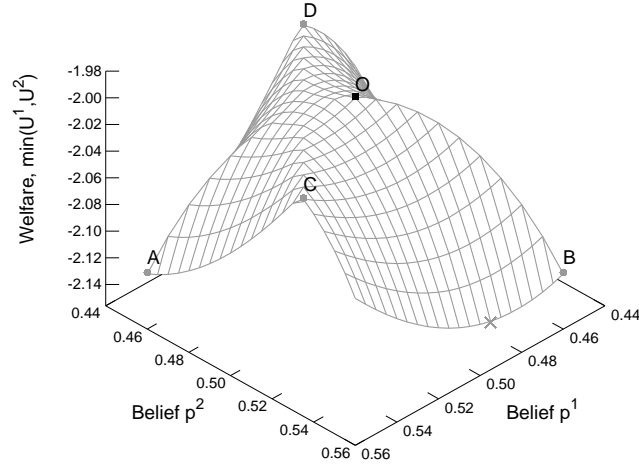of beliefs and this may present an advantage to one of the types.[27]

Figure 12: Welfare in example 1 under the complete markets design (gray). Square point denotes the unconstrained maximum: $(p^1, p^2, W) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5$.

Figure 12 reproduces the welfare surface shown in figure 4. Along arc AOB both types are either optimistic or both pessimistic. The wealth effects for each type offset each other. So welfare is decreasing as we move away from point O because agents disagree more on individual states and accept more volatile consumption. When we perturb beliefs slightly away from the arc, welfare drops. This happens because of the wealth effect. Independently of the direction in which beliefs are perturbed, one type's wealth will be affected negatively, and this reduces both his and society's welfare.

Along arc CD both types are close to agreement, but $p^1 > p^2$. Consider the closer half of arc CD where $p^2 \geqslant 0.5$. Then a type-1 agent is optimistic

---

[27]When beliefs are equally accurate, the direction can be determined by looking at the date-0 consumption level. If the wealth effect impacts both types equally then $c_0^i = 0.5$.

and a type-2 agent is pessimistic. This configuration of beliefs is advantageous to a type-1 agent. (See also our two period example below.) But a type-1 agent also has less accurate beliefs. So he is affected adversely by survival forces. The latter partially offsets the wealth effect and creates a ridge along arc CD.[28]

A.2.1 Wealth effect

It is instructive to study a simple two-period economy. This example demonstrates that an agent with less accurate beliefs can secure a higher objective welfare. The key to this result is a wealth effect.

The period utility function is $u(c) = log(c)$, and future utility is not discounted. The state in period 0 is known, and there are two possible state realizations in period 1. Endowments for the two types are $(0.5, 0.5)$ in period 0. In period 1, they are $(1, 0)$ when the state is 1 and $(0, 1)$ when the state is 2. Under the true probability distribution, both states are equally likely. A type-1 agent's beliefs coincide with the truth. But a type-2 agent believes that $prob(s = 1) = 0.5(1 - \Delta) \neq 0.5$. Depending on whether $\Delta > 0$ or $\Delta < 0$ a type-2 agent is optimistic or pessimistic.

If both types had correct beliefs, in a competitive equilibrium allocation with complete markets, every agent would consume 0.5 in every period and state.

When markets are complete, the optimal consumption plan of a type-2 agent is:

$$c_0^2 = \frac{1}{2 - \Delta}, \quad c_1^2(s = 1) = \frac{1 + 2\Delta}{2 + \Delta}, \quad c_1^2(s = 2) = \frac{1 - 2\Delta}{2 - 3\Delta}. \qquad (16)$$

Two aspects of this equilibrium are important. First, consumption of agent 2 is decreasing on average for all $\Delta \neq 0$:

$$E[c_1^2] = c_0^2 \frac{4 - 4\Delta^2 c_0^2}{4 - \Delta^2 (c_0^2)^2} < c_0^2. \qquad (17)$$

Here the agent with incorrect beliefs is gradually being "driven from the market." Second, if a type-2 agent is optimistic ($\Delta < 0$), then his consumption in period 0 is higher than 0.5. Lastly, the agent with incorrect beliefs may

---

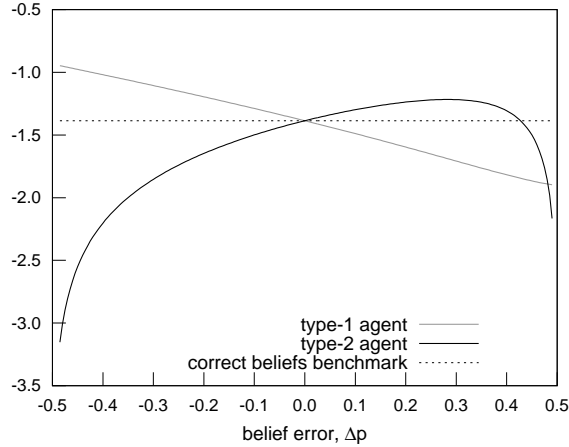[28]Along the more distant half of arc CD, the roles of the two types reverse.

Figure 13: Objective welfare in the two-period complete markets economy

have higher objective welfare:

$$\frac{dW^2(\Delta)}{d\Delta}\Big|_{\Delta=0} = 1 \neq 0, \tag{18}$$

where $W^2(\Delta) \equiv ln(c_0^2) + 0.5ln(c_1^2(s = 1)) + 0.5ln(c_1^2(s = 2))$. Here agent 2 can be better off being an optimist. But $\lim_{p \to 0.5} W^2(p) = -\infty$. Figure 13 plots welfare of the two types of agent. The horizontal dotted line denotes the welfare level in the economy in which beliefs of each agent coincide with the truth. A type-2 agent benefits from being optimistic because of his impact on the equilibrium price system. Optimism increases the relative price of goods delivered in state $s = 2$. This is the wealth effect.

A.3   On choice of preference specification

We made two important assumptions about individual preferences. The first is that preferences are time separable and the second is that the period utility function is unbounded below. Neither is crucial for our analysis.

Suppose that individual preferences have a recursive utility representation as in Epstein and Zin [1989]. When markets are complete and agents have diverse beliefs, some agent types will be driven out of financial markets. The difference is, as Borovicka [2012] shows, that it may not be the agent with the most accurate beliefs who survives as in Blume and Easley [2006]. But

so long as there are agents that could be driven out of financial markets, there is a case for financial regulation. Our arguments could be regarded as becoming more compelling in this case because speculation may impoverish agents with more accurate beliefs.

When the period utility function is bounded from below, survival forces could be stronger because potential financial losses have lower utility cost. We demonstrate this by changing the period utility specification to $u(c) = \sqrt{c}$. Figure 14 plots welfare surfaces for the complete markets design and the complete markets with an exogenous borrowing limit design. Welfare levels under the two financial designs are, respectively, 1.3033 and 1.3762, a difference equivalent to a permanent 5.59% increase in consumption. The welfare effect of imposing the borrowing limit $B = 1$ is less significant than with $u(c) = -1/c$, but the set of beliefs for which the complete markets design is preferred is smaller. Although survival forces are stronger and agents can lose financial wealth more quickly, the welfare effect of losing wealth is less significant.
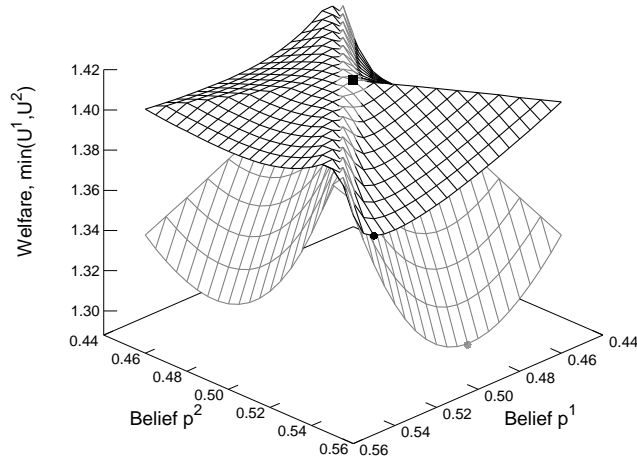

Figure 14: Welfare with bounded below utility for complete markets (gray) and complete markets with a tight borrowing limit (black).
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, p^2 = 0.55, B = 8.$

A.3.1 Effects of time preference

The choice of financial design also depends on the discount factor β. To illustrate the effect of time preference, we fix $p^1 = p^0 = 0.5$ and specify the admissible belief set as $p^2 \in [0.45, 0.55]$. Then we let the common discount factor β vary between 0.8 and 0.99. Figure 15 plots the social welfare surface (again using the Rawlsian aggregator) under the bond-only (black) and complete markets (gray) designs.[29]
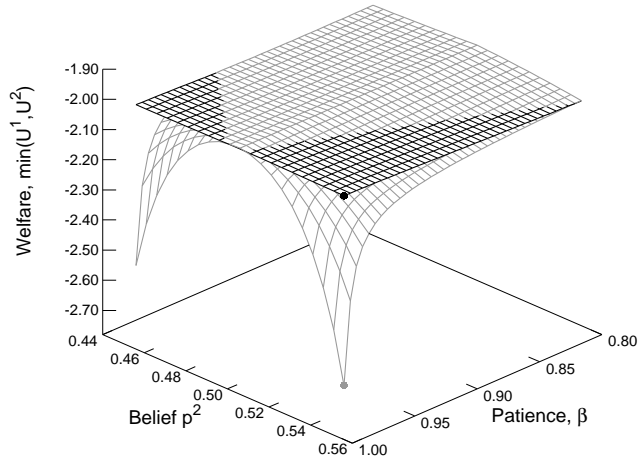

Figure 15: Welfare in example 1: the bond-only (black) vs the complete markets (gray) design. Circle points denote belief assignments that attain the lowest welfare under the corresponding design when $\beta = 0.99$.
Parameters: $e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, B = 2$.

As the discount factor increases, the minimum welfare in the bond economy dominates the complete markets economy on a larger set of belief specifications. This happens because agents care more about the limiting behavior of their consumption plans when they are more patient. So, their welfare can

---

[29]Note that the borrowing limit under the bond-only design was tightened so that we could study preferences with a discount factor as low as 0.8.

be low even when disagreement is small under the complete markets design.[30]
For instance, for $\beta = 0.99$, social welfare is -2.637 and -2.008, respectively,
under the complete markets and the bond-only designs. In this case, restrict-
ing financial markets to allow trade of only a risk-free bond is equivalent to
a permanent 31.3% increase in consumption.

---

[30]Roughly speaking disagreement affects the speed at which an agent with less accurate
beliefs can lose wealth. So, more patient agents care about longer horizons and they can
loose substantial amounts of wealth over long periods of time even if they are losing slowly.