A Case for Incomplete Markets*

Lawrence E. Blume† Timothy Cogley‡ David A. Easley§
Thomas J. Sargent¶ Viktor Tsyrennikov‖

September 20, 2018

Abstract

If two rational agents want to trade and there are no externalities, then trade is Pareto improving. Economists generally oppose restrictions on such trade. Complete markets allocations are Pareto optimal and thus complete markets are generally viewed as good. But when individuals want to trade because of heterogenous beliefs, this standard argument is less compelling. We illustrate this in a standard general equilibrium setting and explore potential social benefits from restrictions on trade that make markets incomplete.

Keywords: social welfare, heterogeneous beliefs, spurious unanimity, speculation, 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

*We thank Jonathan Parker for a helpful and detailed discussion. We also thank two referees for this journal, V.V. Chari, Vincenzo Quadrini, and seminar participants of the 2014 SITE conference, the 2014 ITAM Macroeconomics Workshop, the 2013 NBER Summer Institute, 2013 Econometric Society Australian meetings, 2012 North American Econometric Society meetings, 2012 Cornell/PennState Macroeconomics workshop.

†Cornell University, lb19@cornell.edu
‡New York University, tc60@nyu.edu
§Cornell University, dae3@cornell.edu
¶New York University and Hoover Institution, ts43@nyu.edu
‖Promontory, viktor.tsyrennikov@gmail.com

Pareto optimal and that any optimal allocation can be realized by trade in complete markets with an appropriate lump-sum transfer scheme. So limits that close trading opportunities leave 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 lump-sum transfers are impossible. 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 costly to execute. Consequently, these concerns are typically considered to be second-order.

We offer here a different and perhaps more fundamental critique of complete markets. When markets allocate contingent claims among expected-utility-maximizing agents who have heterogeneous beliefs, a market designer who cares about the realized discounted utility of consumption paths may prefer restrictions on trade to complete markets. As complete markets allocations are Pareto optimal (for the classical economies we study) we necessarily reject Pareto optimality as an appropriate welfare criterion for these economies. Our critique is motivated by two observations in the literature. First, Blume and Easley [2006] show that if traders have heterogenous beliefs, a common discount factor and access to complete markets, only those traders whose beliefs are most nearly correct survive. The consumptions of all other traders converge to zero almost surely. Although the resulting allocation is Pareto optimal the objective impoverishment of traders with more incorrect beliefs seems undesirable. Second, Mongin [2005] shows in his discussion of "spurious unanimity" that, in the presence of heterogeneous beliefs, Pareto optimality is not a compelling welfare criterion. Other more recent literature (Brunnermeier et al. [2014] and Gilboa et al. [2014]) proposes alternative welfare criteria for economies with heterogeneous beliefs. We do not advocate any of these welfare criteria; rather we argue that regardless of how a market designer views welfare, as long as he cares about the realized discounted utility of consumption paths, for some economies he should prefer markets in which trade is restricted to complete markets.

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 actual data generating process were known to an omniscient social planner, Pareto calculations with correct beliefs is an obvious fix. Omniscient social planners do not exist, however, and without them there is no alternative welfare requirement that obviously ameliorates the issues raised in section 3. Section 3 also includes a discussion of alternative welfare criteria offered in the recent literature. We investigate the market design problem through computation of competitive equilibria for simple, classical economies. In sections 4 and 5 we describe the set of market restrictions we consider and how we propose that the market designer views welfare when traders have heterogeneous beliefs. Sections 6, 7 and 8 examine several policy alternatives to complete markets in Markovian instances of the model of financial restrictions developed in section 4, and there we explore the size and location of the set of data generating processes and beliefs for which these policies would lead to a true welfare improvement. We conclude in section 9 with a discussion of the theoretical and the policy implications of our findings. All proofs are provided in the Appendix A.1.

2 The model

We assume that time is discrete and begins at date 0. At each date a state is drawn from the set $\mathcal{S} = \{1, \dots, S\}$. 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. This probability may or may not be known to the market designer. But we, the modelers, use it to describe the actual outcomes of any market design, which of course, are not known to the market designer unless she knows P^0. 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 several random variables of the form $x_t(\sigma)$. All such random variables are assumed to be date-t measurable; that is, their values depend 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, $\sum_i (c_t^i(\sigma) - e_t^i(\sigma)) = 0$.

We assume that consumer i's preferences on consumption plans have a subjective expected utility representation 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 i's 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. However, 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.*

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). While we do certainly agree that tastes for apples and oranges, work and leisure, etc., are to be taken as given, we do not accept the claim that *ex ante* preferences on contingent claims are above dispute.

Following Savage (1954), we represent these ex-ante preferences with time-0 expected utility described by a payoff function and a probability 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. A merit of representation theorems is that they give us access to these objects (payoff functions and beliefs) that we can reason about and work with. When we reason about the correctness of beliefs, and ask whether an individual would prefer one consumption stream to another if only they had correct beliefs, we are indirectly reasoning about preferences. This is standard in the economics profession as is evidenced by the common assumptions of rational expectations in GE and Nash equilibrium in games. But it is important to note that, just as proponents of rational expectations do, we are rejecting some preferences in favor of others.

We do not require market participants to have rational expectations; that is, common, correct beliefs. Instead we allow them to have differing beliefs. When market participants have different beliefs, not all can be correct, and those who are wrong are making decisions that they would regard as suboptimal if only they had correct beliefs. Much of our analysis begins with this fact and asks about the social value of restrictions on trade that would prevent individuals from making decisions that would not be optimal if only they had correct beliefs. It is important to note that we are not asking about restrictions on trade that would prevent *ex post* regret. Rather, in the spirit of *ex ante* Pareto optimality we are asking about *ex ante* welfare improvements with respect to correct beliefs.

It is important to note that we do not restrict payoff functions beyond the restrictions imposed by the standard assumptions of time separability, geometric discounting, risk aversion and strict desirability of consumption. Although the issue does not arise in our setting with a single consumption good, we do not dispute that individuals may have differing tastes for apples and oranges. We do not propose to restrict their abilities to trade these goods within a time period. We also do not dispute individuals' time preferences. We conduct our analysis for an economy with a common discount factor, so trade arising from differences in time preferences does not arise in our economy.

To the extent that constraints on the market restrict trade arising only from differing time preferences or from differing payoff functions, rather than differing beliefs, those constraints are harmful. In a heterogeneous agent economy, trade may arise from some combination of these factors. So there is a tradeoff between restricting beneficial trade and speculative trade. We view our contribution as showing that such a tradeoff actually exists, illustrating its possible magnitude, and providing examples in which simple restrictions are socially beneficial. We leave it to the reader to decide how to weigh the gains and loses arising from the restrictions we illustrate.

3.1 Ex ante welfare economics of contingent claims

If one believes that all individuals have common, correct beliefs then *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 that goes far beyond the notion of rationality embedded in Savage's work. For this reason we do not find *ex ante* Pareto optimality compelling, and we entertain restrictions on asset trade even though individuals might unanimously favor complete markets.

In a different setting Mongin [2005] has also argued that unanimity of preference is not necessarily a good argument for social optimality. He argues that not only preferences, but also the reasons why people hold them, should be considered in making welfare claims and that "spurious unanimity" can arise. We illustrate his point with an example in appendix A.2. In this paper we demonstrate that competitive allocations for economies in which individuals have heterogeneous beliefs are generally suboptimal when they are evaluated using a common, possibly correct, belief.

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 actual data generating process $\mathcal{P}^0$.[1]

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 $P^1, \dots, P^I$, 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, $P^i_t(\sigma) \neq P^j_t(\sigma)$, then no* ex ante *Pareto optimal allocation in which $c^i, c^j \neq 0$ can be true-Pareto optimal for any distribution P^0.*

This proposition, along with the first welfare theorem, implies that complete markets competitive equilibria cannot be true Pareto optimal for any economy with heterogeneous beliefs. What do true Pareto optimal allocations look like? First, each individual must have constant consumption across irrelevant states; that is states with the same aggregate endowment. In any true Pareto optimal allocation (actually in any Pareto optimal allocation for an economy with common beliefs), there is no speculation over irrelevant states.

Corollary 1. *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 $c^i_t(\sigma) = c^i_t(\sigma')$.*

Second, if 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 2. *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, $\liminf_t c^i_t(\sigma) > 0$.*

[1] It is important for our analysis that there is actually a data generating process. We do not assume that any agents in the model know this process, but we the modelers do use data generating processes to describe the distribution of outcomes given various market restrictions.

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 has correct beliefs, proposals that are true-Pareto improvements should gain political support.

Unfortunately, the mechanism design problem depends 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 knowledge of the true distribution, it would be unreasonable to assume that market designers would have any better knowledge.[3] That is, we want to do distribution-independent market design.

Our solution to this problem is to explore the parameter space: possible data generating processes and sets of beliefs. 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 improving, where it is better according to a Rawlsian welfare aggregator, and where it is better according to a Bergson-Samuelson social welfare function in which the welfare weights are those that solve the *ex ante* Pareto optimality problem.

[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.

[3]If agents disagree, then the social planner must disagree with at least some agents independently of P^0. The disagreement between the planner and agents is a necessary consequence of our modeling choices. We argue in section 8 that it has limited impact on our results in the sense that the main force behind our results is the disagreement between agents.

3.2 Spurious unanimity: other approaches

Others have addressed the problem of spurious unanimity in contingent claims allocations. Brunnermeier et al. [2014] introduce *belief-neutral* Pareto optimality. They identify a set of "reasonable beliefs", in the sense of potential true distributions, as 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. [2014] reduce incompleteness by examining partial orders induced by Bergson-Samuelson social welfare functions, formed as weighted averages of each profile of true expected utilities. They also investigate the implications of their criteria for a number of problems including speculative trade.

Gilboa et al. [2014] offer an 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 occurs 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, leaving few welfare comparisons that 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.

Although these approaches provide insights into the difficulties that heterogeneous beliefs create for welfare analysis they do not provide a compelling way actually to undertake welfare analysis when beliefs are heterogeneous. 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 simple economies. 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[4]

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

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).

[4] The borrowing limits are needed because we formulate and solve the agent's problem recursively.

The implied price of a risk-free bond, which we refer to later, is:

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

Definition 1. *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* (2b)*.*

In addition to standard complete markets, we analyze several other designs: complete markets with ad-hoc borrowing limits (CMB), 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 2. *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, we impose an exogenous borrowing limit instead.

4.3 Borrowing limits

Definition 3. *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)

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

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.

Finally, we also analyzed the market design with a transaction tax and reached similar conclusions as in the economy with a borrowing limit. The results of that analysis can be found in the Appendix A.6.

5 Welfare

A designer chooses a market structure $\mathcal{M}$ that, given beliefs $\mathcal{P} = (P^1, ..., P^I)$, induces a competitive equilibrium allocation $(c^1(\mathcal{P}|\mathcal{M}), ..., c^I(\mathcal{P}|\mathcal{M}))$.[5] Individuals evaluate their welfare according to their own beliefs, but as we have argued, when beliefs are heterogeneous social welfare should not be based on individual perception of welfare. Instead the market designer should care about realized welfare; so we evaluate social welfare using some true data generating process. For fixed individual beliefs and data generating process, a market structure $\mathcal{M}$ induces a stochastic process of consumption flows for individuals in the economy. We evaluate individual i's consumption flow, $c^i(\mathcal{P}|\mathcal{M})$, using his utility function and discount factor, and the actual data

[5] We assume that such an incomplete markets equilibrium exists. On existence see Kubler and Schmedders [2003].

generating process (rather than his possibly incorrect beliefs) to get his actual discounted expected utility U_{i,P^0}. The market designer can, for given $\mathcal{P} = (P^1, ..., P^I)$ and P^0, vary $(U_{1,P^0}, ..., U_{I,P^0})$ by varying $\mathcal{M}$.

We do not want to consider the welfare induced by $\mathcal{M}$ for any particular $\mathcal{P}$ or P^0 and so we will consider sets of individual beliefs and possible truths. But before considering how to manage these sets we need to confront the more standard question of evaluating social welfare given any individual beliefs and data generating process. That is, some market structures may improve the realized welfare of some individuals while reducing that of others and the market designer has to decide whether to aggregate individual welfares into some social welfare and, if so, how. First, however, we note that for some economies and some alternative market structures, there is no tradeoff in realized utilities; that is, it may be possible to make everyone better off (according to their realized welfare) by choosing the appropriate market structure. For example, if all traders are homogeneous and have correct beliefs then complete markets Pareto dominate incomplete markets. If agents have no endowment risk but they disagree about irrelevant states, then a bond-only structure may Pareto dominate complete markets that serve only to permit welfare reducing speculation.

In many settings, however, there will be a tradeoff between realized welfares and so we view the market designer as using some utility aggregator to choose over market structures. We see no compelling argument for any particular aggregator, so we instead consider several possibilities: one based on a Rawlsian criterion and another one based on a Bergson-Samuelson criterion. We admit at the outset that this approach requires that we take utility to be interpersonally comparable.

Definition 4. *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

$$\mathcal{W}(U_1, .., U_I) = \sum_{i=1}^{I} \theta^i U_i \tag{7}$$

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

$$\mathcal{W}(U) = \min_i U_i. \tag{8}$$

A designer using (8) 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].[6] A designer using (7) would act similarly to a Pareto planner. But the utility aggregator (7) presents a new degree of arbitrariness: What weights should a designer use? One could choose $\theta_i = 1/I, \forall i$, a choice that is attractive in ex-ante symmetric environments. One could also choose θ to be a vector of "market weights".[7] These two choices are special cases of Bergson-Samuelson social welfare functions.

Fixing either of these utility aggregators and some market structure, we have a measure of social welfare for each configuration of individual beliefs $\mathcal{P} = (P^1, ..., P^I)$ and true data generating process, P^0. Our designer is interested in choosing a market structure that performs best according to the selected social welfare measure. The answer may depend on the individual beliefs and the truth. For example, if Bergson-Samuelson social welfare with market weights is used then complete markets are optimal when each individual's belief coincides with the truth. And, as we will show, complete markets are not optimal if there is enough dispersion in individual beliefs. 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. So we evaluate welfare over a set of possible truths. We also 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

---

[6]Rawls [1971] argues that a fair social choice can be made only 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.

[7]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.

individuals. Therefore we evaluate welfare over a set of individual beliefs.

Rather than relying on an arbitrary aggregator of welfare across the sets of possible beliefs and truths, we instead provide welfare surfaces and the designer (or reader) can decide which surface he prefers. We will display welfare surfaces both for the Rawlsian social welfare measure and for Bergson-Samuelson social welfare with market weights. To help illustrate tradeoffs between market structures, we also provide a numerical representation of social welfare based on one particularly appealing aggregation across sets of possible beliefs and truths.

Definition 5. *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*

$$\min_{P^0 \in \mathcal{B}^0} \min_{\mathcal{P} \in \mathcal{B}^I} \mathcal{W}\left(\left(U_{i,P^0}(c^i(\mathcal{P}|\mathcal{M}))\right)_{i=1}^{I}\right). \tag{9}$$

6 Simple Economies

We present a simple economy that we use to illustrate economic forces operating in economies with heterogeneous beliefs. 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.

Agents share a common utility function

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

where $\gamma = 2$.[8] 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.[9] Endowments are

$$(e_t^1(\sigma), e_t^2(\sigma)) = \begin{cases} (0.5, 0.5) & \text{if } \sigma_t = 0 \\ (e_h, e_l) & \text{if } \sigma_t = 1 \\ (e_l, e_h) & \text{if } \sigma_t = 2 \end{cases}, \quad \forall t, \sigma. \tag{10}$$

[8] Robustness of our results to the specification of preferences is considered in Appendix A.2.

[9] 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.

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

Beliefs are specified as follows:

$$\Pi^i = \begin{bmatrix} 0 & 0.5 & 0.5 \\ 0 & p^i & 1-p^i \\ 0 & p^i & 1-p^i \end{bmatrix}, \tag{11}$$

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

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. These individuals are rational in that they correctly maximize their subjective expected utility; that is, they are Savage (1954) agents. Furthermore, our analysis can be extended to include learning, and our general results carry over to this extension (see Appendix A.3).

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 constant amounts. The competitive equilibrium allocation is:

$$(c_t^1(\sigma), c_t^2(\sigma)) = (0.5 + \beta^2(\mu_e - 0.5), 0.5 - \beta^2(\mu_e - 0.5)), \quad \forall t, \sigma, \tag{12}$$

where $\mu_e \equiv pe_l + (1-p)e_h$ is the expected endowment evaluated using the common beliefs, p, about the probability of state 1. 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 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:

$$\limsup_{t\to\infty} c_t^2(\sigma) = 0 \quad P^0 a.s. \tag{13}$$

Following Blume and Easley [2006], we say that type-2 agents do not survive. The eventual immiseration of agents with incorrect beliefs when markets are complete is the source of our intuition 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 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."[10] 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.

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 -1.948, respectively, for type-1 and type-2 agents. Both agents expect higher welfare than W^*. Both believe that "speculative" financial trades would allow them to accumulate wealth. Objective welfare levels (expected utility of equilibrium consumptions computed using the actual data generating process) are -1.943 and -2.129, for type-1 and type-2 agents, respectively.

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.[11] So welfare

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

[11]Costs of aggregate fluctuations in a standard RBC model are typically found to be

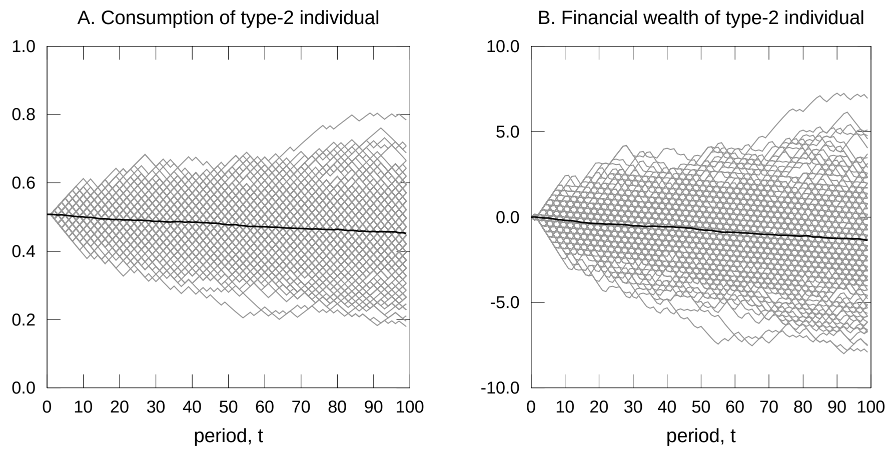

Figure 1: Sample paths of consumption and financial wealth of a type-2 agent in the complete markets economy.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, p^2 = 0.55$.

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 downward trend in his consumption rather than by increased consumption volatility.[12]

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

below 0.1%.

[12]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: $U^1_{P^0} = U^1_{P^1} = -2.002, U^2_{P^0} = -2.063$ and $U^2_{P^2} = -2.058$. 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. Thus, here the welfare loss is attributed mainly to increased consumption volatility. See also section A.5.1.

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: $B^i_t(\sigma) = 8, \forall t, \sigma$.

Continuing with the economy from the previous section, we simulate equilibrium consumption and wealth dynamics in the bond economy. As shown in figure 2, 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 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.

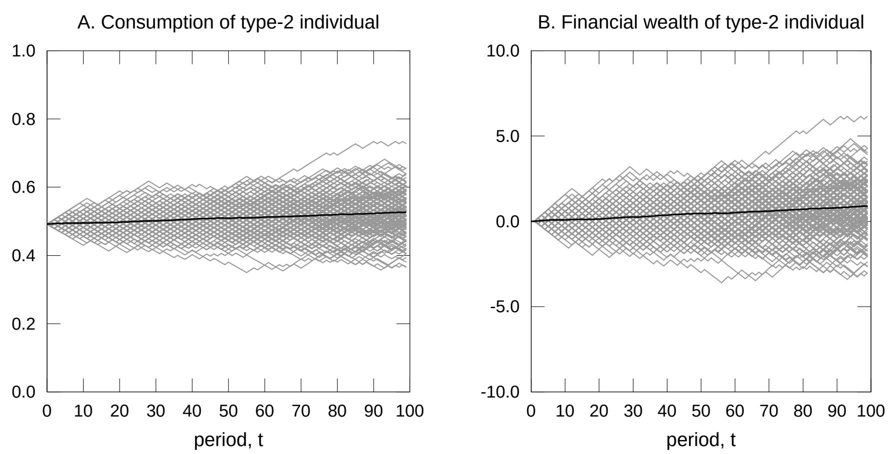

Figure 2: Sample paths of consumption and financial wealth of a type-2 agent in the bond economy.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, p^2 = 0.55, B = 8.$

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, the designer considering complete markets or only bonds would have to decide between giving type-1 agents a lower objective welfare with bonds-only versus giving type-2 agents a lower objective welfare with complete markets. The welfare criterion (9) (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 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 $p^0 = 0.5$.[13] Figure 3 plots the Rawlsian welfare surface, $\min_i[U_{P^0}(c^i(p^1, p^2|\mathcal{M}))]$, for this belief set.[14]

A market designer using the Rawlsian utility aggregator who considers this set of beliefs as possible and who knows the true data generating process would have to choose between the two welfare surfaces in figure 3. The bond-only surface is relatively flat as speculation possibilities are limited without access to state contingent contracts. The complete markets design surface reaches a higher social welfare level than does the bond-only design for beliefs sufficiently close to each other and to the truth; the bond-only design does better when beliefs are sufficiently different or sufficiently far from the truth.

Considering the lowest (Rawlsian) social welfare yielded by each market design is also instructive. 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. The lowest welfare in the complete markets economy is worse, -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 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, both types are pessimistic

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

[14]The shape of this welfare surface is explained in Appendix A.4.

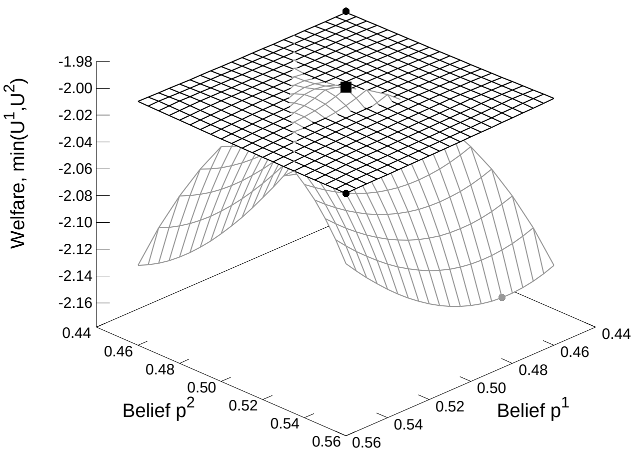

Figure 3: Welfare in example 1: the bond-only (black) vs the complete markets (gray) design. Square point denotes the unconstrained maximum: $(p^1, p^2, W^*) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare under the corresponding design.
Parameters: $\beta =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

about type-1 agent's endowment stream and as a result it is valued less – type-1 agent is subject to a negative wealth effect.

In this example, the welfare criterion (9) (with the Rawlsian aggregator) selects the bond-only design over the complete markets design because

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

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.

6.4 True Pareto Dominance

To this point we have focused on tradeoffs between the objective welfare of type-one agents (those with correct beliefs) and type-two agents (those with incorrect beliefs), but as we noted previously there are belief assignments for which both agents are better off with one market design than with the alternative market design. We say that market design A true-Pareto dominates market design B if the objective welfare of both agents is greater under market design A than under market design B. Clearly if there is a true-Pareto dominant market design then the market designer does not need a welfare aggregator; he should select the true-Pareto dominant design.

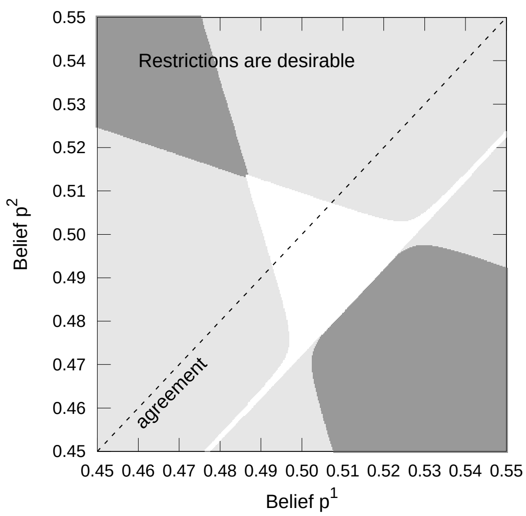

Figure 4: True-Pareto ranking: $\mathcal{B} \succ CM$ (dark gray), $CM \succ \mathcal{B}$ (white), allocations cannot be ranked (light gray).
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 1$.

Figure 4 plots the true-Pareto domination relationship for our example. 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. Restricting financial trade to risk-free bonds effectively shuts down speculation, thereby increasing everyone's utility for sufficiently heterogeneous and sufficiently incorrect beliefs.[15] The light gray area denotes belief assignments under which the two market designs cannot be ranked because one agent gains while another loses. 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.

6.5 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 5 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 lowest 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 society's welfare are both 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

[15] 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.

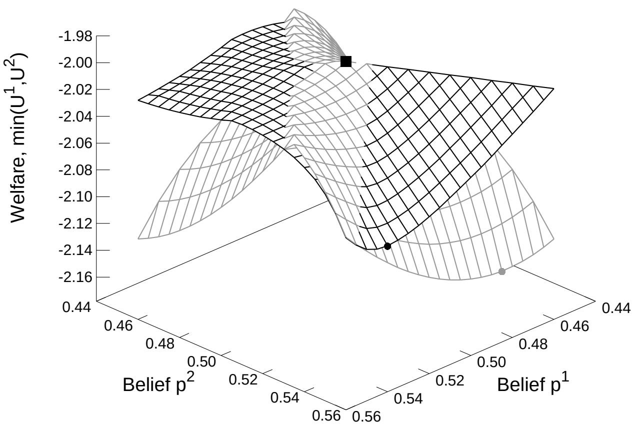

Figure 5: Welfare in example 2: the complete markets with borrowing limits (black) vs the complete markets design (gray). Square point denotes the unconstrained maximum: $(p^1, p^2, W^*) = (0.5, 0.5, -2)$. Circles denote belief assignments that attain the lowest welfare under the corresponding design. Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 1$.

forces. For this example, society's welfare increases from -2.139 to -2.083, using the welfare criterion (9) (with the Rawlsian aggregator), 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, $(p^1, p^2) = (0.50, 0.55)$, and contrast two designs: complete markets with (restrictive) $B = 1$ and (relaxed) $B = 8$ borrowing limits.

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.

7 Bergson-Samuelson criterion

Next, we examine the market design problem using the alternative utility aggregator that makes use of individuals' "market weights." For each 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 utility aggregator is:

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

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 6 plots social welfare $\sum_i [\theta^i(\mathcal{P}) U_{P^0}(c^i(\mathcal{P}|\mathcal{M}))]$ for both the bond-only and complete markets financial markets designs (with $p^0 = 0.5$). Social welfare with complete markets is close to -2 – the maximum achievable under any market design (depicted by the gray square point) – when agents have common beliefs, even if those beliefs are wrong, i.e., on the diagonal with $p^1 = p^2$. This relative insensitivity to common, incorrect beliefs differs from the Rawlsian social welfare graph because these beliefs harm one agent and benefit the other agent, which makes the impact on social, i.e. average, welfare minimal. For this reason, the shape of the Bergson-Samuelson welfare surface is primarily determined by the survival forces, and it declines fastest along the diagonal with the maximum disagreement: $p^1 = 1 - p^2$.

As we move away from the common beliefs diagonal, social welfare stays robustly high under the bond-only design, but declines under the complete markets design. As with the Rawlsian social welfare, the reason for the robust performance of the bond-only design is that it limits survival forces. For this example a market designer using the welfare criterion (9) (with the market-

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

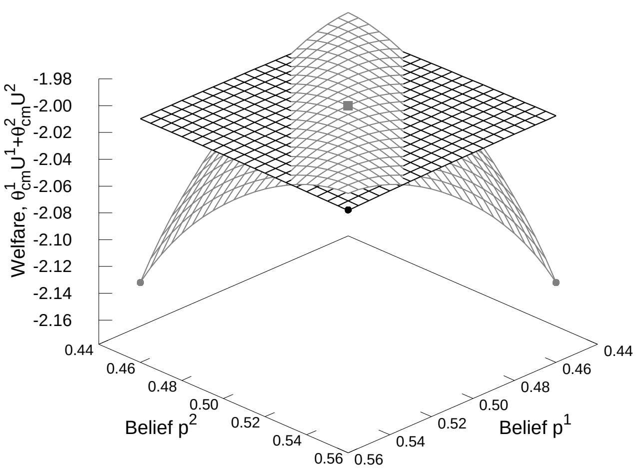

Figure 6: Welfare in example 1: the bond-only (black) vs the complete markets (gray) design. Square point denotes the unconstrained maximum: $(p^1, p^2, W^*) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare under the corresponding design.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 8$.

weighted, Bergson-Samuelson aggregator) would choose the bond-only design as

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

We next compare the complete markets economy with a borrowing limit and the unrestricted complete markets design. As in section 6.5, we impose a borrowing limit of $B = 1$. Figure 7 plots welfare surfaces under the two designs. The shapes of the two welfare surfaces are similar as the two market designs allow the same forces to operate. However, the survival forces are restricted under the design with the borrowing limit, which explains the robust performance of this design off of the common-belief diagonal. A market designer who uses the welfare criterion given in (9) (with the Bergson-Samuelson aggregator and market weights) should choose the design with

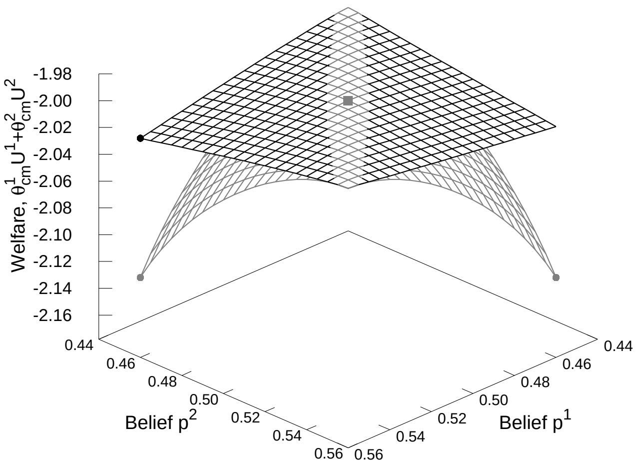

Figure 7: Welfare in example 2: the complete markets with borrowing limits (black) vs the complete markets design (gray). Square point denotes the unconstrained maximum: $(p^1, p^2, W^*) = (0.5, 0.5, -2)$. Circles denote belief assignments that attain the lowest welfare under the corresponding design. Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 1$.

the borrowing limit

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

8 Dependence on P^0

To this point we have analyzed the market design problem for a singleton $\mathcal{B}^0$. In this section, we confront our designer with multiple data-generating processes. Recall that our theoretical results hold for any P^0 and, hence, for any $\mathcal{B}^0$. Our computed equilibria also show that welfare varies more under the complete markets design than under the designs with financial restrictions. In this section, we show that introducing ambiguity about P^0 via expansion of $\mathcal{B}^0$ causes welfare gains from financial restrictions to increase. Here we use the Rawlsian aggregator, but similar results obtain with the Bergson-Samuelson aggregator of Section 7.[17]

P^0	$W_{P^0}(CM)$	$W_{P^0}(B)$ $B=8$	$W_{P^0}(CMB)$ $B=1$
1	2	3	4
0.45	**-2.545**	**-2.084**	**-2.121**
0.46	-2.439	-2.068	-2.113
0.47	-2.347	-2.053	-2.106
0.48	-2.267	-2.039	-2.098
0.49	-2.195	-2.025	-2.090
0.50	-2.139	-2.011	-2.083

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

In constructing the welfare levels reported in Table 1, we assumed that $(p^1, p^2) \in [0.45, 0.55]^2$ and for each choice of $p^0 \in [0.45, 0.50]$ we report welfare using (9) with the Rawlsian aggregator.[18] Columns 2 through 4 present welfare under the unrestricted complete markets design (section 4.1), the bond economy (section 4.2), and the complete markets with ad-hoc borrowing limits (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 $B = 1$ and unrestricted set of securities (CMB) under-performs the bond-only design but still offers a sizeable improvement (equivalent to a permanent 20% increase in consumption) over the unrestricted financial markets.

The worst-case for the complete markets design occurs at $(p^0, p^1, p^2) = (0.45, 0.45, 0.55)$. This point assigns correct beliefs to type-1 agents and maximally wrong beliefs to type-2 agents. This worst-case choice of beliefs

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

[18] 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.

maximizes the strength of survival forces. Type-2 agents have the lowest welfare. For the bond-only design the worst-case occurs at $(p^0, p^1, p^2) = (0.45, 0.55, 0.535)$ where 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 complete markets with a borrowing limit, the worst-case occurs at $(p^0, p^1, p^2) = (0.45, 0.515, 0.55)$. Here 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.

It is the disagreement between agents that drives our results. When agents agree and markets are complete each agent consumes a constant amount in every period and state because the aggregate endowment is constant in our examples. So, the *ex-ante* welfare is independent of P^0. Disagreement between the planner and agents has no impact on the social welfare under complete markets.

To isolate the effect of disagreement with the social planner, consider an environment when agents agree: $\mathcal{B} = \{(P^1, P^2) : P^1 = P^2 \in [0.45, 0.55]\}$. Then the social ranking of the alternative financial market designs would be: $\mathcal{W}(CMB) = \mathcal{W}(CM) = -2.0634 > \mathcal{W}(B) = -2.0840$. Complete markets both with and without the borrowing limit ($B = 1$) achieve maximal social welfare.[19] Imposing borrowing limits does not harm the social welfare. Restricting financial trade to bonds only lowers welfare quantitatively insignificantly.[20]

The larger are $\mathcal{B}^0$ and $\mathcal{B}$, the starker are the welfare differences. Reason-

[19]The borrowing limits are inactive for all considered beliefs.

[20]We deem the reduction in the social welfare insignificant because it is equivalent to a 1.0% permanent reduction in consumption while the same financial market restrictions lead to the welfare gains equivalent to a 22.1% permanent increase in consumption if the agents disagree.

able choices of $\mathcal{B}^0$ and $\mathcal{B}$ can be constructed using error detection probabilities as in Hansen and Sargent [2008].[21]

An important benefit of our approach is that it can be immediately applied to any completely specified economy (payoff functions, discount factors, sets of possible truths and beliefs) to determine optimal financial market restrictions. For example, assuming the Rawlsian welfare aggregator the optimal borrowing limit in the complete markets financial design with borrowing limits is:

$$B^* = \arg\max_B \min_{P^1,P^2,P^0} \min_i W^i_{P^0}(CMB). \tag{14}$$

The optimal borrowing limit B^* is 36% of an average annual income, while in the economy with homogeneous and correct beliefs it is 33%.[22] It is optimal to allow more borrowing than needed to hedge income fluctuations in the homogeneous beliefs economy.

9 Concluding remarks

We propose a framework to evaluate financial market designs for exchange economies in which agents have heterogeneous beliefs. Our analysis illustrates 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. In the economies that we study, financial market designs with simple restrictions like limits on the set of traded assets or borrowing limits 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 numerical simulations and the analysis of learning in Blume and Easley [2006] suggests that our conclusions would not be changed if we considered learning agents. In general, some learners learn faster than others, and the slow learners do not learn fast enough to keep their consumption from becoming negligible. Here too, limiting the bets that traders can take slows down or prevents financial ruin of slow learners.

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

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

Perhaps the most important limitation of our analysis is the absence of incentive effects. In our analysis, restrictions imposed on 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

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, October 2016.

Markus Brunnermeier, Alp Simsek, and Wei Xiong. A welfare criterion for models with distorted beliefs. *Quarterly Journal of Economics*, pages 1753–1797, 2014.

Timothy 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, 2008.

F. Kubler and K. Schmedders. Stationary equilibria in asset-pricing models with incomplete markets and collateral. *Econometrica*, 71:1767–1795, 2003.

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

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.

A Appendix

A.1 Proofs

Proof of proposition 1. 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. 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 true-Pareto optimal, it must be the case that:

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

First-order conditions for optimality on path σ' imply

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

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

Proof of corollary 1. 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. □

Proof of corollary 2. 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. □

A.2 Spurious unanimity

Ithaca NY, the home of two of us, has a pedestrian mall. It is serviceable, but would benefit from renovation. The work 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 persuasive.[23]

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's (1954) subjective expected utility theorem. It is instead a restriction on preferences that goes far beyond the notion of rationality embedded in Savage.

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 [2016] shows, that it may not be the agent with the least accurate beliefs who is driven out of the markets 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 may be stronger because potential financial losses have lower utility cost. We demonstrate this by changing the period utility specification to $u(c) = \sqrt{c}$.

[23] It's also useful to note that this example does not require disagreement about the supports of the distribution on states. Imagine 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 payoff functions and prior beliefs, which are as follows: $u^1(a) = 1$, $u^1(b) = 0$, $\rho^1(E) = 0.99$, $u^2(a) = 0$, $u^2(b) = 1$, and $\rho^2(E) = 0.01$. 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.

Figure 8 plots welfare surfaces for the complete markets design and the complete markets with an exogenous borrowing limit design. Welfare levels (using the welfare function given in equation (10)) 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.

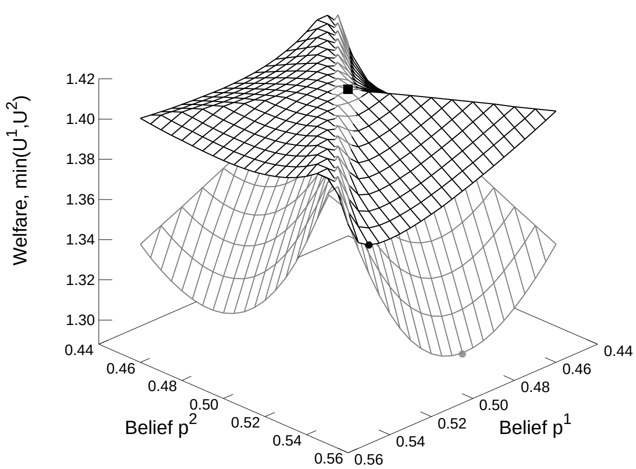

Figure 8: Welfare with bounded below utility for complete markets (gray) and complete markets with a tight borrowing limit (black).
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, p^2 = 0.55, B = 8$.

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 9 plots the social welfare surface (again using the welfare function equation (10) with Rawlsian aggregator) under the bond-only (black) and complete markets (gray) designs.[24]

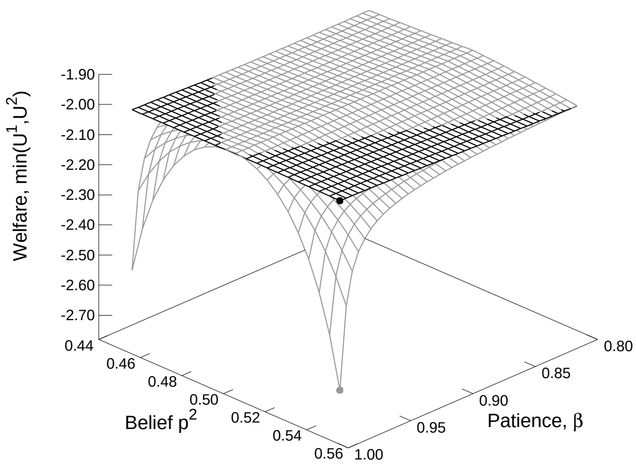

Figure 9: Welfare in example 1: the bond-only (black) vs the complete markets (gray) design. Circle points denote belief assignments that attain the lowest welfare under the corresponding design when $\beta = 0.99$.
Parameters: $e_l = 1/3, e_h = 2/3, p^0 = p^1 = 0.5, B = 2$.

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 be low even when disagreement is small under the complete markets design.[25] 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, restricting financial markets to allow trade of only a risk-free bond is equivalent to

[24]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.

[25]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 lose substantial amounts of wealth over long periods of time even if they are losing slowly.

a permanent 31.3% increase in consumption.

A.4 Learning

Next we extend our analysis to include learning.[26] Consider the two-state, two-agent economy from our examples, but suppose that agents are uncertain about the data generating process. They believe that it is i.i.d., but they do not know the probability distribution. Agent i starts with a Beta prior about the probability of state 1: $\hat{p}^i \sim B(n_0^i, m_0^i)$, where n_0^i and m_0^i represent prior number of realizations of state 1 and 2, respectively. Agents update their beliefs according to Bayes rule. We consider the following set of admissible prior distributions: $(n_0^i, m_0^i) \in \{(8.0, 12.0), (8.4, 11.6), (8.8, 11.2), ..., (12.0, 8.0)\}$. The implied set of admissible prior state-1 probabilities is $\{0.40, 0.42, 0.44, ..., 0.60\}$. Holding the truth fixed at $p_1^0 = 0.50$, we compute society's welfare under all possible assignments of priors. We do computations for two financial market designs: the complete financial markets and the complete financial markets with borrowing limits.

The two welfare surfaces are presented in figure 10.[27] This figure closely resembles figure 5 comparing the same market designs without learning. With the restrictive borrowing limit, speculation is limited and the worst case corresponds to nearly homogeneous priors: $B(11.6, 8.4)$ and $B(12, 8)$ for type-1 and type-2 agents, respectively. Under the first prior assignment, the type-2 agent has lower welfare because both types initially agree that type-2 agents are unlikely to receive high income. Under the complete markets design, the worst case assignment of beliefs is $B(10.4, 9.6)$ and $B(12, 8)$ for type-1 and type-2 agents, respectively. A type-2 agent is impoverished because his prior is less accurate in addition to being believed unlikely to receive high income. So, the same forces operate as in the analogous economy without learning. The designer should opt for imposing the borrowing limit and this would increase welfare by 2.6% compared to 2.7% in the economy without learning.

[26] We do not consider learning from endogenous outcomes, e.g. prices, as doing so requires agents to have a model of how states influence prices. Assuming that agents have a correct model of this relationship (which is an equilibrium relationship that evolves as agents learn) is even more implausible than assuming that they know the process on states. Rather than assume this it would be preferable to just assume common, correct beliefs about the exogenous state process.

[27] Figure 10 plots interpolated data to provide the same resolution as other figures.

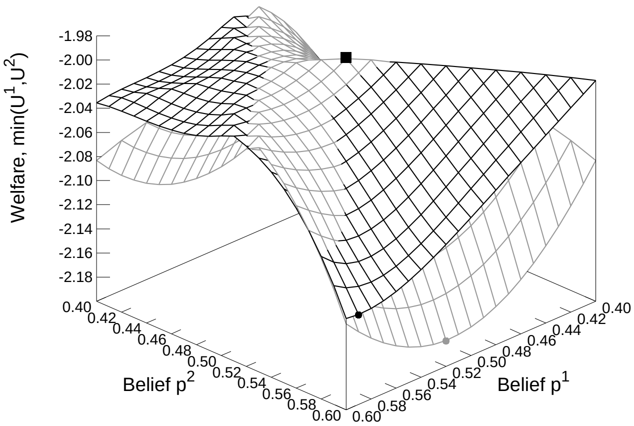

Figure 10: Welfare in example 2 with learning: complete markets with borrowing limits (black) vs complete markets (gray). Square point denotes the unconstrained maximum: $(p^1, p^2, W) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare under the corresponding market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

A.5 Complete markets design

In this section we provide intuition for 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 is likely to have the lowest welfare. The second is the wealth effect: the equilibrium price system is affected by the configuration of beliefs and this may present an advantage to one of the types.[28]

Figure 11 reproduces the complete markets welfare surface shown in figure 3. Along arc AOB both types are either optimistic or both pessimistic. The

[28]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 $c_0^i = 0.5$.

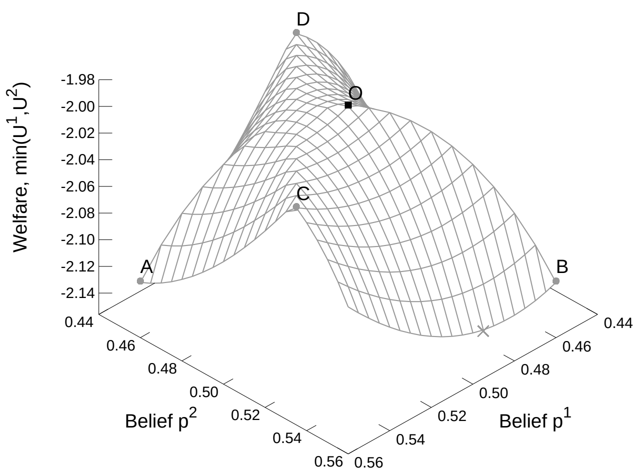

Figure 11: Welfare in example 1 under the complete markets design (gray). Square point denotes the unconstrained maximum: $(p^1, p^2, W) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5$.

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 their speculation induces 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 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.[29]

A.5.1 Wealth effects

To understand the wealth effects that arise in our economy 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 $\text{Prob}(s = 1) = 0.5(1-\Delta) \neq 0.5$. 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.

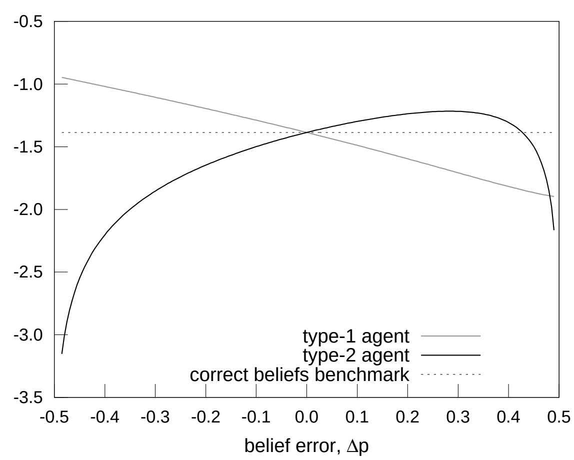

Figure 12: Objective welfare in the two-period complete markets economy

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

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

agent is:

$$c_0^2 = \frac{1}{2-\Delta}, \quad c_1^2(s=1) = \frac{1+2\Delta}{2+\Delta}, \quad c_1^2(s=2) = \frac{1-2\Delta}{2-3\Delta}. \tag{15}$$

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

$$E[c_1^2] = c_0^2 \frac{4 - 4\Delta^2 c_0^2}{4 - \Delta^2 (c_0^2)^2} < c_0^2. \tag{16}$$

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 have higher objective welfare:

$$\left.\frac{dW^2(\Delta)}{d\Delta}\right|_{\Delta=0} = 1 \neq 0, \tag{17}$$

where $W^2(\Delta) \equiv ln(c_0^2) + 0.5ln(c_1^2(s=1)) + 0.5ln(c_1^2(s=2))$. So agent 2 can be better off being an optimist. But $\lim_{p\to 0.5} W^2(p) = -\infty$. Figure 12 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 two. This is the wealth effect.

A.6 Transaction tax

Definition 6. *The complete financial markets with a transaction tax (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 tax that is rebated back to investors as equal lump sums.*

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

$$c_t^i(\sigma) + \sum_{\tilde{\sigma}|\sigma^t} Q_t(\tilde{\sigma}) a_{t+1}^i(\tilde{\sigma}) + \tau \cdot \sum_{\tilde{\sigma}|\sigma^t} [a_{t+1}^i(\tilde{\sigma}) - a_t^i(\sigma)]^2 = a_t^i(\sigma) + e_t^i(\sigma) + T_t(\sigma)/I, \tag{18}$$

where $T_t(\sigma)$ is the total transaction tax revenue. This design assumes that the transaction tax is a quadratic function of security purchases to ensure continuity of demands for securities.

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

Figure 13 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 $\tau = 0.05$. Welfare levels for the first two designs are very

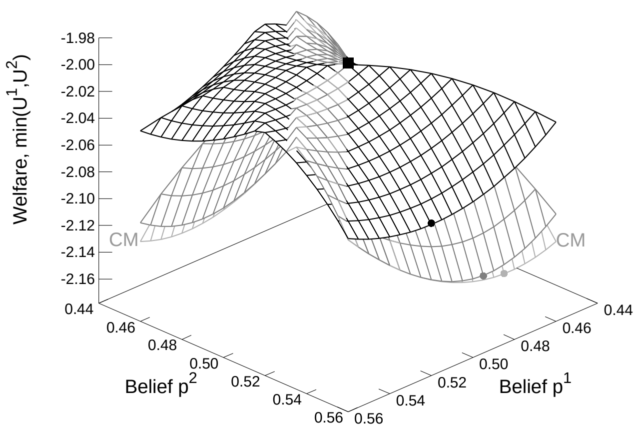

Figure 13: Welfare in example 3: complete markets with borrowing limits and transaction tax (black) vs complete markets with borrowing limit (dark gray) vs complete markets. Square point denotes the unconstrained maximum: $(p^1, p^2, W) = (0.5, 0.5, -2)$. Circle points denote belief assignments that attain the lowest welfare under the corresponding market design.
Parameters: $\beta = 0.96, e_l = 1/3, e_h = 2/3, p^0 = 0.5, B = 8, \tau = 0.05$.

close, as we intended by setting the generous borrowing limit $B = 8$.[30] The impact of the transaction tax can be assessed by comparing the third and second designs.

Imposing a transaction tax on top of the borrowing limit increases society's welfare from -2.134 to -2.079, using the welfare criterion with the Rawlsian aggregator (9), and this welfare gain is equivalent to a permanent 2.6% increase in consumption.

A.7 Heterogeneous Discounting

This section continues the two-period example of the previous section, but allows for heterogeneous discount factors. To facilitate comparisons across the population, life-time utilities are redefined as follows:

$$U^i(c_0, c_1, c_2) \equiv \frac{1}{1+\beta^i} ln(c_0) + \frac{\beta^i}{1+\beta^i}[p_1^i ln(c_1) + p_2^i ln(c_2)]. \tag{19}$$

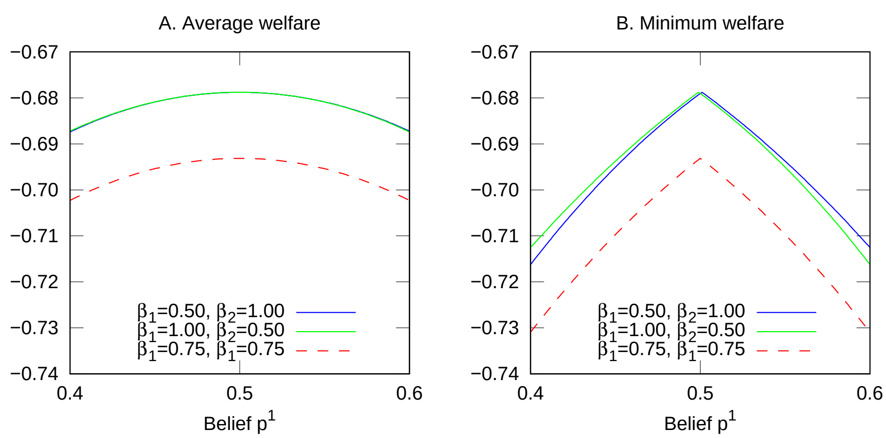

Figure 14: Discounting vs belief heterogeneity.
Parameters: $e_l = 1/3, e_h = 2/3, p^0 = 0.5$.

Figure 14 presents profiles of the social welfare when agents have diametrically opposite beliefs: $p_2 = 1 - p_1$. The case with $(\beta_1, \beta_2) = (0.75, 0.75)$

[30] The natural borrowing limits are difficult to compute in the presence of a transaction cost. So, we impose an exogenous borrowing limit instead.

corresponds to homogeneous discounting benchmark analyzed in the body of the paper. The welfare in the cases with heterogeneous discounting – $(\beta_1, \beta_2) = (1.00, 0.50)$ and $(\beta_1, \beta_2) = (0.50, 1.00)$ – is higher than in the homogeneous discounting economy because of increased intertemporal trade that benefits *everyone*, the patient and impatient. The assumed levels of heterogeneity were meant to be extreme; hence, the welfare impact is substantial.

The main observation emerging from figure 14 is that the effect of discount factor heterogeneity on the society's welfare is best described as a level effect. This suggests that the set of beliefs when the complete markets is dominated by other financial market designs would be largely unaffected by heterogeneous discounting. Said differently, our analysis is robust to the presence of heterogeneity in time preferences.