Doubts, Asymmetries, and Insurance

Job Market Paper *

Anmol Bhandari †

November 2013

Abstract

This paper studies a redistribution problem with heterogeneous beliefs, heterogeneous information and aggregate risk. A Pareto planner assigns consumption to two agents who differ in their initial priors over a common set of forecasting models of aggregate endowment. One of the agents receives a privately observed taste shock. Agents doubt their forecasting models. We use Hansen-Sargent "multiplier-preferences" to represent these doubts by specifying sets of probability distributions close to the forecasting models. To obtain decisions that are robust to model misspecifications, agents construct "ex-post" beliefs that rationalize choices and provide lower bounds to their expected valuations. Ex-post beliefs differ across agents at histories where their wealths differ. Endogenous heterogeneity in ex-post beliefs generates a force that can keep long run inequality bounded in several settings. Wealth inequality affects market price of risk and motives to trade on public information.

*An updated version is available at https://files.nyu.edu/apb296/public/papers.html. I am especially indebted to Thomas Sargent, Jaroslav Borovicka for their advice on this project. I also thank Tim Cogley, Lars Peter Hansen, Andrew Atkeson, Pierre Olivier Weill, Christopher Phelan, V.V. Chari, Larry Jones, Aleh Tsyvinski, Mikhail Golosov, Victor Tsyrennikov, Max Croce, David Evans, Jesse Perla, Chris Tonneti, Gaston Navarro, Isaac Baley, Andres Blanco, Anusha Nath and the participants at the Becker Friedman Institute, SED (Seoul) for helpful comments and suggestions.

†Economics Department; New York University apb296@nyu.edu

1 Introduction

This paper studies a redistribution problem with heterogeneous beliefs, heterogeneous information and aggregate risk. To motivate the question, consider an example of a world with two countries. Both countries value consumption streams of a single commodity, say 'oil,' and one of them is subject to a privately observed stochastic shift in the preference for oil. In such a scenario, what is the efficient way of allocating oil between these countries? The supply of oil depends on many factors and these countries need to make forecasts about future oil production. Arguably, countries can differ in their forecasts. However, the answer to the redistribution question depends on how one thinks about the forecasting problem that these countries face.

In a Bayesian setup, this forecasting problem is reduced to endowing each country with a well specified probability distribution [1] over future contingencies. Risk sharing in environments with heterogeneous beliefs and heterogeneous information has stark implications for how inequality evolves over time. Absent these features, efficiency is characterized by equating relative marginal utilities to initial Pareto weights at all histories. When agents differ in their assessments of future contingencies, efficiency requires a wedge (that corresponds to the time varying relative likelihood ratios) in the inter-temporal evolution of Pareto weights. Similarly, while providing efficient insurance to agents who are privately informed about their insurance needs, optimal incentives cause Pareto weights to fluctuate overtime.

Previous literature [2] has (separately) studied such settings and pointed out that either feature - heterogeneous beliefs or heterogeneous information can generate ever growing cross-sectional inequality. These observations are also referred to as "immiseration" results. Although in stylized economies, these papers rationalize inequality as a necessary condition for efficiency.

This paper provides a qualification to such views by relaxing the assumption that agents' attitudes towards risk are captured by a well specified probability distribution. In the context of our two countries example, the future supply of oil could depend in complex ways on various factors - institutional, technological, political or even climatic. In contrast to the Bayesian approach, these countries specify their views using a parsimonious "approximating" model but they take seriously the possibility that this approximating model is misspecified. We will distinguish between two types of such misspecification concerns - ones that countries can eventually learn about and others that are hard to learn. The enduring misspecification concerns capture low frequency changes, structural breaks, non-linearities that would usually pose an infinite dimensional learning problem to a Bayesian.

We use Hansen-Sargent "multiplier-preferences" to formally represent these attitudes. Given an approximating model, agents entertain a set of probability distributions that are statistically close [3] to it. They use a min-max approach to obtain decision rules that are robust against an entire set of possible models. These decisions are rationalized by "ex-post" beliefs that are located on the boundary of such sets. Constructing these ex-post beliefs provides a lower bound to agents' expected valuations. Since ex-post beliefs depend on the magnitude of fluctuations in continuation values, they differ across agents at histories where their wealth differs. Thus there is a link between inequality in wealth and heterogeneity in ex-post beliefs.

The paper allows for a general specification where each agent describes his approximating model as a prior over a common but finite set of forecasting models for aggregate endow-

[1] In principle this distribution could be in a form of a single prior over a parsimonious set of probability models.

[2] See Blume and Easley (2006) for heterogeneous beliefs and Thomas and Worrall (1990) and Atkeson and Lucas (1990) for heterogeneous information.

[3] In a sense that they cannot be easily distinguished using finite data.

ment. Given some initial (and possibly heterogeneous) priors over these competing specifications, agents update their approximating models through time using observations of realized output. Except in the cases where agents are dogmatic [4], they will eventually focus their doubts around one of these specification.

Heterogeneity in ex-post beliefs affects the planner's trade-offs to redistribute resources. While allocating consumption, the planner internalizes how candidate risk sharing schemes affect ex-post beliefs. To facilitate the understanding of how asymmetries in priors or information combine with doubts to alter the nature of optimal insurance, the analysis is split into two parts. We first focus on the case without asymmetric information and then consider the implications of unobservable taste shocks.

With symmetric information but asymmetric priors the central issues revolve around how Milton Friedman's *market selection hypothesis* is reassessed. This hypothesis states that agents who make systematic errors in forecasting future contingencies are driven out of the market. In our case although some agents may begin with relatively "correct" approximating models, what matters for the market selection is how much agents depart from their approximating models owing to doubts. When utility is bounded from below, relatively poor agents have less to lose from misspecifications and hence distort their approximating models less. It is possible that the ranking of their ex-post beliefs relative to a data generating process flips on paths when either one of the agent is dominating the other in terms of wealth. Enduring model doubts can allow for a wider range of initial asymmetries in priors to be consistent with all agents surviving in the long run. The setup allows us to contrast the cases depending on whether these heterogeneities in priors are permanent or transitory. For example, when both agents are learning, they eventually have a common approximating model and in the limit their Pareto weights are equal.

In absence of asymmetric information, the efficient allocation can be decentralized using Arrow securities. Heterogeneity in ex-post beliefs affects both motives and prices at which trade in securities takes place. In particular, market price of risk depends on inequality in wealth, has cyclical fluctuations and may be sensitive to news shocks. To isolate these effects we study asset pricing moments from a perspective of an outside econometrician using common approximating model shared by both the agents. In the benchmark case when agents trust this common approximating model, market price of risk only depends on the exogenous volatility of aggregate endowment.

With doubts, the link between inequality and market price of risk depends on the intertemporal elasticity of substitution (IES). With IES > 1, concerns for misspecification get magnified when the agents share of aggregate endowments is larger. This leads to higher market prices of risk when inequality is high. The cyclical fluctuations come from state dependent distortions to the priors over the models when learning is active. For e.g., in recession the agent overweight (relative to a Bayesian estimate) the models with higher persistence.

These wealth driven differences in beliefs also open up a new role for "news" shocks. Without concerns for misspecification, complete markets and identical initial priors, these signals about future aggregate endowment are redundant to the agents if they are publicly observed. Concerns for misspecification, however, make the informative value of these shocks different across agents with different wealth. Consequently, consumption and prices of claims to aggregate endowments are sensitive not only to histories of aggregate shocks but also to news shocks. A policy that shuts down markets where people can trade on information is welfare reducing in this environment

These forces extend to settings with private information. To study this, we look at the case

[4] Dogmatic here means an agent with a degenerate initial prior over the set of forecasting model.

when agents have symmetric priors but Agent 2 faces an unobservable taste shock. When agents have complete faith in this common approximating model, Pareto weights fluctuate for incentive reasons. The planner insures this agent against high taste shocks by increasing consumption and lowering continuation values. This causes his Pareto weight to drifts to zero eventually. However, when agents have specification doubts, both the forces (heterogeneity in beliefs and incentive considerations) affect how Pareto weights evolve over time. Along the paths where relative Pareto weight of Agent 2 approaches zero, there are two opposing forces: (a) the fluctuations due to optimal incentives are dampened because the contract converges to the first best and (b) the disagreements across agents about future aggregate endowment, measured with respect to their ex-post beliefs, is maximized. This generates a drift that pushes Agent 2 out of immiseration.

In summary, this paper provides a common tractable framework to analyze the implications of Pareto efficient insurance arrangements in settings where the planner respects agents' heterogeneity in beliefs and information. The underlying theme is that heterogeneity in wealth affects ex-post differences in beliefs and consequently the incentives of the planner to redistribute resources in the future. This in turn affects the evolution of long run inequality.

The paper is structured as follows: we first provide a brief overview of the literature in section 2 and introduce how the attitudes towards uncertainty are modeled in section 3. Sections 4 and 5 will setup the environment for the general risk sharing problem and sections 6 and 7 specialize it to study the dynamics of consumption shares with symmetric and asymmetric information respectively. Finally, sections 8 and 9 study some extensions and outline directions for future research respectively.

2 Literature review

This paper builds on models of decision making under uncertainty developed in Hansen and Sargent (2001, 2005, and 2007). Concerns for misspecification are represented by versions of "multiplier-preferences." Macceroni, Marinacci and Rustichini (2004, 2006a, and 2006b) and Strazalecki (2011) provide axiomatic foundations.

The closest paper to our analysis is Anderson (2005). He sets up a planning problem with agents having similar preferences in an environment with a common approximating model, symmetric information and no learning. His paper identifies points in wealth distribution which are "absorbing" and shows convergence for a class of i.i.d shock process to such points. Our paper extends the analysis by adding heterogeneity in approximating models, private information about taste shocks and learning. With these modifications the stationary wealth distributions generically do not exist. This paper provides conditions under which the long run wealth distribution is bounded away from "immiseration" - a state when all aggregate endowment is consumed by one agent[5].

The asset pricing results complement findings in Hansen and Sargent (2010) [6]. With simpler dynamics for aggregate risk and learning, our paper extends the analysis to heterogeneous agents economies. The key results emphasize the relationship between wealth inequality and market

[5] Colacito and Croce(2010) study the condition under which the long-run equilibria in a two-good economy when agents are endowed with multiplier preferences and differ in the preferences over the two good are non-degenerate.

[6] Other papers that study asset pricing under models of multiple priors in representative agent frameworks include Ju and Miao (2012), Collard, Mukerji and Sheppard (2011) and Epstein and Schneider (2008), Boyarchenko (2012).

price of risk.

The results on long run inequality are related to two groups of papers that study redistribution in economies with heterogeneous agents - ones that have heterogeneous beliefs and others that have heterogeneous information.

With symmetric information Sandrioni (2000), Blume and Easely (2006) [7] study consequences of efficient redistribution in stationary economies with expected utility and heterogeneous beliefs. Their paper provides a benchmark (in our setting this will correspond to agents exhibiting complete trust in their models) against which some of the results in our paper can be contrasted.

Also related are papers by Borovicka (2013), Guerdjikova and Sciubba (2013) that study long run survival of agents with heterogeneous beliefs but preference that depart from expected utility. In Brownian information settings, Borovicka uses Epstein-Zin preferences to emphasizes the role of higher inter-temporal elasticity of substitution (IES) (relative to risk aversion) to guarantee that the agent with possibly incorrect beliefs can "save" their way out of immiseration. Guerdjikova and Sciubba (2013) study economies populated with agents who are either expected utility maximizers or have smooth ambiguity-averse investors, as in Klibanoff, Marinacci and Mukerji (2009). In our settings, IES will matter too. In particular, CRRA utility preferences with IES >1 will provide bounds on utilities which generates lower distortions to approximating models for poor agents.

The setup with heterogeneous information closely follows Phelan (1998), Thomas Worrall (1990) and Atkeson Lucas (1990). In a wide range of economies, these papers derive conditions which imply that long run immiseration is a necessary condition of efficient redistribution in presence of private information. In complementary work, Farhi and Werning (2007) and Phelan (2005) show examples where departing from the Pareto criterion can make efficient insurance compatible with bounded inequality. They interpret these findings as a justification for progressive estate taxation. In our setup, the planner maintains the Pareto criterion but endogenous heterogeneity in ex-post beliefs generates an auxiliary force that prevents agents with unobservable taste shocks be driven to immiseration.

3 Attitudes towards uncertainty

Throughout the paper, a model will refer to probability distributions over sequences of random variables, possibly indexed by a vector of parameters. Environments characterized by Knightian uncertainty correspond to scenarios where agents entertain multiple probability specifications for the exogenous randomness they face. Decision making in such environments often involve a selection rule that captures how they deal with this multiplicity. These can be distinguished from models where agents are assumed to have unique subjective beliefs and all choices over random payoffs are rationalized by a ranking generated via expected utility under those beliefs [8].

Consider a scenario where one begins with a state space Ω whose elements list all possible contingencies ("states of world") and a function $f : \Omega \to \Delta(Z)$ [9] that maps every state of

[7] There is a growing literature that studies survival of irrational traders. For more details see Bhamra and Uppal (2013), Kogan, Ross, Wang, and Westerfield (2011).

[8] Savage (1954) provides axioms on choices over random payoffs that guarantee existence and uniqueness of such beliefs and an associated utility function.

[9] These mappings are called "Anscombe-Aumann acts"

the world to a lottery over a given set of outcomes denoted by Z. The agents are confronted with choices amongst such mappings f. Given that there might not be enough information to ascertain the precise likelihood over the contingencies listed in Ω, should the agents treat mappings that deliver the same lottery for every contingency differently from mappings that deliver a known outcome every contingency[10]? These two kinds of mappings capture the difference between *objective* and *subjective* risks. The literature studying individual decision making under uncertainty, that models departures from the expected utility benchmark, attempts to separate attitudes towards these two types of risks. These departures have been successful in rationalizing facts observed both in experimental settings studying individual decision making and aggregate behavior in asset markets.

We borrow the precise formulation of these attitudes known as "multiplier-preferences " from Hansen Sargent (2005, 2007a, and 2007b). [11].

The decision maker starts with a possibly simple but potentially misspecified description of the world henceforth denoted as an "approximating" model. He surrounds the approximating model with a class of alternative probability specifications that are statistically close to it; and evaluates outcomes under all such models. With multiplier preferences the notion of statistical closeness corresponds to "relative entropy". [12]

Definition 1 *Consider Q and P two absolutely continuous measures over a set X, the relative entropy of Q with respect to P is defined as*

$$\mathcal{E}_{Q,P} = \int_X \log \frac{dQ}{dP} dQ$$

Let $z = \frac{dQ}{dP}$, we can write the above as $\mathbb{E}_P[z \log z]$

This set is hard to discriminate using finite data, but expected valuations can differ across its elements. The decision maker confronts this multiplicity by following a *robust* approach which provides a lower bound that guards him against possible misspecifications.

To elaborate the workings, we use a simple static problem of distributing output from a risky technology that yields $Y \in [\underline{Y}, \bar{Y}]$ amongst K agents who value consumption by $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$. These agents doubt the distribution of y and have $p_i(y)$ as their approximating model. They use multiplier preferences to evaluate the consumption risk implied by an arbitrary distribution of the endowment, $\eta = [\eta_1 \dots \eta_K]$ such that $\sum \eta_i = 1$. Here η_i is the share of Y that agent i consumes. (One can imagine that η represents relative wealth shares)

Given η_i and some $\theta > 0$, we describe agent i's valuations by the following *minimization* problem,

$$V(\eta_i) = \min_{z_i \geq 0; \mathbb{E}^i z_i = 1} \mathbb{E}^i z_i [u(\eta_i Y) + \theta \log(z_i)]$$

This minimization is over functions $z_i(y)$'s that index the set of alternative probability distributions $\hat{p}_i(y)$,

$$\hat{p}_i(y) = z_i(y) p_i(y).$$

[10]For example the Ellesberg paradox captures such thought experiments where people exhibit aversion towards unknown risks that cant be rationalized by unique prior models.

[11]For the sake of completeness we provide a brief summary. For a textbook length discussion refer Hansen and Sargent (2007).

[12]This is also referred to as the *Kullback-Leiblar divergence.* It bounds the (exponential) rate at which the probability error in an likelihood ratio test between two distributions declines with a sample size.

The criterion function has two terms a) $\mathbb{E}^i z_i[u(\eta_i Y)]$ which represents the valuation of agent i's consumption risk under $\hat{p}_i(y)$ and b) $\theta \mathbb{E}^i z_i \log(z_i)$, which is the relative entropy of $\hat{p}_i(y)$ with respect to $p_i(y)$ scaled by a positive constant θ.

As described before, relative entropy measures the distance of the alternative $\hat{p}_i(y)$ from agent i's approximating model $p_i(y)$. It is non-negative [13] and equal to zero if $z_i(y) = 1$. The minimization can be interpreted as a malevolent "alter" ego who chooses the worst case distribution for agent i subject to a cost that scales in the relative entropy.
The choice for z_i^* which solves this problem is

$$z_i^* \propto \exp\left\{ \frac{-\eta_i^{1-\gamma} y^{1-\gamma}}{\theta(1-\gamma)} \right\} \tag{1}$$

Equation (1) is also referred to *exponential twisting* or *statistical Murphy's law*. It states that agents confront doubts about model specification by adjusting probabilities of events inversely to their desirability. The probability distribution $\tilde{p}_i(y) = p_i(y) z_i^*(y)$ is the ex-post belief that rationalizes agent i's actions. The expression for $\tilde{p}_i(y)$ shows how heterogeneity in priors has both the exogenous $p_i(y)$ and endogenous components $(z_i^*(y))$. We may also refer to $\tilde{p}_i(y)$ as agent i's "worst case" model.
The following proposition shows how the relative pessimism depends on η_i and γ or the curvature of the u

Proposition 1 *Let η_i be the agent i's share of Y,*

1. $\lim_{\eta_i \to 0} \tilde{p}_i(\underline{Y}) = p_i(\underline{Y}) \quad \gamma < 1$
2. $\lim_{\eta_i \to 0} \tilde{p}_i(\underline{Y}) = 1 \quad \gamma > 1$

The search for alternatives balances the marginal costs of exploring new models and benefits from recognizing potential losses in expected valuations. Scaling valuations by a factor $k > 1$, will mean that for the same marginal costs, the benefits are larger. Thus agents will explore models in a larger set as measured by relative entropy. With $\gamma < 1$, a decrease a η_i lowers $\mathbb{E}^i z_i[u(\eta_i Y)] - \mathbb{E}^i[u(\eta_i Y)]$ monotonically until the spread reaches zero. With $\gamma > 1$ utility diverges to -∞ as consumption goes to zero and the aforementioned spread increases. This highlights the forces that will affect relative beliefs as wealth shares move around. Also the distinction between γ less than and greater than one plays an important role in how the long run wealth distribution evolves[14].

Using the version of static multiplier preferences as above, attitudes towards randomness are captured by jointly by two parameters : (θ, γ). The typical interpretation of γ as the coefficient of relative risk aversion is thus not alone sufficient for determining agents' willingness to pay for insurance against uncertain events. One could measure them separately by confronting the same agent in (controlled) environment where random events are associated either known probabilities

[13] This follows from Jensen's inequality,

$$-\tilde{\mathbb{E}}^i \log \frac{\tilde{p}_i(y)}{p_i(y)} \leq \log\left(\tilde{\mathbb{E}}^i \frac{p_i(y)}{\tilde{p}_i(y)} \right) = 0$$

.

[14] The dynamic version of $\gamma = 1$ multiplier preference are equivalent to Epstein-Zin preference (with IES=1 and risk aversion $1 - \theta$). This is the only case in which these the two preferences coincide.

or unknown probabilities and study how the certainty equivalent varies. However, in dynamic environments where there is a no uncertainty, γ has the dual interpretation of the inverse of inter-temporal elasticity of substitution (IES). The restriction $\gamma < 1$ then refers to IES greater than one [15]. As an alternative, detection error probabilities (See Anderson, Hansen and Sargent (2003), Barillas, Hansen and Sargent (2009)) can be used to map a pair (θ, γ) into an average odds that a statistician would wrongly conclude the data is generated by the worst case model if it was in fact drawn from the approximating model (or vice versa) based on likelihood ratio tests and a finite sample. In other words, it measures the confidence an agent has in his approximating model. A number close to 0.5 refers to complete trust in approximating model and a number close to zero corresponds to an individual with exceedingly high distrust of the approximating model.

A novel mechanism present when one departs from the expected utility benchmark is the link between *risk sharing* and *risk perceptions*. A risk sharing scheme is a state contingent plan of outcomes for an agent. As mentioned earlier, the agent uses his utility outcomes across various models as an input for selecting a relevant model of risk which rationalizes his actions. Thus how risk is shared affects his ex-post beliefs, but these beliefs in turn affect how much an agent is prepared to pay for the hedging that risk.

4 Setup

There is an exchange economy with stochastic aggregate endowment denoted by $y_t \in \mathcal{Y}$. The demography is described by two types of agents $\mathcal{I} = \{1, 2\}$. Agents of type 2 are subject to unobservable "taste shocks" [16] that take values $s_t \in \mathcal{S}$. The taste shocks will feature a multiplicative adjustment to the utility of consumption at time t. We assume both $\mathcal{Y}$ and $\mathcal{S}$ to be bounded[17].

There is a finite set $\mathcal{M}$ of Markov kernels for endowment risk $\{P_Y(y_{t+1}|y_t, m)\}_{m \in \mathcal{M}}$ with $P_Y(y'|y, m) > 0$ for all (y, y', m). Agents' approximating model over sequences of endowment and taste shocks are described by heterogeneous *initial priors* $\{\pi_{0,i} \in \Delta(\mathcal{M})\}_{i \in \mathcal{I}}$ and a distribution $P_S(s)$ from which the taste shocks are drawn. Under the approximating model, the taste shocks are i.i.d over time and orthogonal to aggregate endowment. The heterogeneity in agents' approximating models comes only from the differences in the initial priors $\pi_{i,0}(m)$ on the specifications for the aggregate endowment.

Given a consumption stream $c = \{c_t(y^t, s^t)\}_t$, let $v_t^1 = u(c_t(y^t, s^t))$ and $v_t^2 = s_t u(c_t(y^t, s^t))$ denote respective agents' felicities along a particular history (y^t, s^t). These felicities are computed using $u : \mathcal{R}_+ \to \mathcal{R}$ that maps consumption into time t utils. We assume u to be strictly concave, satisfy Indada condition and bounded from below. The subjective time discount factor $\delta \in (0, 1)$.

For some $(\pi_{i,0}(m),\ y_0)$, let V_t^i be Agent i's value from c given history (y^t, s^{t-1}). We first describe a Bellman equation that represents these valuations when agents have complete trust in their approximating models and then modify it to capture concerns for misspecification.

[15]There is large literature on whether this value is larger than unity or not. Refer Bansal Yaron (2004) , Hall (1988), Hansen and Singleton (1982) etc.

[16]The specification here assumes that all agents of type 2 get the same realization of the taste shock. The results are unaffected if we assume that there are a continuum of type 2 agents and the tastes shocks are i.i.d across agents. Section 7.2 elaborates on this generalization.

[17]For some results we will further restrict these sets be finite valued.

Under complete trust, valuations are computed by taking expected utility under Agent i's approximating model.

$$V_t^i[c|y^t, s^{t-1}] = \sum_{s_t \in \mathcal{S}} P_S(s_t) \left\{ (1-\delta)\upsilon_t^i(y^t, s^t) + \delta \sum_{m \in \mathcal{M}, y_{t+1} \in \mathcal{Y}} \pi_{i,t}(m) P_Y(y_{t+1}|y_t, m) V_{t+1}^i[c|y^{t+1}, s^t] \right\} \tag{2}$$

The averaging across models in $\mathcal{M}$ using $\pi_{i,t}(m)$ comes from applying Bayes rule,

$$\pi_{i,t}(m) \equiv Pr\{m|y^t\} = \frac{\pi_{i,t-1}(m) P_Y(y_t|y_{t-1}, m)}{\sum_m \pi_{i,t-1}(m) P_Y(y_t|y_{t-1}, m)}$$

We can summarize (2) below,

$$V_t^i[c|y^t, s^{t-1}] = \mathbb{E}\left\{(1-\delta)\upsilon_t^i + \delta \mathbb{E}_t^i V_{t+1}^i[c|y^{t+1}, s^t]\right\}, \tag{3}$$

where $\mathbb{E}_t^i$ takes expectations using

$$P_t^i(y_{t+1}|y^t) = \sum_m \pi_{i,t}(m) P_Y(y_{t+1}|y_t, m)$$

These values are computed before the taste shocks s_t are realized. The "outer" expectation operator in (3) uses $P_S(s)$ to average across the possible realizations of taste shocks.

With misspecification concerns, the Bellman equation representing valuations is given by,

$$V_t^i[c|y^t, s^{t-1}] = \mathbb{T}_{\theta_1}\left\{(1-\delta)\upsilon_t^i + \delta \mathbb{R}_{\pi_{i,t},\theta_2} \mathbb{T}_{\theta_1,m,y_t} V_{t+1}^i[c|y^{t+1}, s^t]\right\} \tag{4}$$

Recursion (4) replaces 'inner" expectation operator $\mathbb{E}_t^i$ with two operators $\mathbb{R}_{\theta_2,t}^i$ and $\mathbb{T}_{\theta_1,m,t}$ and the "outer" expectation operator by $\mathbb{T}_{\theta_1}$. These operators are parametrized by a pair $(\theta_1, \theta_2) \in \Theta \subset \mathcal{R}_+^2$.

We use two operators to distinguish two sources of misspecification to $P_t^i(y_{t+1}, s_t|y^t, s^{t-1}) = \sum_m \pi_{i,t}(m) P_S(s_t) P_Y(y_{t+1}|y_t, m)$.

- Uncertainty about $P_Y(y_{t+1}|y_t, m)$ and $P_S(s_t)$ using $\mathbb{T}_{\theta_1,m,y_t}$ and $\mathbb{T}_{\theta_1}$, respectively
- Uncertainty about $\pi_{i,t}(m)$ using $\mathbb{R}_{\pi_{i,t}\theta_2}$

In our setting, agents begin with a common set of misspecified models $\mathcal{M}$. The concerns about the Markov kernels $P_S(s)$, $P_Y(y_{t+1}|y_t, m)$ capture a large set of alternatives including low frequency changes, structural breaks, non-linearities. This would usually pose an *infinite dimensional learning problem* to a Bayesian[18] and are addressed here using the $\mathbb{T}$ operators.

On the other hand uncertainty about $\pi_{i,t}(m)$ comes from a *finite dimensional learning problem* and is more structured. For example each $m \in \mathcal{M}$ could represent some time invariant parameters that describe the transition matrix of y. The agents' approximating model is described by a initial priors over $\mathcal{M}$ and is updated through time using observations of y^t. The distinction between parameter risks and endowment risk allows us to partition the class of misspecifications into different sets, each centered on $P_Y(y_{t+1}|y_t, m)$. In this generalized setup the

[18] See Sims(1971b), Diaconis and Freedman (1983) for issues related to infinite dimensional learning problems and consistency of Bayes rules.

agents use a two step procedure to confront misspecifications. They first surround each specification with "cloud" of alternatives that are close to it and settle on alternative set of worst case models. Next they address estimation uncertainty by distorting the estimate of $\pi_{i,t}(m)$ [19]

We now describe in detail how these operators are used to construct valuations. These will be the dynamic version of the example studied in section 3.

Fix a model $m \in \mathcal{M}$. Given a bounded random variable $W_{t+1}(y_{t+1}, s_t)$, operator $\mathbb{T}_{m,\theta,y}$ is defined by

Definition 2

$$\mathbb{T}_{\theta_1,m,y_t} W_{t+1} = \min_{z^m_{t,t+1}} \mathbb{E}^m_t z^m_{t,t+1} W_{t+1} + \theta_1 \mathbb{E}^m_t z^m_{t,t+1} \log(z^m_{t,t+1}) \tag{5}$$

s.t

$$\mathbb{E}^m_t z^m_{t,t+1} = 1$$

Operator $\mathbb{T}_{\theta_1,m,y}$ searches over $z^m_{t,t+1}$ which are time t likelihood ratios given by

$$z^m_{t,t+1} = \frac{\tilde{P}^m_t(y_{t+1}|s_t)}{P_Y(y_{t+1}|y_t, m)}$$

Being Radon-Nikodym derivatives, these are constrained to have a mean of unity under the approximating model. The minimizing likelihood ratio associated with $\mathbb{T}_{\theta_1,m,y_t} W_{t+1}$ is given by

$$z^m_{t,t+1}(y_{t+1}, s_t) \propto \exp\left\{-\frac{W_{t+1}}{\theta_1}\right\}$$

and

$$\hat{W}_t(m, s_t) = \mathbb{T}_{\theta_1,m,y_t} W_{t+1}$$

The operator $\mathbb{T}_{\theta_1,m,y_t}$ thus takes conditional expectations of W_{t+1} under the distribution $\tilde{P}_Y(y_{t+1}|y_t, m)$ and adds the relative entropy of $\tilde{P}(y_{t+1}|y_t, s_t, m)$ with respect to $P_Y(y_{t+1}|y_t, m)$ scaled by penalty θ_1. We will denote $\tilde{P}^m_t$ associated with the minimizing $z^m_{t,t+1}$ as agents' "ex-post beliefs."

Note that with $\theta_1 \to \infty$, we approach the limit where the agent completely trusts the specification $m \in \mathcal{M}$ and applies the usual conditional expectation operator $\mathbb{E}^m_t$. In this case the penalty for distorting the approximating model is arbitrarily high.

We now move to how the agent addresses uncertainty associated with $\pi_{i,t}(m)$. Agents apply Bayes rule under the approximating model and express their doubts on the outcome $\pi_{i,t}(m)$.

Definition 3

$$\mathbb{R}_{\pi_{i,t},\theta_2} \hat{W}_t(m, s_t) = \min_{h_t(m)} \sum_m h_t(m)\pi_{i,t}(m)\hat{W}_t(m, s_t) + \theta_2 \sum_m h_t(m)\pi_{i,t}(m)\log(h_t(m)) \tag{6}$$

s.t

$$\sum_m \pi_{i,t}(m) h_t(m) = 1$$

[19] This is obtained using Bayes rule using $\pi_{i,0}$ and $\{P_Y(y_{t+1}|y_t, m)\}$ and the history of signals y^t. Hansen Sargent (2007b) refer to this procedure as "Robustness with no-commitment."

As before h_t is given by

$$h_t(m, s_t) \propto \exp\left\{-\frac{\hat{W}_t(m, s_t)}{\theta_2}\right\}.$$

Let $\bar{W}_t(s_t) = \mathbb{R}_{\pi_{i,t},\theta_2}\hat{W}_t(m, s_t)$. Finally, we can address uncertainty about $P_S(s)$.

Definition 4

$$\mathbb{T}_{\theta_1}\bar{W}_t(s_t) = \min_{z_t^s} \mathbb{E}z_t^s\bar{W}_t(s_t) + \theta_2\mathbb{E}z_t^s\log(z_t^s) \tag{7}$$

s.t

$$\mathbb{E}z_t^s = 1$$

Remark 1 *When $\theta_2 = \theta_1 = \theta$, we can collapse the composite operator: $\mathbb{R}_{\theta,t}\mathbb{T}_{\theta,m,t}$ into $\mathbb{T}_{\theta,t}$ The agent reduces the compound lottery $\{\pi_t(m), P_Y(y_{t+1}|y_t, m)\}_{m\in\mathcal{M}}$ to $P_t(y_{t+1}) = \sum_m \pi_t(m)P_Y(y_{t+1}|y_t, m)$ and explores misspecification around it using*

$$\mathbb{T}_{\theta,t}[W_{t+1}] = \min_{\hat{z}_t(y_{t+1})} \sum_{y_{t+1}} P_t(y_{t+1})\hat{z}_t(y_{t+1})W_{t+1} + \theta\sum_{y_{t+1}} P_t(y_{t+1})\hat{z}_t(y_{t+1})\log(\hat{z}_t(y_{t+1}))$$

s.t

$$\sum_{y_{t+1}} P_t(y_{t+1})\hat{z}_t(y_{t+1}) = 1$$

The Bellman equation (4) uses these operators to iterate on V_t^i. History dependence in valuations translates into history dependence in the distortions and the specification of the ex-post beliefs. Thus, agents explores a rich class of alternatives around the set of misspecified models given observed histories and use fluctuations in valuations to guide which of them matter more to them .

Remark 2 *Both beliefs and taste shocks affect the value from consumption multiplicatively. One can combine these to be interpreted as "new" taste shock with a permanent component that is publicly observable and transitory component which is privately observable. Suppose that the data is generated from some $P_t^0(y^t, s^t)$. One can rewrite Agent 2's preferences sequentially with this combined taste shock. Consider the case when $\theta_1 = \theta_2 = \infty$. Equation (3) that captures the value recursion can be re-written as,*

$$\mathbb{E}^0\left[\sum_{t=0}^{\infty}\beta^t\vartheta_t u_t\right]$$

Denote $\psi_{t,t-1} = \frac{\sum_m \pi_t(m)P_Y(y_t|y_{t-1},m)}{P_t^0(y_t)}$ we have the following law of motion for ϑ_t

$$\log\vartheta_{t+1} = \log\vartheta_t + \log\psi_{t,t+1} + \log s_{t+1}$$

Without concerns for uncertainty the dynamics of ϑ_t are exogenous and independent of the underlying consumption stream. With concerns for uncertainty there is an additional term

$$\zeta_{t,t+1} = \frac{\tilde{P}_t(y_{t+1}, s_{t+1})}{\sum_m \pi_t(m)P_Y(y_{t+1}|y_t, m)P_S(s_{t+1})}$$

such that

$$\log \vartheta_{t+1} = \log \vartheta_t + \log \psi_{t,t+1} + \log \zeta_{t,t+1} + \log s_{t+1}$$

The probabilities $\{\tilde{P}_t\}_t$ *are the Agent 2's ex-post beliefs associated with the minimizations in (5), (6) and (7). The dynamics coming from* $\zeta_{t,t+1}$ *will be endogenous and depend on the equilibrium consumption distribution.*

The restrictions on preferences and technology allow us to give tighter characterizations of the long run dynamics, albeit leaves interesting extensions outside the scope of this paper. Most importantly the boundedness of $\mathcal{Y}$ implies a stationary process for aggregate endowment. Kogan et all (2009) study the implications of heterogeneous beliefs in growing economies where agents have perfect trust in their models. Their paper details the joint restrictions on preferences, growth rate of endowment, and heterogeneity in beliefs that matter for how long run wealth shares evolve. We conjecture that such forces will be relevant even in settings with specification doubts, but leave a detailed analysis for future work. The assumption on the lower bound on the utility function is important for how distortions differ across agents on paths with extreme wealth inequalities. This point alluded to in the simple examples in section 3 and we will return to it in more details later.

5 Efficient allocations

Our economy features non-diversifiable aggregate risk and stochastic taste shocks to Agent 2. This section will describe risk sharing arrangements that are Pareto efficient subject to informational constraints. [20]
We begin with some standard definitions.

Definition 5 (Allocation) *An allocation* $\boldsymbol{c} = \{c_{i,t}(y^t, s^t)\}_{i,t}$ *is a collection of history contingent consumption functions.*

Definition 6 (Feasibility) *An allocation is a said to be feasible, if*

$$\sum_i c_{i,t}(y^t, s^t) \le y_t \quad \forall t$$

In absence of asymmetric information, this will be the only restriction on allocations. Next we define a reporting strategy for Agent 2,

Definition 7 (Reporting strategy) *A reporting strategy* $\sigma = \{\sigma_t(y^t, s^t) \in \mathcal{S}\}$ *is a collection of history contingent reports of contemporaneous taste shocks. Denote* $\sigma^t = \{\sigma_j(y^j, s^j)\}_{j \le t}$ *as the collection of reports up to time* t.

Histories of aggregate shocks y^t are observable. Given a pair $(\boldsymbol{c}, \sigma)$, actual consumption is $\boldsymbol{c}(\sigma) = \{c_{i,t}(y^t, \sigma^t(y^t, s^t))\}_{i,t}$. We denote truth-telling strategies by $\sigma_t^*(y^t, s^t) = s_t$.

Let $V_0^i[\boldsymbol{c_i}(\sigma)]$ be the ex-ante valuations given y_0 from an allocation and strategy pair $(\boldsymbol{c}, \sigma)$ using recursion (4).

[20] Is the Pareto criterion reasonable? Brunnermeier et all (2012) and Blume et all (2013) dig deeper into this issue. Both papers approach the issue in an economy with exogenous heterogeneous beliefs but allow the planner to evaluate social welfare under different criteria that are motivated from taking a normative stance on outcomes like "speculation" or "immiseration" associated with belief heterogeneity under complete markets.

Definition 8 (Efficiency) *An allocation $\boldsymbol{c}$ is (constrained) efficient if there does not exist any other incentive compatible - feasible allocation $\tilde{\mathbf{c}}$ such that*

$$V_0^i[\tilde{c}_i] \geq V_0^i[c_i] \quad i = 1,2 \tag{8}$$

$$V_0^i[\tilde{c}_i] > V_0^i[c_i] \quad \textit{for some } i \tag{9}$$

and

$$\sigma^* = argmax_\sigma V_0^2[\tilde{\boldsymbol{c}}_2(\sigma)] \tag{10a}$$

$$\sum_i \tilde{c}_{i,t}(y^t, s^t) \leq y_t \tag{10b}$$

We can index the efficient allocations by an pair of initial Pareto weights $(\Gamma, 1-\Gamma)$ for Agent 1 and 2 respectively. The optimal allocation $\boldsymbol{c}(\Gamma)$ solves

$$\max_{\boldsymbol{c}(\Gamma)} \Gamma V_0^1[\boldsymbol{c_1}] + (1-\Gamma)V_0^2[\boldsymbol{c_2}] \tag{11}$$

subject to

$$c_{1,t}(y^t, s^t) + c_{2,t}(y^t, s^t) \leq y_t \tag{12a}$$

$$\sigma^* \in \text{argmax}_\sigma V_0^2[\boldsymbol{c}_2(\sigma)] \tag{12b}$$

The valuations V_0^i encode heterogeneity in approximating models and agents doubts regarding them. Equation (12b) respects heterogeneity in information by imposing restrictions on the allocations that induce Agent 2 to truthfully report his taste shock.

In the next section, we will characterize the set of efficient allocations recursively using 'Dynamic programing squared'[21]. In this approach, continuation values are used as state variables to get parsimonious representation of the consumptions dynamics under an efficient plan. In our problem, these continuation values summarize history dependence. As alluded earlier, current beliefs and actions of an agent depend on entire stream of future utilities which are chosen by the planner. Thus at any given history, choices of the planner are constrained by past commitments. A similar force is present in a context of a Ramsey problem with forward looking constraints. [22]

The planner chooses a menu of current and future utilities to maximize the value for Agent 1 while achieving a given promised value for Agent 2. As with the sequential problem (11), the solution to recursive problem characterizes all efficient allocations by tracing out the Pareto frontier.

---

[21] Refer Ljungqvist and Sargent (2013) Chp. 20-24 for a textbook treatment of a wide range of problems with enforcement or information friction that can be analyzed using these methods.

[22] For e.g.. Kydland and Prescott (1980) and Lucas and Stokey (1983), Aiyagari et al. (2002) who study optimal taxation with commitment.

5.1 Recursive formulation of planner's problem

Without loss of generality, we represent choices for consumption c_t with choices for utility levels u_t by inverting the the felicity function $u(c_t) = u_t$. Let $C(u)$ denote this inverse.[23]

Let $\boldsymbol{\pi} = \{\pi_i(m) \in \mathcal{M}\}_i$ and y denote the collection of agents' priors over $\mathcal{M}$ and the current aggregate output respectively. Let $\mathcal{Q}(\boldsymbol{\pi}, v, y)$ be the maximum value the planner can achieve for Agent 1 provided he delivers a level v of promised value to Agent 2 before the taste shocks are realized. The planner does this by choosing a report contingent menu of current and future utilities $\{u_1(s), u_2(s), \bar{v}(s, y^*)\}$ to solve the following Bellman equation,

$$\mathcal{Q}(\boldsymbol{\pi}, v, y) = \max_{u_1(s), u_2(s), \bar{v}(s,y^*)} \mathbb{T}_{\theta_1}\left[(1-\delta)u_1(s) + \delta \mathbb{R}_{\pi_1,\theta_2}\mathbb{T}_{\theta_1,m,y}\mathcal{Q}(\boldsymbol{\pi}^*, \bar{v}(s,y^*), y^*)\right] \tag{13}$$

s.t

(a) **Promise keeping:**

$$\mathbb{T}_{\theta_1}\left[(1-\delta)su_2(s) + \delta \mathbb{R}_{\pi_2,\theta_2}\mathbb{T}_{\theta_1,m,y}\bar{v}(s,y^*)\right] \geq v \tag{14a}$$

(b) **Incentive compatibility:** For all s, s'

$$(1-\delta)su_2(s) + \delta \mathbb{R}_{\pi_2,\theta_2}\mathbb{T}_{\theta_1,m,y}\bar{v}(s,y^*) \geq (1-\delta)su_2(s') + \delta \mathbb{R}_{\pi_2,\theta_2}\mathbb{T}_{\theta_1,m,y}\bar{v}(s',y^*) \tag{14b}$$

(c) **Feasibility:** For all s

$$C(u_1(s)) + C(u_2(s)) \leq y \tag{14c}$$

(d) **Bayes Rule:** For all i

$$\pi_i^*(m) \propto \pi_i(m) P_Y(y^*|y, m) \tag{14d}$$

(e) **Bounds:** For all (s, y^*)

$$\bar{v}(s, y^*) \leq v^{max}(\pi_2^*(y^*), y^*) \tag{14e}$$

Equation (14a) is the promise keeping constraint that makes sure that the value of the optimal contract to Agent 2 is at least higher than what was "promised" to him, v. Equations (14b) are the incentive compatibility constraints that induce Agent 2 to reveal truthfully the realized taste shock. These constraints are ex-post in the sense that they are evaluated for each possible realization of s. The bounds $v^{max}(\pi_2, y)$ solves,

$$v^{max}(\pi_2, y) = \mathbb{T}_{\theta_1}\left[(1-\delta)su(y) + \delta \mathbb{R}_{\theta_2,\pi_2}\mathbb{T}_{\theta_1,m,y}v^{max}(\pi_2^*(y^*), y^*)\right]$$

These bounds represent the value for Agent 2 from consuming the entire aggregate endowment from the next period onwards.

The next proposition establishes existence of $\mathcal{Q}$ under the boundedness restrictions on $\mathcal{Y}$, $\mathcal{S}$ and felicity functions u. The arguments can be extended to handle more general cases [24]

Proposition 2 *There is a unique solution to the planner's problem (13). Further $Q(\boldsymbol{\pi}, v, y)$ is decreasing and concave in v*

[23] It is easy to check that concavity of u is equivalent to convexity of C.

[24] See discussion in Anderson (2005), Kan (1985), Lucas and Stokey (1984).

The Pareto frontier is thus downward sloping. The solution to the Bellman equation (13) gives us a law of motion for the state variables $(\boldsymbol{\pi}, v)$ and policy rules that map $(\boldsymbol{\pi}, v, y)$ to respective consumptions for both the agents. We can use these to construct an efficient allocation allocation $\boldsymbol{c}$.

Suppose we start time 0 with some level of ex-ante promised value v_0 and initial conditions $(\{\pi_{0,i}(m)\}_i, y_0)$. This choice of v_0 will correspond to some Γ in problem (11).

1. **State variables:** v_t and $\{\pi_{i,t}(m)\}_i$ capture history dependence,

$$v_t = \bar{v}(s_{t-1}, y_t | \boldsymbol{\pi}_{t-1}, v_{t-1}, y_{t-1})$$

.

$$\pi_{i,t}(m) = \frac{\pi_{i,t-1}(m) P_Y(y_t | y_{t-1}, m)}{\sum_m \pi_{i,t-1}(m) P_Y(y_t | y_{t-1}, m)}$$

2. **Allocation:** Given by policy rules,

$$c_{2,t}(y^t, s^t) = C[u_2(s_t | \boldsymbol{\pi}_t, v_t, y_t)]$$

$$c_{1,t}(y^t, s^t) = C[u_1(s_t | \boldsymbol{\pi}_t, v_t, y_t)]$$

Consider the Lagrangian for the Planner's problem

$$\mathcal{L}(\pi, v, y, \lambda) \equiv \mathbb{T}_{\theta_1}\left[(1-\delta)u_1(s) + \delta \mathbb{R}_{\pi_1,\theta_2}\mathbb{T}_{\theta_1,m,y}\mathcal{Q}(\boldsymbol{\pi}^*, \bar{v}(s, y^*), y^*)\right] + \lambda\left\{\mathbb{T}_{\theta_1}\left[(1-\delta)su_2(s) + \delta \mathbb{R}_{\pi_2,\theta_2}\mathbb{T}_{\theta_1,m,y}\bar{v}(s, y^*)\right] - v\right\}$$

The multiplier λ plays the role of the relative "Pareto weights" for Agent 2. The concavity of the planner's value function $\mathcal{Q}$ additionally implies a monotonic relationship between λ and v. This follows from the Envelope theorem which implies $\lambda = -\mathcal{Q}_v$

The next proposition establishes the equivalence of the problems (11) and (13). The key step involves showing that the temporary ex-post incentive constraints in (14b) are equivalent to the ex-ante incentive constraints in (12b) by exploiting the one-deviation principle.

Proposition 3 *For every $\Gamma \in (0,1)$ there exists a v_0 such that the solution to recursive problem (13) given v_0 and initial conditions $(\{\pi_{0,i}(m)\}_i, y_0)$ solves the sequential problem (11) for the pair of initial Pareto weights $(\Gamma, 1-\Gamma)$ and conversely.*

Heterogeneity in ex-post beliefs

The planner internalizes the fact that ex-post beliefs [25] change with candidate incentive compatible consumption allocation. Given the optimal allocation $\boldsymbol{c}$ and associated continuation values for each agent $V_t^i(y^t, s^{t-1})$, we can compute as equilibrium ex-post beliefs using,

$$\tilde{P}_t^i(y_{t+1} | s_t, y_t, m) \propto P_Y(y_{t+1} | y_t, m) \exp\left\{-\frac{V_{t+1}^i}{\theta_1}\right\} \tag{15a}$$

[25] The probabilities associated with the minimizations in operators $\mathbb{T}_{\theta_1,m,y}, \mathbb{R}_{\pi_i,\theta_2}$ and $\mathbb{T}_{\theta_1}$.

$$\tilde{\pi}_{i,t}(m|s_t) \propto \pi_{i,t}(m)\exp\left\{-\frac{\mathbb{T}_{\theta_1,m,y_t}V^i_{t+1}}{\theta_2}\right\} \tag{15b}$$

$$\tilde{P}^i_t(s_t) \propto P_S(s)\exp\left\{-\frac{\left[(1-\delta)v^i_t + \delta\mathbb{R}_{\pi_{2,t},\theta_2}\mathbb{T}_{\theta_1,m,y_t}V^i_{t+1}\right]}{\theta_1}\right\} \tag{15c}$$

Dynamic insurance problems in settings with heterogeneity in beliefs or heterogeneous information typically feature time varying Pareto weights λ_t. The heterogeneity of beliefs in our case has two components the exogenous coming from $\{\pi_{i,0}(m)\}_i$. and endogenous coming from the exponential twisting formulae. The next lemma summarizes these dynamics.

Lemma 1 *Let $\tilde{P}^i_t(y_{t+1}|s_t) = \sum_m \tilde{\pi}_{i,t}(m)\tilde{P}^i_t(y_{t+1}|y_t, s_t, m)$, λ_t be the Lagrange multiplier on the PK constraint* (14a) *and $\mu_t(s)$ be the (scaled) multiplier on the incentive constraints* (14b)*. The one period ahead growth rate of λ_t under the optimal allocation are given by,*

$$\frac{\lambda_{t+1}}{\lambda_t} = \underbrace{\frac{\tilde{P}^2_t(s_t)\tilde{P}^2_t(y_{t+1}|s_t)}{\tilde{P}^1_t(s_t)\tilde{P}^1_t(y_{t+1}|s_t)}}_{\textit{Heterogeneous beliefs}} \underbrace{\left[1 + \mu_t(s_t) - \sum_{s'_t \neq s_t} \mu_t(s'_t)\frac{\tilde{P}^2_t(s'_t)\tilde{P}^2_t(y_{t+1}|s'_t)}{\tilde{P}^2_t(s_t)\tilde{P}^2_t(y_{t+1}|s_t)}\right]}_{\textit{Heterogeneous Information}} \tag{16}$$

Proof. Take the F.O.C condition w.r.t. $\bar{v}(s, y^*)$.

$$\begin{aligned}\mathcal{Q}_v(\boldsymbol{\pi}^*, \bar{v}(s, y^*), y^*)\tilde{P}^1(s)\sum_m \tilde{\pi}_1(m)\tilde{P}^1(y^*|y, s, m) =& \lambda\sum_m \tilde{P}^2(s)\tilde{\pi}_2(m)\tilde{P}^2(y^*|y, s, m)\\ &+ \hat{\mu}(s)\sum_m \tilde{\pi}'_2(m)\tilde{P}^2(y^*|y, s, m)\\ &- \sum_{s'\neq s}\hat{\mu}(s)\sum_m \tilde{\pi}'_2(m)\tilde{P}^2(y^*|y, s', m)\end{aligned} \tag{17}$$

Define $\mu(s)$ as

$$\hat{\mu}(s) = \lambda\mu(s)\tilde{P}^2(s)$$

Finally using the Envelope theorem $\lambda = -\mathcal{Q}_v(\pi, v, y)$ with the first order conditions for $\bar{v}(s, y^*)$ to obtain (16)

■

The remaining part of the paper accounts for the differences in consumption patterns due to concerns for model misspecification. Our benchmark will be the case when agents completely trust their approximating models. To focus on the effects of heterogeneous initial priors $\{\pi_{i,0}(m)\}$ and heterogeneous information, the analysis is split in two parts, each of which focuses on one of the sources of heterogeneity. We begin in section 6 by shutting down the heterogeneity in private information by assuming $\mathcal{S} = \{1\}$. Finally in section 7 we show how the main results on long run inequality extend to case with heterogeneous information.

6 Efficient allocation with symmetric information

In absence of taste shocks, the dynamics of Pareto weights in equation (16) simplify and the optimal allocation is characterized by

$$c_{i,t} = c_i(\lambda_t, y_t) \tag{18a}$$

$$\frac{\lambda_{t+1}}{\lambda_t} = \frac{\sum_m \tilde{\pi}_{2,t}(m)\tilde{P}_t^2(y_{t+1}|y_t, m)}{\sum_m \tilde{\pi}_{1,t}(m)\tilde{P}_t^1(y_{t+1}|y_t, m)} \tag{18b}$$

where the probabilities can be computed using (15)

Given a Pareto weight λ_t, the planner uses conditions (18a) to distribute the current aggregate endowment y_t. This combines the first order conditions for c_1, c_2 and the resource constraint. The Lagrange multiplier $\lambda_t = \frac{u_c(c_{1,t})}{u_c(c_{2,t})}$ and measures the time t ratio of marginal utilities. Concavity of u thus implies a strictly monotonic relationship between consumption of Agent 2 and λ.

Equations (18b) show that the planner accounts for the differences in the ex-post beliefs, for example, by by increasing future Pareto weight in states where the Agent 2 is relatively optimistic.

Remark 3 *The allocations are efficient in an "alternative" economy where agents have no misspecification concerns but heterogeneous beliefs given by* $\{\tilde{P}_t^i\}_{i,t} = \sum_m \tilde{\pi}_{i,t}(m)\tilde{P}_t^i(y_{t+1}|y_t, m)$.*In absence of private information, the optimal allocation can be supported by a complete set of one period ahead state contingent securities.*

Suppose agents start with identical initial priors $\pi_{1,0} = \pi_{2,0}$ and the initial promised value v_0 is chosen to satisfy $\mathcal{Q}(\boldsymbol{\pi}_0, \hat{v}(y, \boldsymbol{\pi}_0), y) = \hat{v}(y, \boldsymbol{\pi}_0)$. This is a situation when two agents are identical and the optimal allocation features equal consumption shares forever. It is easy to see that at this candidate allocation, the spread between continuation values is same across agents for all histories and they will agree on the relative assessments of future shocks. The initial condition corresponds to $\lambda_0 = 1$ and the allocation implies $\lambda_t(y^t) = 1$ for all histories y^t.

The next section 6.1 we use the decentralization using Arrow securities to emphasize some properties of asset prices and the role of public information in generating motives for trade.

In section 6.3 we trace how misspecification doubts affect long run inequality due to heterogeneity in initial priors $\{\pi_{i,0}(m)\}_i$. The results will broadly depend on whether differences are dogmatic, e.g., for some $m^* \neq m^{**} \in \mathcal{M}$, $\pi_{1,0}(m^*) = 1$ and $\pi_{2,0}(m^{**}) = 1$ or transient, e.g., $\pi_i(m) > 0$ for all i, m.

6.1 Decentralization

In a decentralized equilibrium with Arrow-securities, agents trade until margin utilities of consumption weighted by their respective beliefs are equalized across states. Heterogeneity in beliefs, thus, affects both motives and prices at which such trade takes place. One can back out the shadow prices for these one-period ahead Arrow securities using the stochastic discount factor implied by the optimal allocation. Using consumption and ex-post beliefs of any agent i one obtains,

$$q(y^*|\pi, v, y) = \delta \left(\frac{\sum_{m\in M} \tilde{\pi}^i(m|\pi, v, y)\tilde{P}^i(y^*|m, \pi, v, y)u_c^i(\lambda^*(y^*), y^*)}{u_c^i(\lambda(\pi, v, y), y)} \right)$$

For the rest of the section, we will assume $\pi_{0,1} = \pi_{0,2} = \pi_0$. This isolates the role of endogenous heterogeneity in beliefs on properties of asset prices and makes it convenient to compute moments from a perspective of an econometrician who shares this common approximating model. Concerns for uncertainty make asset prices more volatile and this volatility is related to the degree of wealth inequality. Further uncertainty about $\pi_t(m)$ can potentially deliver conditional volatilities that are higher in recessions as against booms. For some calculations in this section we assume $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$.

Given λ, equation (18a) gives us consumption for each of the agent

$$c_1(\pi, v, y) = \frac{1}{[1 + \lambda(\pi, v, y)]^{\frac{1}{\gamma}}} y \equiv \eta_1(\pi, v, y)y \tag{19}$$

$$c_2(\pi, v, y) = \frac{\lambda(\pi, v, y)^{\frac{1}{\gamma}}}{[1 + \lambda(\pi, v, y)]^{\frac{1}{\gamma}}} y \equiv \eta_2(\pi, v, z)y \tag{20}$$

Substituting equation (18b) and using the above expressions to compute $q(y^*|\pi, v, y)$, we have

$$q(y^*|\pi, v, y) = \delta \left\{ \sum_i \omega_i^* \left(\sum_{m \in M} \tilde{\pi}^i(m) \tilde{P}^i(y^*|m, \pi, v, y) \right) \right\} \left(\frac{y^*}{y} \right)^{-\gamma} \tag{21}$$

Where ω_i^* are given by

$$\omega_i^*(y^*|\pi, v, y) = \frac{\eta_i^{-\gamma}(\pi^*(y^*), v^*, y^*)}{\sum_{i=1,2} \eta_i^{-\gamma}(\pi, v, y)}$$

Denote $q^{BM}(y^*|y)$ be the asset prices when both concerns for uncertainty are absent $\theta_1 = \theta_2 = \infty$

$$q^{BM}(y^*|\pi, v, y) = \delta \sum_m \pi(m) P_Y(y^*|y, m) \left(\frac{y^*}{y} \right)^{-\gamma}$$

So we have that $\frac{q(z^*|v,z)}{q^{BM}(z^*|z,\pi)} = \zeta(y^*|\pi, v, y)$, where

$$\zeta(y^*|\pi, v, y) = \left\{ \sum_{i=1,2} \omega_i^* \frac{\sum_{m \in M} \tilde{\pi}^i(m) \tilde{P}^i(y^*|y, m)}{\sum_m \pi(m) P_Y(y^*|y, m)} \right\}$$

The asset prices in this economy feature an additional multiplicative adjustment ζ has two components - one coming from the relative likelihood of the worst case model with respect to the approximating model and the other reflecting how consumption shares evolve. Barillas, Hansen and Sargent (2009) point out the importance of the first component in explaining the volatility of asset prices and equity premium related puzzles. When we have $v \neq \hat{v}(y, \pi)$, the point corresponding to equal initial Pareto weights, there is feedback from wealth inequality to asset prices.

Pricing kernel

Define $\rho_t(y)$ as the stochastic discount factor that prices one period ahead state contingent cash-flows $f(y)$ under the common approximating model.

Definition 9 (Pricing kernel)

$$\mathbb{P}_t(f) = \mathbb{E}_t \rho_t(y_{t+1}) f(y_{t+1})$$

Evidently,

$$\rho_t(y_{t+1}) = \delta \frac{u_c(c_{i,t+1})}{u_c(c_{i,t})} \left(\frac{\sum_{m \in M} \tilde{\pi}_{i,t}(m) \tilde{P}_t^i(y_{t+1}|y_t, m)}{\sum_{m \in M} \pi_t(m) P_Y(y_{t+1}|y_t, m)} \right)$$

For CRRA specification for u,

$$\rho_t(y_{t+1}) = \zeta_t(y_{t+1}) \left(\frac{y_{t+1}}{y_t} \right)^{-\gamma} \tag{22}$$

We use conditional variance of the (log) pricing kernel[26] as a measure for summarizing the market price of risk in this economy.

$$\text{MPR}[\pi, v, y] = \text{var}[\log(\rho)|\pi, v, y]$$

The volatility is with respect to the common approximating model i.e $\sum_m \pi(m) P_Y(y^*|y, m)$

Under the benchmark when agent trust their approximating models, MPR corresponds to volatility of aggregate growth rate scaled by γ (risk aversion). It is thus independent of wealth shares.

With concerns for uncertainty MPR is increasing in wealth inequality with IES>1. To see this consider a simple example with $\mathcal{Y} = \{y_l, y_h\}$.

Suppose $y_h > y_l$ and $u(c) = c^{1-\gamma}/(1-\gamma)$ and denote the spread in utilities, $\Delta[\eta] \equiv u(\eta y_h) - u(\eta y_l)$

$$\Delta[\eta] = \eta^{1-\gamma}[u(y_h) - u(y_l)]$$

It is easy to see that with IES > 1 this spread is increasing in η.

With identical initial priors, $v = \hat{v}$ that satisfies $\mathcal{Q}(\pi_0, \hat{v}(y, \pi_0), y) = \hat{v}(y, \pi_0)$. This corresponds to equal initial Pareto weights and an optimal allocation given by

$$c_{1,t} = \frac{1}{2} y_t$$

WLOG, let $v_0 \to v^{min}$, Agent 1 consumes the entire aggregate endowment

$$c_{1,t} = y_t.$$

Thus as we go from $\hat{v}$ and v^{min}, the share of endowment that Agent 1 consumes doubles. With IES>1, the increasing spread in utilities imply that the distortions to the common approximating model are increasing with wealth inequality and vice versa.

Standard consumption based asset pricing models often rely on heteroskedastic innovations to endowment growth in order to generate cyclical properties of market price of risk. Hansen and Sargent (2010) study a representative agent model with unobservable shocks to trend growth and reconcile countercyclical prices of risk using versions of $\mathbb{T}_{\theta_1,m,y}$ and $\mathbb{R}_{\pi_t,\theta_2}$. In a simplified

[26] For a log-normal variable $\sigma[log(x)] = \frac{\sigma(x)}{E(x)}$. Hansen and Jagganathan (1991) show that this is the bound on the maximum Sharpe ratio.

setting we show the relevance of $\theta_2 \neq \theta_1$ for ensuring such cyclical properties and then construct a example with two models that shows how doubts about π_t are capable of generating countercyclical market price of risk.

Proposition 4 *Consider the economy where $\mathcal{Y} = \{y_l, y_h\}$, agents have identical priors π and consider two models $|\mathcal{M}| = 2$ with symmetric transition matrices : $\{\alpha_m = P_Y(y^* = y|y, m)\}_{m \in \mathcal{M}}$*
Let $MPR^{\theta_1,\theta_2}[y, v, \pi]$ be the market price of risk and $\bar{\alpha} = \sum_m \pi(m)\alpha_m$

- $\theta_1 = \theta_2 = \theta$

$$MPR^{\theta,\theta}[\pi, v, y] = \sqrt{\bar{\alpha}(1-\bar{\alpha})}(2g\gamma + \log\zeta^*(y_l|\pi, v, y) - \log\zeta^*(y_h|\pi, v, y))$$

and

$$\lim_{\|\log(\lambda)\| \to \infty} MPR^{\theta,\theta}[\pi, v, y_l] = \lim_{\|\log(\lambda)\| \to \infty} MPR^{\theta,\theta}[\pi, v, y_h]$$

- *With $\theta_2 \neq \theta_1 = \infty$ and $\alpha_2 > \alpha_1 = \frac{1}{2}$ there exists a $\bar{\pi} < 1$*

$$\lim_{\|\log(\lambda)\| \to \infty} MPR^{\infty,\theta_2}[\pi, v, y_l] > \lim_{\|\log(\lambda)\| \to \infty} MPR^{\infty,\theta_2}[\pi, v, y_h] \quad \forall \pi(m=1) > \bar{\pi}$$

Remark 4 *Since the MPR is symmetric with respect to λ, the properties generally hold for all intermediate wealth shares too*

As $\|\log\lambda\| \to \infty$, the asset pricing properties are isomorphic to an economy populated by a single agent. In this case the departures from the benchmark, when agents trusts the approximating model, are primarily driven by how large are the distortions to the common approximating model. When $\theta_1 = \theta_2$, the approximating model is given by a wighted sum of symmetric Markov transition kernels and retains the symmetry. Further the spread between the valuations across states in the next period are invariant to the current aggregate endowment. Thus the limiting market price of risk is acyclic.

However, with $\theta_2 \neq \theta_1$, admitting an approximating model with a small probability of a very persistent model would generate countercyclical MPR. This arises because of the state dependent twisting of priors over models. For e.g.., in recessions, Agent 1 twists $\pi_t(m)$ towards models with more persistence and vice versa. Start from the situation where the common approximating model has a prior π_t that is closer to the IID model. Relative entropy being a convex measure of distance, the departures of ex-post beliefs are smaller in booms than in recessions. This imparts the countercyclical behavior to market price of risk when agents consider a small possibility of persistence. Figure 1 plots the limiting market price of risk as function of π for $\theta_1, \theta_2 = \infty$ and $\theta_1 = \infty, \theta_2 < \infty$ cases. The left panel shows the region is π where the dotted line which plots the market price of risk in booms is lower than that in recessions.

6.2 Uncertainty about $\pi_{i,t}(m)$ and news shocks

The previous section illustrated how distortions to $\pi_t(m)$ that depend on the current aggregate endowment impart cyclical properties to risk prices. This section will show how distortions to $\pi_t(m)$ that depend on the current distribution of wealth can generate a reason to trade on public signals. In particular agents with identical initial priors and information sets will have

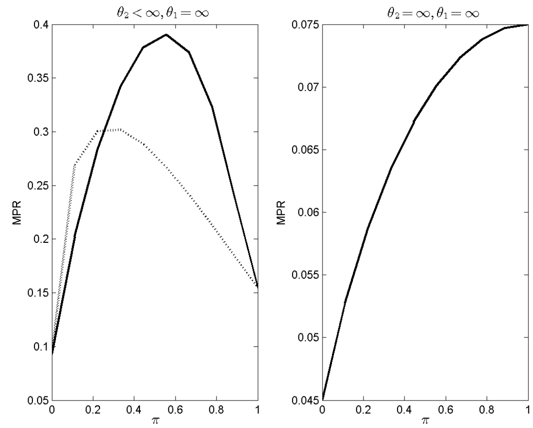

Figure 1: This figure plots the conditional market price of risk as function of π. The solid (dotted) line is $y = y_l(y_h)$. The left panel has $\theta_2 < \infty, \theta_1 = \infty$ and the right panel has both $\theta_1, \theta_2 = \infty$

consumption that fluctuates with "news shocks" (after conditioning for the aggregate endowment). These news shocks will be modeled as public signals that are informative about future distribution of aggregate endowment.

Suppose we augment the economy with news shocks ν_t that provide a noisy signal of the aggregate endowment in the next period.

$$\nu_t = y_{t+1} + \epsilon_t,$$

where ϵ_t is i.i.d over time.

When would these shocks matter for consumption and asset pricing dynamics? With complete markets if agents have identical priors $\pi_i(m) = \pi(m)$ and identical information structure, i.e all agents see $\mathcal{F}_t = \{y^t, \nu^t\}$ it easy to see that in the benchmark case when $\theta_1 = \theta_2 = \infty$ consumption will be exact function of the realized endowment y_t and news shocks are irrelevant. These shocks only affect the information sets and may give agents better forecasting abilities but are symmetric across agents. Thus there is no motive for the planner to vary consumption with these shocks.

However with concerns for misspecification this conclusion is not generally true.

Example 1 *Consider a simple example where the aggregate endowment can take two values and $\mathcal{M}$ or the set of models describing the possible models for the distribution of aggregate endowment has two elements, both describing a symmetric transition matrix with the probability of remaining in the same state is $\alpha_1 \neq \alpha_2$ and $\alpha_1 + \alpha_2 < 1$ and $\theta_1 = \infty$, we have*

$$\lambda_t \neq 1,$$

is necessary and sufficient for

$$c_t(y^t, \nu^t) \neq c_t(y^t, \tilde{\nu}^t) \text{ whenever } \nu^t \neq \tilde{\nu}^t$$

.

Proof.

This follows from the analyzing the FOC of the Planning problem (18b), for all $\nu_{t+1} \neq \tilde{\nu}_{t+1}$

$$\frac{\lambda_t(y_{t+1}, \nu_{t+1})}{\lambda_t(y_{t+1}, \tilde{\nu}_{t+1})} = 1 \qquad \text{iff} \qquad \frac{\sum_m \tilde{\pi}_{i,t}(m)\tilde{P}^i_t(y_{t+1}, \nu_{t+1}|m)}{\sum_m \tilde{\pi}_{i,t}(m)\tilde{P}^i_t(y_{t+1}, \tilde{\nu}_{t+1}|m)} \text{is independent of i}$$

One can simplify the right hand side to

$$\frac{\sum_m \pi_t(m) P_Y(y_{t+1}, \nu_{t+1}|y_t, \nu_t, m) F^i_t(m)}{\sum_m \pi_t(m) P_Y(y_{t+1}, \tilde{\nu}_{t+1}|y_t, \nu_t, m) F^i_t(m)} \tag{23}$$

where $F^1_t(m) = \exp\left\{-\frac{1}{\theta_2}\mathbb{E}_{m,t}Q_{t+1}\right\}$ and similarly for Agent 2, $F^2_t(m) = \exp\left\{-\frac{1}{\theta_2}\mathbb{E}_{m,t}v_{t+1}\right\}$ In the appendix we show how $\lambda_t \neq 1$ and the restrictions on α_i's are sufficient to make (23) vary across agents. ■

Agents trade on news shocks because they have different ex-post assessments of histories with same aggregate endowment but different signals. What is key for this to be true is that a) $Pr_t\{\nu_{t+1}|y_{t+1}\}$ differ across models $m \in \mathcal{M}$ and b) Ex-post weights on models differ across agents.

In the example, the first condition was achieved by the particular signal structure where ν_{t+1} was informative about y_{t+2} and models differed in persistence of aggregate endowment. The ex-post differences in weights on the models came from wealth inequality.

Until agents haven't settled on a single $m \in \mathcal{M}$, learning induces a compound lottery whereby the conditional distribution of the shocks has to *averaged* across models which are themselves distributed according to the some prior $\pi_t(m)$. These preferences implicitly takes a stand on attitudes towards resolution of compound lotteries [27].

In the rest of the section we generalize the structure of the news shocks to see in more detail what patterns of correlation are necessary for consumption to be sensitive to new shocks.

Lemma 2 *Let ν_t be some random process such that $\xi_t = (y_t, \nu_t)$ and the approximating models are modified to have a joint distribution as denoted by $\{P(\xi^*|\xi, m)\}_{m\in\mathcal{M}}$. Consider a partition of $\mathcal{M}$ generated by $F^i_t(m)$, where $M_{i,t}(m) \equiv \{m' \in \mathcal{M} : F^i_t(m) = F^i_t(m')\}$. As long as $\theta_2 \neq \theta_1 = \infty$ and $\lambda_t \neq 1$ a sufficient condition for there to be pair of histories such that $c_{i,t+1}(y^{t+1}, \nu^{t+1}) \neq c_{i,t+1}(y^{t+1}, \nu^t, \tilde{\nu}_{t+1})$ is that $\nexists \quad \kappa^1_t(y_{t+1}, \nu_{t+1}, \tilde{\nu}_{t+1})$ independent of m such that*

$$\begin{aligned} &\sum_{m' \in M_{i,t}(m)} \pi_t(m') P_Y(y_{t+1}|y_t, \nu_t, m') P(\nu_{t+1}|y_{t+1}, y_t, \nu_t, m') = \\ \kappa_t &\sum_{m' \in M_{i,t}(m)} \pi_{i,t}(m) P_Y(y_{t+1}|y_t, \nu_t, m) P(\tilde{\nu_{t+1}}|y_{t+1}, y_t, \nu_t, m) \end{aligned} \tag{24}$$

is satisfied for all m in the support of $\pi_t(m)$

Proof.

Suppose $c_{i,t+1}(y^{t+1}, \nu^{t+1}) = c_{i,t+1}(y^{t+1}, \nu^t, \tilde{\nu}_{t+1})$. Thus

[27] In particular the only kind of preferences that display both indifference to compounding temporal lotteries and ambiguity aversion are Gilboa and Schmeidler's (1989) maxmin expected utility (MEU) preferences. Strzalecki(2011) discusses the connection between these in detail.

$$\frac{\lambda_t(y_{t+1}, \nu_{t+1})}{\lambda_t(y_{t+1}, \tilde{\nu}_{t+1})} = 1$$

Consider first the case when $F_t^1(m) = \iota F_t^2(m)$ for all m in the support of $\pi_t(m)$. This would imply that $Q_{t+1} - v_{t+1}$ is constant across all realizations of (y_{t+1}, ν_{t+1}). The distortions $\tilde{\pi}_t^i$ and $\tilde{P}_t^i(y_{t+1}, \nu_{t+1}|m)$ are same across agents and thus $\lambda_{t+1} = \lambda_t = 1$.

In other cases, the sufficient condition rules of cases where 23 is independent of i.

■

The lemma isolates properties of the joint distribution of the endowment and news shocks that depend on endogenous objects namely continuation values Q_{t+1}, v_{t+1}. Instead of further restricting the class of distribution such that sufficient condition is always satisfied, we give some examples of cases that are ruled out.

Proposition 5 *Without loss of generality, let,*

$$P(\xi^*|\xi, m) = P_Y(y^*|y, \xi, m)P\nu|Y(\nu^*|y^*, \xi, m) \tag{25}$$

Let α_m and $\beta_m(y^*)$ be a vector of parameter that characterize P_Y and $P_{\nu|Y}$ respectively. If the following conditions are true then consumption is only measurable with respect to histories of aggregate shocks in spite of $\theta_2 \neq \theta_1 = \infty$:

1. ***No aggregate risk:*** $|\mathcal{Y}| = 1$
2. ***No correlation:*** $P_{\nu|Y}^m(\nu^*|y^*, \xi) = P_{\nu|Y}^m(\nu^*|\xi)$ *for all m in the support of π_t and* $\mathcal{M} = \{\alpha_m\}_{m\in\mathcal{M}} \times \{\beta_m\}_{m\in\mathcal{M}}$
3. ***Aggregate risk specifications:*** *For all m, m' in the support of π_t, $\alpha_m = \alpha_{m'}$*
4. ***Conditional distributions:*** *For all m, m' in the support of π_t, $\beta_m = \beta_{m'}$*

So far the discussion assumed $\theta_1 = \infty$. This made it easy to derive and study the sufficient conditions for news shocks to matter in terms of the properties of approximating models. However, this restriction is not necessary and with $\theta_1 < \infty$, it is still generally true that news shocks will matter. In the appendix we construct a three period example to highlight how different histories of ν^t's that generate different π_t's require optimal consumption to be sensitive to new shocks. The key insight is that the correlation in the shocks gives a alternative histories ν^t differing *information content* about future distribution of aggregate endowment shocks. The value of this information depends on the difference in expected continuation values (under the model specific ex-post beliefs) across models m as captured by the presence of π in the state variables. With inequality in wealth, this difference in expected continuation values across models can differ across agents, leading to disagreements and trade on histories that only differ in news.

A simple corollary to the analysis is that closing markets where agents can trade on news shocks is suboptimal, a conclusion that would be reversed if one ignores concerns for misspecification. The implications of cross sectional consumption and consequently pricing kernels being sensitive to news shocks can jointly reconcile both higher volatility of prices and larger volume of trade. We leave a quantitative assessment of this channel for future work.

6.3 Long run inequality and market selection

The market selection hypothesis articulated by Friedman (1953) conjectures that agents who make *systematic* errors in evaluation of future risk lose wealth on an average. What is envisioned in this hypothesis is a frictionless economy (i.e complete markets) with agents who have disagreements about how they evaluate future contingencies. It asserts that in such cases, there is a "natural selection" and markets will weed out agents who have incorrect beliefs. Their wealth shares and consequently consumption shares will diminish over time.

In our setting, the differences in initial priors $\{\pi_{i,0}(m)\}_i$ summarize differences in approximating models over time. To categorize the results we will consider two kinds of heterogeneities in $\pi_{i,0}(m)$'s depending on whether agents are "dogmatic" or not. An agent is called dogmatic if his initial prior $\pi_{i,0}$ is degenerate.[28].

The next theorem establishes the long run consequences for Pareto weights when agents trust their approximating model. With bounded aggregate endowment, Inada conditions can be used to map the outcomes on λ_t to consumption shares for each agent.

Proposition 6 *For $\theta_1 = \infty, \theta_2 = \infty$, suppose the data generating process is $P^0_Y(y_{t+1}|y_t) = P_Y(y_{t+1}|y_t, m^*)$,*

- ***Dogmatic initial priors:*** *If $\pi_{1,0}(m^*) = 1$ and $\pi_{2,0}(m^{**}) = 1$ for $m^* \neq m^{**}$*

$$\lambda_t \to 0 \quad P^0_Y - \textit{almost surely}$$

- ***Learning:*** *If $\pi_{1,0}(m^*) > 0$ and $\pi_{2,0}(m^*) > 0$*

$$\lambda_t \to \lambda_0 \frac{\pi_{2,0}(m^*)}{\pi_{1,0}(m^*)} \quad P^0_Y - \textit{almost surely}$$

Furthermore if $\pi_{1,0}(m^*) = 1$

$$\lambda_t \geq \lambda_0 \pi_{2,0}(m^*)$$

Proof.

It should be noted that λ_t is the Lagrange multiplier on the promise keeping constraint and hence non-negative. With $\theta_1, \theta_2 = \infty$ the (18b) simplify to

$$\frac{\lambda_{t+1}}{\lambda_t} = \frac{\sum_m \pi_{2,t}(m) P_Y(y_{t+1}|y_t, m)}{\sum_m \pi_{1,t}(m) P_Y(y_{t+1}|y_t, m)} \tag{26}$$

In the with dogmatic beliefs, λ_{t+1} is a martingale under P^0. Since it is bounded below, Doob's supermartingale theorem implies that it converges almost surely to λ_∞. Suppose for some path $\lambda(y^t) \to \hat{\lambda} > 0$.

$$\frac{P_Y(y_{t+1}|y_t, m^{**})}{P_Y(y_{t+1}|y_t, m^*)} \to 1$$

This is yields a contradiction as $m^* \neq m^{**}$, there would exist y_{t+1} such that given y_t the two models have different likelihood of the shock y_{t+1}.

[28]These were considered in Harrison and Kreps (1978) who studied trading amongst agents with dogmatic priors in presence of short selling constraints.

In the case of Bayesian learning

$$\frac{\lambda_{t+1}}{\lambda_t} = \left(\frac{\sum_m \pi_{2,t}(m) P_Y(y_{t+1}|y_t, m)}{P_Y(y_{t+1}|y_t, m^*)} \right) \left(\frac{P_Y(y_{t+1}|y_t, m^*)}{\sum_m \pi_{1,t}(m) P_Y(y_{t+1}|y_t, m)} \right) \tag{27}$$

Applying Bayes rule,

$$\pi_{i,t+1}(m^*|y_{t+1}) = \frac{\pi_{i,t}(m^*) P_Y(y_{t+1}|y_t, m^*)}{\sum_m \pi_{i,t}(m) P_Y(y_{t+1}|y_t, m)} \tag{28}$$

Thus

$$\sum_m \pi_{i,t}(m) P_Y(y_{t+1}|y_t, m) = \frac{\pi_{i,t}(m^*) P_Y(y_{t+1}|y_t, m^*)}{\pi_{i,t+1}(m^*|y_{t+1})} \tag{29}$$

Combining (27) and (29), we have

$$\lambda_{t+1} \frac{\pi_{2,t+1}(m^*)}{\pi_{1,t+1}(m^*)} = \lambda_t \frac{\pi_{2,t}(m^*)}{\pi_{1,t}(m^*)}$$

One can note that $\pi_{i,t+1}(m^*) \to 1 \quad P^0-$almost surely by Doob's (1949) consistency theorem and $\lambda_{t+1} \to \lambda_0 \frac{\pi_{2,0}(m^*)}{\pi_{1,0}(m^*)}$

If $\pi_{1,0}(m^*) = 1$, then $\pi_{1,t}(m^*) = 1$ for all t. Thus

$$\lambda_{t+1} \geq \lambda_{t+1} \pi_{2,t+1} = \lambda_0 \pi_{2,0}$$

■

The case with dogmatic beliefs and complete trust is studied in Blume and Easley (2001) and we review the main arguments. Consider a static world and an extreme case of belief heterogeneity where some agents incorrectly perceive certain states of the world to be impossible. If the planner had instruments to allocate resources over states, trivially efficiency would require that he allocates zero consumption to these agents in such states. Now consider an infinite horizon world, even if there are no zero probability states of the world in the next period, it is easy to construct events that have zero probabilities over infinite sequences. In our setting, under the benchmark, Agent 2 has complete faith in a Markov kernel that is different from the data generating process m^*, in particular yields a different stationary mean for y_t. Now consider events that predicts a sample mean is equal to the stationary mean under the data generating process $P_Y(y_{t+1}|y_t, m^*)$. By Ergodic theorem, these are zero probability in his mind and as before the planner would allocate (eventually) zero consumption on these events. Unlike the static case, where consumption was zero with some positive probability, here eventually consumption will be zero with probability 1. The previous reasoning isolates absolute continuity of beliefs over infinite sequences as an important criterion for survival[29].

The consequences of survival when agents are Bayesian learners depends on the specific details of the environment. For a setup where the initial priors are absolutely continuous with respect to the Lebesgue measure, the key heterogeneity in initial priors that matters for long run survival is dimension of the support. Under some regularity conditions, Blume and Easley (2006) show that agents with low dimensional priors dominate (provided their priors include the

[29] More generally when $\pi_{i,0}$ have disjoint support(across agents), with Bayesian learning agents' priors converge to model that has the least relative entropy with respect to m^*. Thus studying the dogmatic initial priors captures this case too.

data generating process). Although our setup does not fit the analysis (since the set of models is finite), one can completely characterize how worse off the agents are in the long run as function of their initial differences.

When beliefs are not dogmatic, Proposition 6 says that the limiting relative Pareto weight of Agent 2 is path independent and converges to $\frac{\pi_{2,0}(m^*)}{\pi_{1,0}(m^*)}$ fraction of its initial value. The data generating process is the distribution associated with m^* and $\frac{\pi_{2,0}(m^*)}{\pi_{1,0}(m^*)}$ is measure of the relative heterogeneity in initial priors. A number larger than one implies that Agent 2 puts more mass on the correct model than Agent 1 and consequently has a higher Pareto weight in the limit. The second part of the proposition characterizes a uniform (i.e for all paths) lower bound on the relative Pareto weights of Agent 2. When Agent 1 is both dogmatic and correct but Agent 2 admits the possibility of the truth in his support. Figure 2 shows the simulations of λ_t from a simple example where the aggregate endowment takes two values and $\mathcal{M}$ has two elements, both Markov models with symmetric transition densities parametrized by a α_m that is the probability if remaining in the same state. The data is generated from $m = 1$ and the initial conditions are $(\pi_{1,0}(1), \pi_{2,0}(1)) = (1, \frac{1}{2})$ and $\lambda_0 = 1$.

Consistent with the proposition we see the two properties : a) all paths converge to $\lambda_0\pi_{2,0} = \frac{1}{2}$ and b) The distribution of λ_t at any date t has a support with a lower bound of $\frac{1}{2}$. The red line depicts the mean across paths and the bold black line marks the initial condition. One can see that even though the agent who is learning is worse off in the long run, in the initial periods $Pr\{\lambda_t > \lambda_0\}$ is quite high.

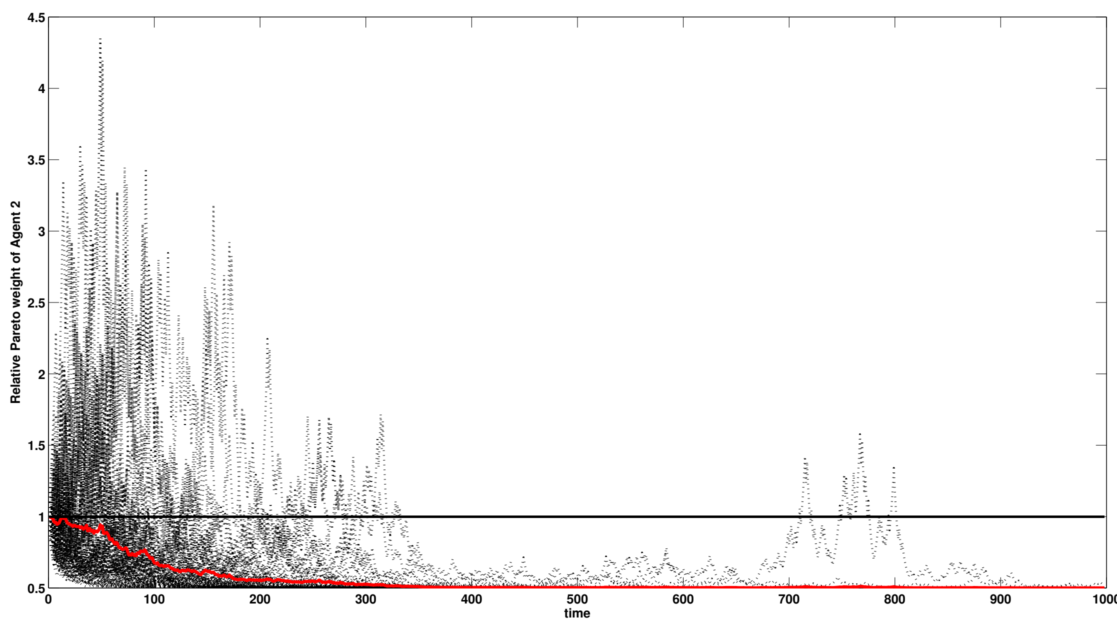

Figure 2: The figure plots paths of λ_t for the case when Agent 1 is both dogmatic and correct while Agent 2 is a Bayesian learner. The red line is the mean across samples and the solid black marks the initial λ_0

The next two sections will describe how the results change when agents admit misspecification concerns. As with the benchmark, we split the discussion in two cases. In section 6.4 we study the situation when agents are dogmatic (in the Harrison Kreps (1978) sense). In section 6.5 we study the case when agents are learning and the differences in their approximating models are transitory.

6.4 Dogmatic beliefs with specification concerns

When agents have different approximating models which they fear are misspecified, what matters for long run inequality is is how far their ex-post beliefs from a data generating process. The proposition below shows how specification concerns expands the set of heterogeneous approximating models that are consistent with none of the agents being driven out.

Let P_Y^i be the $m \in \mathcal{M}$ such that $\pi_i(m) = 1$. Define $\mathbb{I}^{0,i}(y) = \mathbb{E}^0 \log \left(\frac{P_Y^0(y'|y)}{P_Y^i(y'|y)} \right)$ be the conditional relative entropy of Agent i's approximating model with respect to the data generating process.

Proposition 7 *Suppose $u(0) > -\infty$ and y takes more than two values and let Γ_y be the invariant distribution of the Markov chain Y under the data generating process. There exists $\underline{M}$, $\bar{M}$ such that*

$$\mathbb{E}^{\Gamma_y}\mathbb{I}^{0,2}(y) - \mathbb{E}^{\Gamma_y}\mathbb{I}^{0,1}(y) \in (\underline{M}, \bar{M})$$

is sufficient for both agents to survive P^0-almost surely.
For $P_Y^0 = P_Y^1$ we have $\underline{M} = 0, \bar{M} > 0$ and with $P_Y^0 = P_Y^2$ we have $\underline{M} < 0, \bar{M} = 0$.

The survival result exploits two key features : a) presence of non-diversifiable aggregate risk and b) the fact that valuations are bounded from below.

While exploring misspecifications agents are trading off changes in relative entropy and changes in expected future valuations

$$\min_{m:\mathbb{E}^i m=1} \left(\mathbb{E}^i m V^i - \mathbb{E}^i V^i \right) + \theta \mathbb{E} m \log m$$

The search for misspecifications balances the marginal costs of exploring new models and benefits from recognizing potential losses in average valuations. Scaling the valuations by a factor $k > 1$, will mean that for the same marginal costs, the benefits are larger. Thus agents with proportionately larger valuations will explore models in a larger set as measured by relative entropy.

The valuations of the agents roughly scale in individual wealth. Take a hypothetical path such that Agent 2 (who starts with a approximating model that is slightly incorrect) was arbitrarily close to zero wealth. Not only his valuations are low but bounded utility implies that the spread in valuations is also close to zero. As such exploring misspecifications is a rather costly activity for him since the relative changes in expected valuations under alternative models is also close to zero. On the other hand, Agent 1 faces large fluctuations in valuations coming from owning the entire aggregate endowment stream. What matters for the market selection force is not how far the approximating models are, but the gap in terms of worst case models. In cases where Agent 1 might have started with a correct reference point i.e., $P^0 = P^1$, if his doubts switch the relative ranking in terms of distance from the data generating process. Of course for this to work, differences in the initial reference points should be low enough. The interval in proposition 7 gives us a sense of how large these differences can be. The length of the interval shrinks as θ

increase as its more costly for Agent 1 to explore misspecifications for the same set of models [30].

One can interpret this wealth transfer mechanisms by analyzing the differences in perceived insurance needs are across agents. The wealthier agent being more pessimistic "overbuys" recession insurance. This portfolio generates inflows for him in low endowment states against payments booms. Since recessions occur less often than what he perceives, there is on an average a transfer of wealth to the other agent.

With $u(0) = -\infty$, the survival is more delicate. It depends on $(\lambda_0, \pi^1, \pi^2)$. We can construct examples where after distortions, the agent with low wealth is relatively more closer to the data generating process and there is a point in the wealth distribution where the worst case models coincide for both agents.

Characterizing long run distributions: Binary - IID shocks

In a special setting where the aggregate endowment can take two values we can characterize the ergodic distributions of Pareto weights. These results extend Anderson's (2005) findings where he studies a similar risk sharing problem with common approximating models. That paper isolates points in wealth distribution that are absorbing in nature. In our problem with symmetric information and heterogeneous approximating models such "steady states" generally do not exist. However there is a small class of shock process, particularly where aggregate shocks are IID and can can take at most two values under the approximating models of each agent when the ergodic wealth distribution is degenerate. For the rest of the section $\{P_Y^i\}_i$ and the data generating process P_Y^0 satisfy these restriction.

Proposition 8 *Let* $C = \frac{[1-\delta]}{\theta_1}[u(y_h) - u(y_l)]$. *If*

$$\log \frac{P_Y^2(y_h)}{P_Y^1(y_h)} - \log \frac{P_Y^2(y_l)}{P_Y^1(y_l)} \in (-C, C),$$

there exists a constant $\infty > \bar{\lambda} > 0$ *such that*

$$\lambda_T = \bar{\lambda} \implies \lambda_{T+s} = \bar{\lambda} \quad \forall s > 0$$

This $\bar{\lambda}$ is independent of the y and the data generating process. At $\bar{\lambda}$, the agents have the same ex-post beliefs. This requires the heterogeneity in minimizing likelihood ratios (which depend of wealth shares) to offset the heterogeneity in approximating models) for all realizations of aggregate endowments. Given they are "equally" wrong, there are no efficiency gains from changing consumption shares in the future. The IID-binary restrictions to shock process are necessary, otherwise the set of equilibrium conditions consistent with steady states forms a non-square system of unknowns. It is easy to construct examples such that there are generically no roots to such systems.

The steady state is interesting if the economy reaches it from arbitrary initial conditions. The next proposition discusses some mild technical conditions under which $\bar{\lambda}$ is both locally and globally stable.

[30] What is crucial is for the mechanism to work is agents with relatively larger wealth to distort more. It is not necessary that they are pessimistic. One can flip the sign of θ_1, θ_2 to be negative and switch the minimization in operator $\mathbb{T}$'s and $\mathbb{R}$'s to a maximization to obtain a theory of "cautiously" optimistic agents. These agents will have ex-post beliefs that are statistically close to the approximating model but overweight the states with higher valuations.

Proposition 9 *Suppose that* $P^0_Y = P^1_Y \neq P^2_Y$ *and suppose there exists a steady state* $\bar{\lambda}$*. Denote* $\bar{z}_1(y) = \frac{\exp\{-\theta_1^{-1}u(y)\}}{\mathbb{E}^0 \exp\{-\theta_1^{-1}u(y)\}}$*. If*

$$\mathbb{I}^{0,2} \leq -\mathbb{E}^0 \log \bar{z}_1(y) \tag{30a}$$

$$P^2_Y(y_l) \leq P^0_Y(y_l)\bar{z}_1(y_l) \tag{30b}$$

and

$$\tilde{P}^1(y_l|v, y) \geq P^0_Y(y_l) \quad \forall v : \lambda(v, y) \leq \bar{\lambda} \tag{30c}$$

For any $\bar{\lambda} > \lambda_t > 0$ *we have*

$$\lim_s \lambda_{t+s} = \bar{\lambda} \quad P^0 \textit{almost surely}$$

Corollary 1 *Suppose that* $P^0_Y = P^1_Y \neq P^2_Y$ *and suppose there exists a steady state* $\bar{\lambda}$*. Then it is locally stable i.e, there exists a* $\delta > 0$ *such that for all* $\lambda_t \geq \bar{\lambda} - \delta$

$$\lim_s \lambda_{t+s} = \bar{\lambda} \quad P^0_Y \textit{almost surely}$$

Corollary 2 *Assumption* (30c) *is satisfied if* $\frac{P^2_Y(y_l)}{P^2_Y(y_h)} \geq \frac{P^0_Y(y_l)}{P^0_Y(y_h)}$

The restriction to Agent 1 having the correct approximating model is without loss of generality. The set of sufficient conditions will change accordingly, but the structure of proof remains same. We omit the details for brevity. Note that assumption (30c) restricts an endogenous policy rule. Corollary 1 shows that it is satisfied locally around $\bar{\lambda}$ and corollary 2 provides restriction on Agent 2's beliefs that guarantees it holds in the required region $\lambda < \bar{\lambda}$.

The proof starts with the observation that λ_t is a martingale under the Agent 1's ex-post beliefs. The sufficient conditions allow us to argue that it is a sub martingale under the data generating process on the left of the steady state. A helpful lemma used in the proof shows that that $\lambda_{t+1}[y_{t+1}|y, \lambda_t]$ is monotonic in λ and the steady state $\bar{\lambda}$ by definition is a rest point for the its dynamics. The monotonicity property allow us use the steady state as a bound on this sub-martingale. This leaves us two possibilities either it converges to zero or to the steady state. Imposing Assumption (30a) rules out the first case and λ_t converges to the steady state value. The intuition behind this convergence is same as discussed previously, along the paths when Agent 1 is dominating his distortions to the approximating model are large, in particular the ex-post beliefs over estimate (relative to his approximating model) states with low aggregate endowment. Figure 3 depicts the log $\frac{\lambda_{t+1}}{\lambda_t}$ as a function of promised values. The two curves represent growth rates for $y_{t+1} = y_l$ or $y_{t+1} = y_h$. We see that for low λ, the growth rate is positive when aggregate endowment is high (dotted line). This corresponds to the excess returns Agent 2 generates from providing recession insurance to Agent 1. The crossing points are represent the steady states where the growth rates are equalized.

Sample Paths

In this section, we describe sample paths for IID-Binary economy with the following parameters ($\gamma = .5, \theta_1 = 1, \frac{y_h}{y_l} = 1.2, \delta = .9$). The approximating models are $P^1_Y(y_h) = .6, P^2_Y(y_h) = .4$. The plots are constructed to show how convergence rates and patterns differ across different choices

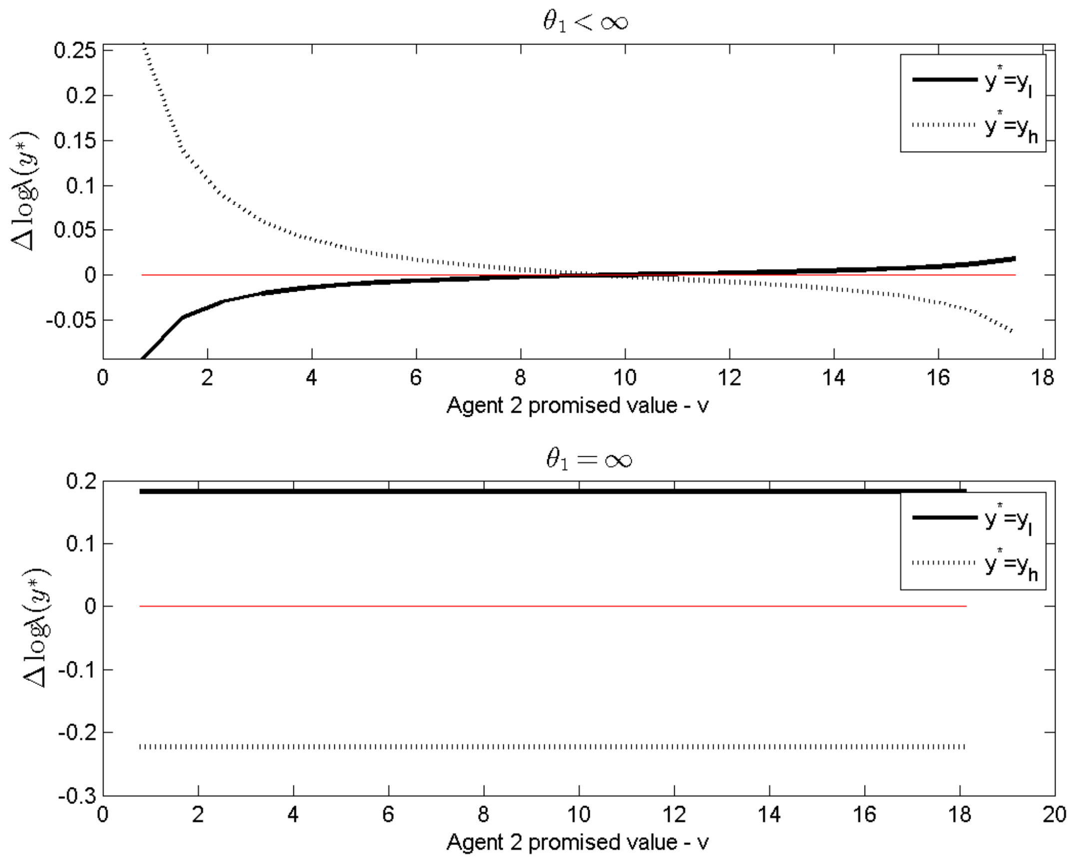

Figure 3: The figure plots the log change in the Lagrange multiplier as a function of the continuation value. The solid (dotted) line refers to$y(y^*) = y_l(y_h)$. The left (right) panel has $\theta < \infty(\theta = \infty)$

for the data generating process. We will discuss two cases a) when P_Y^0 coincides with with the agents approximating models and b) when P_Y^0 coincides with either agents ex-post beliefs.

Figures (4) and (5) plot a typical sample path for consumption share of Agent 1 for shocks drawn from alternative measures and initial conditions. The top panel in figure (4) corresponds to the benchmark when agents trust their approximating models. Depending on whose approximating model we use to draw the shocks, the consumption share of Agent 1 either goes to zero or 1. The bottom panel plots does the same exercise when agents doubt their approximating models. Consistent with proposition 8, Agent 1's consumption share approaches interior $\bar{\lambda}$ (about 0.7) irrespective of the data generating process.

Lastly, figure 5 samples from the ex-post beliefs of the agents. Depending on the initial conditions, we see a very slow convergence to $\bar{\lambda}$.

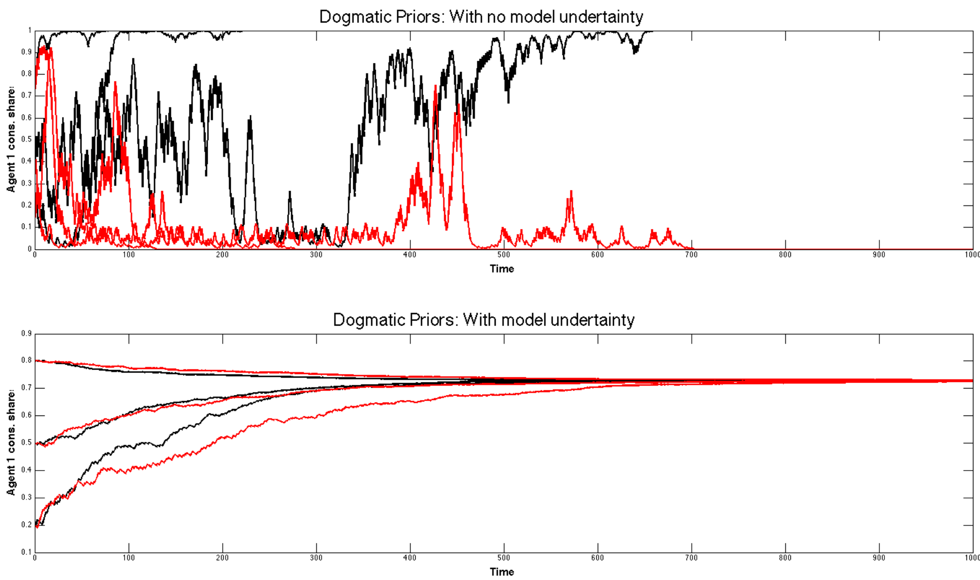

Figure 4: This figure plots a simulated path for the consumption share of Agent 1η_1 starting from 3 different initial conditions. The black (red) line plots outcome under Agent 1 (Agent 2's) reference beliefs

Let $\mathcal{G}_{t,t+1} = \log \frac{\lambda_{t+1}}{\lambda_t}$ be the net growth rate of relative Pareto weights. The outcomes in terms of the convergence rate can depend on the underlying data generating process. Suppose P^0 is either agents' approximating model (or a convex combination of them). At the stationary point the relative entropies of both agent are strictly positive but equal hence $\mathcal{G}_{t,t+1}$ is zero. However, if the data generating process is either agents' ex-post beliefs, at the stationary point both the relative entropies are equal to 0. Since relative entropy is a convex measure of distance, a small perturbation of the system will cause a much larger change in the magnitude of $\mathcal{G}_{t,t+1}$ under the former case than the later. This means that the speed of adjustment will be very slow under the Agent's respective ex-post beliefs.

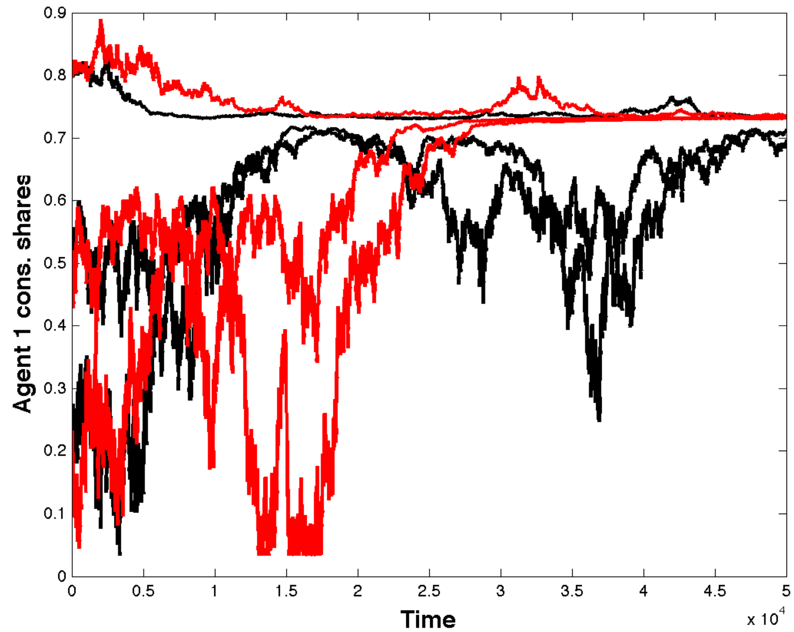

Figure 5: This figure plots a simulated path for the consumption share of Agent $1\eta_1$ starting from 3 different initial conditions. The black (red) line plots outcome under Agent 1 (Agent 2's) worst case beliefs

6.5 Long run inequality under learning

The long run outcomes for Pareto weights when agents are not dogmatic but update their approximating overtime depend on whether they admit specification uncertainty ($\theta_1 < \infty$) or not ($\theta_1 = \infty$). With time the differences in initial priors vanish and the economy converges to situation where both agents have identical posteriors and hence a common approximating model.

Let m^* be the true DGP and π_t^i be Agent i's posterior on $\mathcal{M}$ given history y^t.

Definition 10 *Weak agreement : Agents are said to agree weakly if* $\pi_t^1 = \pi_t^2$. *Let* $\mathcal{T}^w = \inf\{t : \|\pi_{1,t} - \pi_{2,t}\| = 0\}$ *be the time taken to agree weakly*

Definition 11 *Strong agreement : Agents are said to agree strongly if there exist a* m^* *such that* $\pi_{1,t}(m^*) = \pi_{2,t}(m^*) = 1$. *Similarly, let* $\mathcal{T}^s = \inf\{t : \pi_{1,t}(m^*) = \pi_{2,t}(m^*) = 1 \quad \text{for some} m^* \in \mathcal{M}\}$ *to reach this agreement*

The two notions captures the fact that agents starting from possibly heterogeneous initial priors but same information may have similar forecasts even before they actually learn the parameter. Also it is easy to see that weak agreement is an absorbing state. Strong agreement occurs only after all uncertainty is resolved.

When $\theta_1 < \infty$ or agents have enduring specification doubts, the limiting relative Pareto weight of Agent 2 λ_t also depends on the curvature of u. For an iso-elastic utility function $u(c) = \frac{c^{1-\gamma}}{1-\gamma}$ we can extend Anderson's (2005) analysis to show that the limiting properties of

Pareto weight depend on IES. When IES > 1 they converges to 1. This corresponds to a steady state for a setting where both agents are eventually symmetric except for possibly their wealth shares. The intuition for this result comes from the observation that λ_t is a **martingale** under Agent 1's ex-post beliefs

$$\tilde{E}_t^1 \lambda_{t+1} = \lambda_t$$

"Undoing" Agent 1's distortions we get,

$$\mathbb{E}_t^1 \lambda_{t+1} = \lambda_t - Cov_t^1 \left[\lambda_{t+1}, z_{t,t+1}^1\right]$$

or

$$\mathbb{E}_t^1 \lambda_{t+1} = \lambda_t - Cov_t^1 \left[\frac{\tilde{P}_t^2(y_{t+1}|y_t)}{\tilde{P}_t^1(y_{t+1}|y_t)}\lambda_t, \frac{\tilde{P}_t^1(y_{t+1}|y_t)}{P_t^1(y_{t+1}|y_t)}\right]$$

With learning $\pi_{1,0}(m^*) = \pi_{2,0}(m^*) = 1$ for $t > \mathcal{T}^s$

$$\mathbb{E}_t \lambda_{t+1} = \lambda_t - Cov_t \left[\frac{\tilde{P}_t^2(y_{t+1}|y_t)}{\tilde{P}_t^1(y_{t+1}|y_t)}\lambda_t, \frac{\tilde{P}_t^1(y_{t+1}|y_t)}{P_t^1(y_{t+1}|y_t)}\right]$$

Note the two components

1. $\frac{\tilde{P}_t^1(y_{t+1}|y_t)}{P_t^1(y_{t+1}|y_t)}$: Agent 1's pessimism which is countercyclical

2. $\frac{\tilde{P}_t^2(y_{t+1}|y_t)}{\tilde{P}_t^1(y_{t+1}|y_t)}$: Agent 2's *relative* pessimism. This depends on IES and whether Agent 2 is rich or poor.

WLOG suppose $\lambda_t > 1$ or Agent 2 is rich. When IES> 1, utility drops for a given percentage drop in consumption are larger for richer agents. Thus Agent 2 *relative* pessimism or $\frac{\tilde{P}_t^2(y_{t+1}|y_t)}{\tilde{P}_t^1(y_{t+1}|y_t)}$ is countercyclical. This makes the **covariance positive** and imparts a **negative drift** to λ_t. With IES < 1 The covariance flips sign and Agent 2's relative Pareto weights increase on an average. With IES $= 1$, We have Epstein - Zin preferences which are homothetic and the covariance is zero

The next proposition 10 summarizes the observations.

Proposition 10 *Suppose $\pi_{0,i}(m^*) > 0$ for $i = 1, 2$. Let the data generating process $P_t^0(y^t) = P(y_t, m^*)$ where $m^* \in \mathcal{M}$ is i.i.d. The long run Pareto weight depends on IES*

- $\lim_t \lambda_t = 1$ P^0 *almost surely when IES > 1*
- $\lim_t \lambda_t = \lambda_\infty \in (0, \infty)$ *when IES $= 1$*
- $P^0\{\lim_t \lambda_t = 0, \lim_t \lambda_t = \infty\} > 0$ *when IES < 1*

The case when $\theta_1 = \infty$, the concerns for uncertainty vanish as posteriors become degenerate. The transient wealth dynamics persist until either agents strongly agree $\mathcal{T}^s$ or at $T \geq \mathcal{T}^w$ such that $\lambda_T = 1$. This corresponds to a) either both agents having degenerate priors with possibly unequal wealth shares or non-degenerate but equal priors and equal wealth shares.

There is "race" between speed of learning and the survival mechanism associated with heterogeneity in ex-post beliefs. If learning is "slow" the agent with smaller weight on the data generating process has time to catch up. For instance, figure 6 plots sample paths of Agent 1's consumption share after agents agreed weakly on $\pi = \frac{1}{2}$. These transient dynamics come from concerns for model ambiguity and we see an average tendency towards a more equal wealth distribution.

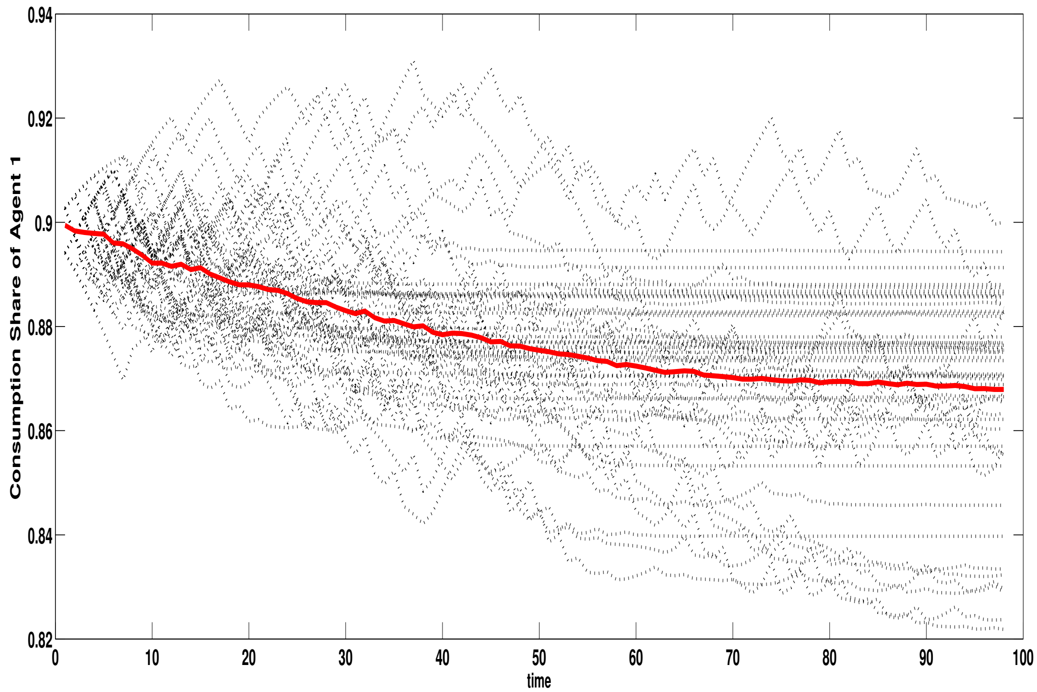

Figure 6: This plots simulated paths of Agent 1's consumption shares after agents weakly agreed on a prior of 0.5.

7 Asymmetric Information

The analysis so far studied the implications of how doubts about asymmetric approximating models under symmetric information arrested the fanning out of consumption shares that is typical in the benchmark without doubts. This section will show how these forces extend to settings with asymmetric information, here coming from the fact that the realizations of the taste shocks are only known to Agent 2. In presence of doubts, the planner will trades-off fluctuating Pareto weights for providing optimal incentives (due to private information) and for heterogeneity in ex-post beliefs (induced by misspecification concerns).

To exclusively focus on the role of heterogeneous information, we abstract from both learning and exogenous heterogeneity in beliefs coming from different approximating models. This amounts to $\pi_i(m^*) = 1$ for all i. Agent 2 now draws privately observed taste shocks from $\mathcal{S} = \{s_l, s_h\}$ with $s_l < s_h$. We first study the case when all individuals of type 2 get the same realization of taste shock and in an extension show that the results carry over to a case when a continuum of individuals of type 2 draw i.i.d taste shocks.

7.1 Planner's problem (II): Recursive version

The recursive planner's problem posed in (13) simplifies and is repeated below.

$$Q(v,y) = \max_{u_1(s),u_2(s),\bar{v}(s,y^*)} \mathbb{T}_{\theta_1}\left[(1-\delta)u_1(s) + \delta \mathbb{T}_{\theta_1,y} Q(\bar{v}(s,y^*),y^*)\right] \tag{31}$$

subject to

$$\mathbb{T}_{\theta_1}\left[(1-\delta)su_2(s) + \delta \mathbb{T}_{\theta_1,y}\bar{v}(s,y^*)\right] \geq v \tag{32a}$$

$$(1-\delta)su_2(s) + \mathbb{T}_{\theta_1,y}\delta\bar{v}(s,y^*) \geq (1-\delta)su_2(s') + \mathbb{T}_{\theta_1,y}\delta\bar{v}(s',y^*) \quad \forall s,s' \tag{32b}$$

$$C(u_1(s)) + C(u_2(s)) \leq y \quad \forall s \tag{32c}$$

$$\bar{v}(s,y^*) \leq v^{max}(y^*) \tag{32d}$$

Simple Contracts

Before discussing the properties of the optimal contract, we will highlight a special class of incentive feasible contracts - *repeated static contracts*[31]. These dynamic contracts are constructed by a sequence of static contracts that deliver a constant share of aggregate endowment to both the agents.

Let $V[\alpha_2, y]$ be the value to agent 2 from a consumption plan that always gives him α_2 fraction of the aggregate endowment. It solves the following recursion.

$$V[\alpha_2, y] = \mathbb{T}_{\theta_1,y}(1-\delta)su(\alpha_2 y) + \delta V[\alpha_2, y^*] \tag{33}$$

Given (v^0, y), we can can solve for $\alpha_2(v^0, y)$ such that $V[\alpha_2(v^0,y), y] = v^0$. The simple contracts are then given by,

1. $u_1^{ss}(s|v^0,y) = u([1-\alpha_2(v^0,y)]y)$,
2. $u_2^{ss}(s|v^0,y) = u(\alpha_2(v^0,y)y)$,
3. $\bar{v}^{ss}(s,y^*|v^0,y) = V[\alpha_2(v^0,y),y^*]$.

Given this allocation, one can obtain the value to the planner by

$$Q^{ss}[v^0,y] = \mathbb{T}_{\theta_1,y}u_1^{ss}(s|v^0,y^*) \tag{34}$$

It is easy to see that these simple contracts are incentive-feasible. The planner essentially provides no insurance against the taste shocks. These contracts are 'absorbing' in nature as the ex-ante promised values are constant and further suboptimal except at $\{(v_{min}, y), (v^{\max}(y), y)\}$. The value of these contracts gives a lower bound to the value function Q^0. Once can construct an upper bound by using the benchmark with no informational frictions. Denote $Q^{cm}(v^0, y)$ as the solution to the planner's problem without imposing constraint 32b.

Thus the optimal value satisfies

$$Q^{ss}[v^0,y] \leq Q^0[v^0,y] \leq Q^{cm}(v^0,y)$$

[31]We call them "static" because they would be optimal if $\delta = 0$.

Optimal Contract

The next proposition summarizes some the key properties of the optimal contract.

Proposition 11 *For all* $v^0 \in (v_{min}, v^{max}(y))$*, and* $s_l < s_h$*, the optimal contract features*

- *Current utility :* $u_2(s_l|v^0, y) < u_2^{ss}(v^0, y) < u_2(s_h|v^0, y)$
- *Future utilities:* $\mathbb{T}_y \bar{v}(s_l, y^*) > \mathbb{T}_y \bar{v}^{ss}(y^*) > \mathbb{T}_y \bar{v}(s_h, y^*)$

The planner provides insurance for periods of high taste shocks by giving consumption higher than the static contracts, but the agent pays for it through lower expected continuation values in the future. This additional lever whereby the (ex-ante) continuation values can be varied relaxes incentives and allows the planner to increase the spread in current utilities. Like before the promised values can be mapped back to effective Pareto weights λ and we can interpret the variation in promised values as variation in Pareto weights. Thus like heterogeneity in beliefs, heterogeneity in information also imparts volatility to Pareto weights.

Inverse Euler Equation

The trade offs that the planner faces while perturbing continuation values are succinctly captured in what has been referred to in literature as an inverse Euler equation. This is a necessary condition for optimality in many environments with private information[32]. In our setting, this equation holds under the worst case beliefs of Agent 1 rather than the common approximating model P^0

Lemma 3 *Let* λ *be the multiplier on the promise keeping constraint* $z_1(s, y^*) = \frac{\tilde{P}^1(s)\tilde{P}^1(y^*|v,y,s)}{P_S(s)P_Y(y^*|y)}$. *The optimal allocation implies,*

$$\tilde{\mathbb{E}}_y^1 \lambda(s, y^*) = \lambda$$

and

$$\mathbb{E}_y \lambda(s, y^*) = \lambda - Cov[\lambda(s, y^*), z_1(s, y^*)|v, y]$$

Proof. The FOC with respect to $\bar{v}(s, y^*)$ gives us

$$\frac{\lambda(s, y^*)}{\lambda} = \left[\frac{\tilde{P}^2(s)\tilde{P}^2(y^*|s)}{\tilde{P}^1(s)\tilde{P}^1(y^*|s)}\right]\left[1 + \mu(s) - \mu(s')\frac{\tilde{P}^2(s')\tilde{P}^2(y^*|y, s')}{\tilde{P}^2(s)\tilde{P}^2(y^*|s)}\right] \tag{35}$$

Multiplying by $\tilde{P}^1(s)\tilde{P}^1(y^*|s)$ and summing over all s, y^* gives us

$$\tilde{\mathbb{E}}_y^1 \lambda(s, y^*) = \lambda$$

The second equality follows from applying expanding $\tilde{\mathbb{E}}_y^1 \lambda(s, y^*) = \mathbb{E}_y z_1(s, y^*)\lambda(s, y^*)$ and the fact that $\mathbb{E}_y z_1(s, y^*) = 1$ ■

The envelope theorem implies that $-Q_v^0(v^0, y) = \lambda$. Thus the multiplier λ captures the marginal cost to the planner of providing an extra unit of promised value to Agent 2. Along any

[32]See Rogerson (1985); Golosov, Kocherlakota, and Tsyvinski (2003) for detailed exposition of the inverse Euler equation and its implication to capital taxation.

history (y^t, s^{t-1}), consider a perturbation (from the optimal contract) that increases continuation values to Agent 2 for all possible reports of taste shocks s_t equally and reduces the ex-ante promised value v_t. This deviation is incentive compatible but would cost the planner $\tilde{\mathbb{E}}_t^1 \lambda_{t+1} - \lambda_t$. These are measured under Agent 1's beliefs as given the resource constraints, the additional future utilities come at lower consumption for Agent 1. The inverse Euler equation captures optimality of the contract by making such perturbations costless. This implies that λ_t is a martingale under the worst case model of Agent 1 but not under P^0. Undoing the distortions one gets the expected law of motion for the Pareto weights under P^0. The proposition suggest that he sign of the covariance between $\lambda(s, y^*)$ and $z_1(s, y^*)$ is important to generate a drift in the dynamics of the Pareto weights over time. In particular a positive covariance for low values of λ can push Agent 2 away from immiseration

Heterogeneous beliefs and Incentives

Asymptotically the agent with taste shocks goes to immiseration if $\theta_1 = \infty$. We already noted that the inverse Euler equation implies that λ_t is a martingale. Since it is also the multiplier of the promise keeping constraint, the KKT necessary conditions require it to be non negative. Proposition 12 shows that the multiplier converges to any positive number, the incentive constraints are violated. However with $\theta_1 < \infty$, we have the both forces - heterogeneity in beliefs and incentive considerations that affect how λ_t evolves over time. Along the paths λ_t approaches zero a) the optimal contract approaches the first best and this dampens the fluctuations due to optimal incentives b) On the other hand, the disagreements across agents about future aggregate endowment, measured with respect to their ex-post beliefs is maximum. This generates a drift that pushes λ_t in the interior.

Proposition 12 *Let λ_t be the Lagrange multiplier on the promise keeping constraint (32a) for problem (31). We have*

- *With $\theta_1 = \infty$, $\lim_t \lambda_t = 0 \qquad P^0 -$ almost surely*
- *With $\theta_1 < \infty$, $\lim_t \lambda_t \neq 0 \qquad P^0 -$ almost surely*

Figure 7 plots the expected growth rate of λ_t for two economies with the shocks are both mutually and over time independent. The black line is the case with $\theta_1 < \infty$ and the red line is the benchmark with no concerns for misspecification. We see that as we go to the left of $\lambda = 1$, there is a strictly positive drift to λ when we activate concerns for misspecification.

Inspecting the forces

One can think of this economy as a combination of two arrangements a) a competitive insurance market with securities contingent on realization of aggregate endowment and b) a bilateral credit market for trading long term loan contracts between these Agent 1 and Agent 2

The multiplier λ is monotonically related to both the promised value to Agent 2 and his average consumption share. This allows us to think of it as a proxy for Agent 2's relative wealth. The next lemma shows that without concerns for misspecification, the Pareto weights are only measurable to the histories of taste shocks.

Lemma 4 *With $\theta = \infty$ we have $\lambda(s, y^*) = \lambda(s, y^{**})$ for all $y^*, y^{**} \in \mathcal{Y}$*

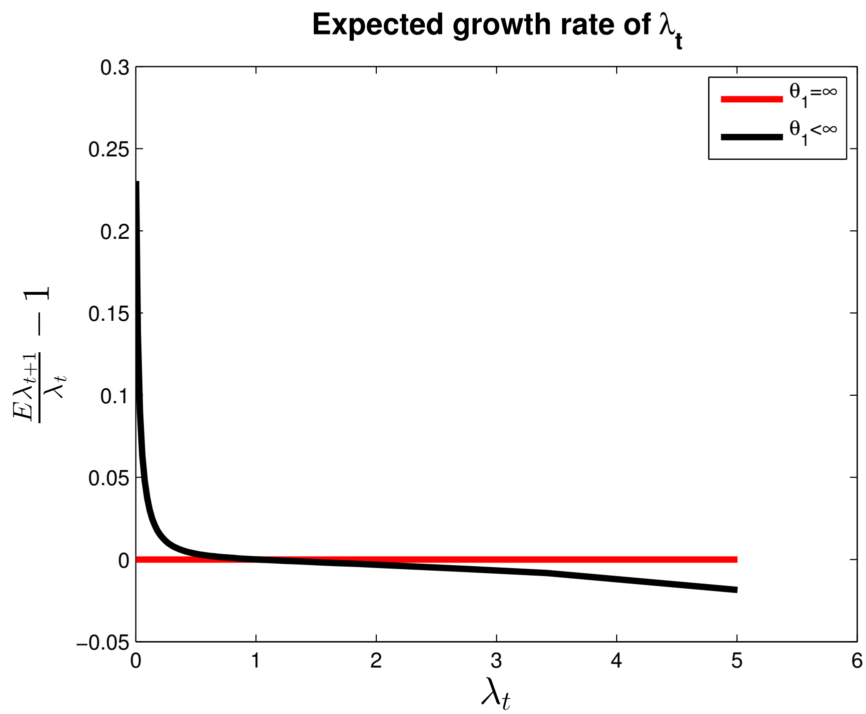

Figure 7: The figure plots the expected one period ahead growth rate of λ_{t+1} under the common approximating model.

Proof. The FOC's with respect to $\bar{v}(s, y^*)$,

$$\frac{\lambda(s, y^*)}{\lambda} = \left[\frac{\tilde{P}^2(s)\tilde{P}^2(y^*|s)}{\tilde{P}^1(s)\tilde{P}^1(y^*|s)}\right]\left[1 + \mu(s) - \mu(s')\frac{\tilde{P}^2(s')\tilde{P}^2(y^*|y, s')}{\tilde{P}^2(s)\tilde{P}^2(y^*|s)}\right] \tag{36}$$

With $\theta_1 = \infty$, $\tilde{P}^i(s) = P_S(s)$ and $\tilde{P}^i(y^*|s) = P_Y(y^*|y, m^*)$. Since the wedges only depend on the reported shock, dividing the expressions for $\lambda(s, y^*)$ and $\lambda(s, y^{**})$ from equations (36) eliminates the terms with $\mu(s)$ and

$$\lambda(s, y^*) = \lambda(y^{**}, s)$$

■

When agents have identical assessments of likelihoods of aggregate shocks, no arbitrage eliminates any "expected" excess returns from trades in the competitive insurance markets.

The bilateral credit market look like "annuities" : Agent 2 borrows from Agent 1 when he has a high taste shock and repays by a forever lower consumption in the future. To see this, consider a simple economy where the informational frictions last only for one period. In time $t = 0$, Agent 2 has a taste shock s_0 which can take a high or low value and then remains constant at $s_t = 1$ for $t \geq 1$. The optimal consumption path for some initial λ^0 (associated with v^0) will feature $\alpha_{2,0}(y, s_h) > \alpha(\lambda_0)$ and $\alpha_{2,t} < \alpha(\lambda_0)$ for $t \geq 1$ where the function $\alpha(\lambda_0)$ is the constant consumption share in an economy without taste shocks. The consumption pattern can be implemented by a long term loan that funds the higher consumption share in $t = 0$ when Agent 2 desires high consumption and is repaid by a permanently lower share in the future. In our problem where these taste shocks are recurrent, they have an permanent effect and eventually drive the agent to immiseration.

When $\theta_1 < \infty$ as λ approaches zero, $\frac{\lambda(y_h,s_h)}{\lambda_0} \approx \frac{P_Y(y_h)}{\tilde{P}^1(y_h)} > 1$. Thus along the paths that threaten to drive Agent 2 towards immiseration, the relative disagreements on aggregate shocks become large. Agent 1 is willing to pay a lot for having consumption is states with low aggregate output. These insurance contracts earn a positive return when $y_t \neq y_l$. The higher returns compensate for possible repayments due to adverse taste shocks and push the economy towards an interior wealth distribution. Figure 8 plots a path for λ_t across two economies that differ in θ but have the same history of the fundamental shocks y^t, s^t drawn from the common approximating model P^0. We see that with $\theta_1 = \infty$, there is an eventual drift towards zero and in the case of $\theta_1 < \infty$ the path eventually clusters around $\lambda_t = 1$.

In related work Farhi and Werning (2007) show that adding a paternalistic planner (whose is more patient than the agents) can also give long run survival with efficiency in presence of asymmetric information.

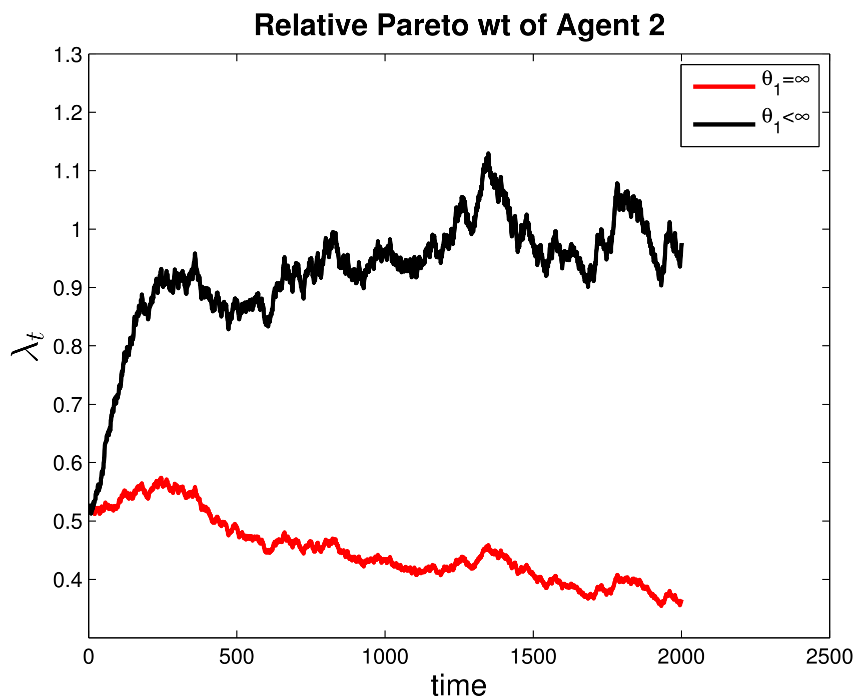

Figure 8: The figure plots a simulation of λ_t for a sequence of shock drawn from the common approximating model

7.2 Continuum of type 2 agents

This section generalizes the setup taste shocks are "idiosyncratic" to individuals to type 2. These are drawn independently across agents from the distribution $P_S(s)$. We setup the recursive planner's problem and re-derive the conditions that characterize the evolution of Pareto weights of any individual of type 2 relative to Agent 1.

Let λ_1 be the Pareto weight on Agent 1 and $\Gamma(\lambda_2)$ be the some distribution of Pareto weights amongst type 2 agents. Let $\mathcal{W}(\lambda_1, \Gamma, y)$ be the value to a planner from the optimal insurance scheme using weights λ_1 and $\Gamma(\lambda_2)$ given the realization aggregate y before the taste shocks are realized.

The planner chooses $\mathbf{\Upsilon} = \{u_1, u_2(\lambda_2, s), \lambda^*(\lambda_2, s), w_1^*(y^*), w_2^*(\lambda_2, s, y^*)\}$ to solve the following Bellman equation[33].

Let

- $V_1[\mathbf{\Upsilon}] = (1-\delta)u_1 + \mathbb{T}_{\theta_1,y}\delta w_1^*(y^*)$
- $V_2[\lambda_2, \mathbf{\Upsilon}] = \mathbb{T}_{\theta_1}\left[(1-\delta)su_2(\lambda_2, s) + \delta\mathbb{T}_{\theta_1,y}w_2^*(\lambda_2, s, y^*)\right]$

The Bellman equation that solves the Planner's problem is described below,

$$\mathcal{W}(\lambda_1, \Gamma, y) = \max_{\mathbf{\Upsilon}} \lambda_1 V_1[\mathbf{\Upsilon}] + (1-\lambda_1)\int d\Gamma(\lambda_2)\lambda_2 V_2[\lambda_2, \mathbf{\Upsilon}]$$

subject to

$$\int \sum_s P_S(s) d\Gamma(\lambda_2) C[u_2(\lambda_2, s)] = y - C[u_1]$$

for all λ_2, s, s'

$$(1-\delta)su_2(\lambda_2, s) + \mathbb{T}_{\theta_1,y}\delta w_2^*(\lambda_2, s, y^*) \geq (1-\delta)su_2(\lambda_2, s') + \mathbb{T}_{\theta_1,y}\delta w_2^*(\lambda_2, s', y^*)$$

for all y^*

$$\mathcal{W}(\lambda_1^*, \Gamma^*, y^*) \geq \lambda_1^* w_1^* + (1-\lambda_1^*)\int \sum_s d\Gamma(\lambda_2) P_S(s) \lambda_2^*(\lambda_2, s, y^*) w_2^*(\lambda_2, s, y^*)$$

for each y^*, Γ^{y^*} is given by

$$\Gamma^{y^*}(\tilde{\lambda_2}|\Gamma, y) = \int \sum_s P_S(s) d\Gamma(\lambda_2) \mathbb{I}_{\lambda^*(\lambda_2, s, y^*) \leq \tilde{\lambda_2}}$$

Notice that the continuation value function for the planner comes in equation (37).

Let $\Gamma'(\lambda_2)(1-\lambda_1)\lambda_2\mu(\lambda_2, s)\tilde{P}^2(\lambda_2, s)$ be the multiplier of the incentive constraint for agent 2 with Pareto weight λ_2 in state s and $\xi(y^*)P_Y(y^*)$ on the constraint in equation (37). Consider the first order necessary conditions with respect to $w_1^*(y^*)$ and $w_2^*(\lambda_2, s, y^*)$

$$\lambda_1 \tilde{P}^1(y^*|y) = \xi(y^*) P_Y(y^*|y) \lambda_1^*(y^*) \tag{38}$$

and

$$(1-\lambda_1)\lambda_2 \tilde{P}^2(s)\tilde{P}^2(y^*|\lambda_2, y)\left[1 + \mu(\lambda_2, s) - \mu(\lambda_2, s')\frac{\tilde{P}^2(s', y^*|\lambda_2, y)}{\tilde{P}^2(s)\tilde{P}^2(y^*|\lambda_2, y)}\right] = \xi(y^*) P_Y(y^*|y) P_S(s)(1-\lambda_1^*(y^*)) \tag{39}$$

Let $\bar{\lambda}_2(\lambda_1, \lambda_2) \equiv \frac{(1-\lambda_1)\lambda_2}{\lambda_1}$ be the relative Pareto weight of type 2 agents. Combining the two we have the following two results

[33]We implicitly impose that all individuals with the same Pareto weight are treated in the same way.

Lemma 5

$$\frac{\bar{\lambda}_2(\lambda_1^*(y^*), \lambda_2^*(\lambda_2, s, y^*))}{\bar{\lambda}_2(\lambda_1, \lambda_2)} = \frac{\tilde{P}^2(s)\tilde{P}^2(y^*|\lambda_2, y)}{\tilde{P}^1(y^*|y)} \left[1 + \mu(\lambda_2, s) - \mu(\lambda_2, s') \frac{\tilde{P}^2(s', y^*|\lambda_2, y)}{\tilde{P}^2(s)\tilde{P}^2(y^*|\lambda_2, y)}\right] \quad (40)$$

and

$$\tilde{\mathbb{E}}\bar{\lambda}_2(\lambda_1^*(y^*), \lambda_2^*(\lambda_2, s, y^*) = \bar{\lambda}_2(\lambda_1, \lambda_2) \quad (41)$$

Thus we have exactly the same dynamics for relative Pareto weights as in equation (16). The ergodic distribution Γ^y that solves

$$\Gamma^{y^*}(\tilde{\lambda_2}|\Gamma^y, y) = \Gamma^{y^*} \quad \forall y^*$$

Remark 5 *Without aggregate risk or $\|\mathcal{Y}\| = 1$, the relative Pareto weights $\bar{\lambda}_{t+1}$ only depend on s_t. Agent 1 faces a deterministic consumption scheme and $z_{1,t+1} = 1$. Thus the covariance term, $Cov_t(\lambda_{t+1}, z_{1,t+1})$ in equation 41 is zero and $\bar{\lambda}_t$ is non negative martingale. Without aggregate risk, the returns from the competitive insurance markets which were that crucial to counter the immiseration force in economies with private information is absent*

8 Extensions

The analysis so far exploited two key features, presence of aggregate risk and the multiplicative nature of taste shocks. In related work Bhandari (2013b), we study implications of model misspecification in environments with no aggregate risk and additive un-insurable idiosyncratic income risk. With complete markets and common approximating models, the risk sharing scheme implies constant consumption for all the agents. In the paper we study the consequences of two forms of market incompleteness.

First we restrict agents such that they can only trade a risk-free bond. The key finding is that in contrast to the results of this paper, relative pessimism is diminishing in wealth, a result that is not dependent on the value of IES. Having a large amount of wealth in assets that yield non-contingent return, lowers the volatility of consumption and consequently concerns for misspecification for richer agents. In Appendix B, we sketch a simple two period example to illustrate this point.

Next we study Pareto optimal risk sharing schemes under the restrictions that individual incomes are private information. Working with ex ante identical but finite agents, we impose an additional restriction that transfers to individuals only depend on the reports of histories of their incomes. This rules out trivially optimal allocations where the planner can impose the first best by using reports from one agent to punish possible misreports by the second agent.

The efficient risk sharing scheme (within the class of such restricted allocations) without the concerns for misspecification has a property that either one of the two agents is driven to immiseration. This comes from the dynamics of continuation values associated with efficient incentives. Since agents linearly aggregate continuation values a mean zero perturbation of continuation values (from a static risk sharing scheme) delivers the same ex ante value but relaxes incentive constraints. The insight in Atkeson Lucas (1992) suggests that such perturbations are always profitable and optimal incentives would imply that continuation values will spread. With enough bad shocks, some agent can drift towards immiseration. However, as we show in Bhandari (2013b), when there are concerns for uncertainty, agents with lower continuation values are relatively more pessimistic and consequently over-estimate the states when they have lower continuation values. The planner alters the risk sharing arrangement by reducing the amount by which continuation values are lowered. This generates a force away from immiseration.

9 Conclusions and future work

In recent years, causes and consequences of growing inequality has been a topic for great debate. The immiseration results in (exogenous) heterogenous beliefs or heterogenous information settings can be interpreted as arguments that there are economic forces associated with efficiency that can *result* in inequality. As such, they caution us from blanket policies that aim at "remedying" the problem of inequality.

This paper re-examined such arguments in settings where the nature of non-diversifiable aggregate risk is complex and a plausible description of agents' attitudes towards risk is captured by a large set of probability distributions. Using a common framework, we analyzed Pareto efficient insurance arrangements where agents centered their misspecification concerns around heterogeneous approximating models and had heterogeneous information. The results contrasted the properties long run inequality and market price of risk with a benchmark where agents had complete trust in their approximating models.

In presence of fluctuations to aggregate endowment, misspecification concerns implied a wedge between approximating models and ex-post beliefs. With bounded utility, the agent with a lower Pareto weight explored a smaller set of misspecification. This enabled him to provide valuable insurance to the other agent whose misspecification fears led to a larger wedge between his approximating model and ex-post beliefs. These asymmetric departures from approximating models prevented immiseration of agents who either had incorrect approximating models or were subject to privately observed taste shocks.

Besides inequality, this paper studied how wealth-driven endogenous heterogeneity in beliefs can alter implications for prices and volume for trades in Arrow-securities. Recent literature studying financial markets has identified heterogeneity in beliefs as a qualitative feature that explains a range of interesting phenomena. For example, Simsek (2013) studies an economy where belief differences of traders interact with endogenous collateral constraints and Buraschi and Jiltsov (2006) who exploit beliefs heterogeneities to explain option prices. Since our analysis in this paper was stylized, a natural next step is a more quantitative exploration of how endogenous heterogeneity in beliefs can account for nature, volume, and prices of a broader set of financial securities.

A Proofs

Proposition 1 :
Proof.

$$z_i^*(y) = \frac{\exp\left\{\frac{-\eta_i^{1-\gamma} y^{1-\gamma}}{\theta(1-\gamma)}\right\}}{\mathbb{E}\exp\left\{\frac{-\eta_i^{1-\gamma} y^{1-\gamma}}{\theta(1-\gamma)}\right\}}$$

Note that at $y = \underline{Y}$

$$\frac{\partial \log z_i(\underline{Y})}{\eta_i} = \frac{\eta_i^{-\gamma}}{\theta}\left[\tilde{\mathbb{E}} y^{1-\gamma} - \underline{Y}^{1-\gamma}\right]$$

Since $\underline{Y}$ is the lower end of the support, we have the term on the RHS of the above derivative to be positive iff $\gamma < 1$

If $\gamma < 1$, $z_i^*(\underline{Y})$ is further bounded below by 1. Thus $\lim_{\eta_i \to 0} \tilde{p}^i(\underline{Y}) = p^i(\underline{Y})$. With $\gamma > 1$, however, $\tilde{p}^i(\underline{Y})$ increases as η_i approaches 0, since $\tilde{p}^i(y) \leq 1$ we have [34]

$$\lim_{\eta_i \to 0} \tilde{p}^i(\underline{Y}) = 1$$

■

Proposition 2 : Proof. Given our assumptions that $\mathcal{Y}$ is finite and $u(0)$ is well defined, the proof is standard. Let $\mathcal{B}(\Delta(\mathcal{M})^2 \times [u(0) \quad v_{max}] \times Y)$ be the set of continuous bounded functions on a compact domain. Define a map $\mathcal{P}$ that solves the problem (13) for a given guess of continuation value function Q.

$$\mathcal{P} : \mathcal{B}(\Delta(\mathcal{M})^2 \times [u(0) \quad v_{max}] \times Y) \to \mathcal{B}(\Delta(\mathcal{M})^2 \times [u(0) \quad v_{max}] \times Y)$$

We first show that $\mathcal{P}$ is indeed a self-map. Under the assumption $u(0)$ is well defined, the resource constraint and bounds for $\bar{v}(s, y^*)$ imply that at any feasible $\{\boldsymbol{u}, \bar{v}(s, y^*)\}$, the objective function is bounded above by

$$\mathcal{Q} = max_{\pi_1 \in (0,1)} \mathbb{T}_{\theta_1}\left[(1-\delta)u(y) + \delta \mathbb{R}_{\pi_1,\theta_2}\mathbb{T}_{\theta_1,m,y}\mathcal{Q}(u(0), y^*)\right]$$

Continuity follows from the Maximum theorem.

Next we show that $\mathcal{P}$ satisfies

1. $\mathcal{P}Q^1 \geq \mathcal{P}Q^2$ for all $Q^1 \geq Q^2$
2. $\mathcal{P}Q + c \leq \mathcal{P}Q + \delta c$ for all $c \geq 0$ and some $\delta \in (0,1)$

The operators $\mathbb{T}_{\theta_1,m,y}$ and $\mathbb{R}_{\theta_2}$ implicit in the map $\mathcal{P}$ are given by

$$\mathbb{T}_{\theta_1,m,y}Q^* = -\theta \log \sum_{y^*} P_Y(y^*|y, m) \exp\left\{\frac{-Q^*}{\theta_1}\right\}$$

$$\mathbb{R}_{\pi_1,\theta_2}\mathbb{T}_{\theta_1,m,y}Q^* = -\theta_2 \log \sum_m \pi^1(m) \exp\left\{\frac{-\mathbb{T}_{\theta_1,m,y}Q^*}{\theta_2}\right\}$$

[34] $\gamma > 1$, z_i^* is not defined at $\eta_i = 0$ as the utility is unbounded in that case.

Its easy to see that $\mathbb{T}_{\theta_1,m,y}$ and $\mathbb{R}_{\pi_1,\theta_2}$ and the outer $\mathbb{T}_{\theta_1}$ are increasing and satisfies the discounting property.

Let $\boldsymbol{u}^k, \bar{v}^{k,*}(s,y^*)$ for $k=1,2$ be the optimal policies associated with Q^1 and Q^2.

$$\begin{aligned}
\mathcal{P}Q^2(\pi,v,y) &= \mathbb{T}_{\theta_1}\left[u_1^2(s) + \mathbb{R}_{\pi_1,\theta_2}\mathbb{T}_{\theta_1,m,y}Q^2(\boldsymbol{\pi}^*(y^*),\bar{v}^{2,*}(y^*),y^*)\right] \\
&\leq \mathbb{T}_{\theta_1}\left[u_1^1(s) + \mathbb{R}_{\pi_1,\theta_2}\mathbb{T}_{\theta_1,m,y}Q^2(\boldsymbol{\pi}^*(y^*),\bar{v}^{1,*}(y^*),y^*)\right] \\
&\leq \mathbb{T}_{\theta_1}\left[u_1^1(s) + \mathbb{R}_{\pi_1,\theta_2}\mathbb{T}_{\theta_1,m,y}Q^1(\boldsymbol{\pi}^*(y^*),\bar{v}^{1,*}(y^*),y^*)\right] \\
&= \mathcal{P}Q^1(\boldsymbol{\pi},v,y)
\end{aligned} \tag{42}$$

The first inequality comes from $\mathbb{R}_{\theta_2}\mathbb{T}_{\theta_1,m,y}$ being increasing and the second inequality comes from the fact that given Q^1 the policies $\{\boldsymbol{c}^2, \boldsymbol{v}^{2,*}(y^*)\}$ are feasible.

The discounting for $\mathcal{P}$ is obvious. Thus $\mathcal{P}$ is a contraction and $\mathcal{B}(\Delta(\mathcal{M})^2 \times [u(0) \quad v_{max}] \times Y)$ is complete under the sup norm. Blackwell's theorem implies that there is a unique $\mathcal{Q}$ that solves the problem.

Fix $\boldsymbol{\zeta} = \{z_i(y^*|s,v,y,\boldsymbol{\pi}), h_i(m|s,v,y,\boldsymbol{\pi}), z_i^s(s|v,y,\boldsymbol{\pi})\}_{i=1,2}$ be some arbitrary likelihood ratios that represent ex post beliefs for agent 1 and agent 2.

1. $\mathcal{I}_i^1 = \theta_1 \sum_{s,m,y^*} P_S(s)\pi_i(m)P_Y(y^*|y,m)h_i(m|s)z_i^s(s)z_i^m(y^*)\log z_i^m(y^*)$
2. $\mathcal{I}_i^2 = \theta_2 \sum_{s,m} P_S(s)z_i^s(s)\pi_i(m)h_i(m)\log h_i(m|s)$
3. $\mathcal{I}_i^3 = \theta_2 \sum_s P_S(s)z_i^s(s)\log z_i^s(s)$

Note that we can write $\mathcal{P}$ as follows

$$\max_{\boldsymbol{u},\bar{v}(s,y^*)}\min_{\zeta_i}\mathbb{E}\left[z_1^s(s)\sum_m \pi_1(m)h_1(m|s)z_1^{m,*}(y^*)\left\{(1-\delta)u_1(s)+\delta Q(\boldsymbol{\pi}^*,\bar{v}(s,y^*),y^*)\right\}\right]+\sum_k \mathcal{I}_1^k(\zeta_1) \tag{43}$$

and

$$\min_{\zeta_2}\mathbb{E}\left[z_2^s(s)\sum_m \pi_2(m)h_1(m|s)z_2^{m,*}(y^*)\left\{(1-\delta)su_2(s)+\delta\bar{v}(s,y^*)\right\}\right]+\sum_k \mathcal{I}_2^k(\zeta_2) \geq v \tag{44}$$

$$\begin{aligned}
&(1-\delta)su_2(s) + \delta\min_{\zeta_2}\left[\sum_m \pi_2(m)h_1(m|s)z_2^{m,*}(y^*)\bar{v}(s,y^*)\right] + \sum_k \mathcal{I}_2^k(\zeta_2) > \\
&(1-\delta)su_2(\tilde{s}) + \delta\min_{\tilde{\zeta}_2}\left[\sum_m \pi_2(m)\tilde{h}_1(m|s)\tilde{z}_2^{m,*}(y^*)\bar{v}(\tilde{s},y^*)\right]
\end{aligned} \tag{45}$$

$$\mathbb{E}^m z_i^m(y^*|y,s) = 1 \quad i=1,2 \text{ and } \forall m \in \mathcal{M} \tag{46a}$$

$$\sum_m h_i(m|s)\pi_i(m) = 1 \qquad i=1,2 \tag{46b}$$

$$\sum_m z_i^s(s)P_S(s) = 1 \qquad i=1,2 \tag{46c}$$

$$\mathbb{E}^m \tilde{z}_2^m(y^*|y,s) = 1 \tag{46d}$$

$$\sum_m \tilde{h}_2(m|s)\pi_2(m) = 1 \tag{46e}$$

and the feasibility constraints on $\{\boldsymbol{c}, \boldsymbol{v}^*(y^*)\}$ and the law of motion for $\boldsymbol{\pi}^*$

Define a sub map $\mathcal{P}_{\boldsymbol{\zeta}}$ that iterates on the Bellman equation given ex-post beliefs that are computed using $\boldsymbol{\zeta}$. We first show that $\mathcal{P}_{\boldsymbol{\zeta}}$ preserves concavity, boundedness and monotonicity. For some $(y, \boldsymbol{\pi})$, let v_1^0, v_2^0, v_3^0 be some initial promised values such that there exist some $\alpha \in (0,1)$ such that

$$v_2^0 = \alpha_1 v^0 + (1-\alpha_1)v_2^0$$

Denote the solutions to $\mathcal{P}_{\boldsymbol{\zeta}}$ associated with $\{v_k^0, y\}_{k=1,3}$ as $\{u_{1,k}(s), u_{2,k}(s), \bar{v}_k(s,y^*)\}$. Construct a candidate solution

$$\hat{u}_i(s) = \alpha u_{i,1}(s) + (1-\alpha)u_{i,3}(s) \quad i = 1,2 \tag{47a}$$

$$\hat{\bar{v}}_2(s,y^*) = \alpha \bar{v}_1(s,y^*) + (1-\alpha)\bar{v}_3(s,y^*) \tag{47b}$$

The promise keeping constraint and incentive constraints are linear in the choices. Under the assumption that $C(u)$ is convex, the resource constraint is also satisfied. This implies that $\{\hat{u}_i(s), \hat{\bar{v}}_2(s,y^*)\}$ is feasible given (v_2^0, y). Thus for a concave $Q(v^0, y)$ we need,

$$\mathcal{P}_{\boldsymbol{\zeta}}(v_2^0, k) \geq \tilde{\mathbb{E}}_y\left[(1-\delta)\hat{u}_1(s) + \delta Q^0(\hat{\bar{v}}(s,y^*), y^*)\right] \tag{48}$$

$$\geq \alpha\left(\tilde{\mathbb{E}}_y\left[(1-\delta)u_{1,1}(s) + \delta\alpha Q^0(\bar{v}_1(s,y^*),y^*)\right]\right) + (1-\alpha)\left(\tilde{\mathbb{E}}_y\left[(1-\delta)u_{1,3}(s) + \delta\alpha Q^0(\bar{v}_3(s,y^*),y^*)\right]\right) \tag{49}$$

$$\geq \alpha Q(v_1^0, y) + (1-\alpha)\alpha Q(v_3^0, y) \tag{50}$$

$$\geq \alpha \mathcal{P}_{\boldsymbol{\zeta}}(v_2^0, k)(v_1^0, y) + (1-\alpha)\alpha \mathcal{P}_{\boldsymbol{\zeta}}(v_2^0, k)(v_3^0, y) \tag{51}$$

Note that the set of feasible policies $\mathcal{F}(\boldsymbol{\pi}, v, y)$ satisfies $\mathcal{F}(\boldsymbol{\pi}, v_1, y) \subset \mathcal{F}(\boldsymbol{\pi}, v_2,)$ for $v_1 \geq v_2$. Thus $Q(\boldsymbol{\pi}, v, y)$ is weakly decreasing.

Thus $\mathcal{P}_{\boldsymbol{\zeta}}$ preserves concavity, monotonicity, boundedness and continuity. We can apply similar arguments as above and show that it is a contraction. Thus $\mathcal{P}_{\boldsymbol{\zeta}}$ has a unique fixed point that is concave. Our argument is complete by evaluating $\mathcal{P}_{\boldsymbol{\zeta}}$ at $\boldsymbol{\zeta}^*$ associated with the fixed point of the previous map $\mathcal{P}$. ■

Proposition 3 :

Proof.

Proposition 2 shows that $\mathcal{Q}$ is concave. The envelope condition implies a increasing monotonic relationship between λ and v. Let v_0 be chosen to satisfy $\lambda_0(\boldsymbol{\pi}, v_0, y_0) = \frac{1-\Gamma}{\Gamma}$. Consider the Lagrangian of the recursive problem,

Consider the Lagrangian for the Planner's problem

$$\mathcal{L}(\pi, v, y, \lambda) \equiv \mathbb{T}_{\theta_1}\left[(1-\delta)u_1(s) + \delta \mathbb{R}_{\pi_1,\theta_2}\mathbb{T}_{\theta_1,m,y}\mathcal{Q}(\boldsymbol{\pi}^*, \bar{v}(s,y^*), y^*)\right] + \lambda\left\{\mathbb{T}_{\theta_1}\left[(1-\delta)su_2(s) + \delta \mathbb{R}_{\pi_2,\theta_2}\mathbb{T}_{\theta_1,m,y}\bar{v}(s,y^*)\right]\right.$$

Notice that Bellman equations in 13 and promise keeping constraint (assuming it holds at equality), imply that at the optimal choice, the Lagrangian

$$\mathcal{L}(\pi, v_0, y, \lambda_0) = \Gamma V_0^1[c_1|y_0] + (1-\Gamma)V_0^2[c_2|y_0]$$

We need to show that the incentive constraints imposed in the recursive problem (13) and time 0 sequential problem (11) imply exactly the same restrictions. Let $c_2 = \{c_{2,t}(y^t, s^t)\}$ be a consumption plan for Agent 2 and let σ^* be the optimal reporting strategy given c_2,i.e.,

$$\sigma^* \in \text{argmax}_\sigma V_0[c_2(\sigma)|y_0]$$

For expected utility, this follows from Blackwell's (1965) one shot deviation principle. The next lemma argues that under some mild technical conditions, this extends to our setting as well.

For any choice of σ we can use the following recursion

$$V_t^i[c|y^t, s^{t-1}] = \mathbb{T}_{\theta_1}\left\{(1-\delta)s_t u(c_2(\sigma_t)) + \delta \mathbb{R}_{\pi_{i,t},\theta_2} \mathbb{T}_{\theta_1,m,y_t} V_{t+1}^i[c_2(\sigma)|y^{t+1}, s^t]\right\} \tag{52}$$

to back out the ex-post beliefs $\tilde{P}_t^2(s^t, y^t)$ that depend on $c_2(\sigma)$ and solve the minimizations implicit in operators $\mathbb{T}_{\theta_1,m,y_t}$, $\mathbb{R}_{\pi_{2,t},\theta_2}$ and $\mathbb{T}_{\theta_1}$. Further given these beliefs, the optimal value can be represented as

$$\tilde{\mathbb{E}}_0 \sum_t (1-\delta)\delta^t u(c_2(\sigma_t)) + \text{terms independent of } \sigma_t$$

When u is concave and bounded [35], one can flip the choices of the ex-post probabilities and σ. Thus for a choice of $\{\hat{P}_t(y^t, s^t)\}$, the choice of $\sigma[\hat{P}]$ solves a standard expected utility optimization. Blackwell (1965) provides mild conditions under which the optimal choice of $\sigma[\hat{P}]$ rules out one shot deviations. The proof follows by evaluating $\hat{P}$ at the solution associated with $\sigma_t^*(y^t, s^t) = s_t$

■

Proposition 4 :

Proof.

With $\theta_1, \theta_2 = \infty$, the marginal rate of substitution, $MRS(y^*|y) = \left(\frac{y^*}{y}\right)^{-\gamma}$. First compute the ingredients for computing the conditional market price of risk. Let $y_h/y_l = \frac{1+g}{1-g}$

1. $y = y_l$

$$\log[MRS(y^*|y_l)] = 0 \quad y^* = y_l \tag{53}$$

$$\log[MRS(y^*|y_l)] = -2\gamma g \quad y^* = y_h \tag{54}$$

1. $y = y_h$

$$\log[MRS(y^*|y_l)] = 2\gamma g \quad y^* = y_l \tag{55}$$

$$\log[MRS(y^*|y_l)] = 0 \quad y^* = y_h \tag{56}$$

$$\sigma[\log(MRS)|v, y] = \sqrt{\bar{\alpha}(1-\bar{\alpha})}(2g\gamma) \tag{57}$$

$$\sigma[\log(MRS)|y_l] = \sigma[\log(MRS)|y_h] = \sqrt{\bar{\alpha}(1-\bar{\alpha})}2\gamma g$$

The market price of risk is independent of v and y when $\theta_2, \theta_1 = \infty$

[35] These are known as known an Isaacs-Bellman conditions. Refer Hansen and Sargent chapter 7 for more details where they study alternative timing protocols for zero-sum games representation of these preferences.

Let $\Delta \log \zeta(v,y) = \log \zeta[y(z^*) = y_l|v,z] - \log \zeta[y(z^*) = y_h|v,z]$. With $\theta_1 = \theta_2$, we expression for market price of risk has a similar expression,

$$\sigma[\log(MRS)|v,z] = \sqrt{\bar{\alpha}(1-\bar{\alpha})}(2g\gamma + \Delta \log \zeta(v,y)) \tag{58}$$

When $v \to 0$, the economy converges to an economy populated by type 1 agents and $\frac{\zeta(y^*=y_l)}{\zeta(y^*=y_h)} \to \exp\{v^{max}(y_h) - v^{max}(y_l)\}$. This is independent of the current state y.

With $\theta_1 = \infty$ and $\theta_2 < \infty$, the cyclical properties of MPR now depend on the conditional volatility of the distorted log likelihood ratio.

Let π denote the probability of $\alpha_m = \frac{1}{2}$ or the common prior weight on the IID model and $\bar{\alpha} = \pi\frac{1}{2} + (1-\pi)\alpha_2$

$$MPR^{\theta_2}(\pi,y) = \sigma\left[\left\{\log \frac{\sum_m \tilde{\pi}[m|y]P_Y(y^*|y,m)}{\sum_m \pi[m|y]P_Y(y^*|y,m)}\right]\right\}|y] + MPR^{\infty}(\pi)$$

Lemma 6 *$MPR^{\theta_2}(\pi,y_l) > MPR^{\theta_2}(\pi,y_h)$ iff $\tilde{\pi}[y_h] + \tilde{\pi}[y_l] < 2\pi \quad \forall\pi$*

Proof. For an arbitrary $\pi, \tilde{\pi}(y_l), \tilde{\pi}(y_h)$ define the function L

$$L[\pi,\tilde{\pi}(y_l),\tilde{\pi}(y_h)] = \sigma\left[\left\{\log \frac{\sum_m \tilde{\pi}[m|y]P_Y(y^*|y=y_l)}{\sum_m \pi[m|y]P_Y(y^*|y=y_l)}\right\}|y=y_l\right] - \sigma\left[\left\{\log \frac{\sum_m \tilde{\pi}[m|y]P_Y(y^*|y=y_h)}{\sum_m \pi[m|y]P_Y(y^*|y=y_l)}\right\}|y=y_h\right]$$

It can be verified that L satisfies the following properties

- $L(\pi,\tilde{\pi}(y_l),\max\{2\pi - \tilde{\pi}(y_l)\}) > 0$ for all $\tilde{\pi}(y_l) \leq \pi$
- $\frac{\partial L}{\partial \tilde{\pi}(y_h)} > 0$

When $y = y_l$ the agent 1 distorts the mixture probability away from the IID model and hence $\tilde{\pi}(y_l) \leq \pi$ ■

Lemma 7 *There exists $\bar{\pi}$ such that $\tilde{\pi}[y_h] + \tilde{\pi}[y_l] < 2\pi \quad \forall\pi > \bar{\pi}$*

Proof.

Let $\Delta(y,\pi) = \mathbb{E}^{NonIID}[Q^{max}(y^*,\pi^*)|y] - \mathbb{E}^{IID}[Q^{max}(y^*,\pi^*)|y]$.

With the IID model we have $\mathbb{E}^{IID}[Q^{max}(y^*,\pi^*)|y=y_h,\pi] = \mathbb{E}^{IID}[Q^{max}(y^*,\pi^*)|y=y_l,\pi]$.

Thus $\Delta(y_l,\pi) = -\Delta(y_h,\pi)$

$$\frac{\tilde{\pi}(y)-\pi}{\pi} = \frac{1-\exp\{-\Delta(y,\pi)\}}{\frac{\pi}{1-\pi} + \exp\{-\Delta(y,\pi)\}}$$

$\bar{\pi}$ solves,

$$\frac{1-\exp\{-\Delta(y,\pi)\}}{\frac{\pi}{1-\pi} + \exp\{-\Delta(y,\pi)\}} = \frac{-1+\exp\{\Delta(y,\pi)\}}{\frac{\pi}{1-\pi} + \exp\{\Delta(y,\pi)\}}$$

■

The hypothesis follows by combining the two lemmas.

■

Proposition 1

Proof.

It is easy to see that $F^i(m)$ will be vary across m's as long as $Q_{t+1}(y_{t+1},\nu_{t+1})$ is not constant across possible values of y_{t+1}.

Further as long as $\lambda_{t+1}\neq 1$ for all shocks, there would be some realization of y_{t+1} where $Q_{t+1}(y_{t+1},\nu_{t+1})\neq v_{t+1}(y_{t+1},\nu_{t+1})$. Since if it was not true v_{t+1} has to equal to the value of consuming half of the aggregate endowment for all $t+j$. This can only happen for $\lambda_t=1$. Thus $F_t^1(m)\neq F_t^2(m)$ if $\lambda_t\neq 1$

Finally

$$P_Y(y_{t+1},\nu_{t+1}|y_t,\nu_t,m)=P_Y(y_{t+1},y_{t+2}+\epsilon_{t+1}|y_t,\nu_t,m)=P_Y(y_{t+1},|y_t,m)P_Y(y_{t+2}+\epsilon_{t+1}|y_{t+1},m)$$

WLOG say $y_{t+1}=y_t=y_l$,

$$P_Y(y_{t+1},\nu_{t+1}|y_t,\nu_t,m)=\alpha_m[\phi(\nu_{t+1}-y_l)\alpha_m+\phi(\nu_{t+1}-y_h)(1-\alpha_m)]$$

Dividing it across the two models $m=1,2$ for any ν_{t+1},

$$\frac{P_Y(y_{t+1},\nu_{t+1}|y_t,\nu_t,m=1)}{P_Y(y_{t+1},\nu_{t+1}|y_t,\nu_t,m=2)}=1 \quad \text{iff } \alpha_1=\alpha_2 \text{ or } \alpha_1+\alpha_2=1+\frac{\phi(v_{t+1}-y_l)}{\phi(v_{t+1}-y_h)}$$

The conditions of α_m rule of these possibilities. Similarly one can show that $P_Y(y_{t+1},\nu_{t+1}|y_t,\nu_t,m)\neq P_Y(y_{t+1},\tilde{\nu}_{t+1}|y_t,\nu_t,m)$ for all m Thus equation (23) will not be independent of i.

and λ_{t+1} and through equation (18a) consumption shares will also be sensitive to news shocks.

■

Proposition 5

Proof.

Suppose not, Let $V^1(\xi^*)=Q(\xi^*,\pi^*,\xi^*)$ and $V^2(\xi^*)=v^*(\xi^*|\xi,v,\pi)$ where $y^*=y^{**}$

The FOC imply that $\lambda[\xi^*]=\lambda[\xi^{**}]$ if and only if

$$\left(\frac{\sum_m \tilde{\pi}_i(m)\tilde{P}^i(\xi^*|\xi)}{\sum_m \tilde{\pi}_i(m)\tilde{P}^i(\xi^{**}|\xi)}\right)$$

is independent of i

This is equivalent to the following expression

$$\frac{\sum_m \pi(m)P_Y(y^*|\xi,m)P_{\nu^*|y^*}(\nu^*|y^*,\xi,m)F_m^i}{\sum_m \pi(m)P_Y(y^*|\xi,m)P_{\nu^*|y^*}(\nu^*|y^*,\xi,m)F_m^i}$$

be independent of i. Where $F_{\xi,m}^i=\left\{\left(\frac{1}{\theta_2}\right)\mathbb{E}_{\xi,m}V^i(\xi^*)\right\}$

The necessity of $\theta_1<\infty$ clear since otherwise $F_{\xi,m}^i$ is independent of m. Similarly, if we have no aggregate risk or more generally have only $\mathcal{M}_\alpha$ to be trivial, consumption being only a function of y^* would imply $V^i(\xi^*)$ to be only a function of $V^i(y^*)$. This would again make $F_{\xi,m}^i$ be independent of m.

Further if $P_{\nu^*|y^*}^m(\nu^*|y^*,\xi,m)=P_{\nu^*|y^*}(\nu^*|y^*,\xi,m) \quad \forall m s.t \pi(m)>0$

$$=\frac{\sum_m \pi(m)P_{\nu^*|y^*}(\nu^*|y^*,\xi)P_Y^m(y^*|\xi)F_m^i}{\sum_m \pi(m)P(\nu^{**}|y^*,\xi)P_Y^m(y^*|\xi)F_m^i} \tag{59}$$

$$=\frac{P(\nu^*|y^*,\xi)}{P(\nu^{**}|y^*,\xi)} \tag{60}$$

When ν^* and y^* are independent and $\mathcal{M}$ is complete i.e $\mathcal{M} = \mathcal{M}_\alpha \times \mathcal{M}_\beta$we can write the ratio as

$$= \frac{\sum_{\alpha_m}\sum_{\beta_m}\pi(\alpha_m)\pi(\beta_m)P_{\nu^*|y^*}(\nu^*|\xi,\beta_m)P_Y(y^*|\xi,\alpha_m)F^i_{\alpha_m}}{\sum_{\alpha_m}\sum_{\beta_m}\pi(\alpha_m)\pi(\beta_m)P_{\nu^*|y^*}(\nu^{**}|\xi,\beta_m)P_Y(y^*|\xi,\alpha_m)F^i_{\alpha_m}} \tag{61}$$

$$= \frac{\left(\sum_{\alpha_m}\pi(\alpha_m)P_Y(y^*|\xi,\alpha_m)F^i_{\alpha_m}\right)\left(\sum_{\beta_m}\pi(\beta_m)P_{\nu^*|y^*}(\nu^*|\xi,\beta_m)\right)}{\left(\sum_{\alpha_m}\pi(\alpha_m)P_Y(y^*|\xi,\alpha_m)F^i_{\alpha_m}\right)\left(\sum_{\beta_m}\pi(\beta_m)P_{\nu^*|y^*}(\nu^{**}|\xi,\beta_m)\right)} \tag{62}$$

$$= \frac{\left(\sum_{\beta_m}\pi(\beta_m)P_{\nu^*|y^*}(\nu^*|\xi,\beta_m)\right)}{\left(\sum_{\beta_m}\pi(\beta_m)P_{\nu^*|y^*}(\nu^{**}|\xi,\beta_m)\right)} \tag{63}$$

which is independent of i

■

Additional details for example in section 6.2

We construct an example to show that news shocks matter even if $\theta_1 < \infty$.

The FOC 18b imply that $\lambda_{t+1}[y^{t+1},\nu^{t+1}] = \lambda_{t+1}[y^{t+1},\tilde{\nu}^{t+1}]$ if and only if

$$\frac{\sum_m \pi(m)P_Y(y_{t+1}|y_t,\nu_t,m)P_{\nu|y}(\nu_{t+1}|y_{t+1},\xi_t,m)F^i_t(m)}{\sum_m \pi(m)P_Y(y_{t+1}|y_t,\nu_t,m)P_{\nu|y}(\tilde{\nu}_{t+1}|y_{t+1},\xi_t,m)F^i_t(m)}\left(\exp\{\frac{V^i_{t+1}(y^{t+1},\nu^{t+1}) - V^i_{t+1}(y^{t+1},\tilde{\nu}^{t+1})}{\theta_1}\}\right)$$

be independent of i. Where $F^i_{t,m} = \left\{\left(\frac{1}{\theta_2} - \frac{1}{\theta_1}\right)\mathbb{T}_{\theta_1,\xi,m}V^i_{t+1}\right\}$

First assume $\theta_1 = \theta_2 = \theta$. This simplifies the previous condition and we require

$$\exp\{\frac{V^i_{t+1}(y^{t+1},\nu^{t+1}) - V^i_{t+1}(y^{t+1},\tilde{\nu}^{t+1})}{\theta}\} \tag{64}$$

to be independent of i

Consider a finite horizon setting with $t = 0, 1, 2$ with a history (y^2,ν^2) such that $\lambda_2(y^2,\nu^2) \neq 1$. This is not vacuous since we are free to choose the initial Pareto weights. The next lemma shows that starting with some initial inequality is sufficient.

Lemma 8 *Suppose news shocks did not matter,* $\lambda_0 \neq 1$ *implies* $\lambda_2(y^2,\nu^2) \neq 1$

Proof. Suppose $\lambda_2(y^2,\nu^2) = 1$. This would imply that the terminal period consumption for both the agents are $\frac{y_2}{2}$. Thus their ex-post beliefs for $y_2|y^1,\nu^1$ will agree for all histories. From equation (18b), we see that $\lambda_1(y^1,\nu^1) = 1$. Repeating the same argument one period backwards, we have $\lambda_0 = \lambda_1 = 1$. ■

Suppose we assume that consumption is not sensitive to news shocks. Thus would imply

$$V^i_1(y^1,\nu^1) - V^i_1(y^1,\tilde{\nu}^1) = \delta\theta[\log(\sum_m \hat{\pi}_1(m)\mathbb{E}^m_{y_1}\exp\{-\theta^{-1}u(c_{i,2}(\lambda_2)\}) - \log(\sum_m \pi_1(m)\mathbb{E}^m_{y_1}\exp\{-\theta^{-1}u(c_{i,2}(\lambda_2)\})]$$

Where $\hat{\pi}_1(m)$ is the prior associated with $\tilde{\nu}^1$. For the difference between the valuations across news shocks to be independent of i

$$\frac{\sum_m \hat{\pi}_1(m)\mathbb{E}^m \exp\{-\theta^{-1}u(c_{i,2}(\lambda_2)\}}{\sum_m \pi_1(m)\mathbb{E}^m_{y_1} \exp\{-\theta^{-1}u(c_{i,2}(\lambda_2)\}}$$

has to independent of λ_2.

The consumption shares will differ across agents if $\lambda_2 \neq 1$ As long the set $\mathcal{M}$ is not trivial, and $\pi(m) \neq \hat{\pi}(m)$ condition 64 will be violated.

Proposition 7 :

Proof. We build the proof through a sequence of lemmas

Lemma 9 *$\lambda(v, y)$ is continuous and monotonic in v*

Proof. The FOC (18a) gives $\lambda(v, y)$ as a continuous function of the allocation. The maximum theorem implies that allocation is continuous in v and hence $\lambda(v, y)$ is a continuous function of v. Monotonicity comes from the fact that the value function is strictly concave and the envelope theorem implies $\lambda_v(v, y) = -Q_{vv}(v, y)$. ■

Lemma 10 *Let $Q^{\max}(y_t)$ be the value of the aggregate endowment to Agent 1. We have $\lambda_t \to 0$ iff $v_t \to v_{min}$ and $Q(v_t) = Q_t^{max}$*

Proof. $Q(v_t, y_t) \leq Q^{\max}(y_t)$ by construction. It is easy to see that $v = v_{min}$, the allocation $c(y) = y$ and $v(y^*) = v_{min}$ is feasible. Thus $Q(v_{min}, y) = Q^{max}(y)$. Again, the maximum theorem implies $Q$is continuous and hence $v_t \to v_{min}$ implies $Q(v_t) \to Q_t^{max}$ pathwise. Evaluating the FOC (18a) at v_{min}, we have $\lambda(v_{min}, y) = 0$.

The previous lemma argues that $\lambda(v, y)$ is strictly increasing in v. Thus we can define a continuous inverse function $\hat{v}(\lambda, y)$ with the property that $\hat{v}(\lambda(v, y), y) = v$ for all $v > v_{min}$. One can extend this function continuously on the positive reals including ∞ by defining $\hat{v}(0, y) = v_{min}$ and $\hat{v}(\infty, y) = v_{max}$.

Thus for any sequence $\lambda_t \to 0$, we have $\lim_t v_t = \lim_t \hat{v}(\lambda_t, y) = v_{min}$

■

Define a matrix valued sequence or random variables $\xi_{i,t}(y^\infty) = \left\{\tilde{P}_t^i(y_{t+1}|y_t)\right\}$ for all possible realizations of (y_t, y_{t+1}). This measures the one period ahead transition matrix of the worst case model (that depends on the history y^t).

Suppose that with strictly positive probability agent 1 dominates agent 2 or $A = \{\omega : \lambda_t \to 0\} > 0$. Restrict the probability space to event A. Along any path that $\lambda_t \to 0$, we can compute the worst case models for both the agents using the exponential twisting formula

$$\lim_t \tilde{P}_t^2(y_{t+1}|y_t) \propto P_Y^2(y_{t+1}|y_t) \exp\{-\frac{v_{min}}{\theta}\}.$$

Since v_{min} is a constant

$$\lim_t \xi_{2,t} = \mathbb{P}_Y^2 \equiv \left\{P_Y^2(y_{t+1}|y_t)\right\} \quad \text{almost surely}$$

For Agent 1,

$$\lim_t \tilde{P}_t^1(y_{t+1}|y_t) \propto= P_Y^1(y_{t+1}|y_t) \exp\{\frac{-Q^{max}(y_t)}{\theta}\}.$$

Let $z_1(y_t)$ be the limiting Radon-Nikodym derivative,

$$\lim_t \xi_{1,t} = \hat{\mathbb{P}}^1 \equiv \left\{ P_Y^1(y_{t+1}|y_t) z_1(y_t) \right\} \quad \text{almost surely}$$

By Egorov's theorem, for every $\epsilon > 0$ there exists a (exceptional) set $E \subset A$ such that $P(E) < \epsilon$ and the almost sure convergence is uniform on the complement set $A - E$. Taking logs of the FOC equation (18b) we have

$$\log(\lambda_{t+1}) - \log(\lambda_t) = \log(\Phi(y_{t+1}, y_t)\xi_{2,t}) - \log(\Phi(y_{t+1}, y_t)\xi_{1,t}) \tag{65}$$

The operator $\Phi(y', y)$ picks out the corresponding element from matrix $\xi_{i,t}$. With uniform convergence on the set $A\ E$, we can arbitrarily choose a $\epsilon_2 > 0$ such that there exists a $T(\epsilon_2)$ that satisfies $t \geq T(\epsilon_2)$

$$\xi_t^2 \geq \mathbb{P}^2(1 - \delta) \tag{66}$$

$$\xi_t^1 \leq \hat{\mathbb{P}}^1(1 + \delta) \tag{67}$$

Where the inequalities are interpreted component-wise. Uniform continuity allows us to have $T(\epsilon_2)$ independent of the particular sample path.

Substituting in equation (65)

$$\log(\lambda_{t+1}) - \log(\lambda_t) \geq \log(\Phi(y_{t+1}, y_t)\mathbb{P}^2) - \log(\Phi(y_{t+1}, y_t)\hat{\mathbb{P}}^1) - C(\epsilon_2) \tag{68}$$

where $C(\epsilon_2) > 0$ Summing up to K terms from $T(\epsilon_2)$ we have

$$\log(\lambda_{T(\epsilon_2)+K}) \geq \sum_{T(\epsilon_2)}^{T(\epsilon_2)+K} \{\log(\Phi(y_{t+1}, y_t)\mathbb{P}^2) - \log(\Phi(y_{t+1}, y_t)\hat{\mathbb{P}}^1)\} - KC(\epsilon_2)$$

we can define an auxiliary Markov process $\chi_t = [Y_t \quad Y_{t-1}]'$ and the associated P_χ^0 using P_Y^0. It is easy to check that χ is Harris recurrent as long as Y is. Further $\Gamma(\chi)$ be the invariant distribution is given by $\Gamma_y(y)P_Y^0(y'|y)$. It is easy to check that $P_Y^0(y'y) > 0$ allows us to use the ergodic theorem,

$$\lim_K \frac{1}{K} \left[\sum_{T(\epsilon_2)}^{T(\epsilon_2)+K} \{\log(\Phi(\chi_{t+1})\mathbb{P}^2) - \log(\Phi(\chi_{t+1})\hat{\mathbb{P}}^1)\} \right] = \mathbb{I}^{0,1} - \mathbb{I}^{0,2} - \mathbb{E}^\Gamma \log z_1 \tag{69}$$

Thus for large K

$$\log(\lambda_{T(\epsilon_2)+K}) \geq K \left\{ \mathbb{I}^{0,1} - \mathbb{I}^{0,2} - \mathbb{E}^\Gamma \log z_1 - C(\epsilon_2) \right\} \tag{70}$$

Let $\underline{\mathrm{M}} = \mathbb{E}^{\Gamma_y} \log z_2(y)$ and $\bar{M} = -\mathbb{E}^{\Gamma_y} \log z_1(y)$.

Since ϵ_2 was arbitrary, the term inside the bracket is positive. Now the left hand side of equation (70) diverges to $-\infty$ and the RHS diverges to $+\infty$. Thus we have a contradiction.

■

Proposition 8:

Proof.

Let $\hat{Q}(\eta, y), \hat{v}(1-\eta, y)$ be the values of Agent 1 and 2 respectively from consuming a constant share of the aggregate endowment.

- $\hat{Q}(\eta, y) = (1-\delta)u(\eta, y) - \delta\theta_1 \log \sum_{y^*} \exp\left\{-\frac{\hat{Q}(\eta,y^*)}{\theta}\right\} P_Y^1(y^*|y)$
- $\hat{v}(1-\eta, y) = (1-\delta)u[(1-\eta)y] - \delta\theta_1 \log \sum_{y^*} \exp\left\{-\frac{\hat{v}(1-\eta,y^*)}{\theta}\right\} P_Y^2(y^*|y)$

The IID shocks imply that $\hat{Q}(\eta, y) - \hat{v}(1-\eta, y) = [(1-\delta)]\left[u(\eta y) - u((1-\eta)y)\right] + D(\eta)$. Using the (18b) we are looking for η such that

$$\frac{P_Y^2(y_h)}{P_Y^1(y_h)} \exp\{\frac{\hat{Q}(\eta, y_h) - v(1-\eta, y_h)}{\theta}\} = \frac{P_Y^2(y_l)}{P_Y^1(y_l)} \exp\left\{\frac{\hat{Q}(\eta, y_l) - v(1-\eta, y_l)}{\theta_1}\right\}$$

The sufficient conditions on probabilities ensure that there exists a root in $\eta = \bar{\eta}$. Now starting with $v = \hat{v}(1-\bar{\eta}, y)$ we have that $Q(\hat{v}, y) = \hat{Q}(\bar{\eta}, y)$ or stationary value function solves the Bellman equation with different approximating models. Thus we can recover the derivatives of $Q(v, y)|_{v=\hat{v}}$ by using the stationary value functions.

$$\frac{\partial Q(v, y)}{\partial v}|_{v=\hat{v}} = \frac{\hat{Q}_\eta(\eta, y)}{\hat{v}_\eta(1-\eta, y)}|_{\eta=\bar{\eta}}$$

This gives $\bar{\lambda} = \frac{\bar{\eta}^{-\gamma}}{(1-\bar{\eta})^\gamma}$

■

Proposition 9:

Proof. The proof will construct arguments that show that $\lambda_{t+s}|\lambda_t$ is a P^0-sub martingale bounded above by $\bar{\lambda}$. The martingale convergence theorem will guarantee that λ_{t+s} converges to λ^*. We then rule of the possibilities that $\lambda^* \neq \bar{\lambda}$. Note that assumption (30c) restricts an endogenous policy rule. It can presumably be weakened by more restrictions on $\frac{y_h}{y_l}$, but for the purpose of this proof, we take it as given. The corollary shows that it is satisfied locally around $\bar{\lambda}$. It can be checked that it is also satisfied as $\lambda \to 0$. We begin with a sequence of lemmas

First we show that $\bar{\lambda}$ exists

Lemma 11 *There exists a $\bar{\lambda}$ such that $\lambda_{t+k} = \bar{\lambda}$ implies $\lim_{t+k} \lambda_{t+k} = \bar{\lambda}$ almost surely.*

Proof. This follows from Proposition 8. ■ We note a useful monotonicity property of the Lagrange multipliers.

Lemma 12 *$\lambda_{t+1}[y_t|\lambda_t]$ is monotonic in λ_t*

Proof. In this proof use the fact that value function $\mathcal{Q}(v, y)$ is concave and decreasing in promised values. For a given v, the envelope condition implies a monotonic relationship between λ and v. Exploiting this they can be used interchangeably and a higher value implies a favorable situation for Agent 2. Given a $\lambda(v, y)$, 18a solves for the current endowment is split between the two agents. The optimal policy for $v^*[y^*|v, y]$ can be described as a solution to the following auxiliary problem Let $\mathbb{T}^i$ represent the $\mathbb{R}^i_{\theta_2}\mathbb{T}_{\theta_1,m,y}$ applied to $\pi^i(m)$ that is degenerate but differs across agents.

$$\mathcal{W}(\lambda) = \max_{v^*(y^*)} \mathbb{T}^1 Q[v^*(y^*), y^*] + \lambda \mathbb{T}^2 v^*(y^*) \tag{71}$$

We first show that $\mathcal{W}$ is convex in λ. Consider a λ', λ'' and $\alpha \in (0, 1)$. Let $v'(y^*), v''(y^*)$ be the associated solutions. Define

$$\lambda^* = \alpha\lambda' + (1-\alpha)\lambda''$$

similarly, $v^*(y^*)$ be the solution to $\mathcal{W}(\lambda^*)$. From the respective optimality, we see that

$$\begin{aligned}\mathcal{W}(\lambda') &\geq \mathbb{T}^1 Q[v^*(y^*), y^*] + \lambda \mathbb{T}^2 v^*(y^*) \\ \mathcal{W}(\lambda'') &\geq \mathbb{T}^1 Q[v^*(y^*), y^*] + \lambda \mathbb{T}^2 v^*(y^*)\end{aligned} \tag{72}$$

Multiplying the above equations by α , $1-\alpha$, we can see that

$$\alpha \mathcal{W}(\lambda') + (1-\alpha)\mathcal{W}(\lambda') \geq \mathcal{W}(\alpha\lambda' + (1-\alpha)\lambda')$$

The envelope theorem applied to this problem implies $\mathcal{W}'(\lambda) = \mathbb{T}^2 v^*$ and convexity implies the later is increasing in λ. Now consider the FOC with respect to v^*

$$\exp\left\{\theta^{-1}\mathbb{T}^1 Q^*\right\}\lambda(y^*) = \lambda \frac{P_Y^2(y^*|y)}{P_Y^1(y^*|y)} \exp\left\{\frac{Q(y^*) - v(y^*)}{\theta}\right\} \frac{1}{\exp\left\{-\theta^{-1}\mathbb{T}^2 v^*\right\}}$$

Suppose $\lambda^*[y^*]$ was decreasing in λ. The last term on the RHS are increasing in λ. A lower $\lambda^*[y^*]$ would imply that spread in value $Q^* - v^*$ would be higher. Note that

$$\mathcal{W}'(\lambda) = \frac{\partial \mathbb{T}^1 Q^*}{\partial \lambda} + \lambda \frac{\partial \mathbb{T}^2 v^*}{\partial \lambda} + \mathbb{T}^2 v^*$$

Substituting $\mathcal{W}'(\lambda) = \mathbb{T}^2 v^*$, we have

$$\frac{\partial \mathbb{T}^1 Q^*}{\partial \lambda} < 0$$

Thus the FOC cannot be satisfied if $\lambda^*[y^*]$ was decreasing in λ. ■

Lemma 13 *The policy rules for λ are ordered $\lambda^*(\lambda, y_l) \leq \lambda \leq \lambda^*(\lambda, y_h)$ for $\lambda \leq \bar{\lambda}$*

Proof. Wee show that as $\lambda \to 0$, the respective policy rules are ordered. Then by uniqueness of steady state in the region $(0, \bar{\lambda})$, we can be sure that this ordering in the limit is uniform on the aforementioned region. If not, the policy rules would switch and there would exist another steady state that is strictly less than $\bar{\lambda}$. In the proof of proposition 7 we argued that,

$$\lim_{\lambda \to 0} \frac{\lambda^*(y)}{\lambda} = \frac{P^2(y)}{P^1(y)\bar{z}_1(y)}$$

The condition (30b) then implies that

$$\lim_{\lambda \to 0} \frac{\lambda^*(\lambda, y_l)}{\lambda^*(\lambda, y_h)} < 1$$

■

Now let $z_{1,t}(y_{t+1})$ be the one period ahead distortion to conditional likelihood of the shock y_{t+1}.

$$z_{1,t}(y_l) = \frac{\tilde{P}^1(y_l|v_t, y_t)}{P^0(y_l)}$$

Lemma 14 *For all $\lambda_t \leq \bar{\lambda}$, $\mathbb{C}_t^0[z_{1,t}(y_{t+1}), \lambda_t(y_{t+1})] \leq 0$*

Proof. we already showed that $\lambda_t(y_{t+1})$ is ordered. By virtue of the restriction (30c), we have $z_{1,t}(y_l) \geq 1$. Since $z_{1,t}$ is a Radon-Nikodym derivative and has mean 1. We also have

$$z_{1,t}(y_l) \geq 1 \geq z_{1,t}(y_l)$$

Thus both $z_{1,t}(y_{t+1})$ and $\lambda_t(y_{t+1})$ are ordered for the possible realization of shocks y_{t+1} and we can sign the covariance. ■

The FOC (18b) imply that λ_{t+1} is a $\tilde{P}_t^1$ martingale. Thus,

$$\tilde{\mathbb{E}}_t^1 \lambda_{t+1} = \lambda_t$$

or

$$\mathbb{E}^0 z_{1,t+1} \lambda_{t+1} = \lambda_t$$

using the formula for covariance, we have

$$\mathbb{E}^0 z_{1,t+1} \mathbb{E}^0 \lambda_{t+1} = \lambda_t - \mathbb{C}_t^0(z_{1,t+1}, \lambda_{t+1})$$

Since $\tilde{P}_t^1(y_l) \geq P^0(y_l)$ and z_1 is ordered.

$$\mathbb{E}^0 z_{1,t+1} \leq \tilde{\mathbb{E}} z_{1,t+1} = 1$$

Combining this with 14 we have

$$\mathbb{E}^0 \lambda_{t+1} \leq \lambda_t$$

we now show that it is bounded above by the steady state value

Lemma 15 $\lambda_{t+s} | \lambda_t \leq \bar{\lambda}$

Proof.

Suppose not, then there exists s such that $\lambda_s < \bar{\lambda} < \lambda_{s+1}$. In lemma 12 we showed that $\lambda(\lambda_t, y)$ is monotonic in λ. It is easy to see that it is also continuous. Thus

$$\lambda^*(\lambda_s, y_{s+1}) \leq \lambda^*(\bar{\lambda}, y_{s+1}) = \bar{\lambda}$$

or

$$\lambda_{s+1} \leq \bar{\lambda}$$

thus we have a contradiction. ■

Now we have a sub martingale that is bounded above by $\bar{\lambda}$. We can appeal to Doob's Convergence theorem and conclude that $\lim_s \lambda_{t+s} = \lambda^*$. Finally we rule out $\lambda^* \neq \bar{\lambda}$

Lemma 16 *The limit* $\lambda^* = \bar{\lambda}$

Proof. The map $\lambda(\lambda_t, y)$ is continuous so the limit λ^* must satisfy $\lambda(\lambda^*, y) = \lambda^*$. The only values $\lambda \leq \bar{\lambda}$ that have this property is $\lambda^* = 0$ and $\lambda^* = \bar{\lambda}$. With proposition 7, restriction 30a is sufficient to rules out $\lambda^* = 0$. ■

The first corollary can be proved by arguing that 30c is satisfied locally near $\bar{\lambda}$. First note that at $\lambda = \bar{\lambda}$,

$$Q(\lambda^*(\lambda, y_l), y_l) < Q(\lambda^*(\lambda, y_h), y_h)$$

Since the shares are constant eventually , as long as the aggregate endowment is strictly ordered, we have the continuation values are ordered too.

$$z_1(\bar{\lambda}, y) \propto \exp\{-\theta^{-1}Q(\bar{\lambda}, y)\}$$

By continuity of $z_1$there exists $\epsilon > 0$ such that

$$z_t(y_l) > 1 > z_t(y_l) \quad \forall \|\lambda_t - \bar{\lambda}\| \leq \epsilon$$

we can apply the previous convergence arguments to conclude that $\lim_k \lambda_{t+k} = \bar{\lambda}$ for all $\bar{\lambda} - \epsilon \leq \lambda_t \leq \bar{\lambda}$ Thus λ_t is locally stable

The second corollary requires us to show that $\frac{P_Y^2(y_l)}{P_Y^2(y_h)} \geq \frac{P_Y^0(y_l)}{P_Y^0(y_h)}$ is sufficient for assumption 30c to be satisfied globally in the required region for v

We observe that assumption 30c is true if we can order the continuation values of Agent 1, in particular $Q[v^*(y_h), y_h] \geq Q[v^*(y_l), y_l]$.

Suppose this is not true and $Q[v^*(y_h), y_h] < Q[v^*(y_l), y_l]$. This would imply

$$Q[v^*(y_l), y_h] > Q[v^*(y_l), y_l] > Q[v^*(y_h), y_h],$$

or $v^*(y_l) < v^*(y_h)$

Lemma 13 shows that $\lambda^*(y_l) < \lambda^*(y_h)$ when v such that $\lambda(v, y) < \bar{\lambda}$. Thus the first order conditions of the planning problem 18b imply that,

$$1 > \frac{\lambda^*(y_l)}{\lambda^*(y_h)} = \left(\frac{P_Y^2(y_l)}{P_Y^2(y_h)}\right)\left(\frac{P_Y^0(y_h)}{P_Y^0(y_l)}\right) \exp\left\{\frac{Q[v^*(y_l), y_l] - Q[v^*(y_h), y_h]}{\theta_1}\right\} \exp\left\{\frac{v^(y_h) - v^*(y_l)}{\theta_1}\right\}$$

Since both the terms involving continuation values are greater than zero, we also have

$$1 > \left(\frac{P_Y^2(y_l)}{P_Y^2(y_h)}\right)\left(\frac{P_Y^0(y_h)}{P_Y^0(y_l)}\right)$$

The sufficient condition in the corollary gives us a contradiction.

■

Proposition 10

Let $\mathcal{T}^s$ be the time such that learning converges and $\pi_{i,t}(m^*) = 1$.

Since $\mathcal{M}$ is finite it is easy to see that $P^0(\mathcal{T}^s < \infty) = 1$. With IES > 1 utilities are bounded below. We can apply the arguments in proof of proposition 7 to show that the one period ahead growth rate of of $\log(\lambda_{t+1})$ is bounded and λ_t cannot converge to zero (or diverge to infinity). Now we show that for $s > \mathcal{T}^s$, $\lambda_{t+s}|\lambda_s$ is a supermartingale if $\lambda_s > 1$ and vice versa.

Summing up the FOCs 18b we get that

$$\tilde{\mathbb{E}}_t^1 \lambda_{t+1} = \lambda_t$$

This implies

$$\mathbb{E}_t \lambda_{t+1} = \lambda_t - Cov_t\left[\lambda_{t+1}, z_{t,t+1}^1\right]$$

The sub or supermartingale nature of λ_t now depends on the sign of the covariance terms. We will sign it in the limit as $\lambda_t \to \infty$ and then argue that it cannot switch signs in the intermediate region $(1,\infty)$. Switching the agents, with symmetry properties switch in the region $(0,1)$ and hence it is without loss of generality to consider what happens in any of this subsets.

The second term in the covariance is the relative pessimism of Agent 1. This is decreasing with respect to y_{t+1}. The next lemma shows that $\lambda_{t+1}(y_{t+1})$ is increasing (decreasing) with respect y_{t+1} for $\lambda_t > 1$ when IES $>$ ($<$)1

Lemma 17 *If IES $>$ ($<$) 1 we have $\lambda_{t+1}(y|\lambda) > (<)\lambda_{t+1}(y'|\lambda)$ if $y > y'$*

Proof. We first show that $\lambda_{t+1}(y|\lambda) = \lambda_{t+1}(y'|\lambda)$ implies $\lambda_{t+1} = 1$. With IID shocks we have

$$Q_{t+1}(y) - Q_{t+1}(y') = u[c_1(\lambda_{t+1}(y), y)] - u[c_1(\lambda_{t+1}(y'), y')](1-\delta)$$

and symmetrically,

$$v_{t+1}(y) - v_{t+1}(y') = u[c_2(\lambda_{t+1}(y)), y] - u[c_2(\lambda_{t+1}(y'), y')](1-\delta)$$

Note that from the FOC 18b,

$$\frac{\lambda_{t+1}(y)}{\lambda_{t+1}(y')} = 1 \text{ iff } \frac{\exp\{\theta_1^{-1}[v_{t+1}(y) - v_{t+1}(y')]\}}{\exp\{\theta_1^{-1}[Q_{t+1}(y) - Q_{t+1}(y')]\}}$$

This implies,

$$u[c_1(\lambda_{t+1}(y), y)] - u[c_2(\lambda_{t+1}(y), y)] = -u[c_1(\lambda_{t+1}(y'), y')] + u[c_2(\lambda_{t+1}(y'), y')]$$

With constant elasticity of substitution, the above expression simplifies to

$$\left(\eta_{1,t+1}(y) - \eta_{2,t+1}(y)\right)\left(u(y) - u(y')\right) = \left(\eta_{2,t+1}(y') - \eta_{1,t+1}(y')\right)\left(u(y) - u(y')\right)$$

If λ_{t+1} is constant across these shocks, we have

$$\eta_{1,t+1}(y) = \eta_{2,t+1}(y) = \frac{1}{2}$$

or equivalently $\lambda_{t+1} = 1$

Thus the policy rules for $\lambda_{t+1}(y)$ can only cross at $\lambda_t = 1$. This implies if we can order them at extremes they will be ordered uniformly on either side of unity. From the FOC,

$$\lim_{\lambda_t \to \infty} \frac{\lambda_{t+1}(y)}{\lambda_{t+1}(y')} = \frac{\exp\{-\theta^{-1} v^{max}(y)\}}{\exp\{-\theta^{-1} v^{max}(y')\}}$$

Now $v^{max}(y) - v^{max}(y') = u(y) - u(y')$ and depending on the IES, we can order $\lim_{\lambda_t \to \infty} \lambda_{t+1}(y)$. The previous discussion allows us to extend the ordering uniformly to $(1,\infty)$. ■

Thus the covariance is positive in the region $\lambda_t > 1$ and the process is a supermartingale. It is bounded below by 1 due to the monotonicity properties i.e $\frac{\partial \lambda_t+1(y_{t+1},\lambda_t)}{\partial \lambda_t} > 0$[36]. Thus by the martingale converge theorem it converges almost surely to 1.

When IES < 1, we can repeat the arguments to show that it is supermartingale in the region $(0,1)$. Thus $\lambda_{\mathcal{T}^s} < 1$ would mean $\lim_t \lambda_{t+\mathcal{T}^s}$ is either 0 or 1 (or by symmetry, if $\lambda_t > 1$ it either

[36]This was proved in Lemma 12

goes to 1 or ∞). The location of λ_{T^s} depends on the sequence of shocks and hence can be either $>$ or <1. Thus we can conclude that $P^0(\lambda_t = 0, \lambda_t = \infty) > 0$.

Let $\mathcal{T}^w$ be the first time when agents agree weakly. As before, $P^0(\mathcal{T}^w < \infty) = 1$. With IES $= 1$, the covariance is zero and Pareto weights are constant when agents priors converge. We now show that the boundaries $\{0, \infty\}$ cannot be approached in finite time unless $\lambda_0 = 0$ or $\lambda_0 = \infty$.

Suppose there exists a history y^T such that $\lambda_T(y^T) = 0$ for $T < \infty$. The continuation allocation from y^T will have $c_{2,T+s} = 0$ and $v_{T+s} = -\infty$ for all $s \geq 0$. With this continuation allocation, the promise keeping at $T-1$ will be violated if $v_{T-1}(y^{T-1}) < -\infty$. By backward induction, this means that $v_0 = \infty$ or $\lambda_0 = 0$.

Proposition 11:

Proof.

We break the proof into two steps. First we show that both constraints cannot bind and then we use this to characterize how the contract varies current and future utilities Let λ be the multiplier of the promise keeping constraint (32a) and $\mu(s)\tilde{P}^2(s)\lambda$ be the multiplier on the incentive compatibility constraint (32b) for state s. We can summarize the first order necessary conditions as follows

$$\frac{\tilde{P}^1(s)C'(u_2(s))}{s\tilde{P}^2(s)C'(u_1(s))} = \lambda\left[1 + \mu(s) - \mu(s')\frac{s'\tilde{P}^2(s')}{s\tilde{P}^2(s)}\right] \tag{73a}$$

$$\frac{\lambda(s, y^*)}{\lambda} = \left[\frac{\tilde{P}^2(s)\tilde{P}^2(y^*|s)}{\tilde{P}^1(s)\tilde{P}^1(y^*|s)}\right]\left[1 + \mu(s) - \mu(s')\frac{\tilde{P}^2(s')\tilde{P}^2(y^*|y, s')}{\tilde{P}^2(s)\tilde{P}^2(y^*|s)}\right] \tag{73b}$$

Presence of asymmetric information introduces a wedge in the allocation as compared to the first best. This measures how much current consumption shares deviate for a given λ from the complete market solution

$$W(s) = \left[\mu(s) - \mu(s')\frac{s'\tilde{P}^2(s')}{s\tilde{P}^2(s)}\right]$$

Lemma 18 *Let $\mu(s)\tilde{P}^2(s)\lambda$ be the multipliers on the IC constraints in state s. We cannot have $\mu(s) > 0$ for both s*

Proof. Suppose there exists (v^0, y) such that both IC's bind. This impels the contract satisfies $u_i(s_l) = u_i(s_h)$ and $\mathbb{T}_{\theta_1,y}\bar{v}(s_l, y*) = \mathbb{T}_y\bar{v}(s_h, y*)$. Next we show that if $\mathbb{T}_{\theta_1,y}\bar{v}(s_l, y*) = \mathbb{T}_{\theta_1,y}\bar{v}(s_h, y*)$, one can without loss of generality choose $\bar{v}(s_l, y^*) = \bar{v}(s_h, y^*)$.

Note that the value of the objective function at this contract can be written as

$$Q(v^0, y) = (1-\delta)u_1 + \delta\mathbb{T}_{\theta_1,y,y^*}\mathbb{T}_{\theta_1,s|y^*}Q(\bar{v}(s, y^*), y^*) \tag{74}$$

Note that,

$$\mathbb{T}_{\theta_1,s|y^*}Q(\bar{v}(s, y^*), y^*) < \mathbb{E}_{\theta_1,s|y^*}Q(\bar{v}(s, y^*), y^*) < Q(\mathbb{E}_{\theta_1,s|y^*}\bar{v}(s, y^*), y^*) < Q(\mathbb{T}_{\theta_1,s|y^*}\bar{v}(s, y^*), y^*) \tag{75}$$

The first inequality comes from the fact that the operator $\mathbb{T}_{\theta_1}$ solves a minimization problem. The second inequality is applying Jensen's inequality. The third inequality comes from the fact that Q is decreasing in v.

Thus, $\bar{v}(s_l, y^*) = \bar{v}(s_h, y^*) = \mathbb{T}_{\theta_1, s|y^*}\bar{v}(s, y^*)$ is feasible and satisfies incentives maintaining the property that both IC's bind.

The FOC's with respect to $u_i(s)$ evaluated at this contract yield

$$\frac{\tilde{P}^1(s)}{s\lambda\tilde{P}^2(s)} = 1 + \mu(s) - \mu(s')\frac{s'\tilde{P}^2(s')}{s\tilde{P}^2(s)} \quad \forall s \tag{76}$$

Similarly the FOCs with respect to $\bar{v}(s, y^*)$ give us

$$\sum_{y^*} \lambda[s, y^*]\tilde{P}^1(s, y^*) = \lambda\tilde{P}^2(s)\left[1 + \mu(s) - \mu(s')\frac{\tilde{P}^2(s')}{\tilde{P}^2(s)}\right] \tag{77}$$

Since $\bar{v}(s_l, y^*) = \bar{v}(s_h, y^*)$ and the value function is concave, $\lambda(s_l, y^*) = \lambda(s_h, y^*)$. Further, the valuations for agent one given s are equal. Hence $\tilde{P}^1(y^*|s_l) = \tilde{P}^1(y^*|s_h)$.

we can substitute for $\lambda\tilde{P}^2(s)$ from equation (76) to obtain

$$s\left[1 + \mu(s) - \mu(s')\frac{s'\tilde{P}^2(s')}{s\tilde{P}^2(s)}\right]\sum_{y^*}\lambda(y^*)\tilde{P}^1(y^*|s) = \left[1 + \mu(s) - \mu(s')\frac{\tilde{P}^2(s')}{\tilde{P}^2(s)}\right] \tag{78}$$

Evaluating and dividing the previous expression for $s = s_l$ by $s = s_h$, we get

$$\frac{\left[1 + \mu(s_l) - \mu(s_h)\frac{s_h\tilde{P}^2(s_h)}{s_l\tilde{P}^2(s_l)}\right]}{\left[1 + \mu(s_h) - \mu(s_l)\frac{s_l\tilde{P}^2(s_l)}{s_h\tilde{P}^2(s_h)}\right]} > \frac{\left[1 + \mu(s_l) - \mu(s_h)\frac{\tilde{P}^2(s_h)}{\tilde{P}^2(s_l)}\right]}{\left[1 + \mu(s_h) - \mu(s_l)\frac{\tilde{P}^2(s_l)}{\tilde{P}^2(s_h)}\right]} \tag{79}$$

However

$$\left[1 + \mu(s_l) - \mu(s_h)\frac{s_h\tilde{P}^2(s_h)}{s_l\tilde{P}^2(s_l)}\right] < [\left[1 + \mu(s_l) - \mu(s_h)\frac{\tilde{P}^2(s_h)}{\tilde{P}^2(s_l)}\right]$$

and and

$$\left[1 + \mu(s_h) - \mu(s_l)\frac{s_l\tilde{P}^2(s_l)}{s_h\tilde{P}^2(s_h)}\right] > \left[1 + \mu(s_h) - \mu(s_l)\frac{\tilde{P}^2(s_l)}{\tilde{P}^2(s_h)}\right]$$

Thus combining both we have a contradiction. ■

The previous lemma states that both constraints cannot bind. Therefore either $IC(s_l) > 0, IC(s_h) = 0$ or $IC(s_h) > 0, IC(s_l) = 0$. Adding IC's for either of the possibilities gives us

$$u_2(s_l)(s_l - s_h) > u_2(s_h)(s_l - s_h)$$

or

$$u_2(s_l) < u_2(s_h)$$

and evaluating the equality gives us

$$\mathbb{T}_{\theta_1, y}\bar{v}(s_l, y^*) > \mathbb{T}_{\theta_1, y}\bar{v}(s_h, y^*)$$

Thus the optimal contract features $u_2(s_l) < u_2(s_h)$ and $\mathbb{T}_{\theta_1, y}\bar{v}(s_l, y^*) > \mathbb{T}_{\theta_1, y}\bar{v}(s_h, y^*)$

■

Proposition 12:

Proof. We first begin with the case with $\theta_1 = \infty$. Lemma 3 implies that λ_t is a P^0-martingale. The KKT necessary conditions imply that $\lambda_t \geq 0$. Thus it is martingale bounded below by zero for t. We can apply Doob's martingale convergence theorem to argue that $\lambda_t \to \lambda^\infty$ almost surely.

Next we argue that it can only converge to its lower bound i.e zero. Suppose not, and for some path $\lambda(y^t, s^t) \to \hat{\lambda} > 0$. Proposition 4 shows that with $\theta_1 = \infty$, $\lambda^*(s, y^*) = \lambda^*(s, y^{**})$. Using the concavity of the value function, we have a monotonic relationship between $\lambda(v^0, y)$ and the promised values such that $v(\lambda(v^0), y) = v^0$

Note that

$$\lambda(y^t, s^t) = \lambda^*[s_t|v(\lambda_{t-1}), y^t)$$

Thus the convergence of λ_t implies $v_t \to \hat{v}(y^t)$ and $\lambda^*[s|\hat{v}(y), y)] = \hat{\lambda}$ for all s.

The FOCs 73b imply that $\lambda^* = \lambda$ imply that the dynamic wedge $[1 + \mu(s) - \mu(s')]$ converges to 1 and the limiting multipliers are zero. The allocation converges to the complete market solution.

The FOC's 73a imply consumption shares for agent 2 vary across s. In particular $u_2(s_l|\hat{v}(y), y) < u_2(s_h|\hat{v}(y), y)$

Now we compute the one period ahead conditional expectation of the continuation values $\mathbb{E}_t\bar{v}(y^{t+1}, s^t)$. In the limit this will converge to the random variable $\mathbb{E}_y\hat{v}$. One can observe that this is independent of s_t.

The ex-post values $v(s, y)$ evidently solve the following equation

$$v(s, y|v^0, y) = (1 - \delta)su_2(s|v^0, y) + \delta\mathbb{E}_y\bar{v}(s, y^*|v^0, y)$$

Let $\Delta(y|v^0, y) = v(s_h, y) - v(s_l, y)$ given by

$$\Delta(y|v^0, y) = (1 - \delta)\left[s_h u_2(s_h) - s_l u_2(s_l)\right] + \delta(\mathbb{E}_y\bar{v}s_h, y^* - \mathbb{E}_y\bar{v}s_l, y^*)$$

Convergence of $\mathbb{E}_t\bar{v}(y^{t+1}, s^t)$ implies

$$\Delta_t = (1 - \delta)\left[s_h u_2(s_h|\hat{v}(y), y) - s_l u_2(s_l|\hat{v}(y), y)\right]$$

Adding $(s_l - s_h)u_2(s_h|\hat{v}(y), y)$ to both sides we gets

$$\Delta_t + (s_l - s_h)u_2(s_h|\hat{v}(y_t), y_t) = (1 - \delta)s_l(u_2(s_h) - u_2(s_l)) > 0$$

This violates the incentive compatibility in state $s = s_l$. Thus we have a contradiction.

Next we turn to the case with $\theta_1 < \infty$ taking logs, the FOC 73b imply that the growth rate of λ_{t+1} is given by

$$\log \frac{\tilde{P}_t^2(y_{t+1}|s_t)}{\tilde{P}_t^1(y_{t+1}|s_t)} + \log \frac{\tilde{P}_t^1(s_t)}{\tilde{P}_t^2(s_t)} + \log(1 + \mu_t(s_t) - \mu_t(\tilde{s}_t)\frac{\tilde{P}^2(s)}{\tilde{P}^1(s)}$$

The proof for survival follows exactly as Proposition 7 if on the paths $\lambda_t \to 0$ we have

- $\lim_t \tilde{P}^i(y_{t+1}|s_t) = \tilde{P}_t^i(y_{t+1})$
- $\lim_t \log \frac{\tilde{P}^1(s_t)}{\tilde{P}_t^2(s_t)} + \log(1 + W_t) = 0$

Following similar arguments as in Proposition 7 one can show that $\lambda_t \to$ iff $v_t \to v_{min}$.

At $v = v_{min}$, the contract $u_2(s) = v_{min}$ and $\bar{v}(s, y^*) = v_{min}$ is incentive compatible and implements the complete markets solution and is hence optimal.

It is easy to see that consumption of agent 1 is independent of s in the limit and hence $\lim_t \tilde{P}^1(y_{t+1}|s_t) = \tilde{P}^1_t(y_{t+1})$.

Let $a(v^0, y)$ be the allocation and associated endogenous beliefs that solves 31. Define a continuous function $\kappa(s|a, \lambda)$ as

$$\kappa(s|a, \lambda) = \frac{C'[u_2(s)]}{\lambda s C'[u_1(s)]} \frac{\tilde{P}^1(s)}{\tilde{P}^2(s)} - 1 \tag{80}$$

Associated with the sequence λ_t that goes to zero, let a^{cm}_t be the complete market allocation and beliefs $a^{cm}(v^(\lambda_t), y_t)$. Continuity with respect to a implies that for every $\epsilon > 0$, there exists $\delta > 0$ such that

$$\|a - a'\| < \delta \implies \|\kappa(s|a, \lambda) - \kappa(s|a', \lambda)\| < \epsilon$$

Since the private information allocation converges to the complete market solution we have $\|a_t - a^{cm}_t\| \to 0$. Thus there exists a T large enough that

$$\|a_t - a^{cm}_t\| < \delta \quad \forall t > T$$

and

$$\|\kappa(s|a_t, \lambda) - \kappa(s|a^{cm}_t, \lambda)\| < \epsilon$$

At $\lambda = \lambda_t$ we have $\kappa(s|a^{cm}_t, \lambda_t) = 0$ since the probability weighted marginal utilities are equal to the time t (implied)Pareto weights on agent 2 λ_t

$$\|\kappa(s|a_t, \lambda_t)\| < \epsilon$$

Using the FOC 73a we have

$$\mu_t < \epsilon$$

The rest of the proof follows from Proposition 7. ■

B Bond Economy

This section describes the bond economy without aggregate risk : $\mathcal{Y} = \{\bar{y}\}$ and normalizing aggregate supply of bonds to zero. The shocks s instead of affecting the tastes of Agent 2 affect the share of endowments. The beliefs of the agents are given by initial priors on a finite set of Markov models for s denoted by $\mathcal{M} = \{P_S(s^*|s, m)\}_m$ and $\pi^0_i \in \Delta(\mathcal{M})$. The endowments of Agent 1 and Agent 2 are ys, $y(1-s)$ respectively[37].

With the zero aggregate supply of bonds and common initial priors - π_0, the sufficient state variables for this economy are (B^1, s, π) :- The assets of Agent 1, current realization of the distributional shock and the common prior over the set of models $\mathcal{M}$

[37]Note that s is not an *idiosyncratic* shock as all agents of type i have the same endowment. This allows us to aggregate symmetric decisions over individual types of agent and keep track only of how wealth is distributed across types. A full extension to a Huggert/ Aiyagari economy is studied in Bhandari (2013b).

Agents problem

Given bond prices q, let $Q^i(b, B^1, s, \pi)$ be the value of Agent i with assets b and aggregate state (B^1, s, π)

$$\mathcal{Q}^1(b, B^1, s, \pi) = \max_{c,b^*} u(c) + \delta \mathbb{R}_{\theta_2} \mathbb{T}^1_{\theta_1,m,s} \mathcal{Q}^1(b^*, B^{1*}, s^*, \pi^*) \tag{81}$$

subject to

$$c + qb^* = ys + b \tag{82a}$$

$$\pi^* \propto \sum_m P_S(s^*|s, m)\pi(m) \tag{82b}$$

$$b^* \geq \underline{\mathrm{b}}^1(s) \tag{82c}$$

Where $\underline{\mathrm{b}}^1(s)$ is the natural debt limit for Agent 1 in state s

Similarly we can describe Agent 2's problem as

$$\mathcal{Q}^2(b, B^1, s, \pi) = \max_{c,b^*} \left[u(c) + \delta \mathbb{R}^2_{\theta_2} \mathbb{T}^1_{\theta_1,m,s} \mathcal{Q}^2(b, B^{1*}, s^*, \pi^*)\right] \tag{83}$$

subject to

$$c + qb^* = y(1 - s) + b \tag{84a}$$

$$\pi^* \propto \sum_m P_S(s^*|s, m)\pi(m) \tag{84b}$$

$$b^* \geq \underline{\mathrm{b}}_2(s) \tag{84c}$$

Where $\underline{\mathrm{b}}^2(s)$ is the natural debt limit for Agent 2 in state s.

Remark 6 *Along equilibrium paths market clearing will impose upper limits on asset positions of individual agents too*

Equilibrium

Given q, the interior solutions to these problems pin down the consumption savings decisions of both agents. Let $\mathcal{B}^i[b, B^1, \pi, s, q]$ be the savings of Agent i.

$$\mathcal{B}^1(b, B^1, s, \pi, q) : qu_c[ys + b - qb^*] - \delta \tilde{\mathbb{E}}^1_s Q^1_b(b^*, B^{1*}, s^*, \pi^*) = 0 \tag{85a}$$

$$\mathcal{B}^2(b, B^1, s, \pi, q) : qu_c[y(1 - s) + b - qb^*] - \delta \tilde{\mathbb{E}}^2_s Q^2_b(b^*, B^{1*}, s^*, \pi^*) = 0 \tag{85b}$$

The expectations are taken with respect to the worst case model averaged marginals $\sum_m \tilde{\pi}_i(m)\tilde{P}^i_S(s^*|s, m)$. Like before $\tilde{\pi}^i$ and $\tilde{P}^i_S$ can be computed using the value functions Q^i for each agent.

B.1 A Minimally Stochastic case

We first analyze the equilibrium under special dynamics for s which reduces the problem essentially to a 2 period version that can be quickly solved and study how wealth differences affect the worst case beliefs of agents. The simple economy is constructed under the following dynamics for $\{s_t\}_{t>0}$ given s_0

1. $s_1|s_0 \sim P_S[s^*|s,m]$
2. $s_{t+1} = s_t$ for $t \geq 1$

This is a *minimally stochastic case* which features an absorbing state for s_t from $t = 1$. The value of the agent can now be computed backwards - Let $\mathcal{Q}^{i*}$ be the value of Agent i from period 1 onwards and $\mathcal{Q}^{i0}$ denote the value after s_0 has been realized. The stationary environment after period 1 implies that

- $\mathcal{Q}^{1*}[b, B^1, s, \pi] = \frac{u[ys+(1-\delta)b]}{1-\delta}$ and $\mathcal{B}^{1*}[b, B^1, s, \pi] = b$
- $\mathcal{Q}^{2*}[b, B^1, s, \pi] = \frac{u[y(1-s)+(1-\delta)b]}{1-\delta}$ and $\mathcal{B}^{2*}[b, B^1, s, \pi] = b$
- $q(B^1, s, \pi) = \delta$

Now we can derive the objects in $t = 0$ using the above as terminal conditions. The following proposition states that there exist an inverse relationship between assets and the weights that agents give to states when they have low income.

Let $z_1(B^1, s, \pi) = \frac{\sum_m \tilde{\pi}^1(m)\tilde{P}^1(s|s_0,m)}{\sum_m \pi(m)P_S(s|s_0,m)}$ be Agent 1's (equilibrium) worst case likelihood ratio.

Proposition 13 *There exists $\bar{y}(b)$ and $\underline{B}^1_{-1,0}[s,\pi]$ such that $\lim_{b\to\underline{B}^1_{-1,0}} \mathcal{B}^{1,0}(b, \underline{B}^1_{-1,0}, s, \pi, q) = -\frac{ys_l}{1-\delta}$ and $\lim_{b\to\underline{B}^1_{-1,0}} \bar{y}[\mathcal{B}(b, \underline{B}^1_{-1,0}, s, \pi, q)] = ys_l$. Further we have,*

$$\frac{\partial z_1(B^1, s, \pi)}{\partial B^1} > 0 \quad \textit{iff} \quad y_1(s) > \bar{y}(b)$$

as long as we have $ys_l < \bar{y}[\mathcal{B}(B^1, s, \pi, q)] < ys_h$

Proof. We first derive some properties of how distortions to priors depend on wealth

Lemma 19 *Suppose consumption was given by $c(y, b) = y + b$ and z solves*

$$V^R(b) = \min_{z, Ez=1} \mathbb{E}z[u(c) + \theta \log(z)]$$

For every b there exists a threshold $\bar{y}(b)$ such that $\frac{\partial m(y,b)}{\partial b} > 0$ iff $y > \bar{y}(b)$

Proof. The choice for z^*

$$z^*(y, b) \propto \exp\left\{\frac{-(y + b)^{1-\gamma}}{\theta(1-\gamma)}\right\}$$

taking logs and differentiating with respect to b we have

$$\frac{\partial \log z^*(y, b)}{\partial b} = -\frac{(y + b)^{-\gamma}}{\theta} + \frac{\mathbb{E}\exp\left\{-\frac{(y+b)^{1-\gamma}}{(1-\gamma)\theta}\right\}(y + b)^{-\gamma}}{\theta_1 \mathbb{E}\exp\left\{-\frac{(y+b)^{1-\gamma}}{(1-\gamma)\theta}\right\}}$$

Define $\tilde{p}(y) = p(y)\frac{\exp\left\{-\frac{(y+b)^{1-\gamma}}{(1-\gamma)\theta}\right\}}{\mathbb{E}\exp\left\{-\frac{(y+b)^{1-\gamma}}{(1-\gamma)\theta}\right\}}$ we have

$$\frac{\partial d\log z^*(y,b)}{\partial b} = -\frac{(y+b)^{-\gamma} - \tilde{\mathbb{E}}(y+b)^{-\gamma}}{\theta}$$

Let $\bar{y}(b)$ be such that the numerator is zero

$$\bar{y}(b) = \left(\tilde{\mathbb{E}}(y+b)^{-\gamma}\right)^{-\frac{1}{\gamma}} - b$$

Since $y + b \geq 0$ as $y > \bar{y}(b)$ we have

$$\frac{\partial d\log z^*(y,b)}{\partial b} > 0$$

■

Lemma 20 *There exists a* $\underline{B}^1_{-1,0}[s,\pi]$ *such that* $\lim_{b\to\underline{B}^1_{-1,0}} \mathcal{B}^{1,0}(b,\underline{B}^1_{-1,0},s,\pi,q) = -\frac{ys_l}{1-\delta}$. *Further we also have*

$$\lim_{b\to\underline{B}^1_{-1,0}} \bar{y}[\mathcal{B}(b,\underline{B}^1_{-1,0},s,\pi,q)] = ys_l$$

As $\mathcal{B}^{1,0}$ approaches $-\frac{ys_l}{1-\delta}$, marginal utility of consumption of Agent 1 in $s^* = s_l$ diverges to ∞. For an interior solution, the FOC would require his current consumption to go to zero as well. This means that $\underline{\mathrm{b}}_{-1,0}[s,\pi]$ will satisfy

$$\underline{\mathrm{b}}_{0,1} \approx q\frac{ys_l}{1-\delta} - ys$$

and from Agent 2's FOC along with market clearing we have that q is

$$q \approx \delta\frac{\tilde{\mathbb{E}}Q^{2*}_b(\frac{ys_l}{1-\delta}, -\frac{ys_l}{1-\delta}, s^*, \pi^*)}{u_c(y)}$$

This suggests $\underline{\mathrm{b}}_{-1,0}[s,\pi] = \delta\left(\frac{\tilde{\mathbb{E}}Q^{2*}_b(\frac{ys_l}{1-\delta},-\frac{ys_l}{1-\delta},s^*,\pi^*)}{u_c(y)}\right)\left(\frac{ys_l}{1-\delta}\right) - ys$

Following steps in lemma 19, the threshold for Agent 1's income to ensure that relative optimism rises with assets satisfies

$$\bar{y}[b^*] = \left(\check{\mathbb{E}}^1_s[ys^* + b^*(1-\delta)]^{-\gamma}\right)^{-\frac{1}{\gamma}} - (1-\delta)b^*$$

Note that the likelihood ratio $m(b^*,s^*) = \frac{\sum_m \tilde{\pi}^1(m)\tilde{P}^1_S(s^*|s,m)}{\sum_m \pi(m)P_Z(s^*|s,m)}$ The numerator can be simplified to

$$\left(\frac{\exp\{\frac{-u[ys^*+b^*(1-\delta)]}{\theta_1}\}}{\sum_m \exp\{\frac{-\delta\mathbb{T}^1_{\theta_1,m}[u(ys^*+b^*(1-\delta))]}{\theta_2}\}}\right)\sum_m \pi(m)P_S(s^*|s,m)F^1(m)$$

and $F^1(m) = \exp\left\{\left(\frac{\theta_2-\delta\theta_1}{\theta_1\theta_2}\right)\mathbb{T}^1_{\theta_1,m}\left[u(ys^*+b^*(1-\delta))\right]\right\}$

The derivative $\frac{\partial\log[z(b^*,s^*)]}{\partial b^*}$ is given by

$$-\left(\frac{1-\delta}{\theta_1}\right)[ys^*+b^*(1-\delta)]^{-\gamma}+(1-\delta)\sum_m \tilde{E}^1_{m,s}[ys^*+b^*(1-\delta)]^{-\gamma}\left(\frac{\delta}{\theta_2}\tilde{\pi}^1(m)+\left(-\frac{\delta}{\theta_2}+\frac{1}{\theta_1}\right)\hat{\pi}^*(m)\right)$$

where $\hat{\pi}^*(m) \propto \pi(m)P_S(s^*|s,m)F^1(m)$

Multiplying by $\frac{\theta_1}{1-\delta}$, we can define $\breve{\pi}^*$ as $\frac{\delta\theta_1}{\theta_2}\tilde{\pi}^1+\left(1-\frac{\delta\theta_1}{\theta_2}\right)\hat{\pi}^*$.

Now $\frac{\partial \log[z(b^*,s^*)]}{\partial b^*} \geq 0$ if and only if $ys_h \geq \bar{y}[b^*] \geq ys_l$

$$\bar{y}[b^*] = \left(\breve{\mathbb{E}}^1_s[ys^*+b^*(1-\delta)]^{-\gamma}\right)^{-\frac{1}{\gamma}} - (1-\delta)b^*$$

Now as $b \to \underline{\mathrm{B}}^1_{0,-1}[s,\pi]$ and $b^* = \mathcal{B}^{1,0}[b,s,\pi] \to -\frac{ys_l}{1-\delta}$

$$\bar{y}\left(\frac{-ys_l}{1-\delta}, s^*\right) = ys_l$$

As long as we have $ys_l < \bar{y}[\mathcal{B}(B^1,s,\pi,q)] < ys_h$ implies a negative association of assets levels and pessimism. ■

References

[1] E W Anderson, Thomas J Sargent, and Lars P Hansen, *A quartet of semigroups for model specification, robustness, prices of risk, and model detection*, Journal of the European Economic Association **1** (2003), no. 1, 68–123.

[2] Evan W Anderson, *The dynamics of risk-sensitive allocations*, Journal of Economic Theory **125** (2005), no. 2, 93–150.

[3] Evan W Anderson, Lars Peter Hansen, and Thomas J Sargent, *Small noise methods for risk-sensitive/robust economies*, Journal of Economic Dynamics & Control **36** (2012), no. 4, 468–500.

[4] A Atkeson and Robert E Lucas, *On efficient distribution with private information*, The Review of Economic Studies **59** (1992), no. 3, 427–453.

[5] DK Backus, BR Routledge, and SE Zin, *Exotic preferences for macroeconomists*, NBER Macroeconomics Annual 2004, ... **19** (2005), no. April.

[6] Ravi Bansal and Amir Yaron, *Risks for the Long Run: A Potential Resolution of Asset*, The Journal of Finance **LIX** (2004), no. 4, 1481–1509.

[7] Francisco Barillas, Lars Peter Hansen, and Thomas J Sargent, *Doubts or variability?*, Journal of Economic Theory **144** (2009), no. 6, 2388–2418.

[8] H Bhamra and R Uppal, *Asset prices with heterogeneity in preferences and beliefs*, AFA 2010 Atlanta Meetings Paper (2010).

[9] Anmol Bhandari, *Risk sharing with model uncertainty in economies without aggregate fluctuations.*

[10] L Blume and D Easley, *If you're so smart, why aren't you rich? Belief selection in complete and incomplete markets*, Econometrica **74** (2006), no. 4, 929–966.

[11] J Borovička, *Survival and long-run dynamics with heterogeneous beliefs under recursive preferences*, (2011).

[12] Nina Boyarchenko, *Ambiguity shifts and the 20072008 financial crisis*, Journal of Monetary Economics (2012).

[13] M Brunnermeier, A Simsek, and W Xiong, *A welfare criterion for models with distorted beliefs*, (2012), no. January.

[14] Andrea Buraschi and Alexei Jiltsov, *Model Uncertainty and Option Markets with Heterogeneous Beliefs*, **61** (2013), no. 6, 2841–2897.

[15] M Cagetti, L P Hansen, T Sargent, and N Williams, *Robustness and pricing with uncertain growth*, Review of Financial Studies **15** (2002), no. 2, SI, 363–404.

[16] Claudio Campanale, *Learning , Ambiguity and Life-Cycle Portfolio*, (2009), 1–54.

[17] Hui Chen, N Ju, and J Miao, *Dynamic asset allocation with ambiguous return predictability*, 2009.

[18] Timothy Cogley, Riccardo Colacito, Lars Peter Hansen, and Thomas J Sargent, *Robustness and US Monetary Policy Experimentation*, Journal of Money and Credit **40** (2008), no. 8, 1599–1623.

[19] Timothy Cogley, T J Sargent, and V Tsyrennikov, *Speculation and wealth when investors have diverse beliefs and financial markets are incomplete*, (2012), 1–22.

[20] Timothy Cogley and Thomas J Sargent, *Diverse Beliefs, Survival and the Market Price of Risk*, The Economic Journal (2009), no. 1963, 1–30.

[21] Timothy Cogley, Thomas J Sargent, and Viktor Tsyrennikov, *Wealth Dynamics in a Bond Economy with Heterogeneous Beliefs*, The Economic Journal (2013), n/a—-n/a.

[22] Riccardo Colacito and M M Croce, *Risks for the long run and the real exchange rate*, Journal of Political Economy **119** (2011), no. 1, 153–181.

[23] ———, *International Robust Disagreement*, American Economic Review **102** (2012), no. 3, 152–155.

[24] Fabrice Collard, S Mukerji, Kevin Sheppard, and JM Tallon, *Ambiguity and the historical equity premium*, (2011), no. 2010.

[25] Scott Condie and JV Ganguli, *Ambiguity and rational expectations equilibria*, The Review of Economic Studies (2011), 1–63.

[26] Persi Diaconis and David Freedman, *On the Consistency of Bayes Estimates*, Annals of Statistics **14** (1986), no. 1, 1–26.

[27] Larry G. Epstein and Jianjun Miao, *A two-person dynamic equilibrium under ambiguity*, Journal of Economic Dynamics and Control **27** (2003), no. 7, 1253–1288.

[28] Larry G Epstein and Martin Schneider, *Ambiguity , Information Quality , and Asset Pricing*, **LXIII** (2008), no. 1, 197–228.

[29] Emmanuel Farhi and Iván Werning, *Inequality and social discounting*, Journal of Political Economy **115** (2007), no. 3, 365–402.

[30] Milton Friedman, *Essays in Positive Economics*, 1953.

[31] Itzhak Gilboa and Massimo Marinacci, *Ambiguity and the Bayesian Paradigm*, 2011.

[32] Mikhail Golosov, *Optimal indirect and capital taxation*, The Review of Economic Studies **70** (2003), no. 3, 569–587.

[33] Robert E. Hall, *Intertemporal Substitution in Consumption*, Journal of Political Economy **96** (1998), no. 2, 339–357.

[34] L P Hansen, P Maenhout, A Rustichini, T J Sargent, and M M Siniscalchi, *Introduction to model uncertainty and robustness*, Journal of Economic Theory **128** (2006), no. 1, 1–3.

[35] L P Hansen and T J Sargent, *Acknowledging misspecification in macroeconomic theory*, Review of Financial Studies **4** (2001), no. 3, 519–535.

[36] ———, *Robust estimation and control under commitment*, Journal of Economic Theory **124** (2005), no. 2, 258–301.

[37] L P Hansen, T J Sargent, G Turmuhambetova, and N Williams, *Robust control and model misspecification*, Journal of Economic Theory **128** (2006), no. 1, 45–90.

[38] L P Hansen, T J Sargent, and N E Wang, *Robust permanent income and pricing with filtering*, Macroeconomic Dynamics **6** (2002), no. 1, 40–84.

[39] Lars Peter Hansen, Ricardo Mayer, and Thomas Sargent, *Robust hidden Markov LQG problems*, Journal of Economic Dynamics & Control **34** (2010), no. 10, 1951–1966.

[40] Lars Peter Hansen and Thomas J Sargent, *Recursive robust estimation and control without commitment*, Journal of Economic Theory **136** (2007), no. 1, 1–27.

[41] ———, *Robustness*, 2007.

[42] ———, *Fragile beliefs and the price of uncertainty*, Quantitative Economics **1** (2010), no. 1, 129–162.

[43] ———, *Robustness and ambiguity in continuous time*, Journal of Economic Theory **146** (2011), no. 3, 1195–1223.

[44] LP Hansen and KJ Singleton, *Generalized instrumental variables estimation of nonlinear rational expectations models*, Econometrica **50** (1982), no. 5, 1269–1286.

[45] JM Harrison and DM Kreps, *Speculative investor behavior in a stock market with heterogeneous expectations*, The Quarterly Journal of Economics (1978).

[46] M James and M Campi, *Nonlinear discrete-time risk-sensitive optimal control*, Journal of Robust and Nonlinear Control **6** (1996), no. April 1993, 1–19.

[47] Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji, *A smooth model of decision making under ambiguity*, Econometrica **73** (2005), no. 6, 1849–1892.

[48] TA Knox, *Foundations for learning how to invest when returns are uncertain*, 2003.

[49] N Kocherlakota and Christopher Phelan, *On the robustness of laissez-faire*, Journal of Economic Theory (2009), no. August.

[50] Leonid Kogan and SA Ross, *The price impact and survival of irrational traders*, The Journal of Finance **LXI** (2006), no. 1.

[51] FE Kydland and EC Prescott, *Dynamic optimal taxation, rational expectations and optimal control*, Journal of Economic Dynamics and Control **2** (1980), 79–91.

[52] Robert E Lucas and Nancy L Stokey, *Optimal growth with many consumers*, Journal of Economic Theory **32** (1984), no. 1, 139–171.

[53] F Maccheroni, M Marinacci, and A Rustichini, *Ambiguity aversion, robustness, and the variational representation of preferences*, Econometrica **74** (2006), no. 6, 1447–1498.

[54] Fabio Maccheroni, Massimo Marinacci, and Aldo Rustichini, *Dynamic variational preferences*, Journal of Economic Theory **128** (2006), no. 1, 4–44.

[55] J Miao and Ju Nengjiu, *Ambiguity, Learning, and Asset Returns*, Econometrica **80** (2012), no. 2, 559–591.

[56] Lubos Pastor and Pietro Veronesi, *Learning in financial markets*, 2009.

[57] Christopher Phelan, *On the long run implications of repeated moral hazard*, Journal of Economic Theory **191** (1998), no. 1998, 174–191.

[58] ———, *Opportunity and social mobility*, Review of Economic Studies (2006), 1–22.

[59] WP Rogerson, *Repeated moral hazard*, Econometrica **53** (1985), no. 1, 69–76.

[60] Alvaro Sandroni, *Do markets favor agents able to make accurate predictions?*, Econometrica **68** (2000), no. 6, 1303–1341.

[61] Leonard J. Savage, *The Foundations of Statistics*, 1954.

[62] J Scheinkman and Wei Xiong, *Overconfidence and speculative bubbles*, Journal of Political Economy **111** (2003), no. 6, 1183–1219.

[63] Ani Sciubba, Emanuela and Guerdjikova, *Survival with Ambiguity*, (2012).

[64] A Simsek, *Belief Disagreements and Collateral Constraints*, Econometrica **81** (2013), no. 1, 1–53.

[65] C Sleet and S Yeltekin, *Misery and Luxury: Long Run Outcomes with Private Information*, (2010), no. February 2009, 1–40.

[66] Christopher Sleet and S Yeltekin, *Social patience, social credibility and long run inequality*, 2005.

[67] Tomasz Strzalecki, *Axiomatic Foundations of Multiplier Preferences*, Econometrica **79** (2011), no. 1, 47–73.

[68] ———, *Temporal Resolution of Uncertainty and Recursive Models of Ambiguity Aversion*, Econometrica **81** (2013), no. 3, 1039–1074.

[69] J Thomas and T Worrall, *Income fluctuation and asymmetric information: An example of a repeated principal-agent problem*, Journal of Economic Theory **390** (1990), 367–390.

[70] Paolo Vanini, Fabio Trojanib, and Paolo Vaninia, *Learning and Asset Prices under Ambiguous Information*, (2004), no. October.