NBER WORKING PAPER SERIES

THE MARKET PRICE OF AGGREGATE RISK AND THE WEALTH DISTRIBUTION

Hanno Lustig

Working Paper 11132
http://www.nber.org/papers/w11132

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2005

This project was started while Lustig was a graduate student at Stanford. I would like to thank my advisors, Robert Hall, Thomas Sargent and Dirk Krueger, for their help and encouragement. Lars Hansen and Stijn Van Nieuwerburgh provided detailed comments. V.V. Chari, John Cochrane, Hal Cole, John Heaton, Narayana Kocherlakota, Lee Ohanian, Pierre-Olivier Weill, Amir Yaron and the members of the Sargent reading group at Stanford, the joint Chicago-Stanford macro workshop and the UCLA macro workshop all helped me along. I would also like to thank seminar participants at Chicago, LBS, Harvard, Minnesota, MIT, Northwestern, NYU, Princeton, UCLA, University of Texas at Austin and on the RES tour. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

The Market Price of Aggregate Risk and the Wealth Distribution
Hanno Lustig
NBER Working Paper No. 11132
February 2005
JEL No. G0

ABSTRACT

I introduce bankruptcy into a complete markets model with a continuum of ex ante identical agents who have power utility. Shares in a Lucas tree serve as collateral. The model yields a large equity premium, a low risk-free rate and a time-varying market price of risk for reasonable risk aversion. Bankruptcy gives rise to a second risk factor in addition to aggregate consumption growth risk. This liquidity risk is created by binding solvency constraints. The risk is measured by one moment of the wealth distribution, which multiplies the standard Breeden-Lucas stochastic discount factor. This captures the aggregate shadow cost of the solvency constraints. The economy is said to experience a negative liquidity shock when this growth rate is high and a large fraction of agents faces severely binding solvency constraints. These shocks occur in recessions. The average investor wants a high excess return on stocks to compensate for the extra liquidity risk, because of low stock returns in recessions. In that sense stocks are "bad collateral". The adjustment to the Breeden-Lucas stochastic discount factor raises the unconditional risk premium and induces time variation in conditional risk premia.

Hanno Lustig
Department of Economics
UCLA
Box 951477
Los Angeles, CA 90095-1477
and NBER
hlustig@econ.ucla.edu

I. Introduction

I develop a model of an exchange economy with a continuum of agents, complete markets, but imperfect enforcement of contracts. Because households can declare themselves bankrupt and escape their debts, they face endogenous solvency constraints that restrain their resort to the bankruptcy option. In the benchmark calibration, the risk associated with these solvency constraints delivers an equity premium of 6 percent, a risk-free rate of one percent and substantial variation in the Sharpe ratio over the business cycle. This variation is driven by shocks to the wealth distribution induced by these solvency constraints.

This paper follows He and Pearson (1991) and Luttmer (1992) in exploring solvency constraints as a device for understanding asset pricing anomalies. I motivate these constraints by the introduction of bankruptcy. The possibility of bankruptcy constrains the price of an individual's consumption claim to exceed the shadow price of a claim to his labor income in all states of the world. The fraction of the economy's endowment yielded by the Lucas tree plays a key role in my economy. If the labor share of aggregate income is one, all wealth is human wealth, the solvency constraints always bind and there can be no risk sharing. As the fraction of wealth contributed by the Lucas tree increases, risk sharing is facilitated.

An economy that is physically identical but with perfect enforcement of contracts forms a natural benchmark with which to compare my model. Because assets only reflect aggregate consumption growth risk in this benchmark representative agent model (Lucas (1978) and Breeden (1979)), two quantitative asset pricing puzzles arise. These puzzles follow from the fact that aggregate consumption growth in the US is approximately i.i.d. and not volatile. First, risk premia are small for plausible levels of risk aversion (Hansen and Singleton (1982) and Mehra and Prescott (1985)), and second, risk premia do not vary much while they do in the data (see e.g. Campbell and Cochrane (1999)). My model produces an additional risk factor that addresses these puzzles.

Beyond risk in the aggregate endowment process, the bankruptcy technology contributes a second source of risk, the risk associated with binding solvency constraints. I call this liquidity risk. In the simplest case all households have power utility with an identical coefficient of risk aversion γ. In the model without solvency constraints households consume a constant share of the aggregate endowment, governed by fixed Pareto-Negishi weights. In the case of limited commitment these weights increase each time the solvency constraint binds. The average of these increases across households contributes a multiplicative adjustment to the standard Lucas-Breeden SDF (stochastic discount factor): the growth rate of the γ^{-1}-th moment of the distribution of stochastic Pareto-Negishi weights. This component reflects the aggregate shadow cost of the solvency constraints. If this growth rate is high, a large fraction of agents is constrained and the economy is said to be hit by a negative liquidity shock. Beyond this "average weight" growth rate, all other features of the wealth distribution are irrelevant for asset prices.

I propose "endogenous" collateral constraints that bound the household's net wealth in each state tomorrow. In the same environment, an "exogenous" constraint on expected net wealth tomorrow contributes a liquidity factor to the Lucas-Breeden SDF that does not depend on the aggregate shock in the next period; the aggregate cost of this type of constraint raises the market price of consumption by the same amount in all states of the world tomorrow. This liquidity factor lowers the risk-free rate, but the risk premia are virtually identical to those in the representative agent economy and they do not vary over time. The state-contingent nature of the collateral constraints is central to my results, not the tightness of the constraints: there is a lot of equilibrium risk sharing.

The wealth distribution dynamics increase the unconditional volatility of the SDF if negative liquidity shocks occur when aggregate consumption growth is low (recessions). Liquidity shocks in recessions emerge from the properties of the labor income process. If the dispersion of idiosyncratic labor income shocks increases in recessions, households would like to borrow against their income in the "high idiosyncratic states" to smooth consumption but they are not allowed to, because they would walk away in the good state. The labor risk channel has support in the data. Storesletten, Telmer, and Yaron (2004) argue that the conditional standard deviation of labor income shocks more than triples in recessions.

Leading asset pricing models cannot generate enough variation in the Sharpe ratio. Lettau and Ludvigson (2003) call this the Sharpe ratio volatility puzzle. The wealth distribution dynamics of my model endogenously generate more time-variation in the conditional volatility of the SDF than competing equilibrium models (Whitelaw (1994), Campbell and Cochrane (1999), and Barberis, Huang, and Santos (2001)). The liquidity shocks are largest when a recession hits after a long expansion. In long expansions, there is a buildup of households in the left tail of the wealth distribution: more agents do not encounter states with binding constraints and they deplete their financial assets because interest rates are lower than in the representative agent economy. When the recession sets in, those low-wealth agents with high income draws encounter severely binding constraints and the left tail of the wealth distribution is erased. After the recession, the conditional market price of risk decreases sharply. If another recession follows shortly thereafter, the mass of households close to the lower bound on wealth is much smaller and so are the liquidity shocks. This lowers the conditional market price of risk sharply after a low aggregate consumption growth shock. This pattern is broadly consistent with the empirical evidence on the cyclicality of the Sharpe ratio.

Zhang (1997a) first endogenized borrowing constraints in a class of incomplete markets models, extending the work of Aiyagari and Gertler (1991), Telmer (1993), Lucas (1994) and Heaton and Lucas (1996). These incomplete market models introduce history dependence in allocations by restricting the menu of traded assets: all of these papers rule out trade in claims that are contingent on labor income realizations. When ruling out trade in these labor-contingent assets in my model, I show the equilibrium SDF to be the highest expected IMRS across all households, obtained by integrating the IMRS for each household over idiosyncratic outcomes tomorrow, while, in the absence of this ban, the SDF is the highest IMRS across all households, in any idiosyncratic state tomorrow. Not surprisingly, the latter tends to be more volatile, even if the actual IMRS are less so, because the effect of the idiosyncratic shock is not averaged out. Instead of taking the largest realized marginal utility growth as the state price of consumption, incomplete market models force it to be the largest expected marginal utility growth. Reducing the span of traded assets is commonly thought to increase the volatility of the SDF; that logic completely breaks down in this environment with endogenous borrowing constraints.

I follow a different route that does not involve exogenous restrictions on this menu of traded assets, but instead it focuses on the restrictions imposed by the lack of commitment. Alvarez and Jermann (2000) decentralize constrained efficient allocations using solvency constraints and make contact with the literature on risk sharing with limited commitment. My model fits in this tradition, but it brings out the importance of collateralizable wealth. Part of the endowment of my economy is yielded by a tradable Lucas tree; the rest of the endowment is labor income. Instead of sending agents into autarky upon default, as Alvarez and Jermann do, I allow agents to file for bankruptcy (Lustig (2000)). When agents declare bankruptcy, they lose their holdings of the Lucas tree, but all of their current and prospective labor income is protected from creditors. Shares in the Lucas tree serve as collateral.[1]

The continuum of agents in my model contributes important differences vis-à-vis the two-agent model of Alvarez and Jermann (2001). In their model, liquidity shocks are larger but much less frequent. To generate a liquidity shock, one of the agents has to switch from a "low" to a "high" state. The persistence of labor income makes these switches rare. Between these switches this economy looks like the benchmark representative agent economy: households consume a

[1] My model outperforms the representative agent asset pricing model only if the collateralizable wealth ratio is small enough. If this ratio exceeds fifteen percent, perfect risk sharing obtains. If this ratio is lower than five percent, the autarkic outcome obtains. My results are for collateralizable wealth ratios within this window.

constant share of the aggregate endowment and the conditional risk premia are constant. In addition, the correlation between the liquidity shocks and stock returns is small and the high volatility of the SDF does not translate into high excess returns. In the benchmark calibration Alvarez and Jermann (2001) report an equity premium of 3.2 percent, while the market price of risk is close to 1. The economy with a continuum is an average of two-agent economies in each period; the liquidity shocks are more frequent and more tightly correlated with the business cycle. My model delivers an equity premium of 6.1 percent in the benchmark calibration at a lower market price of risk of .45. The aggregate cost of the solvency constraints can vary substantially without excessive consumption volatility at the household level in a model with a continuum of agents, but it cannot in a model with two agents.

There is a large literature on heterogeneity and asset pricing, but most authors conclude that heterogeneity contributes little beyond the standard, representative agent model, except for the work by Constantinides and Duffie (1996). They use an insight from Mankiw (1986) to show how a systematic increase in idiosyncratic risk during recessions can deliver a high equity premium. In most models on asset pricing and heterogeneity, assets are essentially being priced off individual consumption processes. In Constantinides and Duffie's model, any agent's intertemporal marginal rate of substitution (IMRS) is a valid SDF for all payoffs in the next state. Similarly, in models with *exogenous borrowing constraints,* (e.g. He and Modest (1995)) the individual IMRS is a valid SDF for excess returns in all states. So is the cross-sectional average of these individual intertemporal marginal rates of substitution. In the continuous time limit the difference between the average marginal utility and the marginal utility of average consumption is absorbed into the drift (Grossman and Shiller (1982)) and the assets can be priced using the Breeden-Lucas SDF. Campbell (2000) concludes this "limits the effects of consumption heterogeneity on asset pricing".

Not so with *endogenous solvency constraints*: the individual IMRS is a valid SDF for payoffs only in those states in which he is unconstrained (Alvarez and Jermann (2000)). Assets can no longer be priced off individual consumption processes and the Lucas-Breeden discounter does not reappear in the continuous-time limit.

To deal with a continuum of consumers and aggregate uncertainty, I extend the methods developed by Atkeson and Lucas (1992,1995) and Krueger (1999). Atkeson and Lucas show how to compute constrained efficient allocations in dynamic economies with private information problems. Krueger computes the equilibrium allocations in a limited commitment economy without aggregate uncertainty, in which households are permanently excluded upon default. These methods cannot handle aggregate uncertainty.

The use of stochastic Pareto-Negishi weights (Marcet and Marimon (1999)) allows me to state an exact aggregation result in the spirit of Luttmer (1992): equilibrium prices depend only on the γ^{-1}-th moment of the distribution of weights and I extend this result to the case of recursive utility. This reduces the problem of forecasting the multiplier distribution -the state of the economy- to one of forecasting a single moment. The exact forecast requires the entire aggregate history or the distribution of weights. I approximate the actual equilibrium by a stationary, truncated-history equilibrium. The state space is reduced to include only the k most recent aggregate events. The allocation errors -the differences between aggregate consumption and the endowment- are equal to the forecast errors. In the simulation results these errors are very small overall when k is 5.

This paper is organized as follows. The second section of the paper describes the environment. The third section discusses the equilibrium allocations prices, using stochastic Pareto-Negishi weights. This section can be skipped by those not interested in the mechanics of the model. The fourth section discusses the calibration and the computation; the fifth section shows the results. All the technical results, including the propositions and the proofs, are in the appendix.

II. Environment and Equilibrium

A. Uncertainty

The events $s = (y, z)$ take on values on a discrete grid $S = Y \times Z$ where $Y = \{y_1, y_2, \ldots, y_n\}$ and $Z = \{z_1, z_2, \ldots, z_m\}$. y is household specific and z is an aggregate event. Let $s^t = (y^t, z^t)$ denote an event history up until period t. This event history includes an individual event history y^t and an aggregate event history z^t. I will use $s^\tau \geq s^t$ to denote all the continuation histories of s^t. s follows a Markov process such that:

$$\pi(z'|z) = \sum_{y' \in Y} \pi(y', z'|y, z) \text{ for all } z \in Z, y \in Y.$$

I assume a law of large numbers holds such that the transition probabilities can be interpreted as fractions of agents making the transition from one state to another. In addition, I assume there is a unique invariant distribution $\pi_z(y)$ in each state z : by the law of large numbers $\pi_z(y)$ is also the fraction of agents drawing y when the aggregate event is z. $(S^\infty, \mathcal{F}, P)$ is a probability space where S^∞ is the set of all possible histories and P is the corresponding probability measure induced by π.

B. Preferences and Endowments

There is a continuum of consumers of measure 1. There is a single consumption good and it is non-storable. The consumers rank consumption streams $\{c_t\}$ according to the following utility function:

$$U(c)(s_0) = \sum_{t=0}^{\infty} \sum_{s^t \geq s^0} \beta^t \pi(s^t|s_0) \frac{c_t(s^t)}{1-\gamma}^{1-\gamma}, \tag{1}$$

where γ is the coefficient of relative risk aversion.

The economy's aggregate endowment process $\{e_t\}$ depends only on the aggregate event history: $e_t(z^t)$ is the realization at aggregate node z^t. Each agent draws a labor income share $\widehat{\eta}(y_t, z_t)$ as a fraction of the aggregate endowment in each period. Her labor income share only depends on the current individual and aggregate event. $\{\eta_t\}$ denotes the individual labor income process $\eta_t(s^t) = \widehat{\eta}(y, z)e_t(z^t)$,with $s^t = (s^{t-1}, y, z)$. I assume $\widehat{\eta}(y_t, z_t) >> 0$ in all states of the world.

There is a Lucas (1978) tree that yields a non-negative dividend process $\{x_t\}$. The dividends are not storable but the tree itself is perfectly durable. The Lucas tree yields a constant share α of the total endowment, the remaining fraction is the labor income share. By definition, the labor share of the aggregate endowment equals the aggregated labor income shares:

$$\sum_{y' \in Y} \pi_z(y')\widehat{\eta}(y', z') = (1 - \alpha), \tag{2}$$

for all z'. An increase in α translates into proportionally lower $\widehat{\eta}(y, z)$ for all (y, z).

Agents are endowed with initial non-labor wealth (net of endowment) θ_0. This represents the value of this agent's share of the Lucas tree producing the dividend flow in units of time 0 consumption. Θ_0 denotes the initial distribution of wealth and endowments (θ_0, y_0).

C. Market Arrangements

Claims to one's entire labor income process $\{\eta_t\}$ cannot be traded directly while shares in the Lucas tree can be traded. Households can write borrowing and lending contracts based on individual labor income realizations. I use $\phi_t(s^t)$ to denote an agent's holdings of shares in the Lucas tree. In each period households go to securities markets to trade $\phi_t(s^t)$ shares in the tree at a price $p_t^e(z^t)$ and a complete set of one-period ahead contingent claims $a_t(s^t, s')$ at prices $q_t(s^t, s')$. $a_t(s^t, s')$ is a security that pays off one unit of the consumption good if the household draws private shock y' and the aggregate shock z' in the next period with $s' = (y', z')$.

$q_t(s^t, s')$ is today's price of that security. In this environment the payoffs are conditional on an individual event history and the aggregate event history rather than just the aggregate state of the economy.

An agent starting period t with initial wealth $\theta_t(s^t)$ buys consumption commodities in the spot market and trades securities subject to the usual budget constraint:

$$c_t(s^t) + p_t^e(z^t)\phi_t(s^t) + \sum_{s'} a_{t+1}(s^t, s')q_t(s^t, s') \leq \theta_t. \tag{3}$$

If the next period's state is $s^{t+1} = (s^t, s')$, her wealth is given by her labor income, the value of her stock holdings -including the dividends issued at the start of the period- less whatever she promised to pay in that state:

$$\theta_{t+1}(s^{t+1}) = \underbrace{\widehat{\eta}(y_{t+1}, z_{t+1})e_{t+1}(z^{t+1})}_{\textit{labor income}} + \underbrace{\left[p_{t+1}^e(z^{t+1}) + \alpha e_{t+1}(z^{t+1})\right]\phi_t(s^t)}_{\textit{value of tree holdings}} + \underbrace{a_{t+1}(s^{t+1})}_{\textit{contingent payoff}}.$$

Incomplete Markets Net financial wealth $\left[p_{t+1}^e(z^{t+1}) + \alpha e_{t+1}(z^{t+1})\right] + a_{t+1}(s^{t+1})$ depends on the realization of y_{t+1}. In much of the literature on asset pricing with heterogenous agents (e.g. Telmer (1993), Lucas (1994) and Heaton and Lucas (1996) and Zhang (1997b)), this is ruled out ex ante, and agents can only trade claims contingent on the aggregate state tomorrow $\{a_t(s^t, z')\}$. In our simple economy, this trading setup is equivalent to trading shares in the Lucas tree and bonds.

D. Enforcement Technology

In this literature, it has been common to assume that households are excluded from financial markets forever when they default, following Kehoe and Levine (1993) and Kocherlakota (1996). I allow agents to file for bankruptcy. When a household files for bankruptcy, it loses all of its asset but its labor income cannot be seized by creditors and it cannot be denied access to financial markets (see Lustig (2000) for a complete discussion).

Bankruptcy imposes borrowing constraints on households, one for each state:

$$\begin{aligned} \left[p_{t+1}^e(z^{t+1}) + \alpha e_{t+1}(z^{t+1})\right]\phi_t(s^t) &\geq -a_{t+1}(s^t, s') \text{ for all } s' \in S, \\ \text{where } s^{t+1} &= (s^t, s'). \end{aligned} \tag{4}$$

These borrowing constraints follow endogenously from the enforcement technology if we rule out borrowing constraints that are too tight (see Alvarez and Jermann (2000)); these constraints only bind when the participation constraint binds. If the agent chooses to default, her assets and that period's dividends are seized and transferred to the lender. Her new wealth level is that period's labor income:

$$\theta_{t+1}(s^{t+1}) = \widehat{\eta}(y_{t+1}, z_{t+1}) e_{t+1}(z^{t+1}).$$

If the next period's state is $s^{t+1} = (s^t, s')$ and the agent decides not to default, her wealth is given by her labor income, the value of her tree holdings less whatever she promised to pay in that state:

$$\theta_{t+1}(s^{t+1}) = \widehat{\eta}(y_{t+1}, z_{t+1}) e_{t+1}(z^{t+1}) + \left[p^e_{t+1}(z^{t+1}) + \alpha e_{t+1}(z^{t+1})\right] \phi_t(s^t) + a_{t+1}(s^{t+1}).$$

This default technology effectively provides the agent with a call option on non-labor wealth at a zero strike price. Lenders keep track of the borrower's asset holdings and they do not buy contingent claims when the agent selling these claims has no incentive to deliver the goods. The constraints in (4) just state that an agent cannot promise to deliver more than the value of his Lucas tree holdings in any state s'.

Bankruptcy and Permanent Exclusion Two key differences between bankruptcy and permanent exclusion deserve mention. First, the bankruptcy constraints in (4) only require information about the household's assets and liabilities. To determine the appropriate borrowing constraints in the case of permanent exclusion, the lender needs to know the borrower's endowment process and her preferences (Alvarez and Jermann (2000)). This type of information is not readily available and costly to acquire. Moreover, the borrower has an incentive to hide his private information. Second, in the case of bankruptcy it is immaterial whether or not the household actually defaults when the constraint binds. The lender is paid back anyhow and the borrower is indifferent as well. Households could randomize between defaulting and not defaulting when the constraint binds.

These collateral constraints are much tighter than the ones that decentralize the constrained efficient allocations when agents can be excluded from trading (see Section A in the Appendix) and they support less risk sharing as a result.

Exogenous Constraints The state-contingent nature of the constraints is central to my results. A simple example of an "exogenous" solvency constraint would be:

$$p_t^e(z^t)\phi_t(s^t) + \sum_{s'} a_{t+1}(s^t, s')q_t(s^t, s') \geq 0. \tag{5}$$

This constraint checks that the agent's consumption promises tomorrow are covered by his collateral in expectation only, but not necessarily state by state, i.e. for each s'. This constraint is simply derived by summing the endogenous collateral constraints across states of the world (weighted by the state price). These constraints obtain in an environment where agents decide today whether to default tomorrow, regardless of the state of the world. Since the price of any excess return tomorrow is zero today, any household can buy a claim to excess returns tomorrow without violating these exogenous constraints. In this environment, any household's IMRS is a valid SDF, and so is the average IMRS as a result. This implies that household consumption growth is conditionally perfectly correlated, as pointed out Luttmer (1992). Later in the paper, I show, not surprisingly, that these constraints fail to deliver interesting asset pricing results.

Equilibrium Default These constraints do not mean I rule out equilibrium default on some traded securities. I only record what the household will actually deliver in each state instead of what it promises to deliver. The borrowing constraints are a simple way of relabelling promises as "sure things" in all states. I can still price defaultable securities -i.e. a collection of promises in different states. What distinguishes this setup from Geanakoplos and Zame (1998) is the fact that only outright default on all financial obligations is allowed, not default on individual obligations. Kubler and Schmedders (2003) introduce collateral constraints in an incomplete markets setting.

E. Sequential Equilibrium

The definition of equilibrium is standard. Each household is assigned a label that consists of its initial financial wealth θ_0 and its initial state s^0. A household of type (θ_0, s^0) chooses consumption $\{c_t(\theta_0, s^t)\}$, trades claims $\{a_t(s'; \theta_0, s^t)\}$ and shares $\{\phi_t(\theta_0, s^t)\}$ to maximize her expected utility:

$$\max_{\{c\},\{\phi\}\{a\}_{s'}} \sum_{t=0}^{\infty} \sum_{s^t \geq s^0} \beta^t \pi(s^t|s_0) \frac{c_t(s^t)}{1-\gamma}^{1-\gamma}$$

subject to the usual budget constraint:

$$c_t(\theta_0, s^t) + p_t^e(z^t)\phi_t(\theta_0, s^t) + \sum_{s'} a_t(s'; \theta_0, s^t) q_t(s^t, s') \leq \theta_t, \tag{6}$$

and a collection of collateral constraints, one for each state:

$$\begin{aligned} \left[p_{t+1}^e(z^{t+1}) + \alpha e_{t+1}(z^{t+1})\right] \phi_t(\theta_0, s^t) &\geq -a_t(s'; \theta_0, s^t) \text{ for all } s' \in S, \\ \text{where } s^{t+1} &= (s^t, s'). \end{aligned} \tag{7}$$

The definition of a competitive equilibrium is straightforward.

Definition 1. *A competitive equilibrium with solvency constraints for initial distribution* Θ_0 *over* (θ_0, y_0) *consists of trading strategies* $\{a_t(s'; \theta_0, s^t)\}$, $\{c_t(\theta_0, s^t)\}$ *and* $\{\phi_t(\theta_0, s^t)\}$ *and prices* $\{q_t(s^t, s'\}$ *and* $\{p_t^e(z^t)\}$ *such that (1) these solve the household problem (2) the markets clear*

$$\int \sum_{y^t} \frac{\pi(y^t, z^t | y_0, z_0)}{\pi(z^t | z_0)} \left(\sum_{y'} a_t(y', z'; \theta_0, y^t, z^t) \right) d\Theta_0 = 0 \textit{ for all } z^t$$

$$\int \sum_{y^t} \frac{\pi(y^t, z^t | y_0, z_0)}{\pi(z^t | z_0)} \phi_t(\theta_0, s^t) d\Theta_0 = 1 \textit{ for all } z^t$$

To prevent arbitrage opportunities in my economy for unconstrained agents in some state tomorrow, the SDF is set equal to the highest IMRS across all agents:

$$m_{t+1} = \max_{(\theta_0, s^t)} \frac{u'(c_{t+1}(\theta_0, y^{t+1}, z^{t+1}))}{u'(c_t(\theta_0, y^t, z^t))}.$$

This follows immediately from the household's first order condition and the observation that same households with positive measure are unconstrained in each node z^{t+1}.

Incomplete Markets Suppose we ban trade in claims contingent on y_{t+1} and only allow trade in claims on aggregate states z_{t+1}, but we keep all other features of the sequential trading setup. Now, the equilibrium SDF is the highest *average* IMRS across all agents, integrated over all of tomorrow's idiosyncratic states y_{t+1}:

$$m_{t+1}^{INC} = \max_{(\theta_0, s^t)} E \left(\frac{u'(c_{t+1}(\theta_0, y^{t+1}, z^{t+1}))}{u'(c_t(\theta_0, y^t, z^t))}) | z_{t+1}, z_t; \theta_0, s^t \right),$$

because we only have to rule out arbitrage opportunities in trades for claims whose payoffs are contingent on aggregate states, not for those that pay off in idiosyncratic states.[2]

If recessions mainly cause an increase in the cross-sectional standard deviation of consumption growth, it is clear that consumption needs to be *more* volatile in the incomplete markets economy, if it is to match the volatility of the SDF in the complete markets economy: the effect of labor income risk is partly averaged out. [3] So, (1) it is not the case that restricting the span produces more volatile equilibrium SDF's in my environment and (2) there is no direct link between the volatility of the SDF and the amount of risk sharing. [4]

The next section considers a more convenient but perhaps less appealing trading arrangement with markets opening only once.

III. Characterizing Equilibrium Prices and Allocations

To facilitate the analysis, I restate the household problem in a time zero trading environment and I define the analogue to Kehoe and Levine (1993) and Krueger (1999)'s equilibrium concept. Pareto-Negishi weights summarize a household's history of shocks. The stochastic discount factor depends on the growth rate of the $1 \setminus \gamma$-th moment of the weight distribution.

This section can be skipped by the reader who wants to get to the asset pricing results.

A. Solvency Constraints

The collateral constraints in the sequential formulation can be restated as restrictions on the price of two claims. $\Pi_{z^t}[\{d\}]$ denotes the price at node z^t in units of z^t consumption of a claim

[2]This follows immediately from the first order condition for consumption and the observation that some households with positive measure need to be unconstrained.

[3]This argument is carefully developed in Krueger and Lustig (2004). They develop the following example. Suppose the constraints are tightened enough to prevent any risk-sharing, and the equilibrium, autarchic SDF is given by:

$$m_{t+1}^{INC} = \beta \left(\frac{e_{t+1}}{e_t}\right)^{-\gamma} \max_{y_t} E\left(\left(\frac{\widehat{\eta}_{t+1}(y_{t+1}, z_{t+1})}{\widehat{\eta}_t(y_t, z_t)}\right)^{-\gamma} | z_{t+1}, z_t; y_t\right),$$

while its autarchic counterpart in the complete markets case is:

$$m_{t+1} = \beta \left(\frac{e_{t+1}}{e_t}\right)^{-\gamma} \left(\frac{\widehat{\eta}_{t+1}(y_1, z_{t+1})}{\widehat{\eta}_t(y_n, z_t)}\right)^{-\gamma}.$$

If the cross-sectional dispersion of labor income shocks is larger in recessions z_{re}, the complete markets SDF is much more volatile than its incomplete markets counterpart.

[4]The complete market model does produce more risk sharing by construction, but not necessarily more volatile SDF's

on $\left\{d_t(s^t)\right\}_{t=0}^{\infty}$. The collateral constraints are equivalent to the following restriction on the price of two claims, one on consumption and one on labor income:

$$\Pi_{s^t}\left[\{c\}\right] \geq \Pi_{s^t}\left[\{\eta\}\right], \text{ for each } s^t. \tag{8}$$

These solvency constraints keep net wealth non-negative in all states of the world. If these constraints are satisfied in all states, households do not wish to exercise their option to default. This is shown in section A in the appendix.

The amount of collateralizable wealth plays a key role. When there is no collateralizable wealth, the solvency constraints bind for all agents in all states of the world and households are in autarky. If the constraint did not bind for one set of households with positive measure, it would be have to be violated for another one with positive measure. Section B in the appendix derives this result.

If there is sufficient collateralizable wealth, then the solvency constraint is satisfied for each (y, z) at perfect-insurance (Breeden-Lucas) prices

$$\Pi_z^*\left[\{e\}\right] \geq \Pi_{y,z}^*\left[\{\eta\}\right] \text{ for all } (y, z),$$

and perfect risk sharing is attainable. If this condition is satisfied, each household can sell a security that replicates its labor income and buy an equivalent claim to the aggregate dividends stream that fully hedges the household.

B. Kehoe-Levine Equilibrium

This section sets up the household's primal problem and defines an equilibrium, when all trading occurs at time zero. Taking prices $\left\{p_t(s^t|s_0)\right\}$ as given, the household purchases history-contingent consumption claims subject to a standard budget constraint and a sequence of solvency constraints, one for each history:

Primal Problem (PP)

$$\sup_{\{c\}} u(c_0(\theta_0, s^0)) + \sum_{t=1} \sum_{s^t \geq s^0} \beta^t \pi(s^t|s_0) u(c_t(\theta_0, s^t)),$$

$$\sum_{t \geq 0} \sum_{s^t \geq s_0} p_t(s^t|s^0)\left[c_t(\theta_0, s^t) - \eta_t(s^t)\right] \leq \theta_0,$$

$$\Pi_{s^t}\left[\left\{c\left(\theta_0, y^t, z^t\right)\right\}\right] \geq \Pi_{s^t}\left[\{\eta\}\right], \text{ for all } s^t \in S^t, t \geq 0.$$

The solvency constraints keep the households from defaulting. The following definition of equilibrium is in the spirit of Kehoe and Levine (1993) and in particular Krueger (1999).

Definition 2. *For given initial state* z_0 *and for given distribution* Θ_0, *an equilibrium consists of prices* $\{p_t(s^t|s_0)\}$ *and allocations* $\{c_t(\theta_0, s^t)\}$ *such that*

- for given prices $\{p_t(s^t|s_0)\}$, the allocations solve the household's problem PP (except possibly on a set of measure zero),
- markets clear for all t, z^t :

$$\sum_{y^t} \int c_t(\theta_0, y^t, z^t) d\Theta_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} = e_t(z_t). \tag{9}$$

In equilibrium households solve their optimization problem subject to the participation constraints and the markets clear.

If interest rates are high enough, the economy with sequential trading is equivalent to an economy in which all trading occurs at time zero subject to these solvency constraints. Alvarez and Jermann (2000) derive the exact condition on interest rates in an economy with finitely many agents and Krueger (1999) adjusts it to an economy with a continuum of agents.

C. Equilibrium Prices

The liquidity risk induced by the wealth distribution shocks interacts with aggregate consumption growth risk to modify the SDF's properties in the right direction. These dynamics are key to understanding the SDF in this model:

$$m_{t+1} = \beta \left(\frac{e_{t+1}}{e_t}\right)^{-\gamma} \left(\frac{h_{t+1}}{h_t}\right)^{\gamma}. \tag{10}$$

The first part is the Breeden-Lucas SDF that emerges in a representative agent economy. The second part is the multiplicative adjustment of the SDF that summarizes the shocks to the wealth distribution induced by the solvency constraints; it is the aggregate shadow cost of the solvency constraints.

The next subsection makes use of Pareto-Negishi weights as a device for characterizing equilibrium allocations and prices. These weights encode the wealth distribution dynamics that are central to my results. I am not solving a planner's resource allocation problem, but I characterize equilibrium allocations and prices from the household's first order conditions. The details are in the appendix in section C.

D. Solvency Constraints and Stochastic Pareto-Negishi Weights

These solvency constraints introduce a stochastic element in the consumption share of each household. The household's wealth at time 0, θ_0, determines its initial Pareto-Negishi weight μ_0. This weight μ_0 governs the share of aggregate consumption allocated to this household in all future states of the world s^t. Φ_0 is the joint measure over initial states and multipliers (μ_0, s_0). When there are no solvency constraints, this share is fixed:

$$c_t\left(\mu_0, s^t\right) = \frac{\mu_0^{1/\gamma}}{E\mu_0^{1/\gamma}} e_t(z^t) \text{ where } s^t = (y^t, z^t), \tag{11}$$

where the constant $E\mu_0^{1/\gamma} = \int \mu_0^{1/\gamma} d\Phi_0$ guarantees market clearing after each aggregate history.

In the presence of solvency constraints, the Pareto-Negishi weights are no longer fixed. I use $\zeta_t(\mu_0, s^t)$ to denote the weight of a household with initial weight μ_0 in state s^t. $\left\{\zeta_t(\mu_0, s^t)\right\}$ is a non-decreasing stochastic process. These weights are constant, unless the household switches to a state with a binding solvency constraint. In these instances the weight increases such that the solvency constraint in (8) is satisfied with equality. Typically, these are states with high labor income realizations. These weights record the sum of all solvency constraint multipliers in history s^t. Section C in the appendix discusses these weight processes in detail.

Consumption is characterized by the same linear risk sharing rule:

$$c_t(\mu_0, s^t) = \frac{\zeta_t^{1/\gamma}(\mu_0, s^t)}{E\left[\zeta_t^{1/\gamma}(\mu_0, s^t)\right]} e_t(z^t), \tag{12}$$

but each household's consumption share is stochastic. Let $h_t(z^t)$ denote this cross-sectional multiplier moment:

$$h_t(z^t) = E\left[\zeta_t^{1/\gamma}(\mu_0, s^t)\right].$$

The average weight process $\left\{h_t(z^t)\right\}$ is a non-decreasing (over time) stochastic process that is adapted to the aggregate history z^t. This process experiences a high growth rate when a large fraction of agents find themselves switching to states with binding constraints. It can be interpreted as the aggregate shadow cost of the solvency constraints.

The risk sharing rule implies that, as long as agents do not switch to a state with a binding solvency constraint, their consumption share drifts downward. When they switch to a state with a binding constraint, their consumption share increases. The rate of decrease is driven by the growth rate of $\{h_t(z^t)\}$ and this growth rate is governed by the wealth distribution dynamics.

In each aggregate state z_{t+1} payoffs are priced off the IMRS of unconstrained agents, whose Pareto-Negishi weight did not change between t and $t+1$. The risk sharing rule for consumption directly implies that his or her IMRS equals the SDF expression in equation (10). The liquidity shocks induced by the solvency constraints are bounded between one and the lowest labor endowment growth rate:

$$1 \leq g_{t+1}(z^{t+1}) \leq \frac{\widehat{\eta}(y_n, z_t)}{\widehat{\eta}(y_1, z_{t+1})}$$

When all households are constrained, the SDF equals the autarchic IMRS of the household switching from the highest to the lowest income state (see section D in the appendix). When none of the households are constrained, their Pareto-Negishi weights are constant. In equilibrium, these liquidity shocks will vary between these bounds depending on the history of aggregate shocks. Obviously, some dispersion in income shares is needed for the liquidity shocks to contribute significant volatility to the SDF.

E. Approximation

A household's Pareto-Negishi weight summarizes its history of private shocks, but obviously not the history of aggregate shocks. In fact, the liquidity shocks depend on the entire history of aggregate shocks.

To compute equilibrium prices and allocations, I propose to keep track of only a truncated version of the aggregate history. This approach is motivated by the limited memory of these economies, if there is sufficient growth in the aggregate weight process. This is borne out by the computations.

Consumption Weights and Cutoff Rule First, I introduce consumption weights as stationary state variables to replace the Pareto-Negishi weights, and, second, I describe the cutoff rule property that characterizes these weights. $g_t(z^t)$ denotes the growth rate of the aggregate weight process h_t/h_{t-1}. For the sake of convenience I renormalize the weights into consumption shares:

$$\omega_t = \frac{\zeta^{1/\gamma}(\mu_0, s^t)}{g_t(z^t)},$$

at the end of each period and store this as this household's state variable. These consumption shares integrate to one by construction and they evolve according to a simple cutoff rule. If the share of a household going into a period is larger than the cutoff value $\underline{\omega}(y', z^t)$, it remains unchanged, else it is increased to its cutoff value:

$$\begin{aligned} \omega'(y', z^t; \omega) &= \omega \text{ if } \omega > \underline{\omega}(y', z^t) \\ &= \underline{\omega}(y', z^t) \text{ elsewhere} \end{aligned} \tag{13}$$

I use Φ_{z^t} to denote the joint measure over (y, ω) in state z^t. Making use of the cutoff rule, the liquidity shock g_{t+1} can be stated as follows:

$$g_t(z', z^{t-1}) = \sum_{y'} \int_{\underline{\omega}(y', z^t)}^{\infty} \omega d\Phi_{z^{t-1}}(dy \times d\omega) \frac{\pi(y', z'|y, z)}{\pi(z'|z)} + \tag{14}$$

$$\sum_{y'} \underline{\omega}(y', z^t) \int_0^{\underline{\omega}(y', z^t)} d\Phi_{z^{t-1}}(dy \times d\omega) \frac{\pi(y', z'|y, z)}{\pi(z'|z)}. \tag{15}$$

It immediately follows that $g \geq 1$. The size of the shock is determined by the mass of households in the left tail.

Using these consumption weights I construct an approximate equilibrium in which agents use only the last k aggregate shocks to forecast g.

Stationary approximating equilibrium. In a stationary equilibrium, there is no probability mass on weights above the highest reservation level. Let L denote the domain for the consumption weights ω. $l(\omega, y', z'; z^k) : L \times Y \times Z \times Z^k \to R$, one for each $(y', z') \in Y \times Z$, gives the new consumption weight for a household entering the period with weight ω, having drawn private shock y' and aggregate shock z'. Its new consumption share is given by:

$$c(\omega, y', z'; z^k) = \frac{l(\omega, y', z'; z^k)}{g^*(z', z^k)},$$

where $g^*(z', z^k)$ is the forecast of the liquidity shock. This consumption share will be stored as the new state variable for this household at the end of the period. The reservation weight policy function $\underline{\omega}(y', z'; z^k) : Z \times Z^k \to R$ and the average weight forecasting function $g^*(z', z^k) : Z^k \to$

R induce the consumption share policy function:

$$\begin{aligned} l(\omega, y', z'; z^k) &= \omega \text{ if } \omega > \underline{\omega}(y', z'; z^k) \\ &= \underline{\omega}(y', z'; z^k) \text{ elsewhere.} \end{aligned} \tag{16}$$

The reservation weights are determined such that the solvency constraints bind exactly. The cost functions $C\left(\omega, y', z'; z^k\right)$ and $C^y\left(y', z'; z^k\right)$ record the price in units of today's consumption of claim to the consumption stream and the labor income stream respectively, scaled by the aggregate endowment today, to keep them stationary. The reservation weights satisfy this functional equation:

$$C\left(\underline{\omega}(y', z'; z^k)\ , y', z'; z^k\right) = C^y\left(y', z'; z^k\right) \text{ for all } \left(y', z'; z^k\right)$$

The optimal forecast when going from state z^k to z' is given by its average for that truncated history:

$$g^*(z', z^k) = E_{z^\infty \subset z^k} g(z', z^\infty), \tag{17}$$

where the actual liquidity shock is given by:

$$g(z', z^\infty) = \sum_{y'} \int l\left(\omega, y', z'; z^k\right) \Phi_{z^\infty}\left(d\omega \times dy\right) \frac{\pi(y', z'|y, z)}{\pi(z'|z)}$$

for each pair (z', z^k). E denotes the expectation operator over all possible histories z^∞ consistent with z^k. The actual measure Φ_{z^∞} depends -possibly- on the entire history of shocks z^∞. The state prices are set using the forecast of the liquidity shock:

$$m(z', z^k) = \beta g^*(z', z^k)^\gamma \lambda(z')^{-\gamma}.$$

Households do not make Euler equation errors, but the markets do not clear exactly. That is the sense in which this equilibrium is approximate. The percentage allocation error is simply the percentage forecast error: $\frac{g(z', z^\infty) - g(z', z^k)}{g(z', z^k)}$. These will turn out to be very small. As $k \to \infty$, the errors tend to zero.

Definition 3. *An approximate stationary equilibrium is fully characterized by a list of functions* $l\left(\omega, y', z'; z^k\right)$, $C\left(\omega, y', z'; z^k\right)$, $C^y\left(y', z'; z^k\right)$ *and* $g(z', z^k)$ *such that (i)* $g(z', z^k)$ *equals the average liquidity shock in* z^k *and (ii)* $l\left(\omega, y', z'; z^k\right)$ *satisfies the optimal policy rule.*

The optimal household consumption policy functions and equilibrium prices are embedded in this information through the risk sharing rule and the expression for the SDF.

Computational Algorithm The algorithm iterates on liquidity shock forecasts:

- The algorithm starts with the perfect insurance growth function $\widehat{g}_1(z^k, z') = 1$ for all (z^k, z') .[5]
- Conditional on this function, I compute the cost functions $C_1(\omega, y', z'; z^k)$, $C_1^y(y', z'; z^k)$ and the policy function $l_1(\omega, y', z'; z^k)$. To do so, I simply determine the cutoff level at which the value of the consumption stream equals the value of the endowment stream: $C_0(\underline{\omega}, y', z'; z^k) = C_0^y(y', z'; z^k)$ for each $(y', z'; z^k)$.
- Next, I simulate a T-period aggregate history $\{z^t\}_{t=0}^{\infty}$ for a cross-section of N agents. I use $T = 10.000$ and $N = 5000$. For each (z^k, z') , I compute the average growth rate $\widehat{g^a}_1(z^k, z')$ implied by the policy function. This provides a new guess $\widehat{g}_2(z^k, z')$ for the weight growth functions.
- Finally, I iterate on the liquidity shock forecasts until $\{\widehat{g}_n(z^k, z')\}$ convergence to $\widehat{g}_*(z^k, z')$. The policy functions and the average weight growth functions characterize a stationary, stochastic equilibrium. The household Euler equations are satisfied exactly by construction. The sup prediction error is exactly the sup percentage allocation error:

$$\varepsilon_k = \sup \left| \frac{g^a(z^k, z') - g_*(z^k, z')}{g_*(z^k, z')} \right| = \sup \left| c^a(z^k, z') - 1 \right| .$$

The allocation error decreases as k is increased.

IV. Calibration and Computation

The first section explains the calibration of the collateralizable wealth fraction and the labor income share dynamics. The second section defines an approximation of the actual equilibrium that can be computed.

[5]This algorithm can be shown to converge as $k \to \infty$. The proof is available upon request.

A. Calibration

The collateralizable wealth share governs the average size of the liquidity shocks, while the labor income dynamics control the time-variation in the size of these shocks.

Collateralizable Wealth The scarcity of collateral is key to generating large and time-varying risk premia in my model. In the next section, I provide a sensitivity analysis by reporting results for $\alpha = 5\%$, $\alpha = 10\%$ and $\alpha = 15\%$.

How do these numbers relate to the US data? The size of the Lucas tree dividend as a share of the total endowment, α, can be matched to the "traded" part of the capital share of net national income. The average labor share of national income in the US between 1946 and 1999 is 70 percent (source, NIPA). An additional 11 percent is proprietor's income derived from farms and partnerships, mainly doctors and lawyers. This should be treated as labor income for the purposes of this exercise. This brings the total labor share to 81 percent. There are two factors that reduce the actual supply of collateralizable wealth. First, a substantial share of the remaining 18 percent are profits from privately held firms. Moskowitz and Vissing-Jorgensen (2002) report that the value of private exceeded the value of public equity in the US until 1995 and part of these profits consists of remuneration for labor services provided by the entrepreneur. In addition, these assets are highly illiquid. Second, bankruptcy exemptions effectively reduce the supply of collateral in the economy.

Aggregate and Idiosyncratic Endowment Risk My calibration is for annual data. The strategy follows Alvarez and Jermann (2001). The moments to be matched are listed in Table I. The first four calibrate the aggregate endowment process. The last six calibrate the labor income process. The standard deviation of the log income shares is .4. The key thing to note is that the cross-sectional dispersion of labor income increases in recessions. This mechanism delivers countercyclical liquidity shocks and Storesletten, Telmer, and Yaron (2004) provide quite some empirical evidence in support of it.

The next subsection discusses the computational details. This can be skipped without loss of continuity.

B. Approximation

To approximate the consumption cost function $C\left(\omega, y', z'; z^k\right)$, I use a Tchebychev polynomial approximation in the consumption weight ω (Judd (1998)). The polynomial is of order 7 and I use

Table I
Alvarez and Jermann Calibration

The first four moments of the four-state Markov process describe US consumption growth dynamics in the 20-th century and are based on Mehra and Prescott (1985). The last six moments in the right panel describe the dynamics of the labor income share process. These combine information from estimation results in Heaton and Lucas (1996) and Storesletten, Telmer and Yaron (2004).

Panel A: Transition matrix				
states (y, z)	l, r	l, e	h, r	h, e
l, r	0.1710	0.8186	0.0040	0.0063
l, e	0.3020	0.5757	0.0172	0.1051
h, r	0.0040	0.0063	0.1710	0.8186
h, e	0.0172	0.1051	0.3020	0.5757
Panel B: Income Shares				
	$\widehat{\eta}_{l,r}$	$\widehat{\eta}_{l,e}$	$\widehat{\eta}_{h,r}$	$\widehat{\eta}_{h,e}$
	0.2048	0.3422	0.7952	0.6578

30 nodes. Table XII in the appendix lists the percentage allocation errors. The approximation works well. The mean of the allocation errors is close to .05 percent for all computations, while the standard deviation is roughly the same size. The low standard deviation of the errors indicates that the errors are tightly distributed around zero. The sup norm is around 2 percent.

V. Asset Pricing Results

This section reports all the asset pricing results. I start with the unconditional moments in the first part; time-variation and predictability in the second part. In the fourth section, I briefly discuss the recursive utility extension. The final part examines the wealth distribution mechanism up close.

Stochastic Discount Factor Absent any arbitrage opportunities, payoffs in state z^{t+1} are priced off the IMRS of the households that are unconstrained in that state. Their consumption share decreases at a rate $g_{t+1} - 1$ between t and $t + 1$. This produces the following SDF:

$$m_{t+1} = \beta \left(\frac{e_{t+1}}{e_t} \right)^{-\gamma} g_{t+1}^{\gamma} \tag{18}$$

This adjustment of the Breeden-Lucas SDF is the aggregate shadow cost of the solvency constraints in that state tomorrow, expressed in units of today's consumption (see section D in the appendix for a derivation). If perfect insurance is feasible, no agent faces a binding solvency constraint and the multiplicative adjustment is 1 exactly. It is always weakly larger than 1 because $\{h_t(z^t)\}$ is a non-decreasing stochastic process and this implies that the market SDF is always weakly larger than the representative agent's valuation of consumption. When the growth rate of the aggregate weight process is larger than average, a large fraction of agents are severely constrained. The multiplicative adjustment of the SDF raises the price of consumption in those states.

The second part g_{t+1} is a function of the entire aggregate history z^{t+1}, including the shock at $t+1$. In the case of exogenous solvency constraints on expected net wealth (see equation 5), a similar aggregation result obtains, but the second part is a function only of z^t (see Luttmer (1992)). In the latter case, each household's IMRS is a valid SDF for excess returns, because these cannot violate the constraints on expected net wealth. The liquidity constraints push up the price of consumption in all states tomorrow. This lowers the risk-free rate, but hardly changes risk premia relative to the full insurance benchmark.

Perfect Insurance Benchmark The model without solvency constraints provides a natural benchmark. Since the consumption shares are constant, the individual intertemporal marginal rates of substitution equal this "aggregate" IMRS. This model is at odds with the data for three main reasons: (1) the model-implied market price of risk is too low relative to observed Sharpe ratios over short holding periods, (2) it is close to constant while the data suggests it should be time-varying and (3) it is even lower relative to Sharpe ratios over longer holding periods. All three failures follow from the properties of aggregate consumption growth. Aggregate consumption growth is not volatile and roughly i.i.d. over time. In addition, the model implies an implausibly large risk-free rate.

A. Results: Equity premium and Risk-free rate

Table VI compares the moments of the data, the representative agent model and the collateral model. My benchmark calibration sets the time discount factor β equal to .95 and γ to 5/7. Equity is a (levered) claim to the aggregate endowment process. The excess return on equity is denoted $R^{l,e}$, while $R^{c,e}$ denotes the excess return on a non-levered claim to the aggregate endowment process. The leverage parameter chosen is four to match the standard deviation of dividend growth of twelve percent. The asset pricing statistics were generated by drawing 10.000 realizations from the model.

The benchmark perfect insurance economy produces a risk-free rate of thirteen percent and an equity premium of three percent, one percent for the non-levered claim to consumption. The

Table II
Benchmark Calibration Results.

$R^{l,e}$ is the return on a leveraged claim to aggregate consumption with a leverage of 4; $R^{c,e}$ is the excess return on a claim to aggregate consumption. The first panel shows moments for the data, the second panel considers the representative agent model and the third panel the collateral model. These moments were generated by averaging 10.000 draws from an economy with 5000 agents. α is 5 percent and β is .95. The CRSP-VW index was used to compute the market return, while the Fama-Bliss risk-free rate was used to compute excess returns.

	$E(R^{l,e})$	$E(R^{c,e})$	$E(r^f)$	$\sigma(R^{l,e})$	$\sigma(R^{c,e})$	$\sigma(r^f)$	$\frac{E(R^{l,e})}{\sigma(R^{l,e})}$	$\frac{E(R^{c,e})}{\sigma(R^{c,e})}$
Sample				*US Data*				
$25-03$	0.078		0.01	0.21		0.04	0.38	
$45-03$	0.084		0.01	0.18		0.03	0.44	
γ				*Representative Agent Economy*				
5	0.03	0.01	0.13	0.17	0.06	0.03	0.19	0.20
7	0.05	0.02	0.15	0.18	0.07	0.05	0.28	0.27
γ				*Collateral Economy*				
5	0.06	0.04	0.03	0.22	0.14	0.10	0.29	0.28
7	0.09	0.06	0.01	0.24	0.15	0.12	0.40	0.39

collateral economy matches the observed US equity premium and misses the risk-free rate by 200 basis points when the fraction of collateralizable wealth is five percent. The liquidity risk induced by the solvency constraints delivers a low risk-free rate of three percent and a high equity premium of about six percent. The compensation per unit of risk is large as well; the Sharpe ratio on the non-levered claim is around thirty-five percent. Increasing the coefficient of risk aversion to seven lowers the risk-free rate to its historical average of around zero percent and increase the Sharpe ratio to thirty-eight percent. The model overstates the volatility of the risk-free rate by a factor of three. This defect will be addressed in section D

Liquidity Shocks To understand these results, we need to understand the effect of these liquidity shocks. First, the liquidity shocks increase the demand for insurance and lower the risk-free rate. This is obvious from the SDF in (18), because $g_t > 1$. The solvency constraints keep the agents from borrowing against their future labor income and the liquidity risk also induces them to save more as a precautionary device.

Second, the liquidity shocks increase the volatility of the SDF if the shocks are negatively correlated with the aggregate consumption growth process. This pattern emerges in equilibrium when a larger fraction of agents is constrained in states with low aggregate consumption growth realization. In general, this need not be the case. In fact, when aggregate consumption growth is i.i.d. and the Markov process for y does not depend on z then the liquidity shocks are constant in the stationary equilibrium. This result is derived in the appendix in proposition 14.

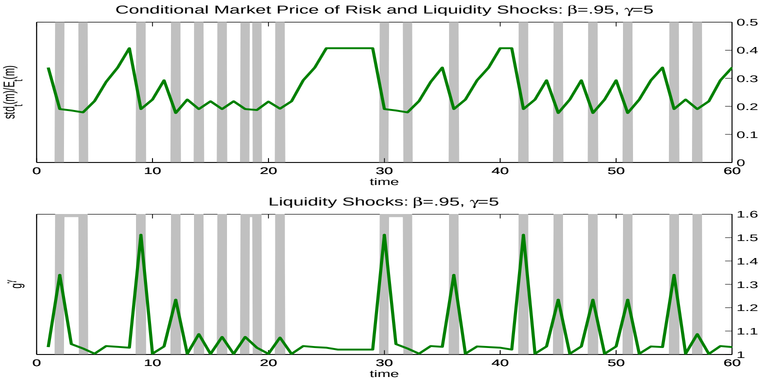

Figure 1. Liquidity Shocks, Market Price of Risk and Aggregate Consumption Growth The upper panel plots the conditional market price of risk and the lower panel plots the liquidity shocks. The shaded area indicates low aggregate consumption growth states. β is .95, γ is 5 and α is .05.

In my calibration liquidity shocks are larger in low aggregate consumption growth states, as shown in the bottom panel of Figure 1, because the increase in the cross-sectional variation of idiosyncratic income shocks increases the size of the consumption increase for households that switch to the high state. During a long series of high aggregate consumption growth realizations, there is a build-up of low wealth households in the left tail of the wealth distribution. Mechanically, this means the mass of agents with weights below the cutoff value is large:

$$\sum_{y'} \int_0^{\underline{\omega}(y',z^t)} d\Phi_{z^{t-1}}\left(dy \times d\omega\right) \frac{\pi(y',z'|y,z)}{\pi(z'|z)} \tag{19}$$

These households have been running down their asset levels as long as they are in low idiosyncratic income states. Their Pareto-Negishi weights remain unchanged throughout. This implies their consumption shares were drifting downwards. When a low aggregate consumption growth state is realized, a larger fraction of households draws a high income state with a high cutoff value $\underline{\omega}(y',z^t)$. This translates into a large liquidity shock as their consumption shares jump up from very low levels (see the definition of the liquidity shock in eq. 15).

Liquidity Premium The increased volatility raises risk premia because returns are low in the low aggregate consumption growth states, when the liquidity shocks are large. I use R^i to denote the return on some risky security. Under joint lognormality of $\Delta\log(e_{t+1}/h_{t+1})$ and

Table III
Summary Statistics for Deterministic Collateral Economy

$R^{l,e}$ is the return on a leveraged claim to aggregate consumption with a leverage of 4; $R^{c,e}$ is the excess return on a claim to aggregate consumption. Aggregate consumption growth in the SDF is deterministic at its mean. The collateralizable share of income α is 5 percent. These moments were generated by averaging 10.000 draws from an economy with 5000 agents .

γ	$E(R^{l,e})$	$E(r^f)$	$\sigma(R^{l,e})$	$\sigma(r^f)$	$\frac{\sigma(m)}{E(m)}$	$\frac{E(R^{c,e})}{\sigma(R^{c,e})}$
5	0.02	0.06	0.20	0.06	0.13	0.08
7	0.03	0.06	0.23	0.10	0.20	0.11

$\log(R^i_{t+1})$ the expected return on asset i is given by:

$$E_t \log R^i_{t+1} - \log r^f_t = \gamma cov_t(\Delta \log(e_{t+1}), R^i_{t+1}) - \gamma cov_t(\log(g_{t+1}), R^i_{t+1})$$

The first part is the standard compensation for consumption growth risk. The second part is the compensation for liquidity risk. This liquidity part accounts for over two thirds of the equity premium in my benchmark calibration.

Deterministic Consumption Growth To highlight the importance of the liquidity shocks, Table III shows the asset pricing moments for an economy without consumption growth uncertainty. Equity is still a claim to the same stochastic aggregate consumption growth process described in the calibration, but the IMRS depends only on the (non-random) average consumption growth rate. Even without consumption growth shocks to the SDF, the equity premium ranges from 2 to 3 percent and the risk-free rate is around 6 percent. This is low considering that we have killed the precautionary effect of aggregate consumption growth risk.

B. Sensitivity Analysis

This section gives on overview of how the asset pricing moments depend on the time discount factor β, the coefficient of risk aversion γ, the collateral share α and the cyclicality of labor income risk. Finally, I briefly contrast these results with those that obtain in different incarnations of my model: a two-agent economy and one with exogenous or non-state-contingent constraints.

Discount Rate Sofar I have kept the annual discount factor β constant at .95. Table IV explores different values of the time discount factor. Lowering the discount rate puts more weight on today's income realization and makes the solvency constraints more binding as a

Table IV
Discount Rate and Risk Aversion

$R^{l,e}$ is the return on a leveraged claim to aggregate consumption with a leverage of 4; $R^{c,e}$ is the excess return on a claim to aggregate consumption. The collateralizable share of income α is .05. These moments were generated by averaging 10.000 draws from an economy with 5000 agents .

γ	β	$E(R^{l,e})$	$E(R^{c,e})$	$E(r^f)$	$\sigma(R^{l,e})$	$\sigma(R^{c,e})$	$\sigma(r^f)$	$\frac{E(R^{l,e})}{\sigma(R^{l,e})}$	$\frac{E(R^{c,e})}{\sigma(R^{c,e})}$
				Panel I: Discount Rate					
5	.85	0.10	0.07	0.05	0.28	0.20	0.17	0.35	0.35
	.90	0.08	0.05	0.04	0.25	0.17	0.13	0.31	0.31
	.95	0.06	0.04	0.03	0.22	0.14	0.10	0.29	0.28
7	.85	0.16	0.13	−0.02	0.33	0.26	0.24	0.49	0.48
	.90	0.13	0.10	−0.00	0.30	0.22	0.19	0.42	0.44
	.95	0.09	0.06	0.01	0.24	0.15	0.12	0.40	0.39
				Panel II: Risk Aversion					
3	.95	0.03	0.01	0.05	0.18	0.09	0.05	0.16	0.15
4		0.04	0.02	0.04	0.20	0.11	0.08	0.22	0.22
5		0.06	0.04	0.03	0.22	0.14	0.10	0.29	0.28
6		0.08	0.05	0.02	0.24	0.16	0.12	0.35	0.33
7		0.09	0.06	0.01	0.24	0.15	0.12	0.40	0.39
8		0.12	0.09	−0.00	0.27	0.19	0.16	0.46	0.45

result. This increases the size of the liquidity shocks considerably. The net effect on the risk-free rate of lowering β is ambiguous, but lower β's always increase the risk premium and the Sharpe ratio at the cost of increasing individual household consumption volatility to less plausible levels. For example, when γ is five and β is .85, the equity premium is 10 percent, but the volatility of household consumption growth is an order of magnitude larger than in my benchmark model. Alvarez and Jermann (2001) choose β equal to .65 in their two-agent model with permanent exclusion from trading.

Risk Aversion Increasing risk aversion lowers the risk-free rate, because the demand for insurance increases and this more than offsets the decreased willingness to substitute intertemporally, as is apparent from Table IV. The market price of liquidity risk increases, raising the risk premium. Returns are more volatile as well, but so is the risk-free rate.

Collateral Share The risk-free rate is highly sensitive to α, the share of collateral, but the risk premia are not, at least not in the zero to 10 percent range. Collateral shares beyond 10 percent produce perfect risk sharing. Table V lists the same moments of returns for different values of the collateralizable share.

Consider the case of γ equal to 5. Increasing α from five to ten percent reduces the equity

Table V
Collateral Share

$R^{l,e}$ is the return on a leveraged claim to aggregate consumption with a leverage of 4; $R^{c,e}$ is the excess return on a claim to aggregate consumption. The subpanels each show the results for different values of the collateralizable share of income α. The coefficient of relative risk aversion γ is 5. These moments were generated by averaging 10.000 draws from an economy with 5000 agents .

γ	α	$E(R^{l,e})$	$E(R^{c,e})$	$E(r^f)$	$\sigma(R^{l,e})$	$\sigma(R^{c,e})$	$\sigma(r^f)$	$\frac{E(R^{l,e})}{\sigma(R^{l,e})}$	$\frac{E(R^{c,e})}{\sigma(R^{c,e})}$
5	.15	0.03	0.01	0.12	0.17	0.06	0.04	0.21	0.20
	.10	0.05	0.02	0.09	0.19	0.09	0.06	0.24	0.24
	.05	0.06	0.04	0.03	0.22	0.14	0.10	0.29	0.28
7	.15	0.06	0.03	0.11	0.19	0.09	0.07	0.31	0.330
	.10	0.09	0.05	0.03	0.23	0.15	0.11	0.37	0.38
	.05	0.09	0.06	0.00	0.24	0.16	0.12	0.39	0.40

premium somewhat, but it mainly increases the risk-free rate to five percent. As α increases to fifteen percent, the solvency constraints bind only infrequently and the asset pricing moments resemble those of the perfect insurance economy. At higher collateral levels, perfect risk sharing obtains in equilibrium.

Cyclicality of Labor Income Risk Table VI reports the results for the calibration without the increase in cross-sectional dispersion of labor income risk in low aggregate consumption growth states: the risk premium is about the same as in the economy without solvency constraints, but the risk-free rate is much lower.

Exogenous Collateral Constraints. The same table also reports the results for the collateral economy with the exogenous constraints of equation 5. As expected, the risk-free rate is even lower than in the collateral economy with endogenous constraints, but the risk premium is about the same as in same economy without solvency constraints.

Two Agent Economy In the two-agent version of this economy, the compensation per unit of risk received by investors is much smaller, because the liquidity shocks are not highly correlated with aggregate consumption growth. Instead, these are driven by the dynamics of the labor income process in a two-agent economy. For example, in the benchmark calibration, for γ is 5 and α is 5 percent, the two agent model delivers a market price of risk of .28 and a Sharpe ratio on a claim to aggregate consumption of only .13. In the continuum economy, the risk price is .32 and the Sharpe ratio is .28. The equity premium on a consumption claim is 3.5 percent and the

Table VI
Labor Income and Collateral Constraints.

$R^{l,e}$ is the return on a leveraged claim to aggregate consumption with a leverage of 4; $R^{c,e}$ is the excess return on a claim to aggregate consumption. These moments were generated by averaging 10.000 draws from an economy with 5000 agents. α is 5 percent and β is .95.

	$E(R^{l,e})$	$E(R^{c,e})$	$E(r^f)$	$\sigma(R^{l,e})$	$\sigma(R^{c,e})$	$\sigma(r^f)$	$\frac{E(R^{l,e})}{\sigma(R^{l,e})}$	$\frac{E(R^{c,e})}{\sigma(R^{c,e})}$
				Representative Agent Economy				
5	0.03	0.01	0.13	0.16	0.058	0.033	0.188	0.196
7	0.05	0.02	0.15	0.18	0.070	0.049	0.277	0.275
				Benchmark Collateral Economy				
5	0.06	0.04	0.03	0.22	0.14	0.10	0.29	0.28
7	0.09	0.06	0.00	0.24	0.15	0.11	0.39	0.40
				A-cyclical Labor Income Risk Collateral Economy				
5	0.03	0.02	0.01	0.15	0.06	0.03	0.22	0.22
7				r^f *too low*				
				Exogenous Constraints Collateral Economy				
5	.03	0.02	0.00	0.15	0.063	0.041	0.18	0.20
7	.04	0.02	0.01	0.15	0.07	0.05	0.26	0.28

risk-free rate is 6 percent, compared to 6 and 3 percent respectively in the continuum economy.

C. Time Variation in Risk Premia

The liquidity shocks introduce time-variation in the conditional volatility of the SDF. The quasi-i.i.d. aggregate consumption growth process endogenously generates a heteroscedastic liquidity shock process. The size of the liquidity shock depends on the history of aggregate shocks, as Figure 1 shows. The liquidity shock is largest when a bad aggregate shock is realized, after a long series of positive aggregate consumption growth shocks. A larger fraction of agents has decumulated its assets and ends up in the left tail of the wealth distribution. This feature generates a substantial amount of time variation in the conditional market price of risk. Under lognormality the conditional market price of risk can be approximated by $\gamma\sigma_t\Delta\log(e_{t+1}/h_{t+1})$, and the market price of risk tracks the conditional volatility of the liquidity shocks. As a result, the conditional volatility of the SDF increases after a long series of high aggregate consumption growth shocks. Once the economy experiences a low growth shock, the conditional volatility of the SDF falls to its lowest level. The left tail of the weight distribution in (19) has been erased.

Figure 1 plots the conditional market price of risk and the liquidity shocks for the benchmark calibration economy. The conditional variance of the liquidity shocks $\sigma_t \log(g_{t+1})$ -and the conditional market price of risk- peak at the end of long expansions. On the other hand, after a series of low aggregate consumption growth realizations, the conditional market price of risk reaches its lowest point. When another bad aggregate growth shock arrives, few households are severely constrained. The left tail of the wealth distribution is completely “erased” and a new low aggregate consumption growth shock does not cause a significant liquidity shock. As a result the conditional market price of risk drops by 50 percent after a string of large liquidity shocks.

Lettau and Ludvigson (2003) argue that leading asset pricing models leave a Sharpe ratio “volatility puzzle ” . I propose a novel mechanism that produces considerable variation in the conditional Sharpe ratio for equity - a claim to aggregate consumption- at moderate levels of risk aversion. The conditional Sharpe ratio is plotted in Figure (2) for the benchmark calibration with γ equal to five and it varies between .4 and .1 For γ equal to eight, the range widens to between .65 and .18. The implied standard deviation of the Sharpe ratio is .075 and .12 respectively.

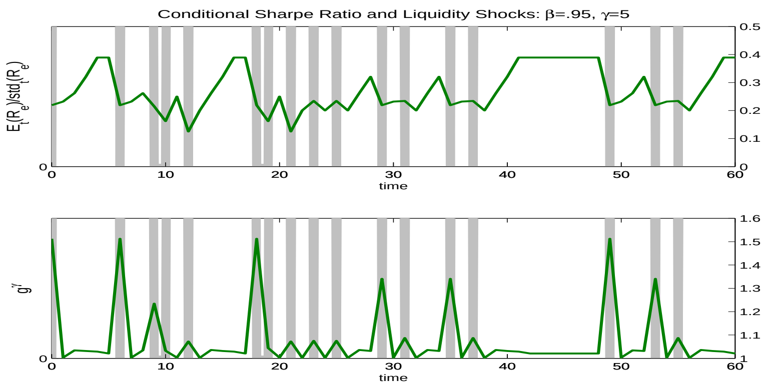

Figure 2. Liquidity Shocks and Conditional Sharpe Ratio The upper panel plots the conditional Sharpe ratio on a levered claim to aggregate consumption and the lower panel plots the liquidity shocks. The leverage parameter is four. The shaded area indicates low aggregate consumption growth states. β is .95, γ is 5 and α is .05.

The Sharpe ratio in my model rises in response to a sequence of high aggregate consumption growth shocks, in anticipation of the liquidity shocks in a recession. A sequence of low aggregate consumption growth shocks pushes the conditional Sharpe ratio to its lowest level. Is this timing consistent with the empirical evidence? The empirical Sharpe ratio estimated by Lettau and Ludvigson (2003) seems to peak around the quarters identified by the NBER as the start of recessions and it drops sharply after the start of the recession (see Figure 3 in Lettau and Ludvigson (2003)). This is broadly consistent with the annual pattern in my model. Recessions are clearly turning points in the data.

Table VII shows regressions of one-year ahead returns on (lagged) consumption growth and

Table VII
Excess returns and lagged consumption growth

The first panel shows the results of regressing log one-year excess return $r^e_{t,t+1}$ on the log price/dividend ratio at t and (lagged) consumption growth. The second panel shows the results of regressing the average monthly standard deviation of the market return in $t+1$ on the log price/dividend ratio at t and (lagged) consumption growth. The t-stats in brackets are computed using the Newel West var/covar matrix with 1 lag. The returns are one-year cum-dividend returns on the VW-CRSP index. The risk-free rate is the average 3-month yield (CRSP's Fama-Bliss risk-free rates). The sample is 1930-2003. Annual data.

Horizon	α_0	α_{pd_t}	α_{c_t}	$\alpha_{c_{t-1}}$	$\alpha_{c_{t-2}}$	R^2
Panel I: Excess Returns						
1	0.07		1.82	−1.76	0.40	0.08
	[1.96]		[2.28]	[−1.15]	[0.57]	
1	0.54	−0.15	1.96	−1.47	0.52	0.15
	[3.00]	[−2.63]	[2.59]	[−0.98]	[0.90]	
Panel II: Volatility of Excess Returns						
1	0.05		−0.32	0.02	0.02	0.10
	[6.07]		[−1.76]	[0.11]	[0.13]	
1	0.06	−0.00	−0.32	0.02	0.52	0.10
	[2.28]	[−0.30]	[−1.76]	[0.13]	[0.90]	

the price/dividend ratio. In the US data, expected excess returns over a one-year horizon *increase* in response to a positive consumption growth shock. This effect is significant and remains so when the price/dividend ratio is included in the regressions. At longer horizons the sign switches as it does in the model (not shown in the table). The conditional volatility of excess returns, proxied by the average standard deviation of market returns, decreases in response to a positive aggregate consumption growth shock. Both results are consistent with an increase in the conditional Sharpe ratio after a positive aggregate consumption growth shock, as predicted by the model.

Predictability The time-varying risk premia impute quite some predictability to returns. Table VIII lists the R-squared and the slope coefficients of Fama-French regressions of log excess returns on log price/dividend ratios. By contrast, there is no predictability in the representative agent economy (not shown in the table). The same pattern was found in the data by numerous authors (see Campbell (2000), p. 1522 for an overview. Table XIII in the appendix shows the same regression results for US postwar data.) When α is fixed, as is the case in the benchmark model, the price/dividend ratio and the consumption/wealth ratio are identical. In other words, regressing returns on consumption/wealth ratios would produce identical results (see Lettau and Ludvigson (2001) for evidence on predictability of returns using the consumption-wealth ratio).

Table VIII
Predictability regressions for Collateral Economy: Excess Returns

Results of regressing log k-horizon returns on the log risk-free rate and the log price/dividend ratio. A simulated sample of 50.000 observations was used. The first panel reports the results for the returns on a claim to aggregate consumption; the second panel reports the results for a levered claim to aggregate consumption. Leverage was set to 4.

	α_0	α_1	R^2	α_0	α_1	R^2
	Risk Aversion 5			*Risk Aversion 8*		
Horizon	*Panel A: Risk-free Rate*					
1	1.07	0.24	0.01	0.07	−0.21	0.02
2	0.11	−1.02	0.17	0.13	−0.36	0.04
3	0.15	−1.33	0.16	0.20	−0.46	0.05
4	0.20	−1.53	0.12	0.25	−0.43	0.03
5	0.22	−1.63	0.11	0.31	−0.42	0.03
6	0.26	−1.68	0.09	0.36	−0.42	0.03
Horizon	*Panel B: Price/dividend Ratio*					
1	1.92	−0.62	0.15	1.82	−0.61	0.23
2	2.68	−0.86	0.19	2.56	−0.84	0.29
3	3.07	−0.97	0.17	2.95	−0.96	0.29
4	3.27	−1.02	0.14	3.14	−1.00	0.23
5	3.29	−1.01	0.13	3.16	−0.99	0.23
6	3.35	−1.02	0.12	3.22	−0.99	0.23

Long-Run Risk The liquidity shocks impute enough persistence to the stochastic discount factor to have the Sharpe ratios on equity rise substantially over longer holding periods, thought not quite as fast as in the data. Table IX reports the simulation results over longer holding periods. The Sharpe ratios on equity rise from .56 to .76 when γ is eight. In the US data, the Sharpe ratio rises from .54 at two years to 1 at eight years see Table XIV in the appendix).

D. Recursive Utility and Long Run Risk

The variation in the conditional volatility of the liquidity shocks imputes too much volatility to the risk-free rate. The mechanism responsible for the time-variation in risk premia also lowers the risk-free rate at the end of a long series of high aggregate consumption growth realizations, in anticipation of a large liquidity shock. This can be mitigated by allowing the intertemporal elasticity of substitution (IES) to differ from γ^{-1}. In this case the market SDF is not necessarily larger than the representative agent's shadow price for consumption in that state. A version of Luttmer's aggregation result extends to the case of recursive utility:

$$m_{t+1} = m^a_{t+1} \left(g_{t+1}\right)^{\rho}, \tag{20}$$

Table IX
Long Horizon Excess Returns

The first panel reports the results for the returns on a levered claim to aggregate consumption in excess of the return on rolling over one-period zero coupons. The leverage parameter is 4. A simulated sample of 50.000 observations was used.

	$E(R^{l,e})$	$\sigma(R^{l,e})$	$\frac{E(R^{l,e})}{\sigma(R^{l,e})}$	$E(R^{l,e})$	$\sigma(R^{l,e})$	$\frac{E(R^{l,e})}{\sigma(R^{l,e})}$
	Risk Aversion 5			*Risk Aversion 8*		
Horizon	*Collateral Economy*					
2	0.08	0.22	0.36	0.17	0.31	0.56
3	0.12	0.30	0.39	0.25	0.41	0.61
4	0.16	0.40	0.40	0.33	0.52	0.63
5	0.19	0.45	0.42	0.40	0.59	0.67
6	0.22	0.51	0.43	0.47	0.66	0.71
7	0.25	0.57	0.44	0.54	0.72	0.74
8	0.28	0.62	0.45	0.61	0.79	0.76
Horizon	*Representative Agent Economy*					
2	0.06	0.19	0.32	0.03	0.18	0.19
3	0.06	0.26	0.22	0.02	0.23	0.12
4	0.02	0.34	0.08	0.00	0.29	0.01
5	−0.03	0.41	−0.09	−0.04	0.34	−0.11
6	−0.14	0.49	−0.29	−0.10	0.40	−0.26
7	−0.30	0.58	−0.52	−0.20	0.47	−0.43
8	−0.51	0.67	−0.75	−0.33	0.53	−0.62

where m^a_{t+1} is the representative agent's SDF derived by Epstein and Zin (1989) and ρ is the inverse IES (see section F in the appendix). In this case g_{t+1} is the growth rate of a moment of the distribution of scaled Pareto-Negishi weights, and these scaled weights change even when the solvency constraints do not bind. They can decrease, and this modifies the properties of the SDF in important ways. In particular, $g_{t+1} < 1$ is possible now. It reduces the volatility of the risk-free rate and it introduces long-run liquidity risk in the SDF.

Results The added flexibility helps to match the moments in the data. Table X summarizes the results. For ρ equal to 3 and γ equal to 6 the model delivers a five percent equity premium on claim to consumption, eight percent on a levered claim. The standard deviation of the risk-free rate has dropped from twelve percent, in the case of additive utility, to eight percent, while the returns and the excess returns are still quite volatile.

Recursive utility obviously creates a role for long-run risk. Bansal and Yaron (2002) rely on a small long-run predictable component in consumption growth that varies over time. My model introduces a role for long-run risk through the solvency constraints. Households are concerned about the risk of depleting their financial wealth in the future.

Table X
Summary Statistics for Collateral Economy with Recursive Utility.

$R^{l,e}$ is the return on a leveraged claim to aggregate consumption with a leverage of 4; $R^{c,e}$ is the excess return on a claim to aggregate consumption. The different panels show the results for different coefficients of relative risk aversion γ. The first column is the inverse of the intertemporal elasticity of substitution ρ. The collateralizable share α is five percent.These moments were generated by averaging 10.000 draws from an economy with 5000 agents.

γ	ρ	$E(R^{l,e})$	$E(R^{c,e})$	$E(r^f)$	$\sigma(R^{l,e})$	$\sigma(R^{c,e})$	$\sigma(r^f)$	$\frac{E(R^{l,e})}{\sigma(R^{l,e})}$	$\frac{E(R^{c,e})}{\sigma(R^{c,e})}$
				Panel I: Additive Utility					
6		0.08	0.05	0.02	0.24	0.16	0.12	0.35	0.33
7		0.09	0.06	0.01	0.24	0.15	0.12	0.40	0.39
8		0.12	0.09	−0.00	0.27	0.19	0.16	0.46	0.45
				Panel II: Recursive utility					
6	3	0.06	0.03	0.01	0.19	0.11	0.08	0.29	0.30
7	4	0.07	0.04	0.00	0.21	0.13	0.09	0.35	0.33
8	5	0.08	0.05	0.01	0.22	0.13	0.10	0.38	0.38

E. Wealth Distribution and Consumption Share Dynamics

The wealth distribution dynamics in this class of models feature limited memory w.r.t. aggregate shocks provided that the solvency constraints bind often.

Consumption Dynamics If the Markov process for y satisfies a standard monotonicity condition (see Stokey, Lucas, and Prescott (1989), p. 267), then the cutoff rules for the consumption weights can be ranked, for each z^t:

$$\widehat{\eta}(y_1, z) \leq \frac{\underline{\omega}(y_1, z^t)}{g_t(z^t)} \leq \ldots \leq \frac{\underline{\omega}(y_n, z^t)}{g_t(z^t)} \leq \widehat{\eta}(y_n, z).$$

The consumption share cutoff always exceeds the lowest income share realization and it never exceeds the highest income share realization. The cutoff values increase monotonically in the value of the income share realization. The appendix contains a proof of these claims in section C. Figure 3 plots the consumption share of a single household against its labor income share. The consumption shares fluctuate between the highest and the lowest income shares. In low income states, the household's consumption share decreases as the household runs down its assets. The largest consumption share increases occur when the household switches from the low state to the high state *after a large string of adverse idiosyncratic shocks*. In the favorable income states, its consumption share increases somewhat when it switches to the highest states. These consumption share increases are larger in recessions (period 45 and 80 in the plot) and produce large liquidity shocks when aggregated across consumers. Recessions are periods when the aggregate show cost of the solvency constraint increases.

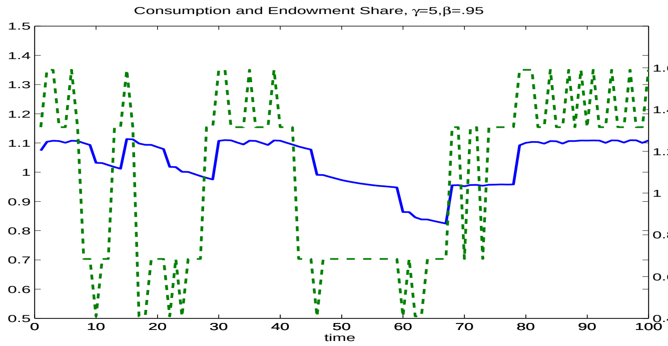

Figure 3. Consumption The full line is the consumption share of a household plotted against the labor income share (dotted line) over a period of 100 years. γ is 5 and β is .95. The y-axis on the left hand side shows the consumption shares; on the right hand side is the labor income share.

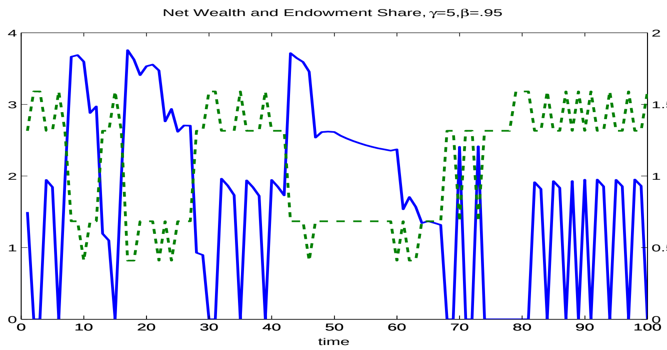

Figure 4. Net Wealth The full line is net wealth of a household scaled by its aggregate endowment, plotted against its labor endowment share (dotted line) over a period of 100 years. γ is 5 and β is .95. The y-axis on the left hand side shows net wealth scaled by the aggregate endowment; on the right hand side is the consumption scale.

Figure (4) plots the consumption share of a single household against its net wealth (net of human wealth) scaled by the aggregate endowment for the same history of shocks. Each time the household switches to a state with a binding solvency constraint, its net wealth position hits zero. This is a feature of complete markets. Net wealth is obviously much more volatile than the consumption. The household's portfolio realizes high returns when bad income shocks arrive and low returns when good income shocks arrive, but the hedge is incomplete because of the collateral constraints.

These dynamics have important ramifications for the distribution of wealth across agents.

Consider a wealthy household with a consumption share $\omega_0 > \widehat{\eta}(y_n, z)$. Its solvency constraint does not bind "today" as long as it stays in the region above the largest income share, but its consumption share is drifting downwards at a rate $g_t(z^t)$. When risk sharing is incomplete and $g > 1$, this wealthy household's consumption share will drop below the highest income share in a finite number of steps. In a "stationary" equilibrium, the consumption shares will fluctuate between the highest and the lowest income shares. Wealthy households chose to run down their assets because interest rates are low. This is the signature of complete markets. It would be inefficient to have some households hold too much financial wealth when collateral is scarce.

As a result, in a stationary equilibrium, all households face at least one binding solvency constraint, the one for the highest income share tomorrow, because their consumption share is -weakly- smaller than $\underline{\omega}(y_n, z^t)$. This explains how this model reconciles fairly smooth individual consumption processes with highly volatile SDF's. This points to a crucial distinction between this model and ex ante incomplete market models (see e.g. Heaton and Lucas (1996)). In these models wealthy agents do not run down their financial wealth holdings, and as a result, may not face any binding solvency constraints at all. In some sense, the stock of scarce collateral is not being used as efficiently in those equilibria.

This feature of a stationary equilibrium also limits the memory of these equilibria. All households face a positive probability of switching to a state with a binding solvency constraint tomorrow. Once they switch, their entire history summarized by the Pareto-Negishi weight is wiped out and reset to a value that depends only on the aggregate history of the economy and that period's private income share draw. The computations will exploit this limited memory feature of the equilibria. To match the empirical wealth distribution, the model obviously needs ex ante heterogeneity in the labor income processes. The homogeneity in my setup allows me to introduce ex ante labor income heterogeneity by scaling up the labor income processes without changing any of the asset pricing results.

The aggregate consumption growth shocks are close to white noise but the wealth dynamics induced by these shocks are not. Figure (5) plots a simulation of the fraction of agents in states with binding constraints and the corresponding liquidity shocks. Both time series peak in the low aggregate consumption growth states. On average the fraction increases to forty-five percent in low aggregate consumption growth states. This is the mechanism that generates high risk premia.

Risk Sharing This economy still manages to sustain a lot of risk sharing. In the benchmark calibration the standard deviation of consumption share growth for households is 4.3 percent, less

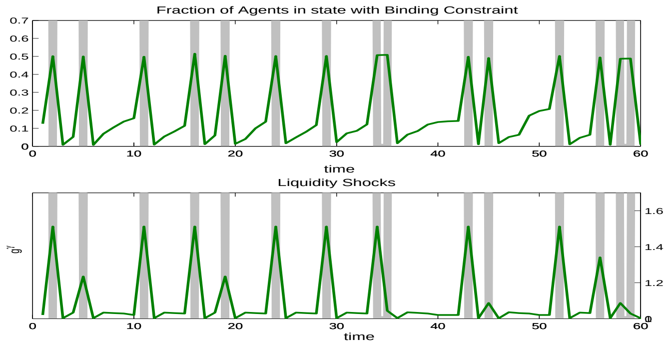

Figure 5. Fraction of Agents in States with Binding Constraints. Results shown for the economy with 5 percent collateralizable income share. This is the low volatility of labor income calibration. γ is 5 and β is .95.

Table XI
Risk Sharing in Collateral Economy

The table lists the standard deviation of household consumption share and endowment share growth. α is .05, γ is 5 and β is .95. These moments were generated by averaging 10.000 draws.

α	.05	.10	.15
$\sigma(\Delta(\log(c)))$	0.0434	.0264	0.0041
$\sigma(\Delta(\log(\eta)))$	0.334	.334	0.334

than twice the standard deviation of aggregate consumption growth, while the standard deviation of endowment share growth is thirty-three percent (see Table XI). These allocations are by no means close to autarky. Not all agents in states with binding solvency constraints experience large shocks to their consumption shares. In the truncated history with the largest liquidity shock, forty-nine percent experience a four percent consumption share drop, thirty-six percent experience an eight percent increase and six percent experience an eleven percent increase. In the history with the smallest liquidity shock (four consecutive low aggregate consumption growth shocks) almost all households have roughly constant consumption shares.

The cyclicality of these liquidity shocks is central to my results. Is there any empirical evidence to support these cyclical shocks to the wealth distribution? Little is known about how the distribution of total wealth, including human wealth, evolves over the business cycle, but my mechanism can also be detected in the consumption growth distribution.

The model predicts that the standard deviation and the skewness of consumption growth covary strongly with the returns on stocks. In particular, the standard deviation increases while the skewness decreases sharply in case of a large liquidity shock. [6] Fig. (6) plots the skewness and the standard deviation of the consumption growth distribution. Most households experience small consumption share decreases. All the households switching to states with binding solvency constraints experience consumption share increases. When the fraction of households switching to these states increases, the skewness (to the left) decreases. The skewness varies from 1 in a recession preceded by a series of expansions to 30 after a series of recessions. The standard deviation ranges respectively from .01 to .09 respectively. The standard deviation of consumption growth across households is perfectly positively correlated and the skewness is perfectly negatively correlated with the liquidity shocks g_{t+1}.

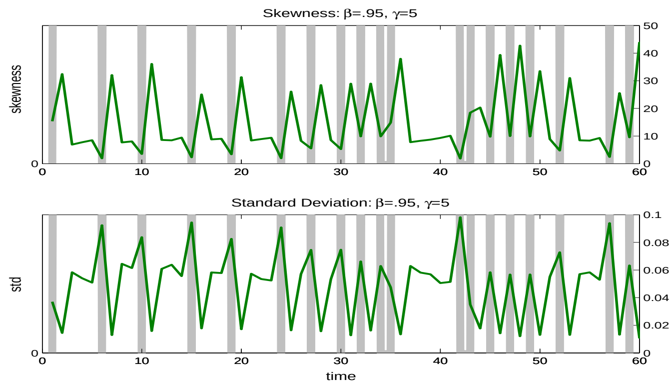

Figure 6. The Distribution of Consumption Growth. The upper panel plots the skewness of the cross-sectional distribution of consumption growth and the lower panel plots the standard deviation. The shaded areas indicate low aggregate consumption growth states. Results for the collateral economy with 5 percent collateralizable income share. γ is 5 and β is .95.

Empirical Evidence Cogley (2002) finds evidence that the cross-sectional dispersion of US household non-durable consumption growth increases when the returns on equity are low, but he argues the covariance is too small to explain the size of the risk premia. Brav, Constantinides, and Gezcy (2002) use the distribution of US household non-durable consumption growth to explain asset returns and argue that the skewness of the consumption growth distribution, in

[6] The centered moments depend only on the size of the left tail of the wealth/consumption-share distribution. The (centered) n-th moment of the consumption growth distribution is:

$$g_t(z^t)^{-n} \left\{ \begin{array}{c} \sum_{y'} \int_0^{\underline{\omega}(y',z^t)} \left(\frac{\underline{\omega}(y',z^t)}{\omega}\right)^n d\Phi_{z^{t-1}}\left(dy \times d\omega\right) \frac{\pi(y',z'|y,z)}{\pi(z'|z)} \\ -\left(\sum_{y'} \int_0^{\underline{\omega}(y',z^t)} \left(\frac{\underline{\omega}(y',z^t)}{\omega}\right) d\Phi_{z^{t-1}}\left(dy \times d\omega\right) \frac{\pi(y',z'|y,z)}{\pi(z'|z)}\right)^n \end{array} \right\}$$

addition to the standard deviation, help in explaining asset returns. They envision an incomplete markets model in which the average IMRS is a valid SDF and use a Taylor expansion to derive an expression for the SDF in terms of the mean, variance and the skewness. Brav, Constantinides, and Gezcy (2002) argue that this SDF can explain the equity premium at smaller levels of risk aversion as the sample is tightened to include only wealthier households. Without the skewness term the Euler equation errors are much larger. In more recent work Kocherlakota and Pistaferri (2004) derive a SDF that shares some features with that in my paper in an environment with private information and use household consumption data to argue that this SDF -the growth rate of the γ-th moment of the consumption distribution- performs at least as well empirically because the consumption distribution has the right cyclical properties.

This paper presents an alternative explanation of why these moments of the consumption growth distribution help explain asset prices and also provides a theory of variation in the skewness and standard deviation over the business cycle. Since the standard deviation and the skewness of consumption growth are perfectly correlated with the liquidity shocks, these moments of the consumption growth distribution can be used as risk factors. In my model individual IMRS are not valid SDF's, while in Brav, Constantinides, and Gezcy (2002) assets could be priced off any individual household's IMRS. There does not seem to be not enough volatility in household consumption growth, even when the sample is confined to stockholders.

VI. Conclusion

Most authors have concluded that heterogeneity cannot help to address the main asset pricing anomalies. This paper offers a dissenting view. While Constantinides and Duffie (1996) had already demonstrated that the right dynamics for the cross-sectional distribution of household consumption risk can generate large risk premia for plausible levels of risk aversion, my model generates this feature endogenously and still allows for considerable risk sharing at the household level, in line with the evidence on household consumption. In my model highly volatile SDF's are consistent with fairly smooth household consumption growth, because the model has a large number of agents and because the solvency constraints are dated $t+1$ instead of t.

The time variation in the distribution of liquidity shocks endogenously delivers more variation in risk premia over the business cycle than competing equilibrium asset pricing models. These risk premia are driven by the co-movement of the aggregate shadow cost of these solvency constraints with the returns on risky assets. Incomplete market models with solvency constraints (e.g. Telmer (1993), Lucas (1994) and Heaton and Lucas (1996)) impose additional measurability constraints on household net wealth across different states of the world and these average out the effect of idiosyncratic risk on the state price of consumption. That explains why, surprisingly perhaps, the risk premia are much smaller in this class of incomplete market models, in spite of the fact these offer less scope for risk sharing than my model.

References

Aiyagari, S. R., and M. Gertler (1991): "Asset Returns with Transaction Costs and Uninsured Individual Risk," *Journal of Monetary Economics*, 27, 311–331.

Alvarez, F., and U. Jermann (2000): "Efficiency, Equilibrium, and Asset Pricing with Risk of Default," *Econometrica*, 68(4), 775–798.

——— (2001): "Quantitative Asset Pricing Implications of Endogenous Solvency Constraints," *Review of Financial Studies*, 14, 1117–1152.

Anderson, E. (1998): "Uncertainty and the Dynamics of Pareto Optimal Allocations," Working Paper University of North Carolina.

Atkeson, A., and R. E. Lucas (1992): "On Efficient Distribution with Private Information," *Review of Economic Studies*, 59, 427–453.

——— (1995): "Efficiency and Equality in a Simple Model of Unemployment Insurance," *Journal of Economic Theory*, 66, 64–85.

Bansal, R., and A. Yaron (2002): "Risks for the Long Run: A Potential Resolution of Asset Prizing Puzzles," Working Paper Duke and Wharton.

Barberis, N., M. Huang, and T. Santos (2001): "Prospect Theory and Asset Prices," *Quarterly Journal of Economics*, 116, 1–53.

Brav, A., G. M. Constantinides, and C. C. Gezcy (2002): "Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence," *Journal of Political Economy*, 110(4), 793–824.

Breeden, D. T. (1979): "An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities," *Journal of Financial Economics*, 7, 265–296.

Campbell, J. Y. (2000): "Asset Pricing at the Millennium," *Journal of Finance*, 55, 1515–1567.

Campbell, J. Y., and J. H. Cochrane (1999): "By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy*, 107(2), 205–251.

Cogley, T. (2002): "Idiosyncratic Risk and The Equity Premium: Evidence from the Consumer Expenditure Survey," *Journal of Monetary Economics*, 49, 309–334.

Constantinides, G. M., and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 104, 219–240.

Epstein, L. G., and S. Zin (1989): "Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework," *Econometrica*, 57, 937–969.

Geanakoplos, J., and W. R. Zame (1998): "Collateral, Default and Market Crashes," Working Paper Yale and UCLA.

Grossman, S., and R. Shiller (1982): "Consumption Correlatedness and Risk Measurement in Economies with Non-Traded Assets and Heterogeneous Information," *Journal of Financial Economics*, 10, 195–210.

Hansen, L. P., and K. J. Singleton (1982): "Generalized Instrumental Variable Estimation of Nonlinear Rational Expectation Models," *Econometrica*, 50, 1269–1286.

He, H., and D. M. Modest (1995): "Market Frictions and Consumption-Based Asset Pricing," *Journal of Political Economy*, 103, 94–117.

He, H., and N. Pearson (1991): "Consumption and Portfolio Choices with Incomplete Markets and Short-Sale Constraints: The Infinite Dimensional Case," *Journal of Economic Theory*, 54, 259–304.

Heaton, J., and D. J. Lucas (1996): "Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing," *Journal of Political Economy*, 104, 668–712.

Judd, K. (1998): *Numerical Methods in Economics*. MIT Press, Cambridge, Mass.

Kehoe, T., and D. K. Levine (1993): "Debt Constrained Asset Markets," *Review of Economic Studies*, 60, 856–888.

Kocherlakota, N. (1996): "Implications of Efficient Risk Sharing Without Commitment," *Review of Economic Studies*, 63, 595–610.

Kocherlakota, N., and L. Pistaferri (2004): "Asset Pricing Implications of Pareto Optimality with Private Information," .

Kreps, D., and E. L. Porteus (1978): "Temporal Resolution of Uncertainty and Dynamic Choice Theory," *Econometrica*, 46, 185–200.

Krueger, D. (1999): "Risk Sharing in Economies with Incomplete Markets," Ph.D. thesis, University of Minnesota.

Krueger, D., and H. Lustig (2004): "(In)-Complete Markets and Asset Pricing: Facts and Fiction," *UCLA*.

Krusell, P., and J. A. Smith (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, 6, 867–896.

Kubler, F., and K. Schmedders (2003): "Stationary Equilibria in Asset Pricing Models with Incomplete Markets and Collateral," *Econometrica*, 71, 1767–1795.

Lettau, M. (2002): "Idiosyncratic Risk and Volatility Bounds," *Review of Economics and Statistics*, 84(2).

Lettau, M., and S. Ludvigson (2001): "Consumption, Aggregate Wealth, and Expected Stock Returns," *Journal of Finance*, 51, 815–849.

Lettau, M., and S. C. Ludvigson (2003): "Measuring and Modeling Variation in the Risk-Return Trade-Off," in *Handbook of Financial Econometrics*, ed. by Y. Ait-Sahalia, and L. P. Hansen.

Lucas, D. (1994): "Asset Pricing with Unidiversifiable Income Risk and Short Sales Constraints: Deepening the Equity Premium Puzzle," *Journal of Monetary Economics*, 34, 325–341.

Lucas, R. E. (1978): "Asset Prices in an Exchange Economy," *Econometrica*, 46(6), 1429–54.

Luenberger, D. G. (1969): *Optimization by Vector Space Methods.*

Lustig, H. (2000): "Secured Lending and Asset Prices," Mimeo Stanford University.

Luttmer, E. (1992): "Asset Pricing in Economies with Frictions," Ph.D. thesis, University of Chicago.

Mankiw, N. G. (1986): "The Equity Premium and the Concentration of Aggregate Shocks," *Journal of Financial Economics*, 17, 211–219.

Marcet, A., and R. Marimon (1999): "Recursive Contracts," Working Paper Universitat Pompeu Fabra.

Mehra, R., and E. C. Prescott (1985): "The Equity Premium: A Puzzle," *Journal of Monetary Economics*, 15., 145–161.

Moskowitz, T., and A. Vissing-Jorgensen (2002): "The Returns To Entrepreneurial Investment: A Private Equity Premium Puzzle?," *American Economic Review*, 92, 745–778.

Stokey, N., R. E. Lucas, and E. C. Prescott (1989): *Recursive Methods in Economic Dynamics.* Harvard University Press., Cambridge, Mass.

Storesletten, K., C. Telmer, and A. Yaron (2004): "Cyclical Dynamics of Idiosyncratic Labor Income Risk," *Journal of Political Economy*, 12(103), 695–717.

Telmer, C. (1993): "Asset-Pricing Puzzles and Inomplete Markets," *Journal of Finance*, 48, 1803–1832.

Whitelaw, R. F. (1994): "Time Variations and Covariations in the Expectation and Volatility of Stock Market Returns," *Journal of Finance*, 49(2), 515–41.

Zhang, H. H. (1997a): "Endogenous Borrowing Constraints with Incomplete Markets," *Journal of Finance*, 52, 2187–2209.

——— (1997b): "Endogenous Short Sale Constraint, Stock Prices and Output Cycles," *Macroeconomic Dynamics*, 1, 228–254.

VII. Tables and Figures

Table XII
Approximation Errors.

The mean, standard deviation and the supremum of the distribution of approximation errors were generated by simulating 10.000 draws from an economy with 5000 agents. The errors are reported in basis points. β is .95

	Collateral : 5 percent			**Collateral : 10 percent**			**Collateral : 15 percent**		
Risk Aversion	$\widehat{E}(x)$	$std(x)$	$\sup\|x\|$	$\widehat{E}(x)$	$std(x)$	$\sup\|x\|$	$\widehat{E}(x)$	$std(x)$	$\sup\|x\|$
3	-.0048	.0033	.012	.0009	.0009	.0047	.0008	.0058	.032
4	-.0002	.0024	.012	.0054	.0088	.0306	.0012	.010	.055
5	.0000	.0026	.015	-.0020	.0031	.0157	.0016	.0124	.073
6	.0005	.0023	.006	-.0009	.0017	.0099	.0018	.0143	.084
7	-.0036	.0037	.011	-.0008	.0017	.0057	.0026	.0151	.080
8	-.0012	.0033	.019	-.0006	.0020	.009	.0022	.0158	.089

Table XIII
Predictability regressions for US excess returns

Results of regressing k-horizon excess returns on the log real risk-free rate and the log price/dividend ratio. The t-stats in brackets are computed using the Newey West var/covar matrix with k lags. The returns are cum-dividend returns on the VW-CRSP index. The risk-free rate is the average 3-month yield (CRSP's Fama-Bliss risk-free rates). Annual data.

	Risk-free rate			*Price/dividend ratio*		
	α_0	α_1	R^2	α_0	α_1	R^2
Horizon			*1945:2003*			
1	0.08	0.89	0.03	0.61	−0.16	0.13
	[3.14]	[1.86]		[3.23]	[−2.71]	
2	0.16	0.83	0.02	0.88	−0.21	0.11
	[4.48]	[1.70]		[3.42]	[−2.63]	
3	0.26	0.26	0.00	1.50	−0.38	0.16
	[4.37]	[0.32]		[3.62]	[−2.83]	
4	0.39	−1.21	0.01	2.26	−0.57	0.19
	[4.36]	[−0.94]		[3.26]	[−2.59]	
5	0.52	−2.13	0.03	3.35	−0.87	0.28
	[4.25]	[−1.56]		[3.37]	[−2.72]	
6	0.66	−2.03	0.02	4.50	−1.19	0.29
	[4.10]	[−1.57]		[3.60]	[−2.92]	

Table XIV
Long Horizon Excess Returns in US Data

The first panel reports the results for the returns on a levered claim to aggregate consumption in excess of the return on rolling over one-period zero coupons. The CRSP-VW index was used and the 3-month average yield on a zero-coupon (Fama risk-free rates). The sample is 1925-2003.

Horizon	2	3	4	5	6	7	8
$E(R^{VW,e})$	0.173	0.270	0.375	0.490	0.622	0.778	0.957
$\sigma(R^{VW,e})$	0.318	0.402	0.500	0.580	0.669	0.785	0.933
$E(R^{VW,e})/\sigma(R^{VW,e})$	0.543	0.671	0.750	0.845	0.930	0.991	1.026

A. Technical Appendix

This appendix is self-contained. The first section (A) starts by deriving the solvency constraints from the participation constraints in an environment where agents cannot be excluded from trading. The second section (B) characterizes the regions of the parameter space where risk sharing can be sustained. In the third section (C), I use stochastic Pareto-Negishi weights to characterize equilibrium prices and allocations. The fourth section (D) examines the properties of the SDF. The final section (E) briefly extends some of my results to the case of recursive utility.

A. Solvency Constraints

First, I show that imposing these solvency constraints is equivalent to imposing participation constraints that prevent default in an environment where agents can default without being excluded from trading. In other words, these solvency constraints are not too tight.

Bankruptcy technology Let Π denote a pricing functional and let $\Pi_{s^t}\left[\{d\}\right]$ denote the price of a sequence of consumption claims $\{d\}$ starting in history s^t in units of s^t consumption:

$$\Pi_{s^t}\left[\{d\}\right] = \sum_{\tau \geq t} \sum_{s^\tau \geq s^t} p_\tau(s^\tau | s^t) d_\tau(s^\tau).$$

This includes the value of today's dividend. Let $\kappa_t(s^t)$ be the continuation utility associated with bankruptcy, conditional on a pricing functional Π :

$$\kappa_t(s^t) = \max_{\{c'\}} U(c)(s^t) \text{ s.t. } \Pi_{s^t}\left[\{c'\}\right] \leq \Pi_{s^t}\left[\{\eta\}\right],$$

and such that the participation constraints are satisfied in all following histories $s^\tau \geq s^t$. Let $U(\{c\})(s^t)$ denote the continuation utility from an allocation at s^t. An allocation is immune to bankruptcy if the household cannot increase its continuation utility by resorting to bankruptcy at any node.

Definition 4. *For given* Π, *an allocation is said to be immune to bankruptcy if*

$$U(\{c(\theta_0, y^t, z^t)\})(s^t) \geq \kappa_t(s^t) \textit{ for all } s^t. \tag{21}$$

These participation constraints can be recast as solvency constraints. I choose solvency constraints that only bind when the participation constraints bind, and hence they are *not too tight*, in the sense of Alvarez and Jermann (2000). These put a lower bound on the value of the household's consumption claim.

Proposition 5. *For given* Π, *an allocation is said to be immune to bankruptcy iff:*

$$\Pi_{s^t}\left[\{c(\theta_0, y^t, z^t)\}\right] \geq \Pi_{s^t}\left[\{\eta\}\right], \textit{ for all } s^t \in S^t, t \geq 0. \tag{22}$$

Proof of Proposition 5: First, I show that the solvency constraints imply that the participation constraints are satisfied:

$$\begin{aligned} U(\{c(\theta_0, y^t, z^t)\})(s^t) &\geq \kappa_t(s^t), \\ \textit{and } U(\{c(\theta_0, y^t, z^t)\})(s^t) &= \kappa_t(s^t) \Longleftrightarrow \Pi_{s^t}\left[\{\eta\}\right] = \Pi_{s^t}\left[\{c(\theta_0, y^t, z^t)\}\right] \end{aligned}$$

and that the participation constraints bind only if the solvency constraints bind. This follows directly from the definition of $\kappa_t(s^t)$. If $\Pi_{s^t}\left[\left\{c\left(\theta_0, y^t, z^t\right)\right\}\right] \geq \Pi_{s^t}\left[\{\eta\}\right]$, then $U(\left\{c\left(\theta_0, y^t, z^t\right)\right\})(s^t) \geq \kappa_t(s^t)$ because

$$U(\left\{c\left(\theta_0, y^t, z^t\right)\right\})(s^t) = \max_{\{c'\}} U(c)(s^t), \tag{23}$$

such that the budget constraint is satisfied $\Pi_{s^t}\left[\{c'\}\right] \leq \Pi_{s^t}\left[\left\{c\left(\theta_0, y^t, z^t\right)\right\}\right]$ and such that the solvency constraints are satisfied in all following histories:

$$U(c)(s^\tau) \geq \kappa_\tau(s^\tau) \text{ for all } s^\tau \geq s^t.$$

The rest of the proof follows from the definition of $\kappa_t(s^t)$:

$$\kappa_t(s^t) = \max_{\{c'\}} U(c)(s^t), \tag{24}$$

such that the budget constraint is satisfied $\Pi_{s^t}\left[\{c'\}\right] \leq \Pi_{s^t}\left[\{\eta\}\right]$ and the solvency constraints are satisfied in all following histories: $U(c)(s^\tau) \geq \kappa_\tau(s^\tau)$ for all $s^\tau \geq s^t$. This shows that the solvency constraints ensure that the participation constraints are satisfied. In addition, the same argument implies that, if the solvency constraints bind, then the participation constraints bind. The solvency constraint is not too tight. Second, the participation constraints imply that the solvency constraints are satisfied. If $U(\left\{c\left(\theta_0, y^t, z^t\right)\right\})(s^t) \geq \kappa_t(s^t)$, then from (23) and (24), it follows that $\Pi_{s^t}\left[\{\eta\}\right] \leq \Pi_{s^t}\left[\left\{c\left(\theta_0, y^t, z^t\right)\right\}\right]$. The second part is obvious.

Sequential Equilibrium and K-L Equilibrium In the text I define a sequential trading equilibrium with collateral constraints. This equilibrium is equivalent with an equilibrium in which all trading occurs at time zero, subject to these solvency constraints.

We assume that interest rates are high enough:

$$\Pi_{s^0}\left[\{\eta\}\right] < \infty \text{ and } \Pi_{z^0}\left[\{e\}\right] < \infty. \tag{25}$$

In the case of a continuum of consumers, it is not sufficient to restrict the value of the aggregate endowment to be finite (as in Alvarez and Jermann (2000)). It is also necessary to restrict the value of labor income to be finite. If the value of the aggregate endowment is finite, then all θ_0 will be finite as well, since these are claims to the aggregate endowment. From the time 0 budget constraint, I know that $\Pi_{s^0}\left[\left\{c\left(\mu_0, s^t\right)\right\}\right] < \infty$. This means I can apply Proposition 4.6 in Alvarez and Jermann (2000) which demonstrates the equivalence between the Arrow-Debreu economy and the economy with sequential trading, provided that there is a ξ such that

$$\frac{c\left(\mu_0, s^t\right)^{1-\gamma}}{1-\gamma} \leq \xi \frac{c_t\left(\mu_0, s^t\right)^{-\gamma}}{1} c_t\left(\mu_0, s^t\right),$$

which is automatically satisfied for power utility.

B. Risk Sharing

This section uses the solvency constraints to characterize the regions of the parameter space where (no) risk sharing can be sustained.

No Collateral If there is no collateral, no risk sharing can be sustained.

Proposition 6. *If there is no outside wealth ($\alpha = 0$), then there can be no risk sharing in equilibrium.*

Proof of Proposition 6: Summing across all of the individual participation constraints at some node z^t:

$$\int \sum_{y^t} \begin{bmatrix} \Pi_{s^t} \left[\left\{c\left(\mu_0, y^t, z^t\right)\right\}\right] \\ -\Pi_{s^t} \left[\{\eta\}\right]) \end{bmatrix} d\Phi_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} \geq 0. \tag{26}$$

Using $p(s^t|s_0) = Q(z^t|z_0)\frac{\pi(y^t,z^t|y_0,z_0)}{\pi(z^t|z_0)}$ -this is w.l.o.g.-, this can be rewritten as:

$$\sum_{z^\tau \succeq z^t} Q(z^\tau|z^t) \left[\int \sum_{y^\tau} \begin{bmatrix} c\left(\mu_0, y^\tau, z^\tau\right) \\ -\widehat{\eta}_\tau(y_\tau, z_\tau) e_\tau(z^\tau) \end{bmatrix} d\Phi_0 \frac{\pi(y^\tau, z^\tau|y_0, z_0)}{\pi(z^\tau|z_0)}\right], \tag{27}$$

with $(z^\tau, y^\tau) \succeq s^t$. To justify the interchange of limits and expectations, I appeal to the monotone convergence theorem. Let $\Pi^n_{s^t}\left[\left\{c\left(\mu_0, y^t, z^t\right)\right\}\right]$ be the value of the claim to the consumption stream until $t+n$ and let $\Pi^n_{s^t}\left[\{\eta\}\right]$ be similarly defined. Then the monotone convergence theorem can be applied for both sequences because for all $n : 0 \leq X_n \leq X_{n+1}$. Let $X = \lim_n X_n$. Then $EX_n \nearrow X$ as $n \to \infty$ (where EX is possibly infinite). This justifies the interchange of limit and the expectation (SLP, 1989, p.187).

The Law of Large Numbers and the definition of the labor share of the aggregate endowment imply that the average labor endowment share equals the labor share:

$$\int \sum_{y^t} \widehat{\eta}_t(y_t, z_t) \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} d\Phi_0 = \sum_{y'} \pi_{z_t}(y_t)\widehat{\eta}_t(y_t, z_t) = (1-\alpha), \tag{28}$$

and the market clearing condition implies that:

$$\int \sum_{y^t} c\left(\mu_0, y^t, z^t\right) \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} d\Phi_0 = e_t(z^t). \tag{29}$$

Plugging eqs. (28) and (29) back into eq. (27) implies the following inequality must hold at all nodes z^t: $\alpha\Pi_{z^t}\left[\left\{e_t(z^t)\right\}\right] \geq 0$. If there is no outside wealth ($\alpha = 0$) in the economy, then the expression is zero at all nodes z^t and eq. (26) holds with equality at all nodes z^t. This implies that each individual constraint binds for all s^t and there can be no risk sharing. Why? Suppose there are some households $\left(\mu_0, y^t, z^t\right) \in A$ at node z^t where A has non-zero measure:

$$\sum \int_A d\Phi_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} > 0,$$

and their constraint is slack: $\Pi_{s^t}\left[\left\{c\left(\mu_0, y^t, z^t\right)\right\}\right] > \Pi_{s^t}\left[\{\eta\}\right]$. Given that eq. (26) holds with equality at all nodes z^t with $\alpha = 0$, there are some households $\left(\mu_0', y^t, z^t\right)$ at node $z^t \in B$ for which

$$\sum \int_B d\Phi_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} > 0,$$

which have constraints that are violated: $\Pi_{s^t}\left[\left\{c\left(\mu_0', y^t, z^t\right)\right\}\right] < \Pi_{s^t}\left[\{\eta\}\right]$. If not, (26) would be violated. But this violates the participation constraints for these agents. So, for $\alpha = 0$, for all households with positive measure:

$$\Pi_{s^t}\left[\left\{c\left(\mu_0, y^t, z^t\right)\right\}\right] = \Pi_{s^t}\left[\{\eta\}\right] \text{ for all } y^t \text{ at } z^t.$$

The same argument can be repeated for all z^t. This implies that the following equality holds for all s^t and for all households with positive measure:

$$\Pi_{s^t}\left[\left\{c\left(\mu_0, y^t, z^t\right)\right\}\right] = \Pi_{s^t}\left[\{\eta\}\right] \text{ for all } s^t,$$

and there can be no risk sharing: $c\left(\mu_0, y^t, z^t\right) = \eta_t(s^t)$ for all s^t and μ_0

Perfect Risk Sharing If there is enough collateral, agents may be able to share risks perfectly. Let Π^* denote the pricing functional defined by the perfect insurance, Lucas-Breeden SDF.

Proposition 7. *If the value of the aggregate endowment exceeds the value of the private endowment at all nodes, perfect risk sharing is feasible:*

$$\Pi^*_{s^t}\left[\{e\}\right] \geq \Pi^*_{s^t}\left[\{\eta\}\right] \text{ for all } s^t.$$

Proof of Proposition 7: If this condition is satisfied: $\Pi^*_{s^t}\left[\{e\}\right] \geq \Pi^*_{s^t}\left[\{\eta\}\right]$ for all s^t, where $\Pi^*_{s^t}$ is the complete insurance pricing functional, then each household can get a constant and equal share of the aggregate endowment at all future nodes. Perfect risk sharing is possible. (q.e.d.)

Permanent Exclusion How does this relate to the Kehoe-Levine-Kocherlakota setup with permanent exclusion? The solvency constraints are tighter in the case of bankruptcy than under permanent exclusion, simply because one could always default and replicate autarky in the economy with bankruptcy by eating one's endowment forever after. The reverse is clearly not true. Let $U(\{\eta\})(s^t)$ denote the continuation utility from autarky.

Proposition 8. *In the economy with permanent exclusion, the participation constraints can be written as solvency constraints as follows:*

$$\Pi_{s^t}\left[\{c\}\right] \geq \Pi_{s^t}\left[\{\eta\}\right] \geq B^{aut}_{s^t}\left[\{\eta\}\right],$$

where $U(\{\eta\})(s^t) = sup_{\{c'\}} U(c')(s^t)$ *s.t.* $\Pi_{s^t}\left[\{c'\}\right] \leq B^{aut}_{s^t}\left[\{\eta\}\right]$ *and s.t. the participation constraint is satisfied at all future nodes .*

Proof of Proposition 8: The value of the outside option at each node s^t is simply the value of autarky: $U(\eta)(s^t)$. The value of bankruptcy has to exceed the value of autarky for any pricing functional, since continuation values are monotonic in wealth:

$$\Pi_{s^t}\left[\{c\}\right] \geq \Pi_{s^t}\left[\{\eta\}\right] \geq B^{aut}_{s^t}\left[\{\eta\}\right],$$

where $U_t(B^{aut}_{s^t}\left[\{\eta\}\right], s^t, c) = U(\{\eta\})(s^t)$.(q.e.d.)

Because this inequality holds for any pricing functional, if perfect risk sharing is feasible in the economy with bankruptcy, it is feasible in the economy with permanent exclusion. Loosely speaking, the Pareto frontier shifts down as one moves from permanent exclusion to bankruptcy (also see Lustig, 2000).

C. Stochastic Pareto-Negishi Weights

This section first describes the household's problem in a time zero trading environment and then defines an equilibrium. In a second step, I characterize these equilibria using stochastic Pareto-Negishi weights.

Taking prices $\left\{p_t(s^t|s_0)\right\}$ as given, the household purchases history-contingent consumption claims subject to a standard budget constraint and a sequence of solvency constraints, one for each history:

Primal Problem (PP)

$$\sup_{\{c\}} u(c_0(\theta_0, s^0)) + \sum_{t=1} \sum_{s^t \geq s^0} \beta^t \pi(s^t|s_0) u(c_t(\theta_0, s^t)),$$

$$\sum_{t\geq 0}\sum_{s^t\geq s_0} p_t(s^t|s^0)\left[c_t(\theta_0,s^t)-\eta_t(s^t)\right]\leq \theta_0,$$

$$\Pi_{s^t}\left[\left\{c\left(\theta_0,y^t,z^t\right)\right\}\right]\geq \Pi_{s^t}\left[\{\eta\}\right], \text{ for all } s^t\in S^t, t\geq 0.$$

The solvency constraints keep the households from defaulting. The following definition of equilibrium is in the spirit of Kehoe and Levine (1993) and in particular Krueger (1999). The definition of equilibrium is in the text. The next section develops an algorithm to solve for equilibrium allocations and prices using stochastic Pareto-Negishi weights.

Solvency Constraints and Stochastic Pareto-Negishi Weights In the complete insurance benchmark model households are assigned Pareto-Negishi weights at time 0 by a social planner and these weights stay fixed throughout. Associated with the equilibrium of my limited commitment economy is a set of Pareto-Negishi weights that are non-decreasing stochastic processes. These keep track of an agent's history. In effect, the Pareto-Negishi weights adjust the value of a household's wealth just enough to prevent it from exercising the bankruptcy option.

I relabel households with initial promised utilities w_0 instead of initial wealth θ_0. The dual program consists of minimizing the resources spent by a consumer who starts out with "promised" utility w_0:

Dual Problem (DP)

$$C^*(w_0,s^0)=\inf_{\{c\}} c_0(w_0,s^0)+\sum_{t=1}\sum_{s^t\geq s^0} p_t(s^t|s^0)c_t(w_0,s^t),$$

$$\sum_{t\geq 0}\sum_{s^t\geq s^0}\beta^t\pi(s^t|s_0)u(c_t(w_0,s^t))=w_0, \tag{30}$$

$$\Pi_{s^t}\left[\left\{c\left(w_0,y^t,z^t\right)\right\}\right]\geq \Pi_{s^t}\left[\{\eta\}\right], \text{ for all } s^t\in S^t, t\geq 0. \tag{31}$$

The convexity of the constraint set implies that the minimizer of DP and the maximizer of PP (the primal problem) coincide for initial wealth $\theta_0=C^*(w_0,s^0)-\Pi_{s^0}\left[\{\eta\}\right]$.(see Luenberger (1969), p. 201.)

To solve for the equilibrium allocations, I make the dual problem recursive. To do so, I borrow and extend some tools recently developed to solve recursive contracting problems by Marcet and Marimon (1999). Let $m_t(s^t|s_0)=p_t(s^t|s_0)/\pi_t(s^t|s_0)$, i.e. the state price deflator for payoffs conditional on event history s^t. $\tau_t(s^t)$ is the multiplier on the solvency constraint at node s^t. I can transform the original dual program into a recursive saddle point problem for household (w_0,s_0) by introducing a cumulative multiplier:

$$\chi_t(w_0,s^t)=\chi_{t-1}(w_0,s^{t-1})-\tau_t(w_0,s^t),\quad \chi_0=1. \tag{32}$$

Let μ_0 denotes the Lagrangian multiplier on the initial promised utility constraint in (30). I will use these to index the households with, instead of promised utilities. It is the initial value of the household's Pareto-Negishi weights. After history s^t, the Pareto-Negishi weight is given by $\zeta_t(\mu_0,s^t)=\mu_0/\chi_t(\mu_0,s^t)$. If a constraint binds $\left(\tau_t(s^t)>0\right)$, the weight ζ goes up, if not, it stays the same. These weight adjustments prevent the value of the consumption claim from dropping below the value of the labor income claim at any node.

Formally, following Marcet and Marimon (1999), I can transform the original dual program into a recursive saddle point problem for household (w_0,s_0) by introducing a cumulative multiplier:

$$D(c,\chi;w_0,s_0)=\sum_{t\geq 0}\sum_{s^t}\left\{\beta^t\pi(s^t|s_0)m_t(s^t|s_0)\begin{bmatrix}\chi_t(s^t|s_0)c_t(w_0,s^t)\\ +\tau_t(s^t)\Pi_{s^t}\left[\{\eta\}\right]\end{bmatrix}\right\}, \tag{33}$$

where $\chi_t(s^t)=\chi_{t-1}(s^{t-1})-\tau_t(s^t),\chi_0=1$. Then the recursive dual saddle point problem facing the household of type

(w_0, s_0) is given by:

$$\inf_{\{c\}} \sup_{\{\chi\}} D(c, \chi; w_0, s_0), \tag{RSDP}$$

such that

$$\sum_{t \geq 0} \sum_{s^t} \beta^t \pi(s^t | s_0) u(c_t(w_0, s^t)) = w_0.$$

Let μ_0 denotes the Lagrangian multiplier on the promise keeping constraint.

The next step is to use those Pareto-Negishi weights and exploit the homogeneity of the utility function to construct a linear consumption sharing rule, as in the benchmark model. Luttmer (1991) derived a similar aggregation result for economies with exogenous wealth constraints and a finite number of agents, without actually solving the model. This allows me to recover allocations and prices from the equilibrium sequence of multipliers $\{\zeta_t(\mu_0, s^t)\}$.

First, consider 2 households having experienced the same history s^t. We know from the first order conditions of the recursive dual saddle point problem for two different households (μ_0', y_0) and (μ_0'', y_0) that the ratio of marginal utilities has to equal the inverse of the weight ratio:

$$\left[\frac{c_t(\mu_0', s^t))}{c_t(\mu_0'', s^t))} \right]^{-\gamma} = \frac{\zeta_t(\mu_0'', s^t)}{\zeta_t(\mu_0', s^t)}. \tag{34}$$

If the constraints never bind, $\zeta_t = \mu_0$ at all nodes and the condition in (34) reduces to condition that characterizes perfect risk sharing.

Second, the resource constraint implies that for all aggregate states of the world z^t consumption adds up to the total endowment:

$$\sum_{y^t} \int c_t(\mu_0, y^t, z^t) d\Phi_0 \frac{\pi(y^t, z^t | y_0, z_0)}{\pi(z^t | z_0)} = e_t(z^t), \tag{35}$$

(34) and (35) completely characterize the equilibrium consumption allocation for a given sequence of multipliers. The objective is to find the risk sharing rule that satisfies these conditions:

$$c_t(\mu_0, s^t) = \frac{\zeta_t^{1/\gamma}(\mu_0, s^t)}{E\left[\zeta_t^{1/\gamma}(\mu_0, s^t)\right]} e_t(z^t). \tag{36}$$

This rule satisfies the condition on the ratio of marginal utilities (34) and it clears the market in each aggregate history z^t. This can be verified by taking cross-sectional averages of the individual consumption rule.

The average weight in the denominator is a non-decreasing stochastic process that is adapted w.r.t. the aggregate history. Let $h_t(z^t)$ denote this cross-sectional multiplier moment: $h_t(z^t) = E\left[\zeta_t^{1/\gamma}(\mu_0, s^t)\right]$. I will refer to this simply as the average weight process.

Cutoff Rule This section characterizes the optimal weight policy and then shows that these weights fully characterize an equilibrium.

First, I will transform this growth economy into a stationary economy with stochastic discount rates (Alvarez and Jermann (2001)) . The aggregate growth rate is a function $\lambda(z_t)$. Let utility over consumption streams be defined as follows:

$$U(\widehat{c})(s^t) = \frac{\widehat{c}_t(s^t)^{1-\gamma}}{1-\gamma} + \widehat{\beta}(z_t) \sum_{s^{t+1}} U(\widehat{c})(s^{t+1}) \widehat{\pi}(s^{t+1} | s_t),$$

where $\widehat{c}$ represents the consumption share of the total endowment and let the transformed transition matrix be given by:

$$\widehat{\pi}(z_{t+1}|z_t) = \frac{\pi(z_{t+1}|z_t)\lambda(z_{t+1})^{1-\gamma}}{\sum_{z_{t+1}} \pi(z_{t+1}|z_t)\lambda(z_{t+1})^{1-\gamma}} \text{ and } \widehat{\beta}(z_t) = \beta \sum_{z_{t+1}} \pi(z_{t+1}|z_t)\lambda(z_{t+1})^{1-\gamma}. \tag{37}$$

The (cum dividend) price-dividend ratio of a dividend stream can be written recursively as:

$$\widehat{\Pi}_{s^t}\left[\left\{\widehat{d}\right\}\right] = \widehat{d}_t(s^t) + \widehat{\beta}(z_t) \sum_{s^{t+1}} \widehat{\Pi}_{s^{t+1}}\left[\left\{\widehat{d}\right\}\right]\left(\frac{h_{t+1}(z^{t+1})}{h_t(z^t)}\right)^{\gamma} \widehat{\pi}(s^{t+1}|s_t), \tag{38}$$

and let $V_{s^t}\left[\left\{\widehat{d}\right\}\right]$ denote the ex-dividend price-dividend ratio (i.e. the previous expression less today's dividend). The equilibrium consumption shares in the stationary economy can simply be scaled up to obtain the allocations in the growth economy. The prices of claims to a dividend stream in the stationary economy are the price-dividend ratio's in the growth economy.

The optimal policy rule has a simple recursive structure. Let $C\left(\mu_0, s^t; \zeta\right)$ denote the continuation cost of a consumption claim derived from a weight policy $\left\{\zeta_t(\mu_0, s^t)\right\}$:

$$C\left(\mu_0, s^t; \zeta\right) = \Pi_{s^t}\left[\left\{c_\tau(\zeta_\tau(\mu_0, s^\tau))\right\}\right],$$

where consumption at each node is given by the risk sharing rule in (36):

$$c_t(\zeta_t(\mu_0, s^t)) = \frac{\zeta_t^{1/\gamma}(\mu_0, s^t)}{h_t(z^t)} e(z^t),$$

and prices of contingent claims are given by the standard expression $p(s^t|s_0) = \pi(s^t|s_0)\widehat{Q}_t(z^t)$. The optimal weight updating rule has a simple structure. I will let $\underline{l}_t(y, z^t)$ denote the weight such that a household starting with that weight has a continuation cost that exactly equals the price of a claim to labor income:

$$C\left(\mu_0, s^t; \zeta\right) = \Pi_{s^t}\left[\{\eta\}\right] \text{ with } \zeta_t(\mu_0, s^t) = \underline{l}_t(y, z^t).$$

A household compares its weight $\zeta_{t-1}(\mu_0, s^{t-1})$ going into period t at node s^t to its cutoff weight and adjusts its weight only if it is lower than the cutoff.

Lemma 9. *The optimal weight updating policy consists of a cutoff rule* $\left\{\underline{l}_t(y, z^t)\right\}$ *where* $\zeta_0(\mu_0, s^0) = \mu_0$ *and for all* $t \geq 1$

$$\begin{aligned} \text{if } \zeta_{t-1}(\mu_0, s^{t-1}) &> \underline{l}_t(y, z^t) \\ \zeta_t(\mu_0, s^t) &= \zeta_{t-1}(\mu_0, s^{t-1}), \\ \text{else } \zeta_t(\mu_0, s^t) &= \underline{l}_t(y, z^t). \end{aligned}$$

Proof of Lemma 9: The sequence of implied weights $\left\{\zeta_t(\mu_0, s^t)\right\}$ satisfies the necessary Kuhn-Tucker conditions for optimality:

$$\left[\zeta_t(\mu_0, s^t) - \zeta_{t-1}(\mu_0, s^{t-1})\right]\left(C\left(\mu_0, s^t; l\right) - \Pi_{s^t}\left[\{\eta\}\right]\right) = 0,$$

and $C\left(\mu_0, s^t; l\right) \geq \Pi_{s^t}\left[\{\eta\}\right]$ for all s^t. The last inequality follows from the fact that $C(\cdot)$ is non-decreasing in μ_0. It is easy to verify that there exist no other weight policy rules that satisfy these necessary conditions. Since the optimal policy is to compare the current weight ζ to the cutoff rule $l_t(y, z^t)$, the continuation cost can be stated as a function of the current weight, the current idiosyncratic state and the aggregate history: $C\left(\mu_0, s^t; l\right) = C_t(\zeta, y, z^t)$.

The household's policy rule $\{\zeta_t(\mu_0, s^t)\}$ can be written recursively as $\{l_t(l, y, z^t)\}$ where $l_0 = \mu_0$ and $l_t(l_{t-1}, y, z^t) = l_{t-1}$ if $l_{t-1} > \underline{l}_t(y, z^t)$ and $l_t(l_{t-1}, y, z^t) = \underline{l}_t(y, z^t)$ elsewhere. The reason is simple. If the constraint does not bind, the weight is left unchanged. If it does bind, it is set to its cutoff value. (q.e.d.)

The following theorem states that an equilibrium is fully characterized by these household weight processes.

Theorem 10. *An allocation $\{\zeta_t(\mu_0, s^t)\}$ for all (μ_0, s^t), state price deflators $\{Q_t(z^t)\}$ and forecasts $\{h_t(z^t|z_0)\}$ define an equilibrium if (i) $\{\zeta_t(\mu_0, s^t)\}_{t=0}^{\infty}$ solves (DP) and (ii) the market clears for all z^t:*

$$h_t(z^t) = \sum_{y^t} \int \zeta_t^{1/\gamma}(\mu_0, y^t, z^t) d\Phi_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)}$$

and (iii) there are no arbitrage opportunities :

$$Q(z^t) = \beta^t \left(\frac{e_t(z^t)}{e_0(z^0)}\right)^{-\gamma} \left(\frac{h_t(z^t)}{h_0(z^0)}\right)^{\gamma}$$

Proof of Theorem 10: $\{\zeta_t(\mu_0, s^t)\}_{t=0}^{\infty}$ and $\{h_t(z^t)\}$ define an allocation $\{c_t(\mu_0, s^t)\}$ through the risk sharing rule

$$c_t(\mu_0, s^t) = \frac{\zeta_t^{1/\gamma}(\mu_0, s^t)}{h_t(z^t)} e_t(z^t).$$

The sequence of Lagrangian multipliers $\{\zeta_t(\mu_0, s^t) - \zeta_{t-1}(\mu_0, s^{t-1})\}$ satisfy the Kuhn-Tucker conditions for a saddle point. The consumption allocations satisfy the first order conditions for optimality (see derivation of risk sharing rule). Market clearing is satisfied because $E\left[\zeta_t^{1/\gamma}(\mu_0, y^t, z^t)\right] = h_t(z^t)$ implies that $E\left[c_t(\mu_0, y^t, z^t)\right] = e_t(z^t)$. Now, let $\theta_0 = C(\mu_0, s^0; l) - \Pi_{s^0}[\{\eta\}]$. The prices implied by $\{m_t(z^t|z_0)\}$ are equilibrium prices by construction and rule out arbitrage opportunities. So, now I can relabel the households as $(\theta_0(\mu_0), s^0)$ and I have recovered the equilibrium allocations $\{c_t(\theta_0, s^t)\}$ and the prices $\{p_t(s^t|s_0)\}$.(q.e.d.)

Useful Properties of the Cutoff Rule I will list two useful properties of these cutoff rules. First, the cutoff rules for the consumption shares are weakly lower than the endowment share. The intuition is simple: the agent consumes less today in exchange for the promise of higher consumption tomorrow.

Lemma 11. *The consumption shares at the cutoff do not exceed the labor endowment shares:*

$$\frac{\underline{l}_t^{1/\gamma}(z^t, y)}{h_t(z^t)} \leq \widehat{\eta}(y, z) \text{ for all } (z^t, y) \tag{39}$$

Proof of Lemma 11: This follows directly from the definition of the cutoff level:

$$C\left(\mu_0, s^t; l\right) = \widehat{\eta}(y, z) + \widehat{\beta}(z_t) \sum_{z'} \left(\frac{h_{t+1}(z^t, z')}{h_t(z^t)}\right)^{\gamma} \widehat{\pi}(z'|z) \sum_{y'} \widehat{\Pi}_{z^{t+1}, y'}\left[\{\widehat{\eta}\}\right] \widehat{\pi}(y', y|z'),$$

where $l_t(\mu_0, s^t) = \underline{l}_t(z^t, y)$. Now since, $C\left(\mu_0, s^{t+1}; l\right) \geq \widehat{\Pi}_{z^{t+1}, y'}\left[\{\widehat{\eta}\}\right]$ for all (y^{t+1}, z^{t+1}), this equality implies that $\frac{\underline{l}_t^{1/\gamma}(z^t, y)}{h_t(z^t)} \leq \widehat{\eta}(y, z)$ for all (y, z).(q.e.d.)

Of course, as the collateralizable share of income decreases, the cutoff consumption shares approach the labor endowment shares; when $\alpha = 0$, equation (39) holds with equality at all nodes.

Second, if the transition matrix satisfies monotonicity, the cutoffs can be ranked and the consumption share in the lowest income state equals the labor endowment share.

Lemma 12. *If the transition matrix satisfies monotonicity, then the cutoff rules can be ranked:*

$$\underline{l}_t(z^t, y_n) \geq \underline{l}_t(z^t, y_{n-1}) \geq \underline{l}_t(z^t, y_{n-2}) \geq \ldots \geq \underline{l}_t(z^t, y_1)$$

and $\frac{\underline{l}_t^{1/\gamma}(z^t, y_1)}{h_t(z^t)} = \widehat{\eta}(y_1, z)$ *for all* z^t.

Proof of Lemma 12: First, I define monotonicity. A transition matrix T is monotone if for any non-decreasing function f on $Y \times Z$, Tf is also non-decreasing. If this condition is satisfied, then I can rank the all of the cutoff weights. Assume the transition matrix of conditional probabilities $\frac{\widehat{\pi}(y', z'|y, z)}{\widehat{\pi}(z'|z)}$ satisfies this condition for all (z', z). Then my claim is that value of the endowment claims can be ranked such that:

$$\widehat{\Pi}_{z^t, y_n}[\{\widehat{\eta}\}] \geq \widehat{\Pi}_{z^t, y_{n-1}}[\{\widehat{\eta}\}] \geq \ldots \geq \widehat{\Pi}_{z^t, y_1}[\{\widehat{\eta}\}], \tag{40}$$

for all z^t. To show this, I start with a truncated version of this economy at $T-1$ I use $\widetilde{\Pi}$ to denote the claims in the truncated version of this economy. By definition, for all z^{T-1} :

$$\widetilde{\Pi}_{z^{T-1}, y}[\{\widehat{\eta}\}] = \widehat{\eta}(y, z_{T-1}) + \widehat{\beta}(z_{T-1}) \sum_{z'} \left(\frac{h_T(z^{T-1}, z')}{h_{T-1}(z^{T-1})}\right)^\gamma \widehat{\pi}(z'|z) \sum_{y'} \eta(y', z') \frac{\widehat{\pi}(y', z'|y, z)}{\widehat{\pi}(z'|z)},$$

and verify that these objects can be ranked:

$$\widetilde{\Pi}_{z^{T-1}, y_n}[\{\widehat{\eta}\}] \geq \widetilde{\Pi}_{z^{T-1}, y_{n-1}}[\{\widehat{\eta}\}] \geq \widetilde{\Pi}_{z^{T-1}, y_1}[\{\widehat{\eta}\}],$$

because $\sum_{y'} \eta(y', z')\widehat{\pi}(y', y|z')$ is non-decreasing in y. This follows immediately from the definition of monotonicity of $\widehat{\pi}(y', y|z')$. Next, I roll the truncated economy back one more period:

$$\widetilde{\Pi}_{z^{T-2}, y}[\{\widehat{\eta}\}] = \widehat{\eta}(y, z_{T-2}) + \widehat{\beta}(z_{T-2}) \sum_{z'} \left(\frac{h_T(z^{T-2}, z')}{h_{T-1}(z^{T-2})}\right)^\gamma \widehat{\pi}(z'|z) \sum_{y'} \widetilde{\Pi}_{z^{t+1}, y'}[\{\widehat{\eta}\}] \frac{\widehat{\pi}(y', z'|y, z)}{\widehat{\pi}(z'|z)},$$

and using the result for $T-1$, one obtains the following ranking:

$$\widetilde{\Pi}_{z^{T-2}, y_n}[\{\widehat{\eta}\}] \geq \widetilde{\Pi}_{z^{T-2}, y_{n-1}}[\{\widehat{\eta}\}] \geq \ldots \geq \widetilde{\Pi}_{z^{T-2}, y_1}[\{\widehat{\eta}\}].$$

By backward induction, for any z^t, the claims in the truncated economy can be ranked such that:

$$\widetilde{\Pi}_{z^t, y_n} \geq \widetilde{\Pi}_{z^t, y_{n-1}} \geq \ldots \geq \widetilde{\Pi}_{z^t, y_1}.$$

Next, I note that the price of a claim in the infinite horizon economy can be stated as:

$$\widehat{\Pi}_{z^t, y_t} = \widehat{\Pi}_{z^t, y_t} + \widetilde{E}_t \beta^{T-t} \left(\frac{h_T}{h_t}\right)^\gamma \widehat{\Pi}_{z^T, y_T},$$

and that $\lim_{T\to\infty} \widetilde{E}_t \beta^{T-t} \frac{h_T}{h_t} \widehat{\Pi}_{z^T, y_T}$ is independent of y_t and converges to some finite x that does not depend on y_t : the transition matrix has no absorbing states, all states y' will be visited infinitely often in the limit and the limit cannot depend on y_t. The limit is finite by assumption. Hence, the results for the truncated economy are valid for the infinite horizon economy. This shows equation (40) holds. Finally, I need to show that this implies a similar ranking for the cutoff weights. When $\zeta_t(\mu_0, s^t) = \underline{l}_t(z^t, y)$, by definition,the following holds:

$$C(\mu_0, s^t; l) = \widehat{\eta}(y, z) + \widehat{\beta}(z_t) \sum_{z'} \left(\frac{h_{t+1}(z^t, z')}{h_t(z^t)} \right)^{\gamma} \widehat{\pi}(z'|z) \left[\sum_{y'} \widehat{\Pi}_{z^{t+1}, y'} [\{\widehat{\eta}\}] \frac{\widehat{\pi}(y', z'|y, z)}{\widehat{\pi}(z'|z)} \right].$$

Since C is monotonically increasing in ζ, I know that for all y' and z^t :

$$\underline{l}_t(z^t, y_n) \geq \underline{l}_t(z^t, y_{n-1}) \geq \ldots \geq \underline{l}_t(z^t, y_1).$$

This result, combined with Lemma 11, implies directly that the consumption share in the lowest state equals the endowment share: $\frac{\underline{l}_t(z^t, y_1)}{h_t(z^t)} = \widehat{\eta}(y_1, z^t)$ for all z^t.(q.e.d.)

I will assume monotonocity is satisfied throughout the rest of the paper.(See Stokey, Lucas, and Prescott (1989) (p. 220, 1986) for the definition of monotonicity.) Naturally, a wealthy household that starts off with an initial weight above the highest cutoff will end up hitting that bound in finite time, unless there is perfect risk sharing. This random stopping time is defined as:

$$\tau = \inf \left\{ t \geq 0 : \frac{\mu_0}{h_t(z^t)} \leq \widehat{\eta}(y_n, z) \right\}$$

The less risk sharing, the smaller τ in expectation for a given μ_0. I will assume this economy has been running long enough such that the agents with weights higher than the highest reservation weight have measure zero:

$$\sum_{y^t} \int_{\underline{l}_t^{1/\gamma}(z^t, y_n)} d\Phi_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} = 0 \text{ for all } z^t.$$

After some finite τ, all of the consumption shares $\omega(\mu_0, s^t)$ are fluctuating between the highest and the lowest endowment shares

$$\widehat{\eta}(y_1, z) \leq \omega(\mu_0, s^t) < \widehat{\eta}(y_1, z) \text{ for all } (\mu_0, s^t) \text{ and } t \geq \tau.$$

This follows directly Lemma (12) and (11). All households face at least one binding solvency constraint -in the highest state tomorrow.

D. Risk Premia

This section first derives an expression for the stochastic discount factor and goes on to derive some properties of the average weight shock.

Stochastic Discount Factor Consider the necessary f.o.c. for optimality in (RSDP):

$$\chi_t(\mu_0', s^t) p(s^t|s_0) = \mu_0 u_c(c_t(\mu_0', s^t)) \beta^t \pi(s^t|s_0).$$

To economize on notation, let $\zeta_t(\mu_0, s^t) = \mu_0 / \chi_t(\mu_0, s^t)$. Consider the ratio of first order conditions for an individual of type (μ_0, s^0) at 2 consecutive nodes (s^{t+1}, s^t):

$$\frac{p(s^{t+1}|s_0)}{p(s^t|s_0)} = \beta \pi(s^{t+1}|s_t) \frac{\zeta_{t+1}(\mu_0, s^{t+1})}{\zeta_t(\mu_0, s^t)} \left[\frac{c_{t+1}(\mu_0, s^{t+1}))}{c_t(\mu_0, s^t))} \right]^{-\gamma},$$

and substitute for the optimal risk sharing rule, noting that the unconstrained investor's weight ζ_{t+1} does not change. Then the following expression for the ratio of prices obtains:

$$\frac{p(s^{t+1}|s_0)}{p(s^t|s_0)} = \beta\pi(s^{t+1}|s_t)\left(\frac{e_{t+1}(z_{t+1})}{e_t(z_t)}\right)^{-\gamma}\left(\frac{h_{t+1}(z^{t+1})}{h_t(z^t)}\right)^{\gamma}.$$

Properties of Liquidity Shock The theory puts upper and lower bounds on the size of these weight shocks that depend only on the primitives of this economy. In the perfect insurance equilibrium, the average weights do not grow. In the autarchic equilibrium, the weights grow at a rate that equals the ratio of the largest and the smallest endowment shares.

Lemma 13. *The equilibrium average weight growth is bounded between the perfect insurance and autarchy values:*

$$1 \leq \frac{h_t(z^{t+1})}{h_t(z^t)} \leq \frac{\widehat{\eta}(y_n, z_t)}{\widehat{\eta}(y_1, z_{t+1})} \text{ for all } (z^t, z)$$

Proof of Lemma 13: First, I prove that $h_{t+1}(z^{t+1})/h_t(z^t) \geq 1$. The definition of h_t implies that:

$$\begin{aligned} h_t(z', z^{t-1}) \quad &= \quad \sum_{y^t}\int_{\underline{l}(y',z^t)}^{\infty}\zeta_{t-1}^{1/\gamma}d\Phi_{z^{t-1}}\,(dy\times d\zeta)\,\frac{\pi(y',z'|y,z)}{\pi(z'|z)} + \\ &\left(\underline{l}(y',z^t)\right)^{1/\gamma}\sum_{y^t}\int_0^{\underline{l}(y',z^t)}d\Phi_{z^{t-1}}\,(dy\times d\zeta)\,\frac{\pi(y',z'|y,z)}{\pi(z'|z)}, \end{aligned}$$

which is obviously larger than:

$$h_{t-1}(z^{t-1}) = \sum_{y^t}\int_0^{\infty}\zeta_{t-1}^{1/\gamma}d\Phi_{z^{t-1}}\,(dy\times d\zeta)\,\frac{\pi(y',z'|y,z)}{\pi(z'|z)}.$$

Second, I prove that the following inequality holds: $h_{t+1}(z^{t+1})/h_t(z^t) \leq \frac{\widehat{\eta}(y_n,z_t)}{\widehat{\eta}(y_1,z_{t+1})}$. If not, this would imply that the highest IMRS satisfies:

$$\max\left(\frac{c_{t+1}(y^{t+1},z^{t+1},\mu_0)}{c_t(y^t,z^t,\mu_0)}/\frac{e_{t+1}(z^{t+1})}{e_t(z^t)}\right)^{-\gamma} > \left(\frac{\widehat{\eta}(y_n,z_t)}{\widehat{\eta}(y_1,z_{t+1})}\right)^{\gamma},$$

which implies that the unconstrained agent is consuming less than her endowment at z^t and more than her endowment at z^{t+1}, but that can be ruled out on the basis of Lemma (11).(q.e.d.)

In an economy with limited income share dispersion, the liquidity shocks cannot be large. This multiplicative adjustment to the SDF will contribute important changes relative to its Breeden-Lucas counterpart.

IID Aggregate Uncertainty In the i.i.d. case, the liquidity shock is constant.

Proposition 14. *: If aggregate uncertainty is i.i.d. and $\pi(y'|y)$ is independent of the aggregate state, then there is a stationary equilibrium in which g^* is constant.*

Proof of Proposition 14: If aggregate uncertainty is i.i.d., the discount rate in the transformed stationary economy is constant:

$$\widehat{\beta} = \beta\sum_{z_{t+1}}\pi(z_{t+1}|z_t)\lambda(z_{t+1})^{1-\gamma}.$$

(see Alvarez and Jermann (2001)). I introduce some notation. I will use ω to denote the consumption share of an agent at the end of the previous period. Let $C(\omega, y)$ denote the cost of the consumption stream for a household in state y. Similarly, I use $C^y(y)$ to denote the cost of the labor endowment stream. Finally, $l(\omega, y)$ denotes the policy rule for the consumption weights. $\omega' = l(\omega, y')/g$ is the new consumption share. The cutoff rule $\underline{l}(y')$ depends only on y because the value of the labor income claim $C_\eta(y)$ does not depend on z^t. The distribution is rescaled at the end of each period (after the cutoff rule is applied) such that growth is eliminated from the consumption weights: $\int \omega\Phi^* (d\omega \times dy) = 1$. This is done simply by dividing all the weights by the growth rate g. The policy rules induce the following growth rate for the average weight: $g^* = \int l(\omega, y') \Phi^* (d\omega \times dy)$. This establishes the equivalence of the economy with i.i.d. aggregate uncertainty and the one without aggregate uncertainty and a twisted transition probability matrix. Given the monotonicity assumptions I have imposed on $\widehat{\pi}$, I know that the consumption weights ω live on a closed domain L because we know that the consumption shares $l(\omega, y)/g \leq \widehat{\eta}(y_n)$ from Lemma 11 and $l(\omega, y)/g \geq \widehat{\eta}(y_1)$. This implies that $\omega \in \left[\underline{l}, \bar{l}\right]$ since g is bounded. If some agent starts with an initial weight $\omega_0 \geq \bar{l}$ their consumption weight drops below $\bar{l}$ after a finite number of steps unless there is perfect risk sharing.

Let $B(L)$ the Borel set of L and let $P(Y)$ be the power set of Y. The policy function l together with the transition function π jointly define a Markov transition function on income shocks and consumption weights: $Q : (L \times Y) \times (\mathcal{B}(L) \times P(Y)) \rightarrow [0, 1]$ where

$$Q(\omega, y, \mathcal{L}, \mathcal{Y}) = \sum_{,y' \in \mathcal{Y}} \pi(y'|y),$$

if $l_h(\omega, y')/h^* \in L$. Next, define an operator on the space of probability measures $\Lambda (L \times Y) \times (\mathcal{B}(L) \times P(Y))$ as

$$T^* \Phi (\mathcal{L}, \mathcal{Y}) = \int Q(\omega, y, \mathcal{L}, \mathcal{Y}) d\Phi.$$

A fixed point of this operator is an invariant probability measure. Let Φ^* denote the invariant measure over the space $(L \times Y) \times (\mathcal{B}(L) \times P(Y))$ that satisfies invariance:

$$T^* \Phi^* (\mathcal{L}, \mathcal{Y}) = \Phi^*.$$

Clearly, if there is unique Φ^*, then there is a unique growth rate that clears the market:

$$g^* = \int \sum_{y'} \widehat{\pi}(y'|y) l_g(\omega, y') d\Phi^*(d\omega \times dy).$$

I can define a stationary equilibrium. A stationary equilibrium consists of cost functions $C(\omega, y), C^y(y)$, shadow discounter Q, updating rules $l(\omega, y)$ and an invariant measure Φ^* such that (i) the recursive updating rule is optimal: $(l(\omega, y') - \omega)$ $(C(\omega, y) - C_\eta(y)) = 0$, (ii) the market clears: $g^* = E[l(\omega, y')]$ and (iii) there is no arbitrage $Q = g^{*\gamma}$, where the expectation is taken w.r.t. Φ^*, the stationary measure over $(L \times Y) \times (\mathcal{B}(L) \times P(Y))$ induced by T^*.

It remains to be shown that this stationary measure exists. This section follows the strategy by Krueger (1999) on p.15 applied to a similar problem. I define an operator on the space of probability measures $\Lambda (L \times Y) \times (\mathcal{B}(L) \times P(Y))$ as

$$T^* \lambda (\mathcal{L}, \mathcal{Y}) = \int Q ((\omega, y), (\mathcal{L}, \mathcal{Y})) d\lambda.$$

A fixed point of this operator is defined to be an invariant probability measure. To show there exists a unique fixed point of this operator, I check condition M in (Stokey, Lucas, and Prescott (1989) p. 348). If this condition is satisfied, I can use Theorem 11.12 in Stokey, Lucas, and Prescott (1989) p. 350. To be perfectly general, let $L = [\underline{l}, l^{\max}]$. There has to be an $\varepsilon > 0$ and an $N \geq 1$ such that for all sets L, Y

$$Q^N ((\omega, y), (\mathcal{L}, \mathcal{Y})) \geq \varepsilon \text{ and } Q^N ((\omega, y), (\mathcal{L}, \mathcal{Y})^c) \geq \varepsilon \ .$$

It is sufficient to show that there exists an $\varepsilon > 0$ and an $N \geq 1$ such that for all $(\omega, y) \in (L, Y)) : Q^N\left((\omega, y), (l_{\max}, y_n)\right) \geq \varepsilon$, but we know that $Q\left((\omega, y), (l_{\max}, y_n)\right) \geq \pi(y_n|y)$. If $l_{\max} \geq \bar{l}$, then define

$$N = \min\left\{n \geq 0 : \frac{l_{\max}}{g^n} \leq \bar{l}\right\},$$

where N is finite unless there is perfect risk sharing. Then we know that $Q^N\left((\omega, y), (l_{\max}, y_n)\right) \geq \varepsilon$ where

$$\varepsilon = \pi(y_n|y) * \left(\pi(y_n|y_n)\right)^{N-1}.$$

If $\bar{l} \geq l_{\max}$, the proof is immediate by setting $\varepsilon = \pi(y_n|y)$. This establishes the existence of a unique, cross-sectional distribution and a unique g^* that clears the market.

$$Tg(\Phi^*) = \sum_{y'} \int_{\underline{l}(y')} \pi(y'|y)\omega d\Phi^* + \sum_{y'} \underline{l}\left(y'\right) \int^{l(y')} \pi(y'|y) d\Phi^*.$$

(q.e.d.)

E. Approximation

This section establishes the existence of a stationary measure over consumption weights and endowment states in the approximating equilibrium.

Let $B(L)$ the Borel set of L and let $P(Y)$ be the power set of Y. The function $l(\cdot)$ together with the transition function π jointly define a Markov transition function on income shocks and "consumption weights": $Q : \left(L \times Y \times Z^k\right) \times \left(\mathcal{B}(L) \times P(Y) \times P(Z^k)\right) \rightarrow [0, 1]$ where

$$Q\left(\left(\omega, y, z^k\right), (\mathcal{L}, \mathcal{Y}, \mathcal{Z})\right) = \sum_{,y' \in \mathcal{Y}, z'.st\ z^{k\prime} \in \mathcal{Z}} \pi(y', z'|y, z)\ if\ l_h(\omega, y', z'; z^k)/g(z^k, z') \in \mathcal{L}.$$

$$= 0 \text{ elsewhere.}$$

Next, define the operator that maps one measure into another on the space of probability measures Λ over $\left(L \times Y \times Z^k\right) \times \left(\mathcal{B}(L) \times P(Y) \times P(Z^k)\right)$ as:

$$T\lambda\left(\mathcal{L}, \mathcal{Y}, \mathcal{Z}\right) = \int Q\left(\left(\omega, y, z^k\right), (\mathcal{L}, \mathcal{Y}, \mathcal{Z})\right) d\lambda.$$

Suppose there exists a unique, invariant measure over weights, endowments and truncated aggregate histories, that is there is a stationary measure λ^* on $(S, S) = \left(L \times Y \times Z^k\right) \times \left(\mathcal{B}(L) \times P(Y) \times P(Z^k)\right)$, such that

$$\lambda^* = T^*\lambda^* = \int Q\left(\left(\omega, y, z^k\right), (\mathcal{L}, \mathcal{Y}, \mathcal{Z})\right) d\lambda^*,$$

where Q is the transition function induced by the policy function and the Markov process. Then the distribution over weights, endowments and histories is unique and stationary, for each $\left(z^{k\prime}, z^k\right) \in Z$ where $z^{k'} = (z', z^k_{k-1})$:

$$\Phi_{z^{k\prime}} = \sum_{z^k} \pi(z^{k\prime}|z^k) \int Q\left(\left(\omega, y, z^k\right), (\mathcal{L}, \mathcal{Y}, \mathcal{Z})\right) \Phi_{z^k}\left(d\omega \times dy\right).$$

If I start off this economy with this measure λ^*, it keeps reproducing itself and I can define a stationary stochastic equilibrium in which the economy moves stochastically between aggregate states and associated wealth/endowment distributions.

The optimal forecast when going from state z^k to z' is given by its unconditional average:

$$g^*(z', z^k) = \sum_{y'} \int l\left(\omega, y', z'; z^k\right) \Phi^*_{z^k}\,(d\omega \times dy)\, \frac{\pi(y', z'|y, z)}{\pi(z'|z)}, \tag{41}$$

To check that a stationary measure exists, it is sufficient to check a mixing condition (Stokey, Lucas, and Prescott (1989), p. 348).

Definition 15. *Condition M: There has to be an* $\varepsilon > 0$ *and an* $N \geq 1$ *such that for all sets* L, Y, Z^k

$$Q^N\left(\omega, y, z^k, \mathcal{L}, \mathcal{Y}, \mathcal{Z}^k\right) \geq \varepsilon \text{ or } Q^N(\omega, y, z^k, \left(\mathcal{L}, \mathcal{Y}, \mathcal{Z}^k\right)^c) \geq \varepsilon\ .$$

The standard argument can be applied. The weights live on a compact set and the upper bound $\max_{(z',z^k)} \frac{\underline{l}(y_n, z'; z^k)}{g^*(z', z^k)}$ will be reached with positive probability provided that π has no zero entries, but convergence will be slower for larger k.

F. Recursive Utility

I will adopt some methods developed by Anderson (1998) to compute Pareto-efficient allocations for recursive utility agents. The main difference with the solution for additive utility specifications is that the Pareto-Negishi weights are stochastic even if the constraints never bind.

I consider the recursive utility formulation due to Kreps and Porteus (1978). The agent's utility at time t is given by a composite index of his utility derived from current consumption and his expected future utility:

$$V_t = \left[(1-\beta)c_t^{1-\rho} + \beta\left(\mathcal{R}_t V_{t+1}\right)^{1-\rho}\right]^{1/(1-\rho)}, \tag{42}$$

where V_t is a continuation utility index and his expected future utility is given by:

$$\mathcal{R}_t V_{t+1} = \left(E\left[V_{t+1}\right]^{1-\gamma}\right)^{1/(1-\gamma)}.$$

Let the utility gradient from a time 0 perspective $M_{0,t}(\mu_0', s^t)$ be defined as $\prod_{s^\tau \preceq s^t} M_\tau(\mu_0', s^\tau)$, where the one-period-ahead gradient is given by: $M_t(\mu_0', s^t) = \left(\frac{V_t(\mu_0, s^t)}{\mathcal{R}_{t-1}V_t}\right)^{\rho-\gamma}$. As before, assume that the optimal sequence of multipliers for each type is known: $\left\{\zeta_t(\mu_0, s^t)\right\}$. I consider 2 households having experienced the same history s^t. I will resort to the dual formulation. This dual formulation is identical to the one in (DP), except for the initial promise keeping constraint:

$$U\left\{(c_t(w_0, s^t))\right\} = w_0, \tag{43}$$

where U is defined by the recursion in (42). I consider two households (μ_0', s^t) and (μ_0'', s^t). First, it is easy to verify that the ratio of marginal utilities has to satisfy this condition. This follows from the necessary first order condition for optimality in the dual recursive saddle point problem:

$$\chi_t(\mu_0', s^t)p(s^t|s_0) = \mu_0(c_t(\mu_0', s^t))^{-\rho}\mathcal{M}_{0,t}(\mu_0', s^t).$$

$$\left(\frac{c_t(\mu_0', s^t)}{c_t(\mu_0'', s^t)}\right)^{-\rho} \frac{\mathcal{M}_{0,t}(\mu_0', s^t)}{\mathcal{M}_{0,t}(\mu_0'', s^t)} = \frac{\zeta_t(\mu_0'', s^t)}{\zeta_t(\mu_0', s^t)}. \tag{44}$$

Let the stochastic Pareto-Negishi weight be defined as before, $\zeta_t(\mu_0, s^t) = \mu_0/\chi_t(\mu_0, s^t)$. Next, I define the scaled weights, using the utility gradient: $\widetilde{\zeta}_t(\mu_0, s^t) = \zeta_t(\mu_0, s^t)\ M_{0,t}(\mu_0, s^t)$. This implies that the scaled weights have a simple recursive structure:

$$\widetilde{\zeta}_t(\mu_0, s^t) = \frac{\zeta_t(\mu_0, s^t)}{\zeta_{t-1}(\mu_0, s^{t-1})} \widetilde{\zeta}_{t-1}(\mu_0, s^{t-1}) M_t(\mu_0', s^t). \tag{45}$$

While $\{\zeta\}$ is a non-decreasing stochastic process, $\left\{\widetilde{\zeta}\right\}$ is not. These scaled Pareto-Negishi weights are stochastic even if the borrowing constraints never bind ($\zeta_t = \zeta_{t-1} = \mu_0$). Only when $\gamma = \rho$ are the scaled weights are constant if the constraints do not bind. That is the standard power utility case. The average scaled weight is defined as $h_t(z^t) = E\left[\widetilde{\zeta}_t^{\frac{1}{\rho}}(z^t)\right]$. The consumption rule for an agent μ_0 in state s^t is identical to the rule in (36), but as a function of the scaled Pareto-Negishi weights:

$$c_t(\mu_0, y^t, z^t) = \frac{\widetilde{\zeta}_t^{\frac{1}{\rho}}(\mu_0, y^t, z^t)}{h_t(z^t)} e_t(z_t).$$

The stochastic discount factor is given by:

$$m_{t,t+1} = \beta \left(\frac{e_{t+1}(z_{t+1})}{e_t(z_t)}\right)^{-\rho} \left(\frac{h_{t+1}(z^{t+1})}{h_t(z^t)}\right)^{\rho}.$$

In the benchmark case of perfect risk sharing at all nodes z^t, it is easy to verify that the weight growth reduces to:

$$\left(\frac{h_{t+1}(z^{t+1})}{h_t(z^t)}\right)^{\rho} = \left(\frac{V_{e,t+1}(\mu_0, s^t)}{\mathcal{R}_t V_{e,t+1}}\right)^{\rho-\gamma},$$

where the last term denotes the utility gradient of a stand-in household that consumes the aggregate endowment.

Computation In the case of recursive utility, the consumption share policy function is modified as follows:

$$\begin{aligned} l(\omega, y', z'; z^k) &= \omega M^{\frac{1}{\rho}}(\omega, y', z'; z^k) \text{ if } \omega > \underline{\omega}(y', z'; z^k) \\ &= \underline{\omega}(y', z'; z^k) \text{ elsewhere,} \end{aligned}$$

The SDF is given by $m(z', z^k) = \beta g^*(z', z^k)^{\rho} \lambda(z')^{-\rho}$. $M(\omega, y', z'; z^k)$ needs to be included in the list of functions to be computed to characterize the approximate equilibrium. Everything else is identical to the additive utility case.