

The Market Price of Aggregate Risk and the Wealth Distribution

Hanno Lustig*

UCLA

January 9, 2007

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. 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. The economy is said to experience a negative liquidity shock when this growth rate is high, a large fraction of agents faces severely binding solvency constraints and the trading volume is low in financial markets. The adjustment to the Breeden-Lucas stochastic discount factor induces time variation in equity, bond and currency risk premia that is consistent with the data.

Keywords: Asset Pricing, Collateral, Liquidity.

JEL: G12,E21

*email: hlustig@econ.ucla.edu. Department of Economics UCLA Box 951477 Los Angeles, CA 90095-1477. <http://www.econ.ucla.edu/hlustig/> 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.

1 Introduction

I develop a model of an exchange economy with a continuum of agents who have power utility with risk aversion coefficient γ , 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 7 percent, a risk-free rate of 1 percent and substantial variation in equity risk premia, as well as an upward sloping yield curve, consistent with the data. This variation in risk premia is driven by shocks to the wealth distribution induced by these solvency constraints. 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 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 in this economy while they do in the data (see e.g. Campbell and Cochrane (1999)). My model produces an additional risk factor that addresses these puzzles.

Since aggregate endowment growth is i.i.d., there are no built-in dynamics in risk premia. Beyond the constant risk in the aggregate endowment process, the bankruptcy technology contributes a second source of time-varying risk, the risk associated with binding solvency constraints¹. I call this liquidity risk. 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 $\beta\lambda_{t+1}^{-\gamma}$ (stochastic discount factor): the growth rate of the γ^{-1} -th moment of the distribution of stochastic Pareto-Negishi weights, denoted g_{t+1} , raised to the power γ :

$$m_{t+1} = \beta\lambda_{t+1}^{-\gamma}g_{t+1}^{\gamma}$$

¹This paper follows He and Pearson (1991) and Luttmer (1992) in exploring solvency constraints as a device for understanding asset pricing anomalies.

This last component reflects the aggregate shadow cost of the solvency constraints. If this growth rate is high, a large fraction of agents is constrained, trading volume is low and the economy is said to be hit by a negative liquidity shock. Figure 1 plots these liquidity shocks g^γ in the top panel and the trading volume in the bottom panel for 55 periods simulated from a calibrated version of my model. The aggregate trade volume in financial markets drops by 20 % when there is a large liquidity shocks, after which it gradually recovers. The shaded areas indicate low aggregate consumption growth states. There is a growing body of evidence that aggregate liquidity risk is priced, both from the cross-section and the time-series variation in stock returns (Pastor and Stambaugh (2003)). This papers delivers a theoretical underpinning for these findings in a model with solvency constraints as the only trading friction.

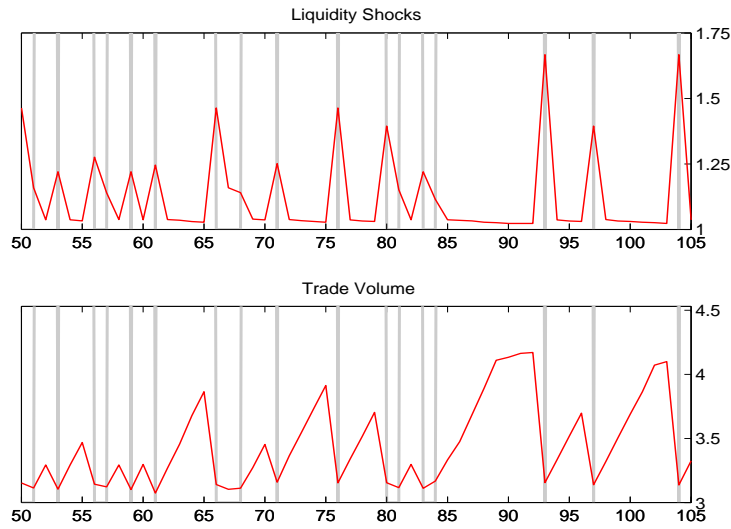


Figure 1: Liquidity Shocks and Trading Volume In the top panel, the full line is the aggregate liquidity shock g^γ over 55 years, generated by simulating the model. In the bottom panel, the full line is the aggregate trading volume in financial markets, $\frac{1}{e_t(z^t)} \int \sum_{y^t} \text{abs}(\Pi_{s^t}\{c(\mu_0, s^t)\} - \Pi_{s^t}\{\eta(s^t)\}) d\Phi_0$ over 55 years. The shaded areas are low aggregate consumption growth states. γ is 7, α is 7.5% and β is .95.

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 when 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 from the contract when that state of the world is realized. 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. 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.

To deal with a continuum of consumers and aggregate uncertainty, I extend the methods developed by and Krueger (1999). Building on work by Atkeson and Lucas (1992,1995), 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: equilibrium state 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.

There is a growing literature on collateral constraints and asset prices. Geanakoplos and Zame (1998)(henceforth GZ) consider an environment in which households can default on their promises at any time, and financial securities are only traded if the promises associated with these securities are backed by collateral. What distinguishes my setup from GZ is the fact that only outright default on all promises is allowed, not default on individual obligations. Kubler and Schmedders (2003) develop a computational algorithm for an infinite horizon version of the GZ economy.

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 results; the fifth section discusses the computation. All the proofs are in the appendix.

2 Environment and Equilibrium

2.1 Uncertainty

The events $s = (y, z)$ take on values on a discrete grid $S = Y \times Z$ where $Y = \{y_1, y_2, \dots, y_n\}$ and $Z = \{z_1, z_2, \dots, z_m\}$. 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 . The transition probabilities for idiosyncratic and aggregate shocks are assumed to be independent.

Condition 2.1. *The transition probabilities can be stated as:*

$$\pi(y', z'|y, z) = \varphi(y'|y)\phi(z'|z)$$

I assume the transition matrix for idiosyncratic events y , $\phi(y'|y)$, satisfies monotonicity and there are no absorbing states, $\phi(y'|y) \gg 0$. Finally, I also assume the aggregate shocks are independent over time:

Condition 2.2. *The aggregate shocks are i.i.d.:*

$$\phi(z'|z) = \phi(z')$$

2.2 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:

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

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 $\hat{\eta}(y_t, z_t)$ 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 $\eta_t(s^t) = \hat{\eta}(y_t, z_t)e_t(z^t)$, with $s^t = (s^{t-1}, y, z)$. I assume $\hat{\eta}(y_{i+1}, z_t) > \hat{\eta}(y_i, z_t)$ and $\hat{\eta}(y_t, z_t) \gg 0$ 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:

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

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

2.3 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:

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

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:

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

2.4 Enforcement Technology

In this literature, it has been common to assume that households can be 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:

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

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:

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

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 3.2) and they support less risk sharing as a result.

2.5 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:

$$c_t(\theta_0, s^t) + p_t^e(z^t) \phi_t(\theta_0, s^t) + \sum_{s'} a_t(s'; \theta_0, s^t) q_t(s^t, s') \leq \theta_t, \quad (5)$$

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

$$\begin{aligned} [p_{t+1}^e(z^{t+1}) + \alpha e_{t+1}(z^{t+1})] \phi_t(\theta_0, s^t) &\geq -a_t(s'; \theta_0, s^t) \text{ for all } s' \in S, \\ \text{where } s^{t+1} &= (s^t, s'). \end{aligned} \quad (6)$$

The appropriate transversality conditions read as

$$\lim_{t \rightarrow \infty} \sum_{s^t} \beta^t \pi(s^t | s_0) u'(\hat{c}_t(s^t)) [a_t(s'; \theta_0, s^t) + [p_{t+1}^e(z^{t+1}) + \alpha e_{t+1}(z^{t+1})] \phi_t(\theta_0, s^t)] = 0$$

The definition of a competitive equilibrium is straightforward.

Definition 2.1. *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*

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

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

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} .

3 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 $1-\gamma$ -th moment of the weight distribution.

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

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

$$\Pi_{s^t}[\{c\}] \geq \Pi_{s^t}[\{\eta\}], \text{ for each } s^t. \quad (7)$$

3.2 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 $\kappa_t(s^t)$ be the continuation utility associated with bankruptcy, conditional on a pricing functional Π :

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

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

$$U(\{c(\theta_0, y^t, z^t)\})(s^t) \geq \kappa_t(s^t) \text{ for all } s^t. \quad (8)$$

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)². These put a lower bound on the value of the household's consumption claim.

Proposition 3.1. *For given Π , an allocation is said to be immune to bankruptcy iff:*

$$\Pi_{s^t} [\{c(\theta_0, y^t, z^t)\}] \geq \Pi_{s^t} [\{\eta\}], \text{ for all } s^t \in S^t, t \geq 0. \quad (9)$$

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 default³.

3.3 Risk Sharing

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

²Zhang (1997) first endogenized borrowing constraints in a class of incomplete markets models.

³Detemple and Serrat (2003) consider an environment in which only a fraction of agents face these constraints. They find small effects on risk premia.

No Collateral 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 have to be violated for another one with positive measure⁴.

Proposition 3.2. *If there is no outside wealth ($\alpha = 0$), then there can be no risk sharing in equilibrium.*

Perfect Risk Sharing When 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 3.3. *If the value of the aggregate endowment exceeds the value of the private endowment at all nodes, perfect risk sharing is feasible:*

$$\Pi_{s^t}^* [\{e\}] \geq \Pi_{s^t}^* [\{\eta\}] \text{ for all } s^t.$$

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

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

$$\Pi_{s^t} [\{c\}] \geq \Pi_{s^t} [\{\eta\}] \geq B_{s^t}^{aut} [\{\eta\}],$$

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

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.

⁴Krueger and Uhlig (2005) derive a similar result in an environment with one-sided commitment, on the part of financial intermediaries.

Loosely speaking, the Pareto frontier shifts down as one moves from permanent exclusion to bankruptcy.

3.4 Kehoe-Levine Equilibrium

This section sets up the household's problem and defines an equilibrium, when all trading occurs at time zero. Taking prices $\{p_t(s^t|s_0)\}$ 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:

$$\begin{aligned}
 & \textbf{Primal Problem} \tag{PP} \\
 & \sup_{\{c\}} u(c_0(\theta_0, s^0)) + \sum_{t=1} \sum_{s^t \geq s^0} \beta^t \pi(s^t|s_0) u(c_t(\theta_0, s^t)), \\
 & \sum_{t \geq 0} \sum_{s^t \geq s_0} p_t(s^t|s^0) [c_t(\theta_0, s^t) - \eta_t(s^t)] \leq \theta_0, \\
 & \Pi_{s^t} [\{c(\theta_0, y^t, z^t)\}] \geq \Pi_{s^t} [\{\eta\}], \text{ for all } s^t \in S^t, t \geq 0.
 \end{aligned}$$

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 3.2. *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 :

$$\sum_{y^t} \int c_t(\theta_0, y^t, z^t) \varphi(y^t|y_0) d\Theta_0 = e_t(z_t). \tag{10}$$

In equilibrium households solve their optimization problem subject to the participation constraints and the markets clear. I assume that the endowments are finitely valued in equilibrium.

Condition 3.1. *Interest rates are high enough:*

$$\Pi_{s^0} [\{\eta\}] < \infty \text{ for all } y_0 \text{ and } \Pi_{z^0} [\{e\}] < \infty. \tag{11}$$

When 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. 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 $\Pi_{s^0} [\{c(\mu_0, s^t)\}] < \infty$. This means I can apply Proposition 4.6 in Alvarez and Jermann (2000)⁵.

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 do not solve a planner's resource allocation problem, but I characterize equilibrium allocations and prices from the household's first order conditions.

3.5 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:

$$c_t(\mu_0, s^t) = \frac{\mu_0^{1/\gamma}}{E\mu_0^{1/\gamma}} e_t(z^t) \text{ where } s^t = (y^t, z^t), \quad (12)$$

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 . $\{\zeta_t(\mu_0, s^t)\}$ 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 (7) is satisfied with equality. Typically, these are states with high labor income realizations. These weights record the sum of all solvency constraint

⁵This proposition demonstrates the equivalence between the Arrow-Debreu economy and the economy with sequential trading, provided that there is a ξ such that

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

which is automatically satisfied for power utility.

multipliers in history s^t .

Consumption is characterized by the same linear risk sharing rule:

$$c_t(\mu_0, s^t) = \frac{\zeta_t^{1/\gamma}(\mu_0, s^t)}{E \left[\zeta_t^{1/\gamma}(\mu_0, s^t) \right]} e_t(z^t), \quad (13)$$

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 $\{h_t(z^t)\}$ 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 -I call this a liquidity shock. $\{h_t(z^t)\}$ can be interpreted as the aggregate shadow cost of the solvency constraints. I will refer to this simply as the *average weight* process.

To derive this consumption sharing rule, 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)

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

$$\sum_{t \geq 0} \sum_{s^t \geq s^0} \beta^t \pi(s^t | s_0) u(c_t(w_0, s^t)) = w_0, \quad (14)$$

$$\Pi_{s^t} [\{c(w_0, y^t, z^t)\}] \geq \Pi_{s^t} [\{\eta\}], \text{ for all } s^t \in S^t, t \geq 0. \quad (15)$$

The convexity of the constraint set implies that the minimizer of *DP* and the maximizer of *PP* (the primal problem) coincide for initial wealth $\theta_0 = C^*(w_0, s^0) - \Pi_{s^0} [\{\eta\}]$ (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 Marcat 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:

$$\chi_t(w_0, s^t) = \chi_{t-1}(w_0, s^{t-1}) - \tau_t(w_0, s^t), \quad \chi_0 = 1. \quad (16)$$

Let μ_0 denotes the Lagrangian multiplier on the initial promised utility constraint in (14). 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 ($\tau_t(s^t) > 0$), 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, I can transform the original dual program into a recursive saddle point problem for household (w_0, s_0) by introducing a cumulative multiplier:

$$D(c, \chi; w_0, s_0) = \sum_{t \geq 0} \sum_{s^t} \left\{ \beta^t \pi(s^t | s_0) m_t(s^t | s_0) \left[\begin{array}{l} \chi_t(s^t | s_0) c_t(w_0, s^t) \\ + \tau_t(s^t) \Pi_{s^t} [\{\eta\}] \end{array} \right] \right\}, \quad (17)$$

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:

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

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. This allows me to recover allocations and prices from the equilibrium sequence of multipliers $\{\zeta_t(\mu_0, s^t)\}$. I will proceed in two steps.

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:

$$\left[\frac{c_t(\mu'_0, s^t)}{c_t(\mu''_0, s^t)} \right]^{-\gamma} = \frac{\zeta_t(\mu''_0, s^t)}{\zeta_t(\mu'_0, s^t)}. \quad (18)$$

If the constraints never bind, $\zeta_t = \mu_0$ at all nodes and the condition in (18) 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:

$$\sum_{y^t} \int c_t(\mu_0, y^t, z^t) \varphi(y^t | y_0) d\Phi_0 = e_t(z^t), \quad (19)$$

(18) and (19) 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:

$$c_t(\mu_0, s^t) = \frac{\zeta_t^{1/\gamma}(\mu_0, s^t)}{E \left[\zeta_t^{1/\gamma}(\mu_0, s^t) \right]} e_t(z^t). \quad (20)$$

This rule satisfies the condition on the ratio of marginal utilities (18) 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.

Cutoff Rule I derive a simple characterization of the optimal weight policy and then I show that these weights fully characterize an equilibrium. The optimal policy rule has a simple recursive structure. Let $C(\mu_0, s^t; \zeta)$ denote the continuation cost of a consumption claim derived from a weight policy $\{\zeta_t(\mu_0, s^t)\}$:

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

where consumption at each node is given by the risk sharing rule in (20). 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:

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

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

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

The following theorem explains that an equilibrium is fully characterized by these Pareto-Negishi weight processes.

Theorem 3.1. *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 :

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

and (iii) there are no arbitrage opportunities :

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

Properties of the Cutoff Rule These cutoff rules have two key properties that will prove useful for understanding the consumption and wealth dynamics inside the model, and for solving the model. 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 3.2. *The consumption shares at the cutoff do not exceed the labor endowment shares:*

$$\frac{l_t^{1/\gamma}(z^t, y)}{h_t(z^t)} \leq \widehat{\eta}(y, z) \text{ for all } (z^t, y) \quad (21)$$

Of course, as the collateralizable share of income decreases, the cutoff consumption shares approach the labor endowment shares; when $\alpha = 0$, equation (21) 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 3.3. *If the transition matrix satisfies monotonicity, then the cutoff rules can be ranked:*

$$l_t(z^t, y_n) \geq l_t(z^t, y_{n-1}) \geq l_t(z^t, y_{n-2}) \geq \dots \geq l_t(z^t, y_1)$$

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

What are the implications for household consumption? Suppose perfect risk sharing cannot be sustained, and $h \gg 1$. 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:

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

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

$$\hat{\eta}(y_1, z) \leq \varpi(\mu_0, s^t) < \hat{\eta}(y_1, z) \text{ for all } (\mu_0, s^t) \text{ and } t \geq \tau \quad (22)$$

This follows directly Lemma (3.3) and (3.2). All households face at least one binding solvency constraint, in the highest state y tomorrow. In this environment, wealthy agents simply run down their wealth, until they reach the region of binding 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. So, if an agent were to start off with a lot of financial wealth at time 0, her consumption share $\varpi(\mu_0, s^t)$ would keep drifting down until she reaches the region in which the solvency constraints start to bind. This is the signature of complete markets: there is no motive for unconstrained households to accumulate wealth. 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. Wealthy households chose to run down their assets because interest rates are low. 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 also points to a crucial distinction between this model and standard incomplete market models. 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. The next subsection derives an expression for the SDF.

3.6 Risk Premia

The structure of the SDF is very revealing. 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 liquidity shock, raised to the power γ .

Proposition 3.5. *The equilibrium SDF is given by:*

$$m_{t+1} = \beta \left(\frac{e_{t+1}}{e_t} \right)^{-\gamma} \left(\frac{h_{t+1}}{h_t} \right)^\gamma. \quad (23)$$

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 (23).

Bounds The theory puts upper and lower bounds on the size of these liquidity 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 3.4. *The equilibrium average weight growth is bounded between the perfect insurance and autarchy values:*

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

When all households are constrained, the SDF equals the autarchic IMRS of the household switching from the highest to the lowest income state. 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.

Why are these *liquidity* shocks? If $g = 1$, then the economy sustains the maximum amount of trading, to implement complete risk insurance. The aggregate volume of trade in node z^t is measured by the average (across households) distance between the consumption and the endowment stream in present discounted value, scaled by the level of the aggregate endowment:

$$\frac{1}{e_t(z^t)} \int \sum_{y^t} \text{abs} (\Pi_{s^t} [\{c(\mu_0, s^t)\}] - \Pi_{s^t} [\{\eta(s^t)\}]) d\Phi_0. \quad (24)$$

This is a direct measure of how far the allocations are from autarchy. The trading volume in financial markets peaks when perfect insurance is implemented. On the other hand, when g hits the upper bound, the trading volume reaches the absolute minimum (zero). So, g is a perfect liquidity indicator. The size of these liquidity shocks is governed by the mass of households in the left tail of the wealth distribution, as explained in the next subsection.

Liquidity Shocks and the Wealth Distribution I use consumption weights as stationary state variables to replace the Pareto-Negishi weights. $g_t(z^t)$ denotes the growth rate of the aggregate weight process h_t/h_{t-1} . At the end of each period, I re-normalize the weights into consumption shares:

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

and I store this as the household's state variable. Φ_{z^t} denotes the joint measure over (y, ω) in state z^t . 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:

$$\begin{aligned} \omega'(y', z^t; \omega) &= \omega \text{ if } \omega > \underline{\omega}(y', z^t) \\ &= \underline{\omega}(y', z^t) \text{ elsewhere} \end{aligned} \quad (25)$$

Making use of the cutoff rule, the liquidity shock g_{t+1} can be stated as follows:

$$g_t(z', z^{t-1}) = \sum_{y'} \int_{\underline{\omega}(y', z^t)}^{\infty} \omega \varphi(y'|y) d\Phi_{z^{t-1}}(dy \times d\omega) + \quad (26)$$

$$\sum_{y'} \underline{\omega}(y', z^t) \int_0^{\underline{\omega}(y', z^t)} \varphi(y'|y) d\Phi_{z^{t-1}}(dy \times d\omega). \quad (27)$$

It immediately follows that $g \geq 1$, because $\sum_{y'} \int \varphi(y'|y) \omega d\Phi_{z^{t-1}}(dy \times d\omega) = 1$ by construction. The size of the liquidity shock is determined by the mass of households in the left tail. In general, the size of these shocks depends on the entire aggregate history z^{t6} . However, if labor income risk is independent of the aggregate shocks, these liquidity shocks are constant. We start by considering this simple case.

Benchmark: Independent Labor Income risk

Condition 3.2. *The labor income shocks are independent of the aggregate shocks if $\hat{\eta}(y_t, z_t) = \hat{\eta}(y_t)$*

In this case, it is easy to show that the cutoff weight $\underline{\omega}(y')$ does not depend on the aggregate history, simply because the price of a claim to labor income relative to the level of the aggregate endowment, $\Pi_{s^t}[\{\eta\}]/e_t(z^t)$, does not depend on z_t . Hence, neither does the cutoff weight $\underline{\omega}(y')$. As a result, after the transitional dynamics have dissipated, the liquidity shock is constant and so is the joint distribution of consumption weights and endowments.

⁶This creates a computational problem that I deal with in section 5.

Proposition 3.6. : *If aggregate uncertainty is i.i.d. and labor income risk is independent of the aggregate state, then there is a stationary equilibrium in which g^* is constant.*

$$g = \sum_{y'} \int_{\underline{\omega}(y')}^{\infty} \omega \varphi(y'|y) d\Phi(dy \times d\omega) + \quad (28)$$

$$\sum_{y'} \underline{\omega}(y') \int_0^{\underline{\omega}(y')} \varphi(y'|y) d\Phi(dy \times d\omega) \quad (29)$$

The mass of households in the left tail is constant over time.

Asset Pricing Implications Absent any arbitrage opportunities, payoffs in state z^{t+1} are priced by:

$$m_{t+1} = \beta \left(\frac{e_{t+1}}{e_t} \right)^{-\gamma} g^\gamma \quad (30)$$

The second part, g^γ , is constant in the case of independent labor income risk. As a result, the liquidity constraints push up the price of consumption in all states tomorrow. This lowers the risk-free rate, but it does not change risk premia relative to the full insurance benchmark. Next, I relax this independence assumption, and I look at a calibrated version of the model.

4 Results

This section starts by explaining the calibration, then we discuss the dynamics of the liquidity shocks and their connection to the wealth distribution and the dynamics of consumption. Finally, I conclude by discussing the asset pricing implications.

4.1 Calibration

We choose a γ of seven and a time discount factor β of .95. These preference parameters allow us to match the collateralizable wealth to income ratio in the data when the collateralizable income share α is 7.5%, as discussed below.

Collateralizable Wealth The average ratio of collateralizable wealth to aggregate income in the US is 3.87 between 1950 and 2005. The wealth measure includes the value of the non-financial corporate sector and the value of residential wealth (Flow of Funds). We exclude government debt. Aggregate income includes tradeable income (payouts to securities owners

of the non-financial, corporate sector, and rental income) and non-tradeable income (labor income and proprietary income). The collateralizable wealth to income ratio is 3.87 in the post-war US data. Tradeable or collateralizable income is 8.3% of total income. In the model, we choose α equal to 7.5 % to match the collateralizable wealth to income ratio of 3.87 in the data⁷. The details are in the appendix in section B. 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.

Aggregate and Idiosyncratic Endowment Risk The Markov process for $\log \eta(y, z)$ is taken from Storesletten, Telmer, and Yaron (2006) (see page 28). We use a 4-state discretization. The conditional variance in recessions and booms is 0.181 and 0.037, and the autocorrelation is 0.89. The elements of the process $\log \eta$ are $\{-2.385, 0.646\}$ in low aggregate consumption growth states and $\{-0.904, 0.467\}$ in high aggregate consumption growth states. Labor income risk doubles in low aggregate consumption growth states. Aggregate consumption growth $\lambda(z_t)$ is i.i.d. This ensures all the dynamics in risk premia flow from the liquidity shocks. The moments for aggregate consumption growth are taken from Mehra and Prescott (1985). The average consumption growth rate is 1.8 %. The standard deviation is 3.15 %. Recessions are less frequent: 27% of realizations are low aggregate consumption growth states. Finally, section 5 explains the computational procedure in detail.

4.2 Liquidity Shocks

While the aggregate consumption growth shocks are i.i.d., the wealth dynamics induced by these shocks are not, as is clear from figure 1. In this calibrated version of the model, the liquidity shocks vary a lot in size depending on the history of aggregate consumption growth shocks.

In this calibrated version of the model, liquidity shocks are larger in low aggregate consumption growth states, because the increase in the cross-sectional variation of idiosyncratic income shocks raises the cutoff values $\underline{\omega}(y', z^t)$, as is clear from inspecting the expression for g in equation (27). In the bottom panel of figure 1, we also plot the trade volume (defined in equation (24)). Large liquidity shocks coincide with equally large drops in trade volume. When these large shocks occur, the fraction of constrained agents increases to forty-five percent, compared to 10 percent in high aggregate consumption growth states.

History of Aggregate Consumption Growth Shocks The model also produces history dependence in these liquidity shocks. Figure 2 plots the liquidity shocks and the moments of

⁷The ratio is 3.79 in the model's benchmark calibration.

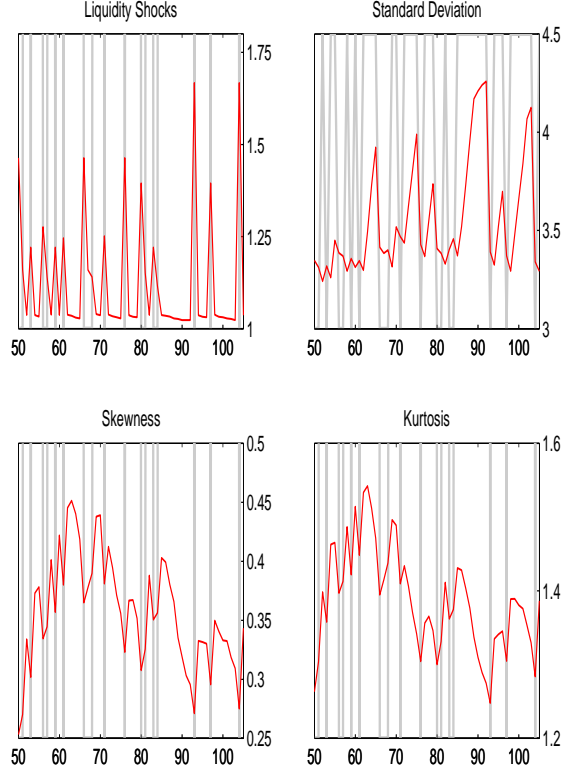


Figure 2: The Distribution of Wealth. The figure plots the liquidity shocks and the centered moments of cross-sectional distribution of wealth over 55 years, generated by simulating the model. Household wealth is defined as: $\frac{1}{e_t(z^t)} (\Pi_{s^t}[\{c(\mu_0, s^t)\}] - \Pi_{s^t}[\{\eta(s^t)\}])$. The shaded areas are low aggregate consumption growth states. γ is 7, α is 7.5% and β is .95.

the wealth distribution; we keep the same draw of aggregate shocks as in figure 1. ‘Household financial wealth’ is defined as

$$\frac{1}{e_t(z^t)} (\Pi_{s^t}[\{c(\mu_0, s^t)\}] - \Pi_{s^t}[\{\eta(s^t)\}]) .$$

I divide by the level of the aggregate endowment, to render it stationary. During a long series of high aggregate consumption growth realizations (say between period 83 and 93 in figure 2), there is a build-up of low wealth households in the left tail of the wealth distribution. The standard deviation of the wealth distribution increases. Mechanically, this means the mass of agents with weights below the cutoff value is large:

$$\sum_{y'} \int_0^{\omega(y', z^t)} \varphi(y'|y) d\Phi_{z^{t-1}}(dy \times d\omega) \quad (31)$$

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, and as a result, 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. 27). The left tail of the wealth distribution is eliminated, and the standard deviation of the wealth distribution drops and so does the skewness and the kurtosis.

4.3 Consumption Dynamics

Risk Sharing This calibrated version of the collateral economy sustains a lot of risk sharing. In the benchmark calibration the standard deviation of consumption share growth for households is 7.5 percent, less than twice the standard deviation of aggregate consumption growth, while the standard deviation of endowment share growth is thirty-three percent. Not all agents in states with binding solvency constraints experience large shocks to their consumption shares. In the 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 (after consecutive low aggregate consumption growth shocks) almost all households have roughly constant consumption shares.

The left panel of figure 3 plots the consumption share as a fraction of the average endowment of a single household against its labor income share. The consumption shares fluctuate between the highest and the lowest income shares. This is what I showed in equation (22). 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 and produce large liquidity shocks when aggregated across consumers. Recessions are periods when the aggregate show cost of the solvency constraint increases.

The right panel of figure 3 plots the 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. 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,




Figure 3: Solvency Constraints In the left panel, the full line is the consumption share of a household plotted against the labor income share (dotted line) over a period of 300 years. The y-axis on the left hand side shows the consumption shares; on the right hand side is the labor income share. In the right panel, 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 55 years. The y-axis on the left hand side shows the endowment shares; on the right hand side is net wealth. γ is 7, α is 7.5% and β is .95.

but the hedge is incomplete because of the collateral constraints. To illustrate the difference, figure 4 plots the consumption share and the wealth of a household not facing any solvency constraints, for the same history of shocks. This household is perfectly hedged, and net wealth is negative in the high y states.

4.4 Asset Pricing Results

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 to match the dynamics of equity and bond risk premia.

Dividend Process Following Bansal and Yaron (2004), dividend growth is a function of aggregate consumption growth and the change in the dividend/consumption ratio q_t :

$$\begin{aligned}\Delta d_{t+1} &= \delta + \phi \Delta c_{t+1} + \Delta q_{t+1} \\ q_{t+1} &= \rho_q q_t + \varphi_d \sigma u_{t+1}\end{aligned}\tag{32}$$

u is white noise with mean zero and variance 1. σ is the standard deviation of aggregate consumption growth. Following Bansal and Yaron (2004), I choose $\rho_q = 0.8$ for the autocor-

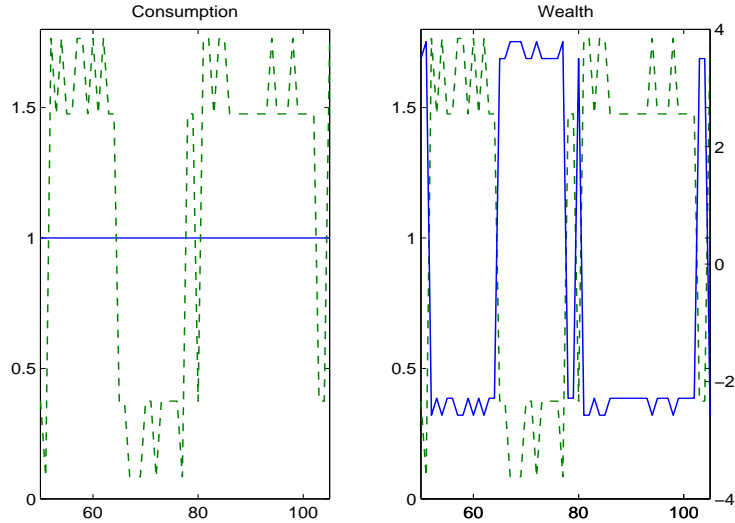


Figure 4: No Solvency Constraints In the left panel, the full line is the consumption share of a household plotted against the labor income share (dotted line) over a period of 300 years. The y-axis on the left hand side shows the consumption shares; on the right hand side is the labor income share. In the right panel, 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 55 years. The y-axis on the left hand side shows the endowment shares; on the right hand side is net wealth. γ is 7, α is 7.5% and β is .95.

relation of the consumption/dividend ratio, the leverage parameter ϕ in the dividend growth process is set to 3, and $\varphi_d = 4.5$. Equity is a claim to this dividend process.

My benchmark calibration sets the time discount factor β equal to .95 and γ to 7. Table 1 compares the moments of the data, the representative agent model and the collateral model. The excess return on equity is denoted R^e , while $R^{c,e}$ denotes the excess return on a non-levered claim to the aggregate endowment process. The asset pricing statistics were generated by drawing 10,000 realizations from the model, simulated with 5000 agents.

Representative agent The benchmark perfect insurance economy produces a risk-free rate of thirteen percent and an equity premium of 2.8 percent, one percent for the non-levered claim to consumption. This is the risk-free rate and the equity premium puzzle. In addition, the model produces a constant conditional market price of risk (second column of Table 1).

Collateral economy In the collateral model, the maximum Sharpe Ratio is .43. The liquidity risk induced by the solvency constraints delivers a low risk-free rate of 19 basis points and a high equity premium of 6.63 percentage points. The compensation per unit of risk is large as well; the Sharpe ratio on the non-levered claim is around 39 percent, compared to 38 percent in the data. The standard deviation of the conditional market price of risk in

Table 1: Benchmark Calibration Results.

R^e is the return on a leveraged dividend claim; $R^{c,e}$ is the excess return on a claim to aggregate consumption. The first panel shows moments for the data, the second panel for the representative agent model and the third panel for the collateral model. These moments were generated by averaging 10,000 draws from an economy with 5000 agents. α is 7.5 percent, γ is 7 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. The sample is 1928-2005.

	$\frac{\sigma(m)}{E(m)}$	$std \frac{\sigma_t(m)}{E_t(m)}$	$E(R^e)$	$std(R^e)$	$\frac{E(R^e)}{std(R^e)}$	$E(R^{c,e})$	$std(R^{c,e})$	$\frac{E(R^{c,e})}{std(R^{c,e})}$	$E(r^f)$	$std(r^f)$
<i>data</i>			7.08	20.0	0.382				0.96	4.80
<i>rep</i>	0.279	0	2.88	13.18	0.218	1.06	4.10	0.260	15.03	0
<i>coll</i>	0.433	0.07	6.63	16.78	0.394	4.18	10.55	0.397	0.193	5.31

the collateral model (second column of Table 1) is 7 percent. Finally, the model overstates the volatility of the risk-free rate relative to the data.

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 (30), 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 because 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, as is the case in this calibrated version of the economy (see subsection 4.2).

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

$$E_t \log R_{t+1}^i - \log r_t^f = \gamma cov_t(\Delta \log(e_{t+1}), \log R_{t+1}^i) - \gamma cov_t(\log(g_{t+1}), \log R_{t+1}^i)$$

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 calibration⁸.

Figure 1 plots liquidity shocks and trading volume for the same sequence of aggregate

⁸In the data, this liquidity premium is large. Pastor and Stambaugh (2003) find that the average return on stocks with high sensitivities to liquidity exceeds that for stocks with low sensitivities by 7.5 % annually.

shocks. These shocks coincide with large negative innovations to trading volume. As is clear from figure 1, since trading volume or liquidity co-moves with all other returns, it also predicts future excess returns in the collateral model. This is consistent with the evidence from the data (see for example Pastor and Stambaugh (2003)).

Bonds The same mechanism increases risk premia on bonds of longer maturity. Since the aggregate shocks are i.i.d., the representative agent model produces a flat yield curve, but the collateral model produces an upward sloping yield curve on average. Let p_t^N denote the log of the price of a N-year zero coupon bond. The yield $y_t^N = -p_t^N/N$ increases from zero to 2.1 % at 10 years. The average excess return $p_t^N - p_{t-1}^{N+1} - r_{t-1}^f$ (in logs) on a 2-year zero coupon bond is 1.28 % and it is 2.78% on a 10-year bond. Longer maturity bonds are riskier:

$$-cov_t(\log(g_{t+1}), p_t^N - p_{t-1}^{N+1})$$

increases as N increases, because the liquidity shocks trigger a persistent, subsequent increase in the short rate, inflicting larger losses on holders of longer maturity bonds. As a result, these zero coupon bond holding period returns become more sensitive to liquidity shocks as the maturity increases.

Time Variation in Risk Premia Recall that, in the representative agent economy, the conditional Sharpe ratio, the conditional risk premium, the conditional volatility, the risk-free rate, the slope of the yield curve and the trading volume are all constant over time, because the aggregate shocks are i.i.d. The collateral model's liquidity shocks generate time varying risk premia.

Figure 5 plots some key asset pricing statistics that illustrate the time variation in risk premia. The first plot shows the liquidity shocks, while the second plot shows the conditional Sharpe ratio on equity, which varies from .3 , after a recession, to .55, after a long series of high aggregate consumption growth shocks. The conditional risk premium on equity (third plot) varies between 5 and 10 percent. After a long series of high aggregate consumption growth realizations, the risk-free rate (fifth plot) drops and the conditional market price of risk increases. The low risk-free rate predicts high excess returns on equity, because it signals large liquidity shocks are likely to occur. This reflects the build-up of households in the left tail of the wealth distribution. At the same time, the conditional volatility of equity returns (fourth plot) increases as well. The slope of the yield curve increases in anticipation of a large liquidity shock, and this partly reflects an increase in the risk premium. After the recession, the conditional market price of risk drops to its lowest level, and the risk-free rate increases sharply, while the yield curve flattens. Interestingly, the negative correlation

between the conditional market price of risk and the riskfree rate is a sufficient condition to explain the forward premium puzzle (Lustig and Verdelhan (2007)).

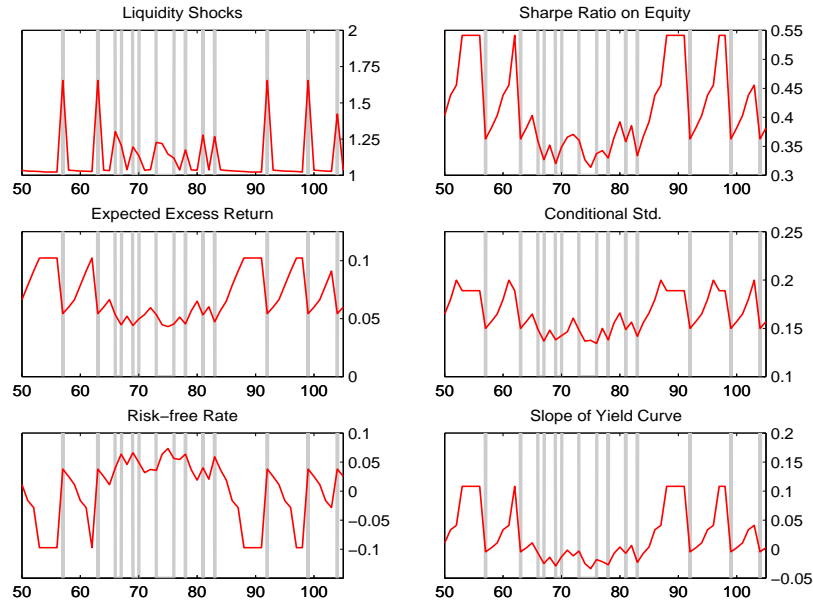


Figure 5: Liquidity Shocks, Market Price of Risk and Aggregate Consumption Growth The shaded area indicates low aggregate consumption growth states. The upper panel plots the liquidity shocks (left) and the conditional Sharpe ratio on equity (right). The middle panel plots the conditional expected excess return on equity (left) and the conditional standard deviation (right). The lower panel plots the risk-free rate (left) and the slope of the yield curve (right). The slope of the yield curve is $y_t^{10} - y_t^1$. β is .95, γ is 7, and α is 7.5%.

Summary What is critical for these result? I list three ingredients: (i) scarcity of collateral, (ii) state-contingent nature of the constraints, and (iii) large number of agents.

- The quantity of collateral: The quantity of collateral was calibrated to match ratio of collateralizable wealth to total income of 3.8 in US data. The collateralizable income fraction α is 7.5 %, compared to 8.3 % in the data. This calibrated version of the collateral economy also matches the equity premium and the risk-free rate. If α is set equal to 10 %, the risk-free rate increases to 1.28 % and the equity premium drops to 6.41 %. When α is 15 %, the equity premium drops to 5.7 % and the risk-free rate increases to 3.91 %. Finally, when α is 30 percent, the collateral economy is identical to the representative agent economy; the solvency constraints no longer bind.
- The state-contingent nature of the collateral constraints, not the tightness of the constraints. In the same environment, an “exogenous” constraint on the value of the net

wealth today contributes a 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 (see Luttmer (1991)). This factor only lowers the risk-free rate, but does not affect risk premia.

- Large number of agents: In the same economy with two agents and i.i.d. aggregate shocks, risk premia are essentially constant over time.

The last section explains the computational procedure in detail.

5 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 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. 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 . $l(\omega, y', z'; z^k) : L \times Y \times Z \times Z^k \rightarrow R$, 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 $\underline{\omega}(y', z'; z^k) : Z \times Z^k \rightarrow R$ and the average weight forecasting function $g^*(z', z^k) : Z^k \rightarrow R$ induce the consumption share policy function:

$$\begin{aligned} l(\omega, y', z'; z^k) &= \omega \text{ if } \omega > \underline{\omega}(y', z'; z^k) \\ &= \underline{\omega}(y', z'; z^k) \text{ elsewhere.} \end{aligned} \tag{33}$$

The reservation weights are determined such that the solvency constraints bind exactly. The cost functions $C(\omega, y', z'; z^k)$ and $C^y(y', z'; z^k)$ 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:

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

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

$$g^*(z', z^k) = E_{z^\infty \subset z^k} g(z', z^\infty), \quad (34)$$

where the actual liquidity shock is given by:

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

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 $k \rightarrow \infty$, the errors tend to zero.

Definition 5.1. *An approximate stationary equilibrium is fully characterized by a list of functions $l(\omega, y', z'; z^k)$, $C(\omega, y', z'; z^k)$, $C^y(y', z'; z^k)$ and $g(z', z^k)$ such that (i) $g(z', z^k)$ equals the average liquidity shock in z^k and (ii) $l(\omega, y', z'; z^k)$ 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') .⁹

- 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 $\hat{g}_1^a(z^k, z')$ implied by the policy function. This provides a new guess $\hat{g}_2(z^k, z')$ for the weight growth functions.
- Finally, I iterate on the liquidity shock forecasts until $\{\hat{g}_n(z^k, z')\}$ convergence to $\hat{g}_*(z^k, z')$. 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:

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

The allocation error decreases as k is increased. To approximate the consumption cost function $C(\omega, y', z'; z^k)$, I use a Tchebychev polynomial approximation in the consumption weight ω (Judd (1998)). The polynomial is of order 7 and I use 30 nodes. 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.

6 Conclusion

There is a growing literature that tries to explain the empirical evidence on liquidity risk (see e.g. Acharya and Pedersen (2005)). My paper shows there is a tight connection between the aggregate volume of trade in securities markets and risk premia in a model with solvency constraints as the only trading friction. The liquidity risk produces a low risk-free rate, a large equity premium, an upward sloping yield curve and substantial time variation in risk premia in a model with i.i.d. aggregate consumption growth innovations and standard power

⁹This algorithm can be shown to converge as $k \rightarrow \infty$. The proof is available upon request.

utility preferences. In related work, Lustig and VanNieuwerburgh (2006) introduce housing into a version of my model and they show the housing collateral dynamics help to match lower frequency variation in risk premia, while Lustig and VanNieuwerburgh (2005) test the empirical predictions of this housing collateral model.

References

- ACHARYA, V. V., AND L. H. PEDERSEN (2005): “Asset Pricing with Liquidity Risk,” *Journal of Financial Economics*, 77, 375–410.
- 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.
- BANSAL, R., AND A. YARON (2004): “Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles,” *The Journal of Finance*, 59, 1481–1509.
- 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., 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.
- DETEMPLE, J., AND A. SERRAT (2003): “Dynamic Equilibrium and Liquidity Constraints,” *Review of Financial Studies*, 16, 597–629.
- GEANAKOPOLOS, J., AND W. R. ZAME (1998): “Collateral, Default and Market Crashes,” Working Paper Yale and UCLA.
- HALL, R. E. (2001): “The Stock Market and Capital Accumulation,” *American Economic Review*, 91, 1185–1202.
- HANSEN, L. P., AND K. J. SINGLETON (1982): “Generalized Instrumental Variable Estimation of Nonlinear Rational Expectation Models,” *Econometrica*, 50, 1269–1286.
- 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.

- 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.
- KRUEGER, D. (1999): “Risk Sharing in Economies with Incomplete Markets,” Ph.D. thesis, University of Minnesota.
- KRUEGER, D., AND H. UHLIG (2005): “Competitive Risk Sharing with Costly Recontracting,” *Journal of Monetary Economics*, Forthcoming.
- KUBLER, F., AND K. SCHMEDDERS (2003): “Stationary Equilibria in Asset Pricing Models with Incomplete Markets and Collateral,” *Econometrica*, 71, 1767–1795.
- 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, 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.
- LUSTIG, H., AND S. VANNIEUWERBURGH (2005): “Housing Collateral, Consumption Insurance and Risk Premia: An Empirical Perspective,” *Journal of Finance*, 60(3), 1167–1219.
- (2006): “Can Housing Collateral Explain Long-Run Swings in Asset Returns?,” *Working Paper UCLA and NYU Stern*.
- LUSTIG, H., AND A. VERDELHAN (2007): “The Cross-section of Foreign Currency Risk Premia and US Consumption Growth Risk,” *American Economic Review*, forthcoming.
- LUTTMER, E. (1992): “Asset Pricing in Economies with Frictions,” Ph.D. thesis, University of Chicago.
- MAR CET, 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.
- PASTOR, L., AND R. F. STAMBAUGH (2003): “Liquidity Risk and Expected Stock Returns,” *Journal of Political Economy*, 3, :642–685.
- 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.
- STORESLETTEN, K., C. I. TELMER, AND A. YARON (2006): “Asset Pricing with Idiosyncratic Risk and Overlapping Generations,” *Working Paper Wharton*.
- ZHANG, H. H. (1997): “Endogenous Borrowing Constraints with Incomplete Markets,” *Journal of Finance*, 52, 2187–2209.

A Proofs

- Proof of Proposition 3.1:

Proof. First, I show that the solvency constraints imply that the participation constraints are satisfied:

$$U(\{c(\theta_0, y^t, z^t)\})(s^t) \geq \kappa_t(s^t),$$

and $U(\{c(\theta_0, y^t, z^t)\})(s^t) = \kappa_t(s^t) \iff \Pi_{s^t}[\{\eta\}] = \Pi_{s^t}[\{c(\theta_0, y^t, z^t)\}]$

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

$$U(\{c(\theta_0, y^t, z^t)\})(s^t) = \max_{\{c'\}} U(c)(s^t), \quad (35)$$

such that the budget constraint is satisfied $\Pi_{s^t}[\{c'\}] \leq \Pi_{s^t}[\{c(\theta_0, y^t, z^t)\}]$ 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 $\kappa_t(s^t)$:

$$\kappa_t(s^t) = \max_{\{c'\}} U(c)(s^t), \quad (36)$$

such that the budget constraint is satisfied $\Pi_{s^t}[\{c'\}] \leq \Pi_{s^t}[\{\eta\}]$ 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 $U(\{c(\theta_0, y^t, z^t)\})(s^t) \geq \kappa_t(s^t)$, then from (35) and (36), it follows that $\Pi_{s^t}[\{\eta\}] \leq \Pi_{s^t}[\{c(\theta_0, y^t, z^t)\}]$. The second part is obvious. \square

- Proof of Proposition 3.2:

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

$$\int \sum_{y^t} \left[\begin{array}{c} \Pi_{s^t} [\{c(\mu_0, y^t, z^t)\}] \\ -\Pi_{s^t} [\{\eta\}] \end{array} \right] \varphi(y^t|y_0) d\Phi_0 \geq 0. \quad (37)$$

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:

$$\sum_{z^\tau \succeq z^t} Q(z^\tau|z^t) \left[\int \sum_{y^\tau} \left[\begin{array}{c} c(\mu_0, y^\tau, z^\tau) \\ -\hat{\eta}_\tau(y_\tau, z_\tau) e_\tau(z^\tau) \end{array} \right] \varphi(y^\tau|y_0) d\Phi_0 \right], \quad (38)$$

with $(z^\tau, y^\tau) \succeq s^t$. To justify the interchange of limits and expectations, I appeal to the monotone convergence theorem. Let $\Pi_{s^t}^n [\{c(\mu_0, y^t, z^t)\}]$ be the value of the claim to the consumption stream until $t+n$ and let $\Pi_{s^t}^n [\{\eta\}]$ 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 $n \rightarrow \infty$ (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:

$$\int \sum_{y^t} \hat{\eta}_t(y_t, z_t) \varphi(y^t|y_0) d\Phi_0 = \sum_{y^t} \pi_{z_t}(y_t) \hat{\eta}_t(y_t, z_t) = (1 - \alpha), \quad (39)$$

and the market clearing condition implies that:

$$\int \sum_{y^t} c(\mu_0, y^t, z^t) \varphi(y^t|y_0) d\Phi_0 = e_t(z^t). \quad (40)$$

Plugging eqs. (39) and (40) back into eq. (38) implies the following inequality must hold at all nodes z^t : $\alpha \Pi_{z^t} [\{e_t(z^t)\}] \geq 0$. If there is no outside wealth ($\alpha = 0$) in the economy, then the expression is zero at all nodes z^t and eq. (37) 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 $(\mu_0, y^t, z^t) \in A$ at node z^t where A has non-zero measure:

$$\sum \int_A \varphi(y^t|y_0) d\Phi_0 > 0,$$

and their constraint is slack: $\Pi_{s^t} [\{c(\mu_0, y^t, z^t)\}] > \Pi_{s^t} [\{\eta\}]$. Given that eq. (37) holds

with equality at all nodes z^t with $\alpha = 0$, there are some households (μ'_0, y^t, z^t) at node $z^t \in B$ for which

$$\sum \int_B \varphi(y^t|y_0) d\Phi_0 > 0,$$

which have constraints that are violated: $\Pi_{s^t} [\{c(\mu'_0, y^t, z^t)\}] < \Pi_{s^t} [\{\eta\}]$. If not, (37) would be violated. But this violates the participation constraints for these agents. So, for $\alpha = 0$, for all households with positive measure:

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

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:

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

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

- Proof of Proposition 3.3:

Proof. If this condition is satisfied: $\Pi_{s^t}^* [\{e\}] \geq \Pi_{s^t}^* [\{\eta\}]$ 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. □

- Proof of Proposition 3.4:

Proof. 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:

$$\Pi_{s^t} [\{c\}] \geq \Pi_{s^t} [\{\eta\}] \geq B_{s^t}^{aut} [\{\eta\}],$$

where $U_t(B_{s^t}^{aut} [\{\eta\}], s^t, c) = U(\{\eta\})(s^t)$. □

- Proof of Lemma 3.1:

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

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

and $C(\mu_0, s^t; l) \geq \Pi_{s^t}[\{\eta\}]$ 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: $C(\mu_0, s^t; l) = C_t(\zeta, y, z^t)$.

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 $l_{t-1} > l_t(y, z^t)$ and $l_t(l_{t-1}, y, z^t) = l_t(y, z^t)$ 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. \square

- Proof of Theorem 3.1:

Proof. $\{\zeta_t(\mu_0, s^t)\}_{t=0}^\infty$ 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 $E[\zeta_t^{1/\gamma}(\mu_0, y^t, z^t)] = h_t(z^t)$ implies that $E[c_t(\mu_0, y^t, z^t)] = e_t(z^t)$. 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 $\{p_t(s^t|s_0)\}$. \square

- Proof of Lemma 3.2:

Proof. 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:

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

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

$$\widehat{\phi}(z_{t+1}) = \frac{\phi(z_{t+1})\lambda(z_{t+1})^{1-\gamma}}{\sum_{z_{t+1}} \phi(z_{t+1})\lambda(z_{t+1})^{1-\gamma}} \text{ and } \widehat{\beta} = \beta \sum_{z_{t+1}} \phi(z_{t+1})\lambda(z_{t+1})^{1-\gamma}. \quad (41)$$

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

$$\widehat{\Pi}_{s^t} [\{\widehat{d}\}] = \widehat{d}_t(s^t) + \widehat{\beta}(z_t) \sum_{s^{t+1}} \widehat{\Pi}_{s^{t+1}} [\{\widehat{d}\}] \left(\frac{h_{t+1}(z^{t+1})}{h_t(z^t)} \right)^\gamma \widehat{\pi}(s^{t+1}|s_t), \quad (42)$$

and let $V_{s^t} [\{\widehat{d}\}]$ 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.

Second, the lemma itself follows directly from the definition of the cutoff level:

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

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

- Proof of Lemma 3.3:

Proof. Since $\varphi(y'|y)$ satisfied monotonicity, I can rank the cutoff weights, because the value of the endowment claims can be ranked such that:

$$\widehat{\Pi}_{z^t, y_n} [\{\widehat{\eta}\}] \geq \widehat{\Pi}_{z^t, y_{n-1}} [\{\widehat{\eta}\}] \geq \dots \geq \widehat{\Pi}_{z^t, y_1} [\{\widehat{\eta}\}], \quad (43)$$

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} :

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

and verify that these objects can be ranked:

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

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

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

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

$$\tilde{\Pi}_{z^{T-2}, y_n} [\{\hat{\eta}\}] \geq \tilde{\Pi}_{z^{T-2}, y_{n-1}} [\{\hat{\eta}\}] \geq \dots \geq \tilde{\Pi}_{z^{T-2}, y_1} [\{\hat{\eta}\}].$$

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

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

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

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

and that $\lim_{T \rightarrow \infty} \tilde{E}_t \beta^{T-t} \frac{h_T}{h_t} \hat{\Pi}_{z^T, y_T}$ 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 (43) holds. Finally, I need to show that this implies a similar ranking for the cutoff weights. When $\zeta_t(\mu_0, s^t) = L_t(z^t, y)$, by definition, the following holds:

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

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

$$L_t(z^t, y_n) \geq L_t(z^t, y_{n-1}) \geq \dots \geq L_t(z^t, y_1).$$

This result, combined with Lemma 3.2, implies directly that the consumption share in

the lowest state equals the endowment share: $\frac{l_t(z^t, y_1)}{h_t(z^t)} = \widehat{\eta}(y_1, z^t)$ for all z^t . (q.e.d.) \square

- Proof of Proposition 3.5:

Proof. 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:

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

\square

- Proof of Lemma 3.4:

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

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

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 (3.2). \square

- Proof of Proposition 3.6:

Proof. In this case, in the transformed economy, the z shocks have disappeared altogether, since $\hat{\eta}$ does not depend on z . I will use ω to denote the consumption share of an agent at the end of the previous period. Let $\hat{C}(\omega, y)$ denote the cost of the consumption stream for a household in state y . Similarly, I use $\hat{C}^y(y)$ 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 $\hat{C}_\eta(y)$ does not depend on z . The proof proceeds in two steps. First, I assume that there exists a stationary equilibrium characterized by the following condition:

$$\frac{h_{t+1}(z^{t+1})}{h_t(z^t)} = g^* \text{ for all } z^{t+1}$$

I compute g^* . Second, I show that for given g^* , there exists a stationary distribution of consumption weights ω .

First, 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 :

$$\hat{C}_\eta(y) = \hat{\Pi}_y[\{\eta(y)\}] = \hat{\eta}(y) + \hat{\beta} \sum_{y'} \hat{\Pi}_{y'}[\{d\}] (g^*)^\gamma \varphi(y'|y)$$

and neither does the value of the consumption claim $C(\omega, y)$:

$$\hat{C}(\omega, y) = l(\omega, y')/g^* + \hat{\beta} \sum_{y'} \hat{C}(\omega', y') (g^*)^\gamma \varphi(y'|y),$$

where the next period's weight is discounted: $\omega' = l(\omega, y')/g$.

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 φ , 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 3.2 and $l(\omega, y)/g \geq \widehat{\eta}(y_1)$. This implies that $\omega \in [L, \bar{l}]$ 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. Second, we establish the existence of a stationary equilibrium. 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

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

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:

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

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 $(\omega, y) \in (L, Y)$: $Q^N((\omega, y), (l_{\max}, y_n)) \geq \varepsilon$, but we know that $Q((\omega, y), (l_{\max}, y_n)) \geq \varphi(y_n|y)$. 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 $Q^N((\omega, y), (l_{\max}, y_n)) \geq \varepsilon$ where

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

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

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

□

B Data Appendix

The computation of firm value returns is based on Hall (2001). The data to construct our measure of returns on firm value were obtained from the Federal Flow of Funds ¹⁰. We use the (seasonally not adjusted) flow tables for the non-farm, non-financial corporate sector, in UTABS 102D. I calculate the value of all securities as the sum of financial liabilities (144190005) plus the market value of equity (1031640030) less financial assets (144090005), adjusted for the difference between market and book for bonds. I correct for changes in the

¹⁰at <http://www.federalreserve.gov/RELEASES/z1/current/data.htm>

market value of outstanding bonds by applying the Dow Jones Corporate Bond Index to the level of outstanding corporate bonds at the end of the previous year.

The flow of pay-outs to securities holders is measured as dividends (10612005) plus the interest paid on debt (from the NIPA Table on Gross Product of non-financial, corporate business) less the increase in net financial liabilities (10419005), which includes issues of equity (103164003). I obtain compensation of all employees (line 2, Table 1.12), proprietary's income (line 9, Table 1.12) and rental income (NIPA line 12 Table 1.12) from NIPA. Finally, I obtain the value of residential housing wealth from the Flow of Funds Tables (FoF-FL155035015.Q).

C Technical Appendix

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": $Q : (L \times Y \times Z^k) \times (\mathcal{B}(L) \times P(Y) \times P(Z^k)) \rightarrow [0, 1]$ where

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

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

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

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

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

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

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 $(z^{k'}, z^k) \in Z$ where $z^{k'} = (z', z_{k-1}^k)$:

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

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:

$$g^*(z', z^k) = \sum_{y'} \int l(\omega, y', z'; z^k) \Phi_{z^k}^* (d\omega \times dy) \varphi(y'|y), \quad (44)$$

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

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

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

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