The Market Price of Aggregate Risk and the Wealth Distribution

Hanno Lustig*
Stanford University

November, 2001

Abstract

Bankruptcy brings the asset pricing implications of Lucas's (1978) endowment economy in line with the data. I introduce bankruptcy into a complete markets model with a continuum of ex ante identical agents who have CRRA utility. Shares in a Lucas tree serve as collateral. The model yields a large equity premium, a low risk-free rate and a time-varying market price of risk for reasonable risk aversion γ. 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 the growth rate of a particular moment of the Pareto-Negishi weight 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 and a large fraction of agents faces severely binding solvency constraints. These shocks occur in recessions. The average investor wants a high excess return on stocks to compensate for the extra liquidity risk, because of low stock returns in recessions. In that sense stocks are "bad collateral". The adjustment to the Breeden-Lucas stochastic discount factor raises the unconditional risk premium and induces time variation in conditional risk premia. This explains why stock returns are highly predictable over longer holding periods.

Keywords: Asset Pricing, Wealth Heterogeneity, Limited Commitment.

*Graduate student at Dept. of Economics, email: hlustig@stanford.edu. I would like to thank my advisors, Robert Hall, Thomas Sargent and Dirk Krueger, for their help and encouragement. I would also like to thank V.V. Chari, John Cochrane, Hal Cole, Felix Kubler, Pascal Maenhout, Lasse Pedersen, Martin Schneider, Neng Wang, Pierre-Olivier Weill, Barry Weingast, the members of the Sargent reading group at Stanford, the joint Chicago-Stanford macro workshop and the UCLA macro workshop for comments. Monika Piazzesi and Stijn Van Nieuwerburgh provided detailed comments. I am grateful for financial support from the William and Sarah Kimball Stanford Graduate Fellowship and the Flemish NSF. The most recent version and the Computational Appendix are downloadable at www.stanford.edu/~hlustig.

1 Introduction

I develop a model of an endowment economy with a continuum of agents, complete markets, but imperfect enforcement of contracts. Because households can declare themselves bankrupt and escape their debts, they face endogenous solvency constraints that restrain their resort to the bankruptcy option. In a calibrated general equilibrium version of the model, the risk associated with these liquidity constraints delivers an equity premium of 6 percent, a risk-free rate of .8 percent and a time-varying market price of risk.

An economy that is physically identical but with perfect enforcement of contracts forms a natural benchmark with which to compare my model. If markets are complete, if contracts are perfectly enforceable, and if agents have CRRA utility with risk aversion γ, then individual consumption is perfectly correlated across households. Assets can be priced by assuming that there is a single investor who consumes the aggregate endowment (Rubinstein, 1974). The SDF (stochastic discount factor) that prices payoffs is this representative agent's intertemporal marginal rate of substitution (Lucas, 1978 and Breeden, 1979).

Because assets only price aggregate consumption growth risk in this benchmark representative agent model, three quantitative asset pricing puzzles arise. These puzzles follow from the fact that aggregate consumption growth in the US is approximately i.i.d. and smooth. First, the covariance of stock returns and consumption growth is small, implying that implausibly high risk aversion (γ) is needed to match the observed equity premium of 6.2 percent in annual post-war US data. This is Hansen and Singleton's (1982) and Mehra and Prescott's (1985) equity premium puzzle. Second, because consumption growth is approximately i.i.d., the risk premia implied by the model are roughly constant, while in the data aggregate stock returns are predictable and the volatility of returns varies systematically over time, a symptom of time-varying risk premia.[1] Third, the observed mean reversion of returns makes stocks less risky over longer holding periods than over shorter ones. At the same time, long streaks of low aggregate consumption growth are unlikely, because consumption growth is i.i.d. The equity premium puzzle worsens over longer holding periods.

This paper addresses each of these puzzles within an economy that is phys-

[1] See Campbell's 2000 survey for an overview.

ically identical to the benchmark economy of Lucas (1978), but follows Alvarez and Jermann (2000a, 2001a) in relaxing the assumption that contracts are perfectly enforceable. Part of the endowment of my economy is yielded by a tradable Lucas tree; the rest of the endowment is labor income. Instead of sending agents into autarky upon default, as Alvarez and Jermann do, I allow agents to file for bankruptcy (Lustig, 1999). When agents declare bankruptcy, they lose their holdings of the Lucas tree, but all of their current and prospective labor income is protected from creditors. Shares in the Lucas tree serve as collateral. It is important that there is a continuum of agents in my economy, in contrast to the two agents in Alvarez and Jermann's economy, because the wealth distribution dynamics drive the variation in risk premia.

The possibility of bankruptcy constrains the price of an individual's consumption claim to exceed the shadow price of a claim to his labor income in all states of the world. The fraction of the economy's endowment yielded by the Lucas tree plays a key role in my economy. If the labor share of aggregate income is one, all wealth is human wealth, the solvency constraints always bind and there can be no risk sharing. As the fraction of wealth contributed by the Lucas tree increases, risk sharing is facilitated.

Beyond risk in the aggregate endowment process, the bankruptcy technology introduces a second source of risk, the risk associated with binding solvency constraints. I call this liquidity risk.[2] In the benchmark model 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: the growth rate of the γ^{-1}-th moment of the distribution of stochastic Pareto-Negishi weights. If this growth rate is high, a large fraction of agents is constrained and the economy is said to be hit by a negative liquidity shock. Beyond this "average weight" growth rate, all other features of the wealth distribution are irrelevant for asset prices.

If negative liquidity shocks occur when aggregate consumption growth is low (recessions), then the liquidity shocks raise the unconditional volatility of the SDF. Liquidity shocks in recessions emerge from two sources. If the dispersion of idiosyncratic labor income shocks increases in recessions, households would

[2] I would like to thank Lasse Pedersen for suggesting this term to me.

like to borrow against their income in the "high idiosyncratic state" to smooth consumption but they are not allowed to, because they would walk away in the good state. If the capital share of income shrinks in recessions, there is "less collateral" available and more agents are constrained as a result. Both the labor risk and the collateral risk channel have support in the data.

The wealth distribution dynamics of the model generate time-variation in the conditional market price of risk. 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, this causes those low-wealth agents with high income draws to hit severely binding constraints and the left tail of the wealth distribution is erased. After the recession, the conditional market price of risk gradually decreases. If another recession follows shortly thereafter, the mass of households close to the lower bound on wealth is much smaller and so are the liquidity shocks. This lowers the conditional market price of risk.

The continuum of agents in my model contributes important differences vis-à-vis the two-agent model of Alvarez and Jermann (2001a). In their model, liquidity shocks are larger but much less frequent. To generate a liquidity shock, one of the agents has to switch from a "low" to a "high" state. The persistence of labor income makes these switches rare. Between these switches this economy looks like the benchmark representative agent economy: households consume a constant share of the aggregate endowment and the conditional risk premia are constant. In addition, the correlation between the liquidity shocks and stock returns is small and the high volatility of the SDF does not translate into high excess returns. In the benchmark calibration Alvarez and Jermann (2001a) report an equity premium of 3.2 percent, while the market price of risk is close to 1. The economy with a continuum is an average of two-agent economies in each period; the liquidity shocks are more frequent and more tightly correlated with the business cycle. My model delivers an equity premium of 6.1 percent in the benchmark calibration at a lower market price of risk of .41.

There is a large literature on heterogeneity and asset pricing. Constantinides and Duffie (1996) use an insight from Mankiw (1986) to show how a systematic increase in idiosyncratic risk during recessions can deliver a high equity premium. Cogley (1998) finds some evidence of correlation between the dispersion of consumption growth and returns that is of the right sign but of small magnitude. Lettau (2001) remarks that even if agents do not share any risks and just eat their labor income, high risk aversion is still needed to put the SDF's inside the Hansen-Jagannathan bounds.

In most models on asset pricing and heterogeneity, assets are essentially being priced off individual consumption processes. In the Constantinides and Duffie model, any agent's intertemporal marginal rate of substitution (IMRS) is a valid SDF for all payoffs in the next state. Similarly, in models with *exogenous borrowing constraints,* (e.g. He and Modest, 1995) the individual IMRS is a valid SDF for excess returns in all states. So is the cross-sectional average of these individual intertemporal marginal rates of substitution. In the continuous time limit the difference between the average marginal utility and the marginal utility of average consumption is absorbed into the drift (Grossman and Shiller, 1982) and the assets can be priced using the Breeden-Lucas SDF. Campbell (2000) concludes this "limits the effects of consumption heterogeneity on asset pricing".

Not so with *endogenous solvency constraints*: the individual IMRS is a valid SDF for payoffs only in those states in which he is unconstrained (Alvarez and Jermann, 2000a). Assets can no longer be priced off individual consumption processes.[3] The Lucas-Breeden discounter does not reappear in the continuous-time limit. Nevertheless, there is a representative agent in this model with a stochastic pseudo-habit. The habit features time-varying sensitivity to aggregate consumption growth shocks, in the spirit of Campbell and Cochrane (1999).

To deal with a continuum of consumers and aggregate uncertainty, I extend the methods developed by Atkeson and Lucas (1992,1995) and Krueger (1999). Atkeson and Lucas show how to compute constrained efficient allocations in dynamic economies with private information problems. Krueger computes the equilibrium allocations in a limited commitment economy without aggregate uncertainty, in which households are permanently excluded upon default.

The use of stochastic Pareto-Negishi weights (Marcet and Marimon, 1999) allows me to state an exact aggregation result in the spirit of Luttmer (1991): equilibrium prices depend only on the γ^{-1}-th moment of the distribution of

[3] Lettau's (2001) criticism does not apply here. The autarchy SDF enters the H-J bounds for low values of γ.

weights. This reduces the problem of forecasting the multiplier distribution - the state of the economy- to one of forecasting a single moment. In Krusell and Smith (1998) agents have to forecast the next period's capital stock; the agents in my model have to forecast this growth rate in every state tomorrow. The exact forecast requires the entire aggregate history or the distribution of weights. I approximate the actual equilibrium by a stationary, truncated-history equilibrium. The state space is reduced to include only the k most recent aggregate events. The allocation errors are exactly zero on average for each truncated history. In the simulation results the errors are fairly small overall.

This paper is organized as follows. The second section of the paper describes the environment. The third section discusses the benchmark representative agent model and its empirical failure. The fourth section introduces the bankruptcy technology. The fifth section derives the policy functions for the stochastic Pareto-Negishi weights. The forces driving asset prices are discussed in the sixth section; the final section reports the simulation results and then I conclude by summarizing my findings. All Proofs are in the Appendix and so are the Figures and Tables.

2 Environment

2.1 Uncertainty

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

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

I will 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.[4] In addition, I will assume there is a unique invariant distribution

[4] Hammond and Sung (2001) show how to solve the measurability problem for a continuum of independent random variables using a Monte Carlo based approach.

$\pi_z(y)$ 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. $(Z^\infty, \mathcal{F}, P)$ is a probability space where Z^∞ is the set of possible aggregate histories and P is the corresponding probability measure induced by π.

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:

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

where γ is the coefficient of relative risk aversion.

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

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

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

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

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

2.3 Market Arrangements

Claims to the labor income process $\{\eta_t\}$ cannot be traded while shares in the Lucas tree can be traded. I use ω to denote an agent's holdings of shares in the

Lucas tree. In each period households go to securities markets to trade ω 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 (see Krueger, 1999).[5]

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:

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

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

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

2.4 Enforcement Technology

In this literature, it has been common to assume that households are excluded from financial markets forever when they default, following Kehoe and Levine (1993) and Kocherlakota (1996). I will 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, 1999 for a complete discussion).

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

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

[5] In an economy with a finite number of agents, the payoffs of the contingent claims needed to complete the markets would depend on the cross-product of private event histories (see Kehoe, Levine and Prescott, 2001, for more on this).

These borrowing constraints follow endogenously from the enforcement technology. 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:

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

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:

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

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'.[6]

Table 1: **Enforcement Technologies**

	Permanent Exclusion	Bankruptcy
renegotiation proof	no	yes
information required	assets, endowments and preferences	assets
equilibrium default	no	yes

Three key differences between bankruptcy and permanent exclusion deserve mention. First, in my model all borrowing is fully collateralized and the lender recovers the full amount in case of default. In case of permanent exclusion, the lender actually loses money and the borrower is denied access to financial markets forever.[7] Default causes some loss of joint surplus and both parties obviously have a strong incentive to renegotiate (see Table 1).

Second, 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

[6] The Appendix contains the definition of a sequential equilibrium on p. 48.

[7] Kocherlakota (1996) points out that there are payoff-equivalent strategies that are renegotiation proof, but these are not consistent with a decentralized market equilibrium in the sense of Alvarez and Jermann (2000a).

the borrower's endowment process and her preferences (Alvarez and Jermann, 2000a). This type of information is not readily available and costly to acquire. Moreover, the borrower has an incentive to hide his private information. Finally, 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 constraints do not mean I rule out equilibrium default on some traded securities. Assume households issue J securities, traded in each period, promising to pay $b_J(s^t)$ in s^t. Households will simply declare bankruptcy whenever their constraint binds. Actual payouts $a_{t+1}(s^t, s')$ are given by the min over the promised payoffs and the *cum dividend* value of securities:

$$-a_{t+1}(s^t, s') = \min \left[\omega_t(s^t) \left[p^e_{t+1}(z^{t+1}) + \alpha e(z_{t+1}) \right], \sum_{j=1}^{J} b_j(s^{t+1}) \right]. \quad (5)$$

Here I only record what the household will actually deliver in each state -the right hand side of eq. (5)- instead of what it promises to deliver. The borrowing constraints are a simple way of relabeling promises as "sure things" in all states. I can still price defaultable securities -i.e. a collection of promises in different states.

What distinguishes this setup from Geneakoplos and Zame (1998) is the fact that only outright default on all financial obligations is allowed, not default on individual obligations. Kubler and Schmedders (2001) introduce collateral constraints in an incomplete markets setting.

The next section abstracts from limited commitment and discusses the empirical failure of the Lucas-Breeden stochastic discounter that emerges in equilibrium (Hansen and Singleton, 1982 and Mehra and Prescott, 1985). It turns out that prices and allocations are easier to derive in an equivalent version of this economy with all trading at time zero.

3 Perfect Enforcement

The absence of perfect enforcement plays a key role in governing the behavior of the distribution of wealth and asset prices. To set the stage for my model, I consider a version of this economy with perfect enforcement. This is a natural benchmark. All trading occurs at time 0 after having observed s^0. A complete set of contingent claims is traded at price $p_t(s^t|s_0)$. This is the price of a security that delivers one unit of the consumption good at t conditional on an event history $s^t = (y^t, z^t)$. Households choose an allocation $\{c_t(\theta_0, s^t)\}$ to maximize expected utility subject to the time 0 budget constraint. A household of type (θ_0, y_0) faces one budget constraint at time 0:

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

3.1 Risk Sharing Rule

In this frictionless environment households equalize their IMRS. Since households have power utility, all of the household consumption processes are growing at the same rate in equilibrium. This follows directly from dividing the first order condition for one household (θ_0', y_0) by that for an *arbitrary* other household (θ_0'', y_0):

$$\left[\frac{c_t(\theta_0', s^t)}{c_t(\theta_0'', s^t)}\right]^{-\gamma} = \frac{\lambda_0'}{\lambda_0''}, \tag{6}$$

where λ_0' and λ_0'' denote the budget constraint multipliers of two different households. To derive the risk sharing rule, I can assign Negishi weights $\mu_0 = \lambda_0^{-1}$ to each household (θ_0, y_0) and allocate consumption on the basis of these fixed weights in all future states of the world. Let Φ_0 denote the initial joint distribution over weights and endowments (μ_0, y_0) induced by the wealth distribution Θ_0. Conjecture a linear risk sharing rule that assigns a constant fraction of the aggregate endowment in all future nodes s^t :

$$c_t\left(\mu_0, s^t\right) = \frac{\mu_0^{1/\gamma}}{E\mu_0^{1/\gamma}} e_t(z^t) \text{ where } s^t = (y^t, z^t), \tag{7}$$

where the constant $E\mu_0^{1/\gamma} = \int \mu_0^{1/\gamma} d\Phi_0$ guarantees market clearing after each aggregate history. This risk sharing rule satisfies the condition for the ratio of marginal utilities in (6). Since the Negishi weights are constant, the individual intertemporal marginal rates of substitution equal this "aggregate" IMRS:

$$m_{t,t+1} = \beta\left(\frac{e_{t+1}}{e_t}\right)^{-\gamma}. \tag{8}$$

Preferences aggregate. This result is due to Rubinstein (1974); this economy also fits into Wilson 's (1968) theory of syndicates.[8] Following Rubinstein (1974)

[8] See Browning, Hansen and Heckman, 2000 for a discussion.

and Hansen and Singleton (1982), I can construct a pricing operator $P_t(x_{t+1}) = E_t(m_{t,t+1}x_{t+1})$ which assigns a price at time t to random payoffs x_{t+1} that are measurable w.r.t $\mathcal{F}_{t+1}$, using the Breeden-Lucas SDF in (8).[9]

The benchmark model has strong predictions about what drives asset prices and what does not. The wealth distribution is fixed throughout and it plays no role whatsoever in pricing payoffs. In addition, the nature of individual labor income risk has no bearing on asset prices; investors are compensated only for aggregate consumption risk because all of the idiosyncratic risk. This implication of the model is hard to reconcile with the empirical evidence.

3.2 Aggregate Consumption Risk

The benchmark model is at odds with the data for three main reasons: (1) the model-implied market price of risk is too low relative to observed Sharpe ratios over short holding periods, (2) it is close to constant while the data suggests it should be time-varying and (3) it is even lower relative to Sharpe ratios over longer holding periods.

Hansen and Jagannathan (1991) derive a theoretical upper bound on the Sharpe ratio, i.e. the ratio of the expected excess return R^e over its standard deviation:

$$\left|\frac{E[R^e]}{\sigma[R^e]}\right| \leq \frac{\sigma(m_{t,t+1})}{E(m_{t,t+1})}. \tag{9}$$

The right hand side is called the market price of risk.

Mehra and Prescott (1985) use annual US data from 1889 to 1979 on stock returns, T-bills and consumption growth to study the equity premium. Figure 1 plots the Hansen-Jagannathan bounds implied by these data (Figures and Tables are in the Appendix). The slope of the dotted line is the market price of risk on the left hand side in (9). This bound only uses the information contained in the excess returns. The bounds that mark the cup-sized region use the restrictions imposed by stock returns and T-bill returns: for a given mean Em, this bound gives the minimum standard deviation of the SDF consistent with these returns. Figure 1 also plots the model-implied sample moments of the SDF in (8) using actual consumption growth over the sample (stars in Figure 1), for γ starting at 4 and increasing to 28.

[9] $\{\mathcal{F}_t\}$ is the sequence of increasing $\sigma-$algebras induced by z^t on $(Z^\infty, \mathcal{F}, P)$.

The Sharpe ratio in the Mehra Prescott data is .35. It takes value of γ in excess of 12 to bring the SDF across the market price of risk line - for the stars to cross the dotted line. Aggregate US consumption growth is fairly smooth; its standard deviation is 3.5 percent in the Mehra-Prescott data and it takes a large γ to get a sufficiently volatile SDF. This is the *equity premium puzzle* (Mehra and Prescott, 1985).[10] Still, even these combinations of Em and σm lie outside the cup-sized region because the mean Em is too low. Households want to borrow heavily against their future labor income in a growing economy to flatten their consumption profile. This makes for high risk-free rates or low Em. This is the *risk-free rate puzzle* (Weil, 1989).

If this benchmark model is to be reconciled with the data, agents need to be extremely risk averse. Many authors have interpreted this as evidence in favor of an alternative source of curvature. The work by Abel (1990) and Campbell and Cochrane (1999) on habits and the work by Hansen, Sargent and Tallarini (1999) on robustness fit in this category.

Higher curvature does not solve all of this model's problems. There is considerable evidence of time-varying conditional Sharpe ratios: returns are predictable and the variation in expected returns $E_t[R^e]$ is not offset by changes in the conditional standard deviation $\sigma_t[R^e]$.[11] Since the conditional mean of the SDF equals the risk-free rate, which is very stable, this implies the stochastic discounter has to be heteroskedastic:

$$\frac{E_t[R^e]}{\sigma_t[R^e]} = -\frac{1}{\rho_t(m_{t,t+1}, R^e_{t+1})} \frac{\sigma_t(m_{t,t+1})}{E_t(m_{t,t+1})}. \tag{10}$$

The benchmark model is not equipped to deliver time-varying Sharpe ratios unless there is time-varying consumption risk $\sigma_t(\Delta \log e)$. The data offer little evidence in support of this.[12] Cochrane (2001) calls this the *conditional equity premium puzzle.*

The equity premium puzzle worsens over longer holding periods. Let $P_{t,t+k}(x_{t+k})$ $= E_t\left[m_{t,t+k}x_{t+k}\right]$ denote the period t pricing operator for random payoffs at $t+k$. A similar upper bound applies for returns over k holding periods:

$$\left|\frac{E[R^e_{t\to t+k}]}{\sigma[R^e_{t\to t+k}]}\right| \leq \frac{\sigma(m_{t,t+k})}{E(m_{t,t+k})}. \tag{11}$$

[10] See Grossman and Shiller (1981) for an earlier empirical application of the Lucas-Breeden SDF.

[11] See Cochrane, 2001, p. 463.

[12] See Campbell, Lo and MacKinlay, 1997, p. 311 for a survey of the evidence.

Over longer holding periods, returns are predictable by dividend/price ratios and other variables (Fama and French, 1988). Predictability and mean-reversion slow the growth of the conditional standard deviation of returns relative to the case of i.i.d. returns. The left hand side of (11) grows faster than the square root of the horizon as a result, but aggregate consumption is roughly i.i.d. and therefore its standard deviation grows at a rate equal to the square root of the horizon. Mean-reversion worsens the equity premium at longer holding periods (see Figure 2). This is the *long run equity premium puzzle.*

This dichotomy between asset prices and risk factors other than aggregate consumption growth conflicts with the data. My paper follows Alvarez and Jermann (2000a, 2001a) in introducing limited commitment into the benchmark model. It differs in that (1) my model allows agents to file for bankruptcy and (2) my model considers a large number of agents. The first change yields a more appealing decentralization and improves the model's overall asset pricing implications. I show that the second is essential to generate the right SDF dynamics: the solvency constraints introduce time variation in the wealth distribution that endogenously activates a version of Mankiw's (1986) recipe for high excess returns. It also delivers time-varying and possibly persistent shocks to the SDF.

4 Limited Commitment

The benchmark model assumes that households can commit to repaying their debt, regardless of their individual labor income history or the aggregate history of the economy. This section relaxes this assumption by endowing agents with a bankruptcy technology.

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

$$\Pi_{s^t}[\{d\}] = \sum_{\tau \geq t} \sum_{s^\tau \geq s^t} p(s^\tau | s^t) d_\tau(s^\tau). \tag{12}$$

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

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

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

$$U(\{c(\theta_0, y^t, z^t)\})(s^t) \geq \kappa_t(s^t) \text{ for all } s^t. \tag{13}$$

4.1 Solvency Constraints

The participation constraints in (13) can be reinterpreted as solvency constraints.[13] These put a lower bound on the value of the household's consumption claim.

Proposition 4.1 *For given* Π, *an allocation is said to be immune to bankruptcy iff:*

$$\Pi_{s^t}\left[\{c(\theta_0, y^t, z^t)\}\right] \geq \Pi_{s^t}\left[\{\eta\}\right], \text{ for all } s^t \in S^t, t \geq 0. \tag{14}$$

Default leaves the household with the claim to its labor income stream. If this claim is worth more than the claim to consumption, then it is optimal to default. If not, the household prefers keeping the assets that finance its future consumption to default.

Collateral is essential. Without it, there can be no risk sharing in this economy.

Proposition 4.2 *If there is no outside wealth* ($\alpha = 0$)*, then there can be no risk sharing in equilibrium.*

If there is no collateral, the constraint in (14) holds with equality for all households, regardless of the pricing functional. If some household's consumption claim were worth more than its labor income claim, another household's claim would have to be worth less by definition of the aggregate endowment and the linearity of the pricing functional, but then it would violate its solvency constraint.

[13] In independent work, Detemple and Serrat (1999) exogenously posit similar constraints in a continuous-time model for only one of 2 agents. There is no collateral. They conclude that its asset pricing effects are limited. Lustig (1999) shows that it is key to have all agents face these constraints.

If there is enough collateral, agents may be able to share risks perfectly. Let Π^* denote the pricing functional defined by the perfect insurance, Lucas-Breeden SDF in eq. (8).

Proposition 4.3 *If the value of the aggregate endowment exceeds the value of the private endowment at all nodes, perfect risk sharing is feasible:*

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

If this condition is satisfied, each household can sell a security that replicates its labor income and buy an equivalent claim to the aggregate dividends stream that fully hedges the household. The benchmark model's results apply. Increases in the supply of collateral facilitate risk sharing: if perfect risk is sharing is sustainable in an economy with α', then it is also sustainable in an economy with $\alpha'' > \alpha'$. That follows immediately from Proposition 4.3.

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

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

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

Because this inequality holds for any pricing functional, if perfect risk sharing is feasible in the economy with bankruptcy, it is feasible in the economy with permanent exclusion. Loosely speaking, the Pareto frontier shifts down as one moves from permanent exclusion to bankruptcy (also see Lustig, 1999).

4.2 Equilibrium

This section sets up the household's primal problem and defines an equilibrium. 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:

Primal Problem (PP)

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

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

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

The solvency constraints keep the households from defaulting. The following definition of equilibrium is in the spirit of Kehoe and Levine (1993) and in particular Krueger (1998).

Definition 4.2 *For given initial state z_0 and for given distribution Θ_0, an equilibrium consists of prices $\{p_t(s^t|s_0)\}$ and allocations $\{c_t(\theta_0, s^t)\}$ such that*

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

$$\sum_{y^t} \int c_t(\theta_0, y^t, z^t) d\Theta_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} = e_t(z_t). \tag{15}$$

In equilibrium households solve their optimization problem subject to the participation constraints and the markets clear. Provided that interest rates are high enough, the set of allocations that can be supported as equilibria coincide with the sequential markets equilibria.[14] The next section develops an algorithm to solve for equilibrium allocations and prices using stochastic Pareto-Negishi weights.

5 Solving for Equilibrium Allocations

The objective is to solve for the equilibrium allocations and prices. The computation proceeds in two stages. The first stage derives a new linear risk sharing rule which takes stochastic Pareto-Negishi weights as inputs. Using this rule, I

[14] The proof is a version of Alvarez and Jermann, 2000a; see p. 50 in the Appendix.

derive the equilibrium SDF as a function of the γ^{-1}-th moment of the Pareto-Negishi weight distribution. In the second stage I compute these weights using a simple cutoff rule.

5.1 Stochastic Pareto-Negishi Weights

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

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

Dual Problem (DP)

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

$$\sum_{t\geq 0} \sum_{s^t} \beta^t \pi(s^t|s_0) u(c_t(w_0, s^t)) = w_0, \tag{16}$$

$$\Pi_{s^t} \left[\left\{ c\left(w_0, y^t, z^t\right) \right\} \right] \geq \Pi_{s^t} \left[\{\eta\} \right], \text{ for all } s^t \in S^t, t \geq 0. \tag{17}$$

The convexity of the constraint set implies that the minimizer of DP and the maximizer of PP (the primal problem) coincide for initial wealth $\theta_0 = C^*(w_0, s^0) - \Pi_{s^0}[\{\eta\}]$.[15]

To solve for the equilibrium allocations, I make the dual problem recursive. To do so, I borrow and extend some tools recently developed to solve recursive contracting problems by Marcet and Marimon (1999). Let $m_t(s^t|s_0) = p_t(s^t|s_0)/\pi_t(s^t|s_0)$, i.e. the state price deflator for payoffs conditional on event history s^t. $\tau_t(s^t)$ is the multiplier on the solvency constraint at node s^t. I can transform the original dual program into a recursive saddle point problem for household (w_0, s_0) by introducing a cumulative multiplier:[16]

$$\chi_t(\mu_0, s^t) = \chi_{t-1}(\mu_0, s^{t-1}) - \tau_t(s^t), \quad \chi_0 = 1. \tag{18}$$

[15] See Luenberger, p. 201, 1969.

[16] The **recursive dual** saddle point problem is stated in the Appendix on p. 52.

Let μ_0 denotes the Lagrangian multiplier on the initial promised utility constraint in (16). 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.

5.2 Risk Sharing Rule

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.[17] This allows me to recover allocations and prices from the equilibrium sequence of multipliers $\{\zeta_t(\mu_0, s^t)\}$.

First, consider 2 households having experienced the same history s^t. We know from the first order conditions of the recursive dual saddle point problem[18] for two different households (μ_0', y_0) and (μ_0'', y_0) that the ratio of marginal utilities has to equal the inverse of the weight ratio:

$$\left[\frac{c_t(\mu_0', s^t))}{c_t(\mu_0'', s^t))}\right]^{-\gamma} = \frac{\zeta_t(\mu_0'', s^t)}{\zeta_t(\mu_0', s^t)}. \tag{19}$$

If the constraints never bind, $\zeta_t = \mu_0$ at all nodes and the condition in (19) reduces to condition (6) in the benchmark model.

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

$$\sum_{y^t} \int c_t(\mu_0, y^t, z^t) d\Phi_0 \frac{\pi(y^t, z^t | y_0, z_0)}{\pi(z^t | z_0)} = e_t(z^t), \tag{20}$$

(19) and (20) completely characterize the equilibrium consumption allocation for a given sequence of multipliers. The objective is to find the risk sharing

[17] This approach builds on the work by Negishi (1960). Luttmer (1991) derived a similar aggregation result for economies with exogenous wealth constraints and a finite number of agents, without actually solving the model.

[18] See p. 52 for the dual program.

rule that satisfies these conditions. The natural analogue to the rule in the benchmark model eq. (7) is given by:

$$c_t(\mu_0, s^t) = \frac{\zeta_t^{1/\gamma}(\mu_0, s^t)}{E\left[\zeta_t^{1/\gamma}(\mu_0, s^t)\right]} e_t(z^t). \tag{21}$$

This rule satisfies the condition on the ratio of marginal utilities (19) and it clears the market in each z^t. This can be verified by taking cross-sectional averages of the individual consumption rule.

The average weight in the denominator is not a constant, but moves over time and across states. Let $h_t(z^t)$ denote this cross-sectional multiplier moment:

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

I will refer to this simply as the *average weight* process. This average tracks changes in the distribution of weights over time and provides a sufficient statistic for the distribution. It is a random variable that is a function only of aggregate histories.[19] The *individual weights* ζ_t track an agent's individual history y^t. Every time the agent enters a state with a binding constraint, his weight is increased. If an agent would like to transfer resources from tomorrow to today by borrowing against his future income but cannot do so because of a binding constraint, he is compensated for low consumption today by an increase in his "weight" tomorrow. The *average weight* h_t tracks the economy's history. It is also a non-decreasing stochastic sequence. If its growth rate is high, a large fraction of households is currently constrained. If it does not grow, nobody is currently constrained.

The log of individual consumption growth is pinned down by the growth rate of the individual weight relative to the average weight. As long as the household does not run into binding constraints, the household's consumption share decreases as it runs down its assets because the average weight pushes down interest rates (see eq. 23 below). When the constraint binds and its private multiplier grows faster than the average, its wealth and consumption share rise. Finally, if perfect risk sharing can be sustained, all of the weights are constant and agents consume a constant fraction of the aggregate endowment:

$$c_t = \frac{\mu_0^{1/\gamma}}{E\mu_0} e_t(z^t) \text{ and } h_t(z^t) = E\mu_0,$$

[19] Formally, $\left\{\left(h_t, \mathcal{F}_t^z\right)\right\}$ is a stochastic process on $(Z^\infty, \mathcal{F}, P)$.

for all z^t. This brings us back to the perfect insurance model. The average weight process also delivers the stochastic discount factor.

5.3 Stochastic Discount Factor

The solvency constraints prevent all agents from equalizing their IMRS. After any history z^{t+1}, only the subset of agents who are unconstrained at this node equalize their intertemporal marginal rates of substitution. If this set of agents has measure zero, there can be no risk sharing. Apart from this case, there is some set of agents at some node $s^{t+1\prime} = (z^{t+1}, y^{t+1\prime})$ with intertemporal marginal rates of substitution equal to:

$$\beta\pi(s^{t+1}|s_t)\left(\frac{e_{t+1}(z^{t+1})}{e_t(z^t)}\right)^{-\gamma}\left(\frac{h_t(z^{t+1})}{h_t(z^t)}\right)^{\gamma}, \tag{22}$$

and this has to be true for each z^{t+1}. The right hand side of (22) only depends on y^{t+1} through the transition probabilities.[20] This implies I can define a pricing operator $P_{t,t+1}(x_{t+1}) = E_t[x_{t+1}m_{t,t+1}]$ for random payoffs x_{t+1} that are measurable w.r.t. $\mathcal{F}^z_{t+1}$, i.e. a function of the aggregate history, using the following SDF:

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

This SDF consists of two parts. The first part is the Breeden-Lucas SDF. The second part is a multiplicative liquidity adjustment. I will refer to the growth rate of the average moment, $h_{t+1}/h_t = g_t$, as a liquidity shock. If this growth rate is higher than average, a large fraction of households is severely constrained. This is a negative liquidity shock. Payoffs in states with large liquidity shocks are valued more highly, because they relax the constraints of the average household. Note that interest rates $(1/E_t m_{t,t+1})$ are lower than in the perfect insurance model, lowering the price of today's consumption.

5.4 Equilibrium Weights

Prices and allocations can easily be recovered from the stochastic Pareto-Negishi weights using the risk sharing rule. This section describes the individual weight policy functions and sets up an algorithm for computing the average weight policy functions. The weight policy functions have a cutoff rule property. This makes the computations relatively simple.

[20] The Appendix contains a formal derivation of this result on p. 53.

5.4.1 Characterizing the weight policy function

Agents face a forecasting problem for two objects: the average weights $\left\{\widehat{h}_t(z^t)\right\}$ and the state price deflators $\left\{\widehat{Q}_t(z^t)\right\}$. Armed with these forecasts, they compute the prices of the consumption and labor income claims. Let $C(\mu_0, s^t; l)$ denote the continuation cost of a consumption claim derived from a weight policy $\{\zeta_t(\mu_0, s^t)\}$:

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

where consumption at each node is given by the risk sharing rule in (21):

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

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

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

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

Lemma 5.1 *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 household's policy rule $\{\zeta_t(\mu_0, s^t)\}$ can be written recursively as $\{l_t(l, y, z^t)\}$ where $l_0 = \mu_0$ and $l_t(l_{t-1}, y, z^t) = l_{t-1}$ if $l_{t-1} > \underline{l}_t(y, z^t)$ and $l_t(l_{t-1}, y, z^t) = \underline{l}_t(y, z^t)$ elsewhere. The reason is simple. If the constraint does not bind, the weight is left unchanged. If it does bind, it is set to its cutoff value.

5.4.2 Algorithm

An equilibrium is characterized by two conditions. (i) Aggregating these individual policy rules has to reproduce the average weight forecasts:

$$h_t(z^t) = E_t\left[l_t^{1/\gamma}(\mu_0, s^t)\right] = \widehat{h}_t(z^t), \tag{24}$$

and (ii) the state price deflators have to be “correct”:

$$\widehat{Q}_t(z^t) = \beta^t \left(\frac{e_t(z^t)}{e_0(z^0)}\right)^{-\gamma} \left(\frac{h_t(z^t)}{h_0(z^0)}\right)^{\gamma}. \tag{25}$$

This last condition imposes *no arbitrage*. The first condition imposes rational expectations and *market clearing*. Together with the individual weights that satisfy the constraints, these two conditions completely characterize an equilibrium.[21]

These last two conditions are the basis for an iterative algorithm. There are two operators: for given state price deflators, the first one maps average weight moment forecasts into actual implied weight moments:

$$T\left\{\widehat{h}_t^i(z^t)\right\} = \left\{E_t\left[\zeta_t^{1/\gamma}\right]\right\} = \left\{\widehat{h}_t^{i+1}(z^t)\right\}.$$

This is the inside loop. The second one iterates on the deflators. It takes the fixed point of the inside loop $\{h_t(z^t)\}$ to construct a new guess for state price deflators $\left\{\widehat{Q}_t^j(z^t)\right\}$ .The algorithm starts out by conjecturing the complete insurance state price deflators $\left\{\widehat{Q}_t^0(z^t)\right\}$ for all aggregate histories. Next, I iterate on the operator T mapping forecasted average weights into actual average weights until a fixed point is reached. If I start off with the perfect insurance average moments for all histories, this will converge to a fixed point. This fixed point then provides me with a new guess for the state price deflator $\left\{\widehat{Q}_t^1(z^t)\right\}$. The inside loop of the algorithm starts over. This algorithm is shown to converge to a fixed point $\{h_t^*\}$ that satisfies (24) and (25). This establishes the existence of an equilibrium.[22] It also provides the basis for the computational algorithm, explained in Section 7.1.

6 The Wealth Distribution and Liquidity Risk

In my model liquidity risk is the risk associated with binding solvency constraints. The binding constraints cause agents to move within the wealth distribution and, possibly, induce cyclical changes in the wealth distribution itself induced by aggregate endowment growth shocks. These liquidity shocks create a second asset pricing factor, liquidity risk. The first part of this section derives

[21] Theorem 9.1 in the Appendix on p. 54 establishes this result formally.

[22] There is a separate Computational Appendix which derives and explains this result. It is available at *www.stanford.edu/~hlustig*.

some comparative dynamics and the second part describes exactly how liquidity risk affects asset prices.

Growth in this economy is taken to be of the form: $e_t(z^t) = e_{t-1}(z^{t-1})\lambda(z_t)$. I follow Alvarez and Jermann (2001) in transforming the growth economy into a stationary one with stochastic discount rates $\widehat{\beta}(z)$ and a transformed transition matrix $\widehat{\pi}$. Under perfect risk sharing pricing payoffs is equivalent to computing expected values under the distorted transition law $\widehat{\pi}$; this is the perfect-risk-sharing risk-neutral measure. $\widehat{\Pi}$ denotes the pricing functional associated with $\widehat{\pi}$.[23]

6.1 Risk sharing regimes

Households are tempted to default when they draw high income shares. Suppose there are two states (y_1, y_2) and $\widehat{\eta}(y_1, z) \leq \widehat{\eta}(y_2, z)$ for all z. If the labor income process is persistent, i.e. the diagonal elements of $\widehat{\pi}_z\,(y'|y)$ are larger than the off-diagonal ones, then the cutoff weight is always larger in the "good" state y_2. This means the constraint binds in y_2 if it does in y_1 but not vice-versa.[24]

If the price of a claim to aggregate consumption at perfect insurance prices is higher than the shadow price of a claim to labor income in y_2, for all states z, then perfect risk sharing is feasible.

Lemma 6.1 *: Perfect risk sharing is feasible if*

$$\widehat{\Pi}^*_z\,[\{e\}] \geq \widehat{\Pi}^*_z\,[\{\eta\}] \;\; \textit{for all } (y, z). \tag{26}$$

In case of monotonicity, this condition need only be checked for y_n, the "best state".[25] The effect of collateral is straightforward. Lowering the supply of collateral or increasing the labor share $(1-\alpha)$ pushes up both terms on the right hand side proportionally in all states (y, z) and makes perfect risk sharing harder to sustain.

In the bankruptcy economy, the i.i.d. component of labor income risk does not affect the price/dividend ratio of the labor income claim, i.e. the right hand side of (26). Only the risk that is correlated with the aggregate consumption growth affects this price since all the other risks can be traded away even after

[23] The details are in the Appendix on p. 54.
[24] See Lemma 11.10 in the Computational Appendix.
[25] let f denote a n-dimensional vector with $f_n \geq f_{n-1} \geq \ldots f_1$. $\widehat{\pi}_z$ satisfies monotonicity if $\widehat{\pi}_z f$ is nondecreasing in n.

defaulting. Let ${}^{ex}PD^*_z$ denote the ex-dividend price/dividend ratio. If the labor income process is i.i.d., only the current labor income share and the p/d ratio of a claim to labor income enter into the condition for perfect risk sharing. The other properties of the labor income process have no effect.

Lemma 6.2 *: If the idiosyncratic shocks are i.i.d, perfect risk sharing is possible if*

$$ {}^{ex}PD^*_z\left[\{e\}\right] \geq \frac{(\max \eta(\widehat{y}, z) - 1)}{\alpha} \text{ for all } (z). \tag{27} $$

More generally, even if perfect risk sharing is not sustainable, only those increases in risk that are correlated with aggregate liquidity shocks affect the value of the outside option.

In an economy with permanent exclusion the condition for perfect risk sharing (Alvarez and Jermann, 2001a) is:

$$ U(e)(z) \geq U(\eta)(y, z) \text{ for all } (y, z). \tag{28} $$

The continuation utility from consuming the aggregate endowment has to exceed the value of autarky. These participation constraints can be restated as solvency constraint by using individual-specific autarchic prices:

$$ \Pi^{aut}_z\left[\{e\}\right] \geq \Pi^{aut}_{(y,z)}\left[\{\eta\}\right] \text{ for all } (y, z), \tag{29} $$

where payoffs are priced off the individual's IMRS in autarky:

$$ m^{aut}_{t,t+k} = \widehat{\beta}^t \left(\frac{\widehat{\eta}(y_{t+k}, z_{t+k})}{\widehat{\eta}(y_t, z_t)} \right)^{-\gamma}. $$

At these prices the agent does not want to trade and she eats her labor income. Here white noise risk does matter. An increase in the volatility of i.i.d. income shocks can decreases the value of autarky, generate more risk sharing and be welfare-improving (see Krueger and Perri, 2000). In other words, an increase in white noise risk lowers the "shadow" price/dividend ratio of the labor income claim in autarky in (29) because the agent only trades with others facing the exact same history of private shocks. This creates a volatility paradox in this class of models. More income volatility may not translate into more consumption volatility because of the perverse effect on the value of autarky. This holds *a fortiori* at higher levels of γ. But consumption volatility and risk aversion are key to generating substantial risk premia. Alvarez and Jermann (2001a) deal with this problem in a two-agent economy by lowering β to .65 in annual data ($\gamma = 4$). Myopic agents put more weight on the current income draw. For standard values of β and γ, perfect insurance, or something close to it, obtains for reasonable calibrations of the income process (see section 7.3). In the continuum-of-agents setup with permanent exclusion, quasi-perfect risk sharing obtains even for $\beta = .65$ if $\gamma = 4$.

In the bankruptcy economy, an increase in risk aversion only induces more risk sharing to the extent that it lowers the shadow price of the labor income claim *more* than the price of the consumption claim.[26] This motivates the introduction of a bankruptcy technology.

6.2 Asset Prices

The wealth distribution is time-varying and agents switch positions within the distribution. These changes add a second multiplicative factor to the standard Breeden-Lucas SDF:

$$m_{t,t+1} = m^r_{t,t+1} g^{\gamma}_{t+1}, \tag{30}$$

where $m^r_{t,t+1}$ was defined in eq. (8) and $g_{t+1} = h_{t+1}/h_t$ is the liquidity shock.

The growth rate of the average moment is a sufficient statistic for the changes in the underlying weight distribution; it is not necessary to determine exactly how wealth is allocated across agents in each aggregate state. This aggregation result applies to any model of endogenously incomplete markets, regardless of the exact nature of the solvency constraints. A similar result can be derived for CARA utility.

Alvarez and Jermann's (2001a) crucial insight is that models with *endogenous* solvency constraints have a better chance of generating volatile SDF's than other heterogeneous agent models because assets are not priced off of a single agent's IMRS. Rather, different agents are constrained in different states tomorrow and the composite IMRS that emerges is more volatile.

This is not the case with *exogenous* constraints on the expected value of a household's portfolio. If there is a complete set of contingent claims, the SDF with exogenous wealth constraints of the form $E_t[\theta_{t+1}] \geq \kappa_t$ can be written as:

$$m_{t,t+1} = m^r_{t,t+1} g^{\gamma}_{t}, \tag{31}$$

[26] In fact, if labor income risk is i.i.d., an increase in γ makes perfect risk sharing harder to sustain because it lowers the ex dividend price of the aggregate consumption claim in (27).

where g_t is measurable w.r.t. $\mathcal{F}_t^z$ (Luttmer, 1992). This follows from the fact each agent's Euler equation has to be satisfied for any set of excess returns because these have price zero and do not violate the constraints. Note that the conditional market price of risk is the one of the benchmark model and time-varying market prices of risk are not attainable. Since aggregate consumption growth is close to i.i.d., this also rules out significant correlation between the two components and lowers the potential unconditional volatility of the SDF.[27]

The next sections show that liquidity risk in models with endogenous solvency constraints can deliver (1) a higher market price of risk, (2) time-variation in the market price of risk and possibly (3) persistence.

6.3 Changes in the Wealth Distribution and the Market Price of Risk

Only the risk associated with changes in the wealth distribution affects excess returns. I refer to this as aggregate liquidity risk. The risk due to changes in the relative wealth position affects the risk-free rates but leaves excess returns unchanged. To distinguish between aggregate and idiosyncratic liquidity risk, I consider the case in which aggregate uncertainty is i.i.d. and the conditional distribution of y' is independent of z'.

The next proposition delivers a generic "perfect" aggregation result for economies with limited commitment and i.i.d. aggregate uncertainty.[28]

Proposition 6.1 *: If aggregate uncertainty is i.i.d. and $\pi(y'|y)$ is independent of the aggregate state, then there is a stationary equilibrium in which g^* is constant.*

The multipliers grow at the same rate in each period. If we rescale all of the weights by $g_{t+1}(z^{t+1}) = h_{t+1}(z^{t+1})/h_t(z^t)$ after every period, then the joint distribution over these "discounted" weights and endowments is stationary. This implies the wealth distribution itself is stationary and there is no aggregate

[27]A similar logic applies to models with wealth constraints in which the traded assets do not span the space.

[28]The proof uses some insights by Alvarez and Jermann (2001) for two-agent economies and by Krueger (1999) on the existence of a stationary measure in a similar environment without aggregate uncertainty.

liquidity risk. Interest rates are lower than in the benchmark model:

$$r = 1/\left(\beta \sum_{z'} \pi(z') e^{-\gamma} g^{*\gamma}\right),$$

but the market price of risk is unchanged from its benchmark model value: $\sigma(e^{-\gamma})/E\left(e^{-\gamma}\right)$. The representative agent for this economy has the same risk aversion as the agents but a higher $\widetilde{\beta} = \beta g^{*\gamma}$. The higher mean for the discount factor may help resolve the equity premium puzzle some by allowing higher risk aversion with lower and constant risk-free rates -it pushes the stars in Figure 1 to the right and closer to the cup-sized region by increasing Em, but it cannot go all the way because perfect risk sharing obtains if γ is set too high. More importantly, it cannot deliver cyclicality of the market price of risk or predictability.[29] Cyclical changes in the wealth distribution are essential, either via shocks to the dispersion of income or via shocks to the collateral share.

Mankiw effect In an *exogenously incomplete markets setting*, Mankiw (1986) shows how excess returns do compensate for idiosyncratic consumption risk if the dispersion of consumption growth is negatively correlated with returns. In the case of *endogenously incomplete markets*, the excess returns are driven by the covariance with changes in the γ^{-1}-th moment of the weight distribution. Suppose there is no aggregate consumption risk. The expected excess returns are given by:

$$ER^e = -r^f cov\left(g_{t+1}^{\gamma}, R^s\right). \quad (32)$$

This is the endogenously incomplete markets version of the Mankiw mechanism. My model activates this version of the Mankiw mechanism if labor income risk is countercyclical or if the capital share is countercyclical. Both of these have some empirical support.

Labor Income Risk Households do seem to face more idiosyncratic risk during recessions. Storesletten, Telmer and Yaron (1998) estimate that the conditional standard deviation of idiosyncratic shocks more than doubles when the economy goes into a recession while Attanasio and Davis (1996) have found the distribution of consumption to be very sensitive to relative wage changes.

Consider an example with two idiosyncratic states and aggregate states and let y_2 be the good state. Suppose that the cross-sectional distribution fans out

[29] This result also applies to economies with permanent exclusion, with lower g^*.

in recessions and that recessions are short-lived. This has two effects on the relative value of human wealth for a household currently in the good state: (1) the current high income share is larger in recessions:

$$\widehat{\eta}(y_2, z_{re}) \geq \widehat{\eta}(y_2, z_{ex}), \qquad if \ \pi_{re}(y_2) = \pi_{ex}(y_2).$$

(2) the price of the claim to future labor income is higher; the recessions are short-lived but the labor income process is persistent and this household is likely to find itself in a high state in an expansion in the next state (y_2, z_{ex}). Its expected future labor income is large and less risky. This increases $\Pi_{y_2, z^{t-1}, z_{re}}[\{\eta\}]$ relative to $\Pi_{z^{t-1}, z_{re}}[\{e\}]$. As a result the cutoff rule in the good state is much higher in recessions, $\underline{l}(y_2, z_{re}, z^{t-1}) > \underline{l}(y_2, z_{ex}, z^{t-1})$ regardless of z^{t-1}. This in turn generates larger liquidity shocks in recessions, $g(\ldots, z_{ex}) < g(\ldots, z_{re})$, regardless of the history z^{t-1}.[30] The volatility of the discount factor increases because the correlation of λ and g is negative:

$$m_{t,t+1} = \beta \lambda (z_{t+1})^{-\gamma} g_t(z^{t+1})^{\gamma}.$$

In addition, its mean increases as well. This moves the moments of the discounter in the right direction relative to the benchmark model (see Figure 1). At the same time, returns on aggregate consumption claims are low in recessions. This delivers the right covariance between returns and liquidity shocks to boost excess returns via the mechanism in (32). At least in theory the model can deliver a higher market price of risk and higher excess returns.

Collateral Risk The labor income risk channel is one way of activating the Mankiw mechanism for higher market prices of risk and excess returns, similar to Constantinides and Duffie (1996). The collateral mechanism offers a second way. Up to now, I have been keeping the capital share α constant, but in the US it has experienced several large shocks. For example, between 1929 and 1932 the *capital share of national income* dropped from 41 to 28 percent and between 1966 and 1975 the capital share of national income dropped from 35 percent to 25 percent (source: NIPA). The capital share hovered around this level until the 90's. In particular, corporate payouts to securities holders as a share of GDP are subject to large swings (Hall, 2001). If the stream of corporate payouts drops persistently, non-labor wealth is destroyed and sharing risks becomes harder.

[30] see p.57 in the Appendix for a derivation of the size of the liquidity shock.

Recall that the solvency constraint can be stated as:

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

Consider the simplest case of 2 aggregate states. If $\alpha(z_{re}) < \alpha(z_{ex})$, then all the income shares in the recession are scaled down and this constraint is more likely to bind in recessions. By the same logic as before, liquidity shocks will be larger in recessions: $g(\ldots, z_{ex}) < g(\ldots, z_{re})$. This collateral effect amplifies the labor market effect, boosting the volatility of the discounter and delivering the right covariance between returns and liquidity shocks.

6.4 History Dependence and Time-variation of the Market Price of Risk

The aggregate history plays a key role in the economy with a continuum of agents. As before, suppose the cross-sectional distribution fans out and recessions are short. In a recession all of the agents who draw the high income share face binding constraints and get a high consumption share. Their "discounted" Pareto-Negishi weight $\zeta_t/h_t(z^t)$ reaches its maximum value and after that their weights drift downwards in the distribution. Agents deplete their assets in light of the low interest rates. This process continues until the next recession hits and they draw the high income share again. Now, if the last recession was a couple of periods ago, a large fraction of agents will have fairly low "discounted" weights. These households have little assets left and are severely constrained in the recession if they draw the high state y_2 : they have no collateral to borrow against. The liquidity shock induced by a recession is large. By contrast, if the last recession was last period, their weights are still fairly close to the cutoff level and the liquidity shock is small. As an example, this would imply the following:

$$g(.., z_{ex}, z_{ex}, z_{ex}, z_{re}) > g(.., z_{re}, z_{re}, z_{re}, z_{re}). \tag{33}$$

The market price of liquidity risk, $\sigma_t(g_{t+1}(z^{t+1})^\gamma)/E(g_{t+1}(z^{t+1})^\gamma)$, reaches a peak after a long string of good shocks, not because recessions are more likely but rather because of the large liquidity shock that results when a recession does set in. That is when its conditional standard deviation is highest. This delivers a time-varying market price of risk and the history dependence may also help to deliver persistence.

This asymmetric history-dependence is absent in the two-agent version of this economy, analyzed (with permanent exclusion) by Alvarez and Jermann (2001a) or bankruptcy (Lustig, 1999). Only one of two things can happen. If no one is constrained or the same agent as last period is equally constrained, the state variable -the relative consumption weight- stays the same. In this case the perfect insurance SDF applies.[31] If a different agent is constrained, the relative weight hits a new upper or lower bound. This is a large liquidity shock, but it resets the initial condition for this economy.[32] The state variable stays nearly unchanged for long periods of time since labor income is persistent: the "low wealth" agent needs to move to the "good" state to generate a liquidity shock but these switches are infrequent if labor income is persistent. In the continuum economy these switches occur all the time.[33]

6.5 Pseudo-habit

The distribution of cumulative weights maps into a wealth distribution. A priori, one would expect asset prices to depend on the entire wealth distribution but it turns out that one moment of the distribution is sufficient.

Another way of stating this aggregation result is that there is a representative agent with preferences that depend on a single moment of the wealth distribution. This agent ranks stochastic aggregate consumption streams according to:

$$U(C) = E_0 \sum_{t=1}^{\infty} \frac{(C_t - X_t)^{1-\gamma}}{1-\gamma},$$

where $X_t = C_t \left(1 - h_t^{-1}\right)$ and $h_0 = 1$. This agent has a stochastic habit X_t which grows over time, even if consumption itself does not. Whenever a large fraction of households are constrained -i.e. h_{t+1}/h_t is large-, the habit jumps up, pushing up the marginal utility of consumption.

The curvature of the utility function at each point in time is given by γh_t. As more people are constrained, this representative agent's habit approaches her consumption from below and she becomes less willing to take on consumption

[31] If the same agent is constrained and the constraint binds more than last period (e.g. when moving from expansion to recession), his consumption share increases some, but this is a small shock.

[32] See Appendix for a formal statement of this result on p. 58.

[33] DenHaan (2001) provides a detailed comparison of two-agent and large-number-of-agents economies within an incomplete markets setting.

gambles. (h_{t+1}/h_t) can be interpreted as a recession proxy: it is high in bad times and low in good times.

The risk associated with changes in the wealth distribution or simply liquidity risk reconciles low risk aversion of individual agents with high implied risk aversion at the aggregate level while the growth of the habit over time mitigates the agent's desire to flatten the consumption profile. The first effect delivers a large and time-varying market price of risk, the latter delivers low interest rates.

Campbell and Cochrane (1999) endow the representative agent with an external habit to obtain time-varying curvature of the utility function. They manage to deliver a large equity premium, time-varying market prices of risk and predictability of returns. In their model the process for the habit is set up to generate the right SDF dynamics. In my model, the sensitivity of the pseudo-habit to aggregate shocks depends endogenously on the history z^t. In particular, the habit responds more to a negative consumption growth shock after a string of good shocks (see eq. (33)). The persistence of the habit depends on the persistence of liquidity shocks.

7 Quantitative Results

This section calibrates a parsimonious version of the model and compares the moments of asset prices generated by the model to those in the data.

7.1 Computation

In this economy agents can (1) keep track of the distribution of weights or (2) the entire history of the economy z^∞. If the first approach is taken, the distribution can be approximated on the computer using a finite vector of moments and this vector can be used to forecast g_{t+1}, in the spirit of Krusell and Smith (1998). I will take the second approach and work around the curse of dimensionality by truncating aggregate histories. Let z_t^k denote $(z_t, z_{t-1}, \ldots, z_{t-k})$. I will define a stationary stochastic equilibrium in which the joint distribution over weights and endowments conditional on a history z^k is invariant on average. This means that in a sample of aggregate history simulations the average distribution over weights and endowments in z^k equals this stationary measure. This stationary measure exists and is unique.[34]

[34] The details are stated in the Computational Appendix on p. 78.

This approximation clears the market on average in each truncated history z^k. Agents do not keep track of the entire aggregate history, only a truncated version. Their policy functions (which have the cutoff rule property) and liquidity shock forecasting functions map truncated aggregate histories into new **consumption** weights and forecasts[35]: $l(\vartheta, y', z'; z^k)$ yields the new weight for an agent entering the period with consumption weight ϑ, private shock y' and aggregate shock z', conditional on a truncated history z^k. The liquidity shock forecasting function is given by $g^*(z', z^k)$ and consumption is given by $c = \vartheta / g^*(z', z^k)$. The optimal forecast when going from state z^k to z' is given by its unconditional average:

$$g^*(z', z^k) = \sum_{y'} \int l\left(\vartheta, y', z'; z^k\right) \Phi^*_{z^k}\left(d\vartheta \times dy\right) \frac{\pi(y', z'|y, z)}{\pi(z'|z)}, \qquad (34)$$

for each pair (z', z^k), where Φ_{z^k} denotes the stationary measure over weights and endowments in z^k. Prices in this economy will be set on the basis of (34) the following SDF:

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

Households know this SDF perfectly and as a result they make no Euler equation mistakes but the Walrasian auctioneer does make mistakes in setting those prices. In Krusell and Smith (1998), households make Euler equation errors but the markets clear. In my approximation, markets only clear on average in each aggregate state and truncated history (z', z^k). Let $\widetilde{\Phi}_{z^k}$ be the actual distribution in that state (which depends on the entire history z^{∞}). On average in node z^k, when the next shock is z', the markets will clear, because aggregate consumption equals:

$$c(z', z^k) = \frac{\sum_{y'} \int l\left(\vartheta, y', z'; z^k\right) \widetilde{\Phi}_{z^k}\left(d\vartheta \times dy\right) \frac{\pi(y', z'|y, z)}{\pi(z'|z)}}{g^*(z', z^k)} e(z').$$

On average, $\widetilde{\Phi}_{z^k} = \Phi^*_{z^k}$ by the definition of a stationary measure and the market clears: $E_{z^\infty \subseteq z^k} c(z', z^k) = e(z')$. For any given realization the actual growth rate differs from the average one because the distribution over weights and endowments differs from the average one. If these errors tend to be small, this is a good approximation. I report statistics for the percentage allocation error:

$$x = \left| \frac{c(z', z^k) - e(z')}{e(z')} \right|$$

[35] The consumption weights ϑ are equal to $\zeta^{1/\gamma}$.

as a measure of the closeness of the approximation. Recall that $\sup|x| = 0$ for $k = 0$ if aggregate uncertainty is i.i.d.[36] As $k \to \infty$, the approximation approaches the actual equilibrium.

7.2 Calibration

Endowment process To make the results comparable to Alvarez and Jermann (2001a), I follow their calibration approach. The endowment process has two different values for the income share and two different values for the growth rates. The subscripts h and l indicate a high income and low income share. The subscripts ex and re indicate an expansion and a recession.

The first four moments of the four-state Markov process describe US output dynamics in the 20-th century and are based on Mehra and Prescott (1985). The remaining 6 moments characterize the household's labor income and its relation to aggregate growth. These combine information from estimation results in Heaton and Lucas (1996) and Storesletten, Telmer and Yaron (1997). The moments are listed in Table 7. The cross-sectional dispersion of labor income shares in state s is defined as

$$v^2(z) = .5 \sum_{y=y_l,y_h} [\widehat{\eta}(y,z) - .5]^2 = 1,$$

and the relative standard deviations of individual shares conditional on current and past realizations of the aggregate shocks are defined as (M8):

$$\frac{\sigma_{r'e}}{\sigma_{e'e}} = \frac{std(\ln \widehat{\eta}(s_{t+1})|\lambda_{t+1} = \lambda_r, \lambda_t = \lambda_e)}{std(\ln \widehat{\eta}(s_{t+1})|\lambda_{t+1} = \lambda_e, \lambda_t = \lambda_e)}.$$

Finally, the relative standard deviations of individual shares conditional on past realizations of the aggregate shocks (M10) are defined as:

$$\frac{\sigma_{r'}}{\sigma_{e'}} = \frac{std(\ln \widehat{\eta}(s_{t+1})\lambda_t = \lambda_r)}{std(\ln \widehat{\eta}(s_{t+1}), \lambda_t = \lambda_e)}.$$

The key thing to note is that the cross-sectional dispersion of labor income fans out in recessions. This mechanism helps deliver a countercyclical market price of risk.

[36] See exact aggregation result in Proposition 6.1.

Share of Collateral The scarcity of collateral is key to generating large risk premia in my model. In the next section, I provide a sensitivity analysis by reporting results for $\alpha = 10\%$, $\alpha = 15\%$ and $\alpha = 20\%$.

How do these numbers relate to the US data? The size of the Lucas tree dividend as a share of the total endowment, α, can be matched to the capital share of net national income. The average labor share of national income in the US between 1946 and 1999 is 70 percent (source, NIPA[37]; see Table 6). An additional 11 percent is proprietor's income derived from farms and partnerships (mainly doctors and lawyers). This should be treated as labor income for the purposes of this exercise. This brings the total labor share to 81 percent.

There are two facts that could further reduce the supply of collateralizable wealth. First, a substantial share of the remaining 18 percent are profits from privately held firms. Jorgensen and Moskovotiz (2000, p. 45) report that on average 19 percent of the value of corporate equity stems from privately held companies. A substantial part of these profits consist of remuneration for labor services provided by the entrepreneur. In addition, these assets are highly illiquid.

Second, bankruptcy exemptions effectively reduce the supply of collateral in the economy. An economy with collateral share α and a zero percent exemption is formally equivalent to an economy with $(1-\alpha)/(1-\phi)$ labor share where ϕ is the proportional exemption. To see why, note that the solvency constraints in this economy are given by:

$$\Pi_{s^t}\left[\{c\}\right] \geq \frac{\Pi_{s^t}\left[\{\eta\}\right]}{1-\phi}.$$

All of the labor income is scaled up by $1/(1-\phi)$. Exemptions lower the effective capital share in the economy. Assuming there is a proportional exemption rate of 10%, the effective supply of collateral in the economy is only 9% (1-.81/.90).[38]

[37] NIPA does not provide a consistent treatment of the household sector in that it make no attempt to include the flow of services of non-residential consumer durables. (see Cooley and Prescott, p. 19, 1995).

[38] In the US a husband and wife would have a \$30,000 homestead (house) exemption under the 1994 Federal Amendments. Below the 75th net worth percentile households hardly have any financial assets and the most important asset is clearly the house: 66 percent of households own a house and its median value is \$ 100.000 (SCF, 1998).

7.3 Results

The asset pricing statistics were generated by drawing 20.000 realizations from the model. In the benchmark calibration the discount factor β is set to .95 and the share of collateral α is 10 percent. Table 8 in the Appendix lists the percentage allocation errors for the benchmark calibrations where the percentage error $x = \frac{c(z',z^k)-e(z')}{e(z')}$. For $k = 4$ the mean of the allocation errors is close to .5 percent for all computations, while the standard deviation is roughly the same size. Increasing k to 5 lowers the mean allocation errors to around .05 percent but the standard deviation is roughly the same order of magnitude. The low standard deviation indicates that the errors are tightly distributed around zero. The sup norm decreases to 2 percent when increasing k from 4 to 5.

The key statistics for this benchmark calibration are reported in Table 9. The first line gives the statistics for the Mehra Prescott data. The returns and price/dividend ratios all pertain to a claim to aggregate consumption. I have also reported the corresponding numbers for the representative agent economy. Table 10 reports the same statistics for $k = 5$. The results are quasi-identical. Increasing k increases the precision but does not change the results.

7.3.1 Equity premium and Risk-free rate

Table 2 in the text summarizes the results for the benchmark calibration: $\gamma = 7$ and $\beta = .95$. The equity premium exceeds 6 percent, while the risk-free rate is .8 percent. These numbers match the data. The market price of risk is .41 and exceeds the Sharpe ratio of .35 in the Mehra Prescott data.

Table 2: **Summary of stats.**

$\beta = .95$	r^f	ER^e	$\sigma m/Em$	$\exp(p/d)$	$\sigma(r^f)$	$\sigma(R^s)$
US (MP)	.80	6.1	.37	25	5	17
γ=**7**	**.8**	**6.1**	**.41**	**30**	**13**	**19**
rep. agent	11.6	2.3	.28	8.3	6.8	11

How does this model manage to match both the equity premium and the risk-free rate? The risk-free rate is pushed down by 1000 basis points relative to its perfect insurance value. Two forces drive down the interest rate. The constraints prevent agents from borrowing against their growing labor income. This is the direct liquidity effect. Moreover, the risk-free asset provides a hedge against liquidity shocks and earns a lower return in equilibrium. This is the precautionary liquidity effect. To see these forces at work, suppose $log(e_t/h_t)$ follows a random walk with drift $E_t\lambda_{t+1} - E_t g_{t+1}$ and variance $\sigma_t \Delta \log(e_{t+1}/h_{t+1})$, and assume conditional normality. Then the risk-free rate is given by this expression:

$$
\begin{aligned}
r_t \quad = \quad & \underbrace{-\ln\beta + \gamma E_t \ln\lambda_{t+1} - \frac{\gamma^2}{2}\sigma_t^2 \Delta \log e_{t+1}} \quad \textit{perfect insurance} \\
& -\gamma E_t \ln g_{t+1} - \frac{\gamma}{2}^2 \left(\sigma_t^2 \Delta \log(e_{t+1}/h_{t+1}) - \sigma_t^2 \Delta \log e_t\right), \qquad (35)
\end{aligned}
$$

The first term in (35) is the perfect insurance risk-free rate, the second term captures the direct liquidity effect and the third term captures the additional precautionary effect. This last term is positive because liquidity shocks and consumption growth are negatively correlated; this implies $\sigma_t \Delta \log(e_{t+1}/h_{t+1}) > \sigma_t \Delta \log(e_{t+1})$. That drives the risk-free rate below its perfect insurance counterpart.

The average investor wants to be compensated for taking on aggregate liquidity risk by holding equity because returns are low in recessions, when the liquidity shocks are large:

$$
ER^e = -\beta r^f cov\left(\lambda_{t+1}^{\gamma} g_{t+1}^{\gamma}, R^s\right). \qquad (36)
$$

This explains why equity premia are high.

The model can pass the market price of risk test but what about the Hansen-Jagannathan bounds implied by the returns on stocks and bonds? Figure 3 plots model-generated moments of the SDF against the H-J bounds for $\gamma = 7$. The increase in Em and σm that result from liquidity risk pushes the SDF moments well inside the bounds for a large range of $\beta's$. As β is lowered, agents share less risks because today's income draw is weighted heavily in the price of the claim on labor income and consumption has to be more correlated with income to keep agents from defaulting. This generates more aggregate liquidity risk and the market price of risk increases. Figure 4 plots the equity premium for the same range of parameters. The equity premium rises as γ increases and as β decreases. For $\beta = .85$ it reaches values of up to 8 percent.

The effect of changes in risk aversion is twofold. First, more risk averse agent have lower outside options and share more risks in equilibrium. This decreases the volatility of the liquidity shocks. Second, at higher γ, the liquidity shocks have a bigger impact on the volatility of the discounter σm. The second effect dominates for the range of parameters I consider. This accounts for the increase in the market price of risk as γ increases (see Table 9). Figure 5 plots the Sharpe ratio against γ and β.

For large γ, the excess returns are much more volatile than in the representative agent model. This brings the model closer to the data. However, the time-variation in the expected liquidity shocks imputes the risk free rate with too much volatility relative to the data (see $\sigma(r^f)$ in Table 9) . This is a feature of the link between the intertemporal elasticity of substitution and risk aversion in the case of power utility: highly risk averse agents are reluctant to substitute consumption intertemporally and this gives rise to large changes in interest rates in response to changes in $E_t \ln g_{t+1}$; this is obvious from $-\gamma E_t \ln g_{t+1}$ in eq. (35). Introducing Epstein-Zin preferences or a production technology should alleviate this problem.

The stochastic discounter is the "composite" IMRS of the unconstrained agent in each state and low individual consumption volatility is consistent with large SDF volatility. In fact, individual consumption is smooth even though the composite IMRS is volatile. Figure 6 plots the risk sharing coefficient

$$std(h_{t+1}/h_t)/std(h_{t+1}^{aut}/h_t^{aut}).$$

This number is bounded between zero (perfect insurance) and 1 (autarchic equilibrium). It remains below 1.5 percent across the parameter grid. Clearly, there is substantial risk sharing going on.

Collateral share The results are highly sensitive to α, the share of collateral. Table 11 lists the same means and standard deviations for $\alpha = 15$ percent. The equity premium drops to 5.1 percent while the risk-free rate increases to 4 percent. When $\alpha = 20$ percent, the numbers are much closer to those for the full insurance economy (see Table 12). Figure 7 plots the Sharpe ratios for both economies. The ones in the high collateral economy lie uniformly below the ones in the low collateral economy.

In the data, payouts by businesses to securities holders are an order of magnitude more volatile than aggregate consumption. What happens if α drops in recessions? Table 3 in the text shows the effect collateral share volatility: the collateral share declines in recessions to 10 percent and 5 percent respectively.

This lowers the risk-free rate from 4 percent to 1 percent, increases the excess return by 100 basis points and doubles the volatility of excess returns.

Table 3: **Collateral**

$\beta = .95, \gamma = 7$	r^f	ER^e	$\sigma m/Em$	$\sigma(r^f)$	$\sigma(R^e)$
$\alpha = 15\%$	4	5.1	.36	11	13
$\alpha_{ex} = 15\%; \alpha_{re} = 10\%$	2	5.6	.40	11	23
$\alpha_{ex} = 15\%; \alpha_{re} = 5\%$	1	6	.41	12	28

Permanent Exclusion For the range of parameters considered here, the value of autarky for highly risk-averse agents is so low that the constraints rarely bind in the economy with permanent exclusion and the aggregates resemble those in the benchmark economy closely. Figure 8 plots the Sharpe ratios for the continuum-of-agents economy with permanent exclusion. Even for $\beta = .7$ the SDF is not nearly volatile enough: liquidity shocks are much smaller than in the two-agent economy considered by Alvarez and Jermann (2001a).

Two-agent economy There is less risk sharing in the two-agent economy but excess returns are lower. Table 4 in the text compares the AJ-benchmark calibration to its continuum-of-agents and/or bankruptcy alternatives. The volatility of the SDF is lower in the continuum economy because household risk is truly idiosyncratic. With permanent exclusion, the AJ-benchmark generates allocations very close to perfect risk sharing. The market price of risk drops from 1 with two agents to .16 with a continuum.

Table 4: **Comparison**

$\beta = .65, \gamma = 4$	Permanent Exclusion	$\alpha = 0\%$	Bankruptcy	$\alpha = 10\%$
# agents	2 *(AJ benchmark)*	continuum	2	continuum
$\sigma m/Em$	1	.16	4.3	.60
ER^e	3.19	1.9	9	12

Figure 9 plots the liquidity shocks in the two-agent economy. These shocks are large but infrequent. They are large because the individual shocks are not independent but -by definition- perfectly negatively correlated. Figure 10 shows how the market price of risk is much higher in the two-agent economy, but this does not translate into higher returns because the shocks are not as strongly correlated with the business cycle as in the economy with a continuum of agents (see Table 4). Alvarez and Jermann (2001a) report an equity premium of only 3.4 percent even though the market price of risk is one.[39] Table 14 lists the two-agent results for my benchmark calibration with bankruptcy. Excess returns are still lower than in the continuum economy.

7.3.2 Time Variation and Predictability

The model generates a substantial amount of time variation in the conditional market price of risk. Under lognormality the conditional market price of risk can be approximated by $\gamma\sigma_t\Delta\log(e_{t+1}/h_{t+1})$. Liquidity shocks are largest in recession preceded by a large number of small liquidity shocks. This causes the conditional variance of the liquidity shocks and the conditional market price of risk to peak at the end of long expansions. On the other hand, after a series of recessions, the conditional market price of risk is low. Even if another recession hits, few households will be severely constrained. Figure 11 plots the conditional market price of risk and the liquidity shocks for the benchmark calibration economy, keeping the same sequence of shocks $\{z^t\}$; the stars indicate expansions. Note that the conditional market price of risk drops by 50 percent after a string of large liquidity shocks. The left tail of the wealth distribution is completely "erased" and a new recession does not cause a significant liquidity shock. Table 15 contains the size of the liquidity shocks for different histories for the benchmark calibration and Table 16 reports the conditional market price of risk. Note that the shocks are much larger if the recession is preceded by expansions. Figure 13 shows how the distribution of weights grows more skewed to the left during long expansions.

Are the shocks persistent enough for the standard deviation of the l-period SDF to grow quickly as the horizon l increases? Figure 15 plots the moments of the discounters for the benchmark calibration of the model for $l = 1$ to 7. The standard deviation of the perfect insurance discounter hardly increases and decreases after $l = 4$. The moments of the model with bankruptcy both grow quickly enough to stay inside or at least close to the cup-sized region. This is

[39]Introducing bankruptcy in the two-agent economy at these low discount rates ($\beta = .65$) leaves little room for risk sharing and the market price of risk explodes to 4.

delivered by the autocorrelation in the shocks g_t : small shocks tend to persist for a while and are followed by a sequence of larger shocks. Table 17 reports these moments for the benchmark calibration. The time-variation in risk premia induces highly volatile excess returns (see Table 9).

Finally, Table 18 lists the R-squared and the slope coefficient of Fama-French regressions of log stock returns on log price/dividend ratios. The growth rate of dividends on stocks here was calibrated to have an unconditional standard deviation of 15%. (The results for aggregate consumption claims are identical at longer holding periods but the R-squared are higher at shorter holding periods). Note how the R-squared dramatically increases as l increases. The same pattern was found in the data by numerous authors.[40]

One distinct advantage of the two-agent setup is that the second part of the stochastic discounter is almost non-stationary for large γ; its conditional standard deviation explodes over longer horizons, because of the size of the shocks.[41] Table 5 in the text provides a summary of the main differences between these two setups.

Table 5: **Summary**

	two-agent	continuum
$corr(m, R^e)$	smaller	larger
$\sigma m/Em$	larger	smaller
$\sigma_t m/E_t m$	mostly constant	volatile
history dependence	no	yes
predictability	yes, for large samples	yes
non-stationary m_t	yes, for large γ	no
	Permanent Exclusion	Bankruptcy
large ER^e	low β, low γ in two-agent economy	high β, higher γ
	not in continuum	

[40] See Campbell, 2000, p. 1522 for an overview.

[41] Campbell and Cochrane (1998) and Alvarez and Jermann (2001b) argue it is important to have a non-stationary SDF.

8 Conclusion

The introduction of bankruptcy brings the asset pricing implications of Lucas' s (1978) endowment economy in line with the data. The liquidity risk created by binding solvency constraints doubles the unconditional market price of aggregate risk, provided that large liquidity shocks coincide with low aggregate consumption growth. This pattern emerges in equilibrium when the capital share of income drops in recessions or when the cross-sectional dispersion of labor income fans out in recessions. The model can also match conditional moments. The endogenous build-up of low-wealth households at the end of expansions generates large and persistent changes in the conditional risk premia. As a result, the returns on stocks are highly predictable. However, the model does generate too much interest rate volatility relative to the data. This defect could be mitigated by severing the link between risk aversion and the intertemporal elasticity of substitution, or by introducing a production technology.

The analysis uncovers a new link between equilibrium changes in the wealth distribution and risk premia, and the calibration suggests that the effects on risk premia are large enough to close at least part of the gap between standard, consumption-based theories of asset prices and the data. The representative agent in my model has a stochastically growing pseudo-habit, even though the households have power utility; the shocks to the habit are driven by changes in the wealth distribution. When a large fraction of agents is constrained, the habit approaches aggregate consumption from below and the implied risk aversion increases. If these results continue to hold in a model with capital accumulation, Krusell and Smith's (1998) irrelevance of the wealth distribution does not survive the introduction of endogenous solvency constraints.

This paper extends the methods developed by Atkeson and Lucas (1992, 1995) and Krueger (1999) by showing how to compute prices and allocations in limited commitment economies with a continuum of agents and aggregate uncertainty, using a version of Marcet and Marimon's (1999) cumulative multiplier approach. The computations are simple because the weight updating policy has a cutoff rule property. A version of this algorithm can be applied to economies with private information and aggregate uncertainty.

References

[1] Abel, A. (1990), Asset Prices under Habit Formation and Catching up with the Joneses, American Economic Review, vol. 80, 38-42.

[2] Abreu, D., Pearce, D. and E. Stacchetti (1986), Optimal Cartel Equilibria with Imperfect Monitoring, Journal of Economic Theory, vol. 39, 251-69.

[3] Alvarez, F. and U. J. Jermann (2000 a), Efficiency, Equilibrium, and Asset Pricing with Risk of Default , Econometrica, vol. 68, 4, 775-798.

[4] Alvarez, F. and U. J. Jermann (2000 b), Using Asset Prices to Measure the Marginal Cost of Business Cycles, mimeo, University of Chicago and Wharton.

[5] Alvarez, F. and U. J. Jermann (2001a), Quantitative Asset Pricing Implications of Endogenous Solvency Constraints, Review of Financial Studies, vol. 14., 1117-1152.

[6] Alvarez, F. and U. J. Jermann (2001b), The Size of the Permanent Component of Asset Pricing Kernels, Discussion Paper, mimeo, University of Chicago and Wharton.

[7] Atkeson, A. and R. E. Lucas, Jr. (1992), On Efficient Distribution with Private Information, Review of Economic Studies, vol. 59, 427-453.

[8] Atkeson, A. and R. E. Lucas, Jr. (1995), Efficiency and Equality in a Simple Model of Unemployment Insurance, Journal of Economic Theory, vol. 66., 64-85.

[9] Attanasio, O. and S. Davis, Relative Wage Movements and Distribution of Consumption, Journal of Political Economy, vol. 104(6), 1227-1262.

[10] Breeden, D. T. (1979), An Intertemporal Asset Pricing Model with Stochastic Consumption and Investment Opportunities, Journal of Financial Economics, vol. 7(3), 265-296.

[11] Browning, M., L. P. Hansen, and J. J. Heckman (2000), 'Micro data and General Equilibrium Models', in J. Taylor and M. Woodford (eds.), Handbook of Macroeconomics. Amsterdam: North Holland.

[12] Campbell, J., A. Lo and A. MacKinlay (1997), The Econometrics of Financial Markets, Princeton: Princeton University Press.

[13] Campbell, J. (1999), Asset Prices, Consumption and the Business Cycle, in J. Taylor and M. Woodford (eds.), vol. 1, Handbook of Macroeconomics, Amsterdam: North Holland.

[14] 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, vol. 107, 205-251.

[15] Campbell, J. (2000), Asset Pricing at the Millennium, Journal of Finance, vol. 55 (4), 1515-1567.

[16] Cogley, T. (1999), Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey', mimeo, Arizona State University.

[17] Cochrane, J. H. and L. P. Hansen (1992), Asset Pricing Explorations for Macroeconomics, NBER Macroeconomics Annual, 115-64.

[18] Cochrane (2001), J. H. , Asset Pricing, Princeton: Princeton University Press.

[19] Constantinides, G. M. (1982), Intertemporal Asset Pricing with Heterogeneous Consumers and without Demand Aggregation, Journal of Business, vol. 55(2), 253-267.

[20] Constantinides, G. M. and D. Duffie (1996), Asset Pricing with Heterogeneous Consumers, Journal of Political Economy, vol. 104, 219-240.

[21] Cooley, T. F. and E. C. Prescott (1995), Economic Growth and Business Cycles, in Cooley (ed.): Frontiers of Business Cycle Research, Princeton: Princeton University Press.

[22] Den Haan, W. J. (2001), Understanding Equilibrium Models with a Small and a Large Number of Agents, Journal of Economic Dynamics and Control , vol. 25(5), 721-746.

[23] Detemple, J. and A. Serrat (1999), Dynamic Equilibrium and Liquidity Constraints, mimeo, Université de Montreal.

[24] Dubey, P., J. Geneakoplos and M. Shubik (1990), Default and Inefficiency in a General Equilibrium Model with Incomplete Markets, Discussion Paper, Cowles Foundation.

[25] Epstein, L. G. and S. Zin (1989), Substitution, Risk Aversion and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework, Econometrica, vol. 57(4), 937-969.

[26] Fama, E. F. and K. French (1988), Dividend Yields and Expected Stock Returns, Journal Of Financial Economics, vol. 25, 23-49.

[27] Geanakoplos, J. and W. R. Zame (1998), Collateral, Default and Market Crashes, mimeo, Yale and UCLA.

[28] Green, E. J. (1987), 'Lending and the Smoothing of Uninsurable Income', in Prescott, E.C. and N. Wallace (eds.), Contractual Arrangements for Intertemporal Trade., Minnesota Studies in Macroeconomics series, vol. 1, Minneapolis: University of Minnesota Press, 3-25.

[29] Grossman, S. and R. Shiller (1981), The Determinants of the Variability of Stock Market Prices, American Econmic Review, vol. 71(2), 222-227.

[30] Grossman, S. and R. Shiller (1982), Consumption Correlatedness and Risk Measurement in Economies with Non-traded Assets and Heterogeneous Information, Journal of Financial Economics, vol. 10, 195-210.

[31] Hall, R. E. (2001a), The Stock Market and Capital Accumulation, forthcoming in American Economic Review.

[32] Hall, R. E. (2001b), Dynamics Of Business Earnings, mimeo, Stanford University.

[33] Hammond, P. and Y. Sun (2001), Monte Carlo Simulation of Macroeconomic Risk with a Continuum of Agents: The Symmetric Case, forthcoming in Economic Theory.

[34] Hansen, L. P. and K. J. Singleton (1982), Generalized Instrumental Variable estimation of Nonlinear Rational Expectation Models, Econometrica, vol. 50(5), 1269-1286.

[35] Hansen, L. P. and K. J. Singleton (1983), Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns, Journal of Political Economy, vol. 91(2), 249-65.

[36] Hansen, L. P. and R. Jagannathan (1991), Implications of Security Markets Data for Models of Dynamic Economies, Journal of Policitical Economy, vol. 99(2), 252-262.

[37] Hansen, L. P., T. J. Sargent and T. Tallarini (1999), Robust Permanent Income and Pricing, Review of Economic Studies, vol. 66(4), 873-907.

[38] He, H. and D. M. Modest (1995), Market frictions and Consumption-based Asset Pricing, Journal of Political Economy, vol. 103, 94-117.

[39] Heaton, J. and D. Lucas (1996), Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing, Journal of Political Economy, vol. 104, 668-712.

[40] Heaton, J. and D. Lucas (1997) , Portfolio Choice and Asset Prices; The importance of Entrepreneurial Risk, Working Paper Northwestern University.

[41] Kehoe, T. and D. K. Levine (1993), Debt Constrained Asset Markets, Review of Economic Studies, vol. 60, 856-888.

[42] Kehoe, T., D. K. Levine and E. Prescott (2001), Lotteries, Sunspots and Incentive Constraints, mimeo, University of Minnesota, forthcoming in JET.

[43] Kehoe, P. and F. Perri (2000), International Business Cycles with Endogenous Incomplete Markets, mimeo, University of Minnesota and Stern School of Business.

[44] Kiyotaki N. and J. Moore (1997), Credit Cycles, Journal of Political Economy, 105, 211-248.

[45] Kocherlakota, N. (1996), Implications of Efficient Risk Sharing without Commitment, Review of Economic Studies, vol. 63, 595-609.

[46] Krueger, D. (1999), Risk Sharing in Economies with Incomplete Markets, Ph.D. Thesis, University of Minnesota.

[47] Krueger, D. and H. Uhlig (2000), Competitive Risk Sharing Contracts with One-Sided Commitment, mimeo, Stanford University.

[48] Krueger, D. and F. Perri (2001), Private Insurance Markets or Redistributive Taxes?, mimeo, Stanford University and Stern School of Business.

[49] Krusell, P. and A. Smith, Jr. (1998), Income and Wealth Heterogeneity in the Macroeconomy, Journal of Political Economy, vol. 6, 867-896.

[50] Kubler, F. and K. Schmedders (2001), Stationary Equilibria in Asset-Pricing Models with Incomplete Markets and Collateral, mimeo, Stanford University and Kellogg.

[51] Lettau, M. (2001), Idiosyncratic Risk and Volatility Bounds, Federal Reserve Bank of New York Staff Report No. 130 (forthcoming in the Review of Economics and Statistics.)

[52] Ligon, T. and T. Worall (1997), Informational Insurance Arrangements in Village Economies, mimeo, Berkeley.

[53] Ljungqvist, L. and T. Sargent (2000), Recursive Macroeconomic Theory, Cambridge: The MIT Press.

[54] Lucas, R. E. (1978), Asset Prices in an Exchange Economy, Econometrica, vol. 46, 1429-45.

[55] Lustig, H. (1999), Secured Lending and Asset Prices, mimeo, Dept. of Economics, Stanford University.

[56] Luttmer, E. (1991), Asset Pricing in Economies with Frictions, Ph.D. Thesis, University of Chicago.

[57] Luttmer, E. (1996), Asset Pricing in Economies with Frictions, Econometrica, vol. 64, 1439-1467.

[58] Luenberger, D. G., Optimization by Vector Space Methods, 1969.

[59] Mankiw (1986), N. G., The Equity Premium and the Concentration of Aggregate Shocks, Journal of Financial Economics, vol. 17, 211-219.

[60] Marcet, A. and R. Marimon (1999), Recursive Contracts, mimeo, Universitat Pompeu Fabra.

[61] Marcet, A. and K. Singleton (1999), Equilibrium Asset Prices and Savings of Heterogeneous Agents in the Presence of Incomplete Markets and Portfolio Constraints, Macroeconomic Dynamics, vol. 3(2), 243-277.

[62] Mehra, R. and E. C. Prescott (1985), The Equity Premium: A Puzzle, Journal of Monetary Economics, vol. 15, 145-61.

[63] Monge A. (2000), Financial Markets, Creation and Liquidation of Firms and Aggregate Dynamics, mimeo, Northwestern University.

[64] Moskowitz, T. and A. Vissing-Jorgensen (2001), The Private Equity Premium Puzzle, mimeo, University of Chicago.

[65] Negishi, T. (1960), Welfare Economics and Existence of an Equilibrium for a Competitive Economy, Metroeconomica, vol. 12, 92-97.

[66] Phelan, C. and R.M. Townsend (1991), Computing Multi-period, Information-Constrained Optima, Review of Economic Studies, vol. 58(5), 853-881.

[67] Rubinstein, M. (1974), An Aggregation Theorem for Security Markets, Journal of Financial Economics, vol. 1, No 3, 225-244.

[68] Schneider, M. (1999), Essays on Banking, Asset Pricing and Learning, Ph.D. Thesis, Dept. of Economics at Stanford University.

[69] Shiller, R. (1982), Consumption, Asset Markets and Macroeconomic Fluctuations, Carnegie-Rochester Conference on Public Policy, vol. 17, 203-238.

[70] Stokey, N. L. and R.E. Lucas, Jr. with E. Prescott (1989), Recursive Methods in Economic Dynamics. Cambridge: Harvard University Press.

[71] Storesletten, K. , C. Telmer. and A. Yaron (1998), Persistent Idiosyncratic Risk and Incomplete Markets, mimeo, Carnegie Mellon University and Wharton School.

[72] Townsend, R. M. (1995), Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies, Journal of Economic Perspectives, vol. 9, 83-102.

[73] Uhlig, H. (2001), Harald Uhlig on Dynamic Contracts, Economic Dynamics Newsletter, vol. 2(2).

[74] Weil, Philippe (1989), The Equity Premium Puzzle and the Risk-free Rate Puzzle, Journal of Monetary Economics, vol. 24, 401-421.

[75] Wilson (1968), The Theory of Syndicates, Econometrica vol. 36(1), 119-132.

9 Appendix

- **Sequential Markets Equilibrium**:

This section defines a sequential markets equilibrium. 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

$\{\omega_t(\theta_0, s^t)\}$ to maximize her expected utility:

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

subject to the usual budget constraint:

$$c_t(\theta_0, s^t) + p_t^e(z^t)\omega_t(\theta_0, s^t) + \sum_{s'} a_t(s'; \theta_0, s^t) q_t(s^t, s') \leq \theta_t, \tag{37}$$

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

$$\begin{aligned} \left[p_{t+1}^e(z^{t+1}) + \alpha e(z_{t+1})\right] \omega_t(\theta_0, s^t) &\geq -a_t(s'; \theta_0, s^t) \text{ for all } s' \in S, \\ \text{where } s^{t+1} &= (s^t, s'). \end{aligned} \tag{38}$$

The definition of a competitive equilibrium is straightforward.

Definition 9.1 *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* $\{\omega_t(\theta_0, s^t)\}$ *and prices* $\{q_t(s^t, s'\}$ *and* $\{p_t^e(z^t)\}$ *such that (1) these solve the household problem (2) the markets clear*

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

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

- **Proof of Proposition 4.1:**

The solvency constraints imply that

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

and

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

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

$$U(\left\{c\left(\theta_0, y^t, z^t\right)\right\})(s^t) = \max_{\{c'\}} U(c)(s^t) \tag{39}$$

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

$$\kappa_t(s^t) = \max_{\{c'\}} U(c)(s^t) \tag{40}$$

such that the budget constraint is satisfied $\Pi_{s^t}\left[\{c'\}\right] \leq \Pi_{s^t}\left[\{\eta\}\right]$ and the solvency constraints are satisfied in all following histories: $U(c)(s^\tau) \geq \kappa_\tau(s^\tau)$ for all $s^\tau \geq s^t$. Second, if $U(\{c\left(\theta_0, y^t, z^t\right)\})(s^t) \geq \kappa_t(s^t)$, then from (39) and (40), it follows that $\Pi_{s^t}\left[\{\eta\}\right] \leq \Pi_{s^t}\left[\{c\left(\theta_0, y^t, z^t\right)\}\right]$. The second part is obvious.

- **Equivalence between sequential trading equilibria and Kehoe-Levine equilibria**

Assumption 1 : *Interest rates are high enough:*

$$\Pi_{s^0}\left[\{\eta\}\right] < \infty \text{ and } \Pi_{z^0}\left[\{e\}\right] < \infty \tag{41}$$

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

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

which is automatically satisfied for power utility.

- **Proof of Proposition 4.2:**

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

$$\int \sum_{y^t} \begin{bmatrix} \Pi_{s^t}\left[\{c\left(\mu_0, y^t, z^t\right)\}\right] \\ -\Pi_{s^t}\left[\{\eta\}\right]) \end{bmatrix} d\Phi_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} \geq 0 \tag{42}$$

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:

$$\sum_{z^\tau \succeq z^t} Q(z^\tau|z^t)\left[\int \sum_{y^\tau}\left[\begin{array}{c} c\left(\mu_0, y^\tau, z^\tau\right) \\ -\eta_\tau(y_\tau, z_\tau)e_\tau(z^\tau) \end{array}\right] d\Phi_0 \frac{\pi(y^\tau, z^\tau|y_0, z_0)}{\pi(z^\tau|z_0)}\right], \qquad (43)$$

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

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

$$\int \sum_{y^t} \eta_t(y_t, z_t)\frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} d\Phi_0 = \sum_{y'} \pi_{z_t}(y_t)\eta_t(y_t, z_t) = (1-\alpha) \qquad (44)$$

and the market clearing condition implies that:

$$\int \sum_{y^t} c\left(\mu_0, y^t, z^t\right)\frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} d\Phi_0 = e_t(z^t) \qquad (45)$$

Plugging eqs. (44) and (45) back into eq. (43) implies the following inequality must hold at all nodes z^t: $\alpha\Pi_{z^t}\left[\{e_t(z^t)\}\right] \geq 0$. If there is no outside wealth ($\alpha = 0$) in the economy, then the expression is zero at all nodes z^t and eq. (42) 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:

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

and their constraint is slack: $\Pi_{s^t}\left[\{c\left(\mu_0, y^t, z^t\right)\}\right] > \Pi_{s^t}\left[\{\eta\}\right]$. Given that eq. (42) 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

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

which have constraints that are violated: $\Pi_{s^t}\left[\{c\left(\mu_0', y^t, z^t\right)\}\right] < \Pi_{s^t}\left[\{\eta\}\right]$. If not, (42) would be violated. But this violates the participation constraints for these agents. So, for $\alpha = 0$, for all households with positive measure:

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

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

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

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

- **Proof of Proposition 4.3:**

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

- **Proof of Proposition 4.4:**

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

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

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

- **Dual Recursive Saddle Point Problem:**

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

$$D(c, \chi; w_0, s_0) = \sum_{t\geq 0}\sum_{s^t}\left\{\beta^t \pi(s^t|s_0)\left[\begin{array}{c} m_t(s^t|s_0)\chi_t(s^t|s_0)c_t(w_0, s^t) \\ +\tau_t(s^t)\Pi_{s^t}\left[\{\eta\}\right]\end{array}\right]\right\} \quad (46)$$

where:

$$\chi_t(s^t) = \chi_{t-1}(s^{t-1}) - \tau_t(s^t), \chi_0 = 1. \quad (47)$$

Then the **recursive dual** saddle point problem facing the household of type (w_0, s_0) is given by:

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

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.

- **Stochastic Discount Factor**

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

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

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

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

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

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

- **Computing the weights**

Proof of Lemma 5.1: The optimal weight updating policy consists of a cutoff rule $\{\underline{l}_t(y, z^t)\}_{y\in Y, z^t}$ where $l_0(\mu_0, s^0) = \mu_0$ and for all $t \geq 1$

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

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

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

and $C\left(\mu_0, s^t; l\right) \geq \Pi_{s^t}\left[\{\eta\}\right]$ for all s^t. The last inequality follows from the fact that $C()$ is non-decreasing in μ_0. It is easy to verify that there exist no other weight policy rules that satisfy these necessary conditions.

Since the optimal policy is to compare the current weight ζ to the cutoff rule $l_t(y, z^t)$, the continuation cost can be stated as a function of the current weight, the current idiosyncratic state and the aggregate history: $C\left(\mu_0, s^t; l\right) = C_t(\zeta, y, z^t)$. I am ready to define an equilibrium in terms of the individual weights and the average weights.

Theorem 9.1 *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*:*

$$h_t(z^t) = \sum_{y^t} \int \zeta_t^{1/\gamma}(\mu_0, y^t, z^t) d\Phi_0 \frac{\pi(y^t, z^t|y_0, z_0)}{\pi(z^t|z_0)} \tag{48}$$

and (iii) there are no arbitrage opportunities :

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

Proof: $\{\zeta_t(\mu_0, s^t)\}_{t=0}^{\infty}$ 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 $\left\{\zeta_t(\mu_0, s^t) - \zeta_{t-1}(\mu_0, s^{t-1})\right\}$ satisfy the Kuhn-Tucker conditions for a saddle point. The consumption allocations satisfy the first order conditions for optimality (see derivation of risk sharing rule). Market clearing is satisfied because $E\left[\zeta_t^{1/\gamma}(\mu_0, y^t, z^t)\right] = h_t(z^t)$ implies that

$$E\left[c_t(\mu_0, y^t, z^t)\right] = e_t(z^t)$$

Now, let $\theta_0 = C(\mu_0, s^0; l) - \Pi_{s^0}\left[\{\eta\}\right]$. 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 $\left(\theta_0(\mu_0), s^0\right)$ and I have recovered the equilibrium allocations $\{c_t(\theta_0, s^t)\}$ and the prices $\{p_t(s^t|s_0)\}$.

- **Transformation of growth economy** (Alvarez and Jermann, 2001a):

The aggregate growth rate is a function $\lambda(z_t)$. Let utility over consumption streams be defined as follows:

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

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

$$\widehat{\pi}(z_{t+1}|z_t) = \frac{\pi(z_{t+1}|z_t)\lambda(z_{t+1})^{1-\gamma}}{\sum_{z_{t+1}} \pi(z_{t+1}|z_t)\lambda(z_{t+1})^{1-\gamma}} \text{ and } \widehat{\beta}(z_t) = \beta \sum_{z_{t+1}} \pi(z_{t+1}|z_t)\lambda(z_{t+1})^{1-\gamma} \tag{49}$$

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

$$\Pi_{s^t}\left[\left\{\widehat{d}\right\}\right] = \widehat{d}_t(s^t) + \widehat{\beta}(z_t) \sum_{s^{t+1}} \Pi_{s^{t+1}}\left[\{d\}\right] \left(\frac{h_{t+1}(z^{t+1})}{h_t(z^t)}\right)^{\gamma} \widehat{\pi}(s^{t+1}|s_t) \tag{50}$$

and let $V_{s^t}\left[\left\{\widehat{d}\right\}\right]$ denote the ex-dividend price-dividend ratio (.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.

- **Proof of Lemma 6.1**:

Let $\widehat{\Pi}^*_z$ denote the perfect insurance functional in the stationary economy. Pricing is equivalent to taking expectations under the risk neutral measure $\widehat{\pi}$. This is the sufficient condition for perfect risk sharing in the stationary economy:

$$\widehat{\Pi}^*_z\left[\{1\}\right] \geq \widehat{\Pi}^*_{(y,z)}\left[\{\widehat{\eta}\}\right] \text{ for all } (y,z) \tag{51}$$

where the aggregate endowment is normalized to unity. This follows immediately from Prop. (3), by exploiting the Markov structure of the shocks. It need not be checked for all s^t because the complete insurance prices only depend on the current state of the economy. It is sufficient to check it for all (y,z). Obviously, if this condition is satisfied, then perfect risk sharing is feasible: each household can consume a constant fraction of the aggregate endowment.

Second, note that the price of a claim to aggregate consumption in the stationary economy equals the price/dividend ratio in the growth economy:

$$\widehat{\Pi}^*_z\left[\{1\}\right] = \Pi^*_{z^t}\left[\{e\}\right]/e_t(z^t)$$

I can restate this condition in terms of price/dividend ratios for the growth economy:

$$PD^*_z\left[\{e\}\right] \geq \widehat{\eta}\left(y,z\right) PD^*_{(y,z)}\left[\{\eta\}\right] \text{ for all } (y,z) \tag{52}$$

where $PD^*_{(y,z)}\left[\{\eta\}\right]$ is the shadow price/dividend ratio on a claim to labor income. To see why, note that by definition the p/d ratio is the ratio of the claim in the stationary economy to the labor income share:

$$PD^*_{(y,z)}\left[\{\eta\}\right] = \widehat{\Pi}^*_{(y,z)}\left[\{\widehat{\eta}\}\right]/\widehat{\eta}\left(y,z\right)$$

- **Proof of Lemma 6.2**:

I proceed by checking the sufficient condition in Lemma 5:

$$1+\widehat{\beta}(z)\sum_{s'}\widehat{\Pi}^*\left[\{1\}\right](z')\widehat{\pi}(z'|z) \geq \widehat{\eta}(y,z)+\widehat{\beta}(z)\sum_{s'}\widehat{\Pi}^*_{y',z'}\left[\{\widehat{\eta}\}\right]\widehat{\eta}(y',z')\widehat{\pi}(y',z'|y,z) \tag{53}$$

If the idiosyncratic shocks are iid: $\widehat{\pi}(y',z'|y,z) = \widehat{\pi}_{z'}(y')$, then the r.h.s. of (53) equals:

$$\begin{aligned} &= \widehat{\eta}(y,z) + \widehat{\beta}(z)\sum_{z'}\widehat{\pi}(z'|z)\left(\sum_{y'}\widehat{\Pi}^*_{y',z'}\left[\{\widehat{\eta}\}\right]\widehat{\pi}_{z'}(y')\right) \\ &= \widehat{\eta}(y,z) + (1-\alpha)\widehat{\beta}(z)\sum_{z'}\widehat{\Pi}^*_{z'}\left[\{1\}\right](z')\widehat{\pi}(z'|z) \end{aligned} \tag{54}$$

because $\sum_{y'}\widehat{\pi}_{z'}(y')\widehat{\eta}(y',z') = (1-\alpha)e(z')$ and the pricing functional is linear. So, the last part is the ex dividend price of a claim to the labor share of the aggregate endowment:

$$\begin{aligned} \sum_{y'}\widehat{\Pi}^*_{y',z'}\left[\{\widehat{\eta}\}\right]\widehat{\pi}_{z'}(y') &= (1-\alpha)e(z') + (1-\alpha)V^*_{z'}\left[\{1\}\right] \\ &= (1-\alpha)\widehat{\Pi}^*_{z'}\left[\{1\}\right] \end{aligned}$$

Eq. (54) becomes:

$$\widehat{\eta}(y,z) + (1-\alpha)\widehat{\beta}(z)\sum_{s'}\widehat{\Pi}^*_{z'}\left[\{1\}\right]\widehat{\pi}(z'|z)$$

This implies perfect risk sharing is possible if

$$\alpha\widehat{\beta}(z)\sum_{z'}\widehat{\Pi}^*_{z'}\left[\{1\}\right]\widehat{\pi}(z'|z) \geq \left(\widehat{\eta}(y,z)-1\right) \text{ for all } (y,z)$$

which can be restated in terms of the ex dividend value:

$$\alpha\widehat{V}^*\left[e\right](z) \geq \left(\widehat{\eta}(y,z)-1\right) \text{ for all } (y,z)$$

Lemma 9.1 *: If the idiosyncratic shocks are highly persistent, perfect risk sharing is not feasible*

Proof of Lemma 9.1:

Let $\widehat{\pi}_{z',z}(y'|y) = \widehat{\pi}(y, z'|y, z)/\widehat{\pi}(z'|z)$ and let $\widetilde{\Pi}_{z',z} = \delta I + (1-\delta)\widehat{\Pi}_{z,z}$. Return to equation (53). As $\delta \to 1$, the second term converges pointwise (for each z) to $\widehat{\eta}(y, z)\left[1 + \widehat{V}_z\left[\{1\}\right]\right]$. This implies a contradiction for $\widehat{\eta}(y, z) > 1$:

$$\left[1 + \widehat{V}_z\left[\{1\}\right]\right] \geq \widehat{\eta}(y, z)\left[1 + \widehat{V}_z\left[\{1\}\right]\right]$$

Remark 1 *As* $min_z\widehat{\beta}(z_t) \to 1$*, perfect risk sharing is feasible for any* $\alpha > 0$.

Proof of Remark: Assume $\widehat{\pi}$ has a single ergodic class and let $\widehat{\pi}^*$ denote the ergodic distribution. Then $\widehat{\pi}(z'|z)$ converges exponentially to $\widehat{\pi}^*(\cdot)$

$$1 = \sum_z \widehat{\pi}^*(z) \geq \sum_z \widehat{\pi}^*(z) \sum_y \widehat{\pi}^*_z(y)\widehat{\eta}(y, z) = (1 - \alpha)$$

where the last equality follows from the definition of the aggregate endowment.

- **Proof of Proposition 6.1**:

If aggregate uncertainty is i.i.d., then the discount rate in the transformed stationary economy is constant

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

(see Alvarez and Jermann, 2001). The (cum dividend) price/dividend ratio of the labor endowment stream can be written as:

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

Only the $\{h_t(z^t)\}$ process depends on z^t. The natural question to ask is whether there exists a stationary equilibrium in which $\frac{h_{t+1}(z^{t+1})}{h_t(z^t)} = g^*$. Suppose $\frac{h_{t+1}(z^{t+1})}{h_t(z^t)} = g^*$. First I assume that there exists a stationary distribution Φ^* over endowments and weights and analyze the iterations on the 2 operators. Second, I show there exists such a stationary distribution. The complete proof is provided in the Comp. Appendix on p. 74.

- Derivation of liquidity shock:

Assume w.l.o.g. that the aggregate moment of the end-of-period weights has been normalized to one. Then the average weight growth rate is given by

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

where $\Phi_{z^{t-1}}$ is the joint distribution over weights and endowments. $\underline{l}(y_2, z_{re}, z^{t-1}) > \underline{l}(y_2, z_{ex}, z^{t-1})$ implies $g_t(z_{re}, z^{t-1}) > g_t(z_{ex}, z^{t-1})$.

- **2-agent Economy**

In the 2-agent economy the SDF still has the same generic with h_t

$$h_t = \left(\zeta_t^1\right)^{1/\gamma} + \left(\zeta_t^2\right)^{1/\gamma}.$$

In each state s there will be a range $[\underline{x}(s), \overline{x}(s)]$ for the relative weights

$$x = \frac{(\zeta^1)^{1/\gamma}}{(\zeta^1)^{1/\gamma} + (\zeta^2)^{1/\gamma}}$$

If agent one faces a new binding constraint, x hits the upper bound, if agent two faces a new binding constraint, x hits the lower bound, else x remains unchanged. The SDF can be restated as:

$$m_{t+1} = \beta \left(\frac{e_{t+1}}{e_t}\right)^{-\gamma} \left(\min\left(\frac{x_{t+1}}{x_t}, \frac{1 - x_{t+1}}{1 - x_t}\right)\right)^{-\gamma}$$

Let g_t^γ equal the second term. The interpretation is similar to the continuum case but not identical. If none of the households is constrained or the same household is equally constrained, $h_{t+1} = h_t$, else $h_{t+1} \geq h_t$. Only one of both can be constrained in any given state. The computational method is discussed in Lustig (1999). This setup does not allow for history dependence. Only two things can happen. (i) no one is constrained or the same household is equally constrained, and x is unchanged: the complete insurance discounter is applied to payoffs in this state. (ii) A new constraint does bind and x hits the upper or lower bound. Since all of the variables are Markov in the consumption share x, the liquidity sequence is probabilistically "restarted" whenever x hits the lower/upper bound and this two-agent economy consists of a sequence of i.i.d. economies.

10 Tables and Figures

Table 6: **Composition of National Income**

US 1946-1999	./National Inc.	1929-1999
labor income	0.70	.70
corporate profits.	0.11	.105
proprietary income	0.11	.12
rental income	0.05	.03
interest income	0.025	.0105

Table 7: **Alvarez and Jermann Calibration**

	e_{t+1}			$\widehat{\eta}_t(y_t, z_t)$	
M1	$\rho(\lambda)$	-.14	M5	$\sigma(\ln \widehat{\eta}(s))$	.296
M2	Pr(ex)/Pr(re)	2.65	M6	$\rho(\ln \eta(s))$	.75
M3	$E(\lambda)$	1.83	M7	v_r/v_e	1
M4	$\sigma(\lambda)$	3.57	M8	$\frac{\sigma_{r'e}}{\sigma_{e'e}}$	1.88
			M9	$\frac{\sigma_{r'r}}{\sigma_{e'r}}$	1.88
			M10	$\frac{\sigma_{r'}}{\sigma_{e'}}$	.90

Table 8: **Approximation errors in percentages**

	k=4				k=5		
$\beta = .95$	$\widehat{E}\,(x)$	$std\,(x)$	$\sup\lvert x\rvert$	$\beta = .95$	$\widehat{E}\,(x)$	$std\,(x)$	$\sup\lvert x\rvert$
γ=2	.0034	.0036	.11	γ=2	-.0009	.0025	.021
γ=3	-.0051	.0062	.08	γ=3	-.0002	.0025	.023
γ =4,	-.0045	.0051	.20	γ =4,	-.0006	.0028	.020
γ =5	-.0022	.0032	.023	γ =5	-.0008	.0026	.002
γ =6	-.0025	.0032	.25	γ =6	-.0006	.0019	.024
γ =7	-.0061	.0065	.16	γ =7	-.0000	.0024	.026

Table 9: **10 percent collateral with k=4**

$\beta = .95$	r^f	ER^e	$\sigma m/Em$	$\exp(p/d)$	$\sigma(r^f)$	$\sigma(R^e)$	$\sigma(R^s)$
US (MP)	.80	6.1	.37	25	5	16	17
γ=2	5	.08	.10	23	3.2	6	6.7
rep.agent	7	.00	.11	12	1.8	.04	10
γ=3	4.8	1.1	.17	23	5	7.7	8.9
rep.agent	9.5	1.0	.11	12	2.7	6.1	11
γ=4	4.5	1.7	.21	22	6.1	10	10
rep. agent	9.8	.9	.15	11.	3.7	5.94	10
γ=5	4.4	2.4	.27	20	8.2	12	13
rep. agent	10.5	1.3	.20	9.8	4.7	6.5	10
γ=6	2.4	3.9	.33	26	9.7	14	15
rep. agent	11	2.3	.24	8.8	5.7	7.33	11
γ=7	.8	6.1	.41	30	13	18	19
rep. agent	11.6	2.3	.28	8.3	6.8	7.8	11
γ=8	.8	5.5	.45	28	13	15	20
rep. agent	12	2.6	.32	7.8	7.9	8.8	13

Table 10: **10 percent collateral with k=5**

$\beta = .95$	r^f	ER^e	$\sigma m/Em$	$\exp(p/d)$	$\sigma(r^f)$	$\sigma(R^e)$	$\sigma(R^s)$
US (MP)	.80	6.1	.37	25	5	16	17
γ =2	5.9	.8	.10	23.	5	7.7	6.7
γ =3	4.8	1.1	.16	23	5	7.8	9
γ =4	4.05	2.0	.21	21	6.7	9	11
γ =5	41	2.4	.27	22	8.8	10	14
γ =6	2.8	3.6	.33	25	10	13	16
γ =7	.002	6.5	.40	30	12	14	18
γ =8	.01	5.3	.44	27	14	16	20

Table 11: **15 percent collateral with k=4**

$\beta = .95$	r^f	$Er^e - r^f$	$\sigma m/Em$	$\exp(p/d)$	$\sigma(r^f)$	$\sigma(R^e)$
γ=2	6.7	.00	.09	18	2.42	5.2
γ=3	6.2	1.3	.14	18	4	7
γ=4	7.3	1.6	.20	18.5	6	8
γ=5	5.4	2.1	.25	17	7	10
γ=6	5.2	3.0	.31	17	9	12
γ=7	4	5.1	.36	18	11	13
γ=8	2.8	5.3	.43	19	13	14
γ=9	3.5	6.1	.46	15	14	14

Table 12: **20 percent collateral with k=4**

$\beta = .95$	r^f	$Er^e - r^f$	$\sigma m/Em$	$\exp(p/d)$	$\sigma(r^f)$
US (MP)	.80	6.1	.37	25	5
γ =4	7.25	1.4	.18	14	5
γ =5	7.5	1.0	.23	13	6
γ =6	6.5	3.1	.29	14	2
γ =7	9.1	2.9	.309	10.0	2

Table 13: **5 percent collateral with k=4**

$\beta = .95$	r^f	ER^e	$\sigma m/Em$	$\exp(p/d)$	$\sigma(r^f)$
US (MP)	.80	6.1	.37	25	5
γ =4	2.88	1.0	.24	38	7
γ =5	1.5	3.4	.30	37	9
γ =6	2.2	4.0	.34	27	10
γ =7	-1.1	6	.42	65	13

Table 14: **Two-agent economy: 10 percent collateral**

$\beta = .95$	r^f	ER^e	$\sigma m/Em$	$\exp(p/d)$	$\sigma(r^f)$
US (MP)	.80	6.1	.37	25	5.3
γ =4	2.44	2.36	1.18	41	8.3
γ =5	2.8	2.7	1.24	31	8.3
γ =6	3.0	3.51	1.69	26	10
γ =7	3.9	4.22	.80	17	9.3

Table 15: **Liquidity Shocks**

$z_t = re$	$z_{t-3} = re$	$z_{t-4} = ex$	$z_t = re$	$z_{t-3} = re$	$z_{t-4} = ex$
$z_{t-1} = re$	$g(z^k)$		$z_{t-1} = ex$	$g(z^k)$	
$z_{t-3} = re$	1.0029	1.0081	$z_{t-3} = re$	1.0305	1.0389
$z_{t-4} = ex$	1.0035	1.0094	$z_{t-4} = ex$	1.0353	1.05

Table 16: **Cond. Market Price of Risk**

$z_t = re$	$z_{t-2} = re$	$z_{t-1} = ex$	$z_t = ex$	$z_{t-2} = re$	$z_{t-1} = ex$
$\frac{\sigma_t(m)}{E_t(m)}$			$\frac{\sigma_t(m)}{E_t(m)}$		
$z_{t-2} = re$	.24	.29	$z_{t-2} = re$	38	.40
$z_{t-1} = ex$	.25	.28	$z_{t-1} = ex$	.39	.45

Table 17: **k-horizon Stochastic Discounter**

$\beta = .95$	$Em_{t,t+l}$	$\sigma m_{t,t+l}$	$Em_{t,t+l}$	$\sigma m_{t,t+l}$
$\gamma = 7$			rep.	agent
l=1	.96	.46	.85	.31
l=2	.93	.48	.75	.32
l=3	.91	.53	.66	.32
l=4	.89	.56	.58	.31
l=5	.86	.59	.52	.30
l=6	.84	.62	.46	.29
l=7	.82	.65	.40	.28

Table 18: **Price/Dividend Ratio Regressions**

$\beta = .95$	R^2	$\widehat{\beta}$		R^2	$\widehat{\beta}$
$\gamma = 7$					
$l = 1$	.03	-1.48	$l = 6$	.57	-1.04
$l = 2$	.52	$-.92$	$l = 7$	.56	-1.02
$l = 3$	.49	-1.06	$l = 8$	.59	-0.98
$l = 4$	.52	-1.03	$l = 9$	.60	-1.03
$l = 5$	.55	-1.04			

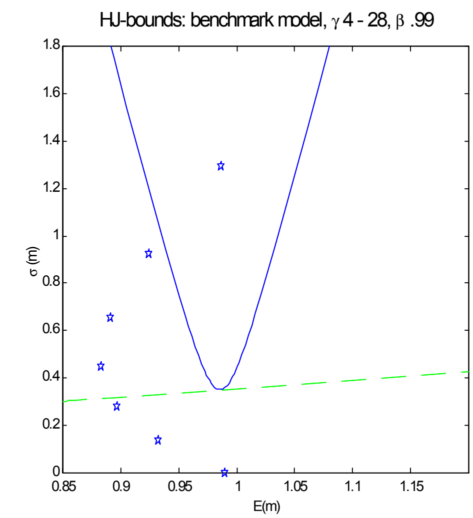

Figure 1:

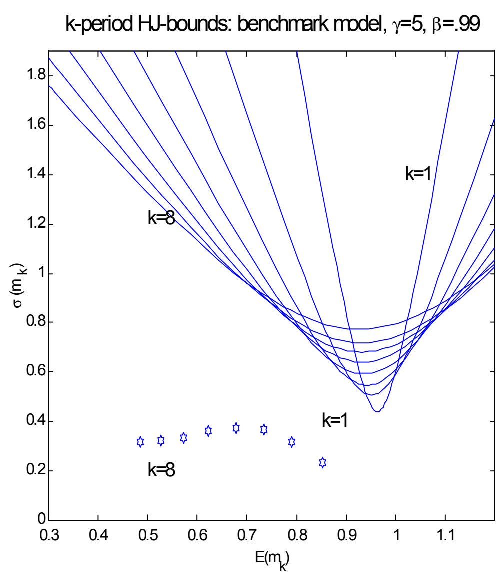

Figure 2:

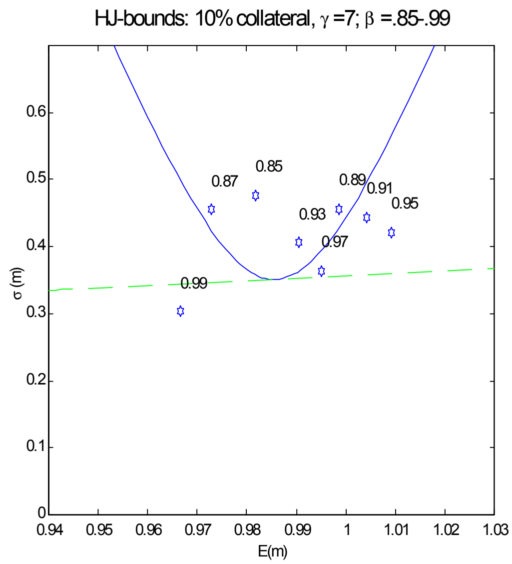

Figure 3:

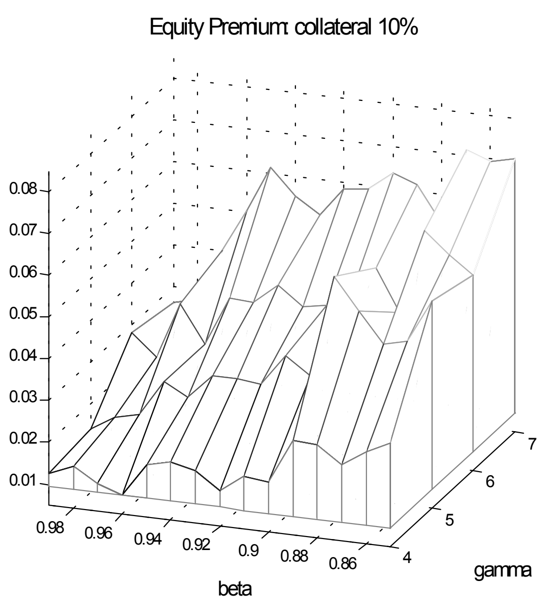

Figure 4:

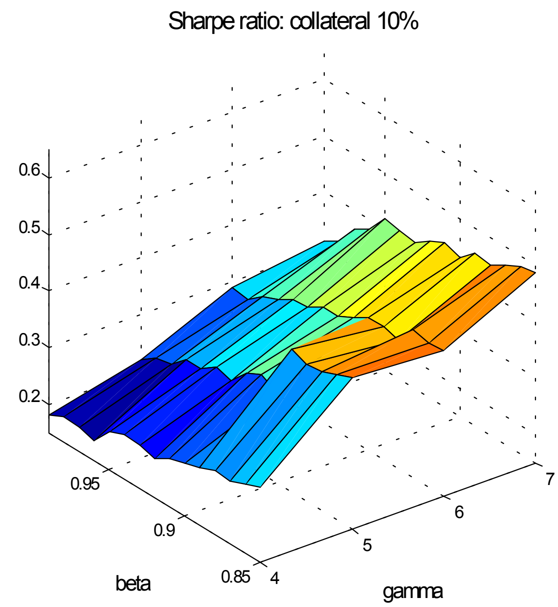

Figure 5:

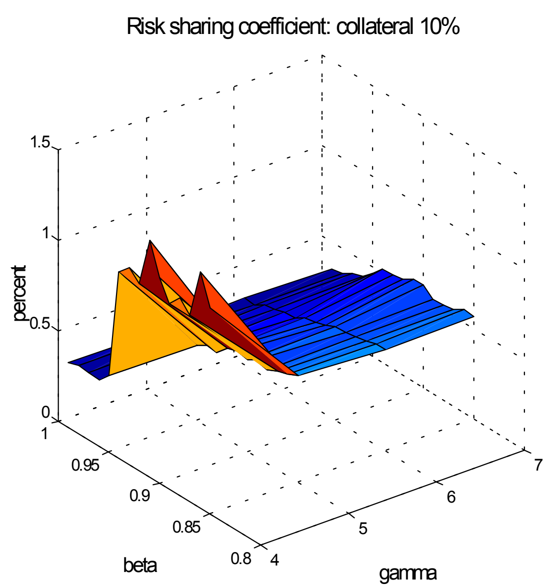

Figure 6:

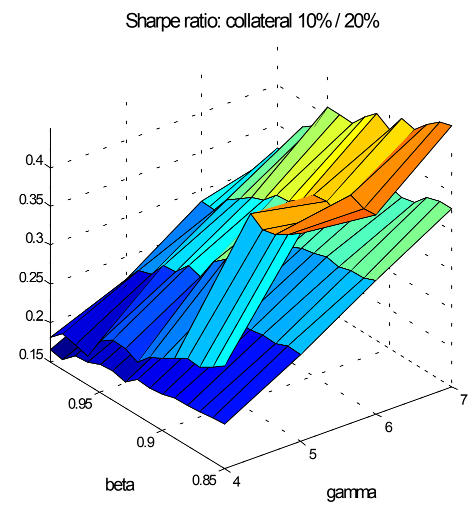

Figure 7:

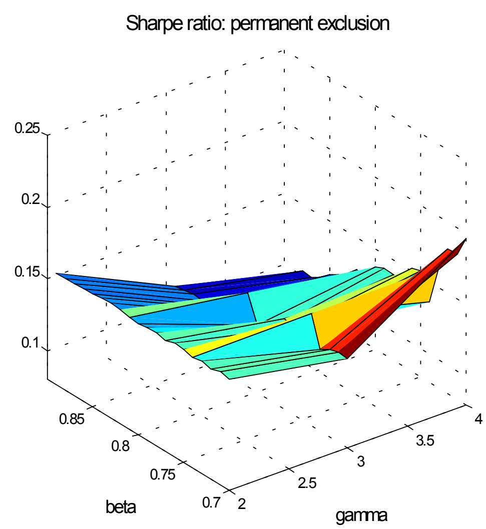

Figure 8:

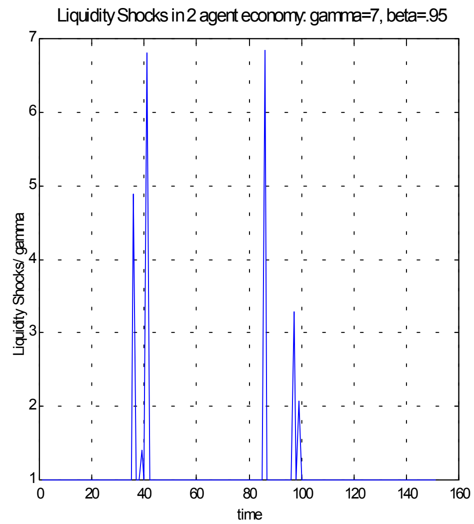

Figure 9:

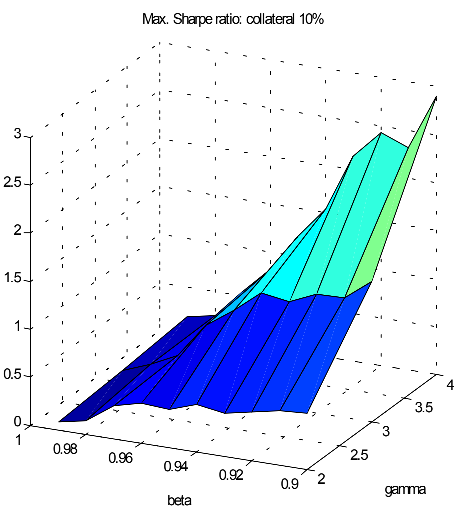

Figure 10:

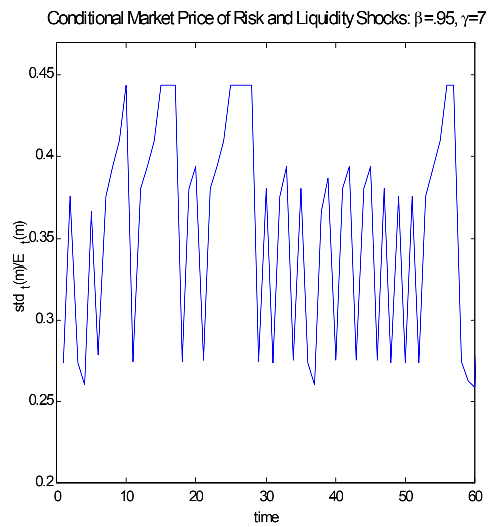

Figure 11:

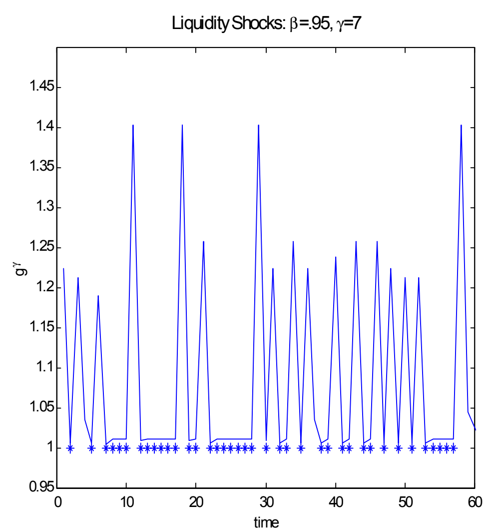

Figure 12:

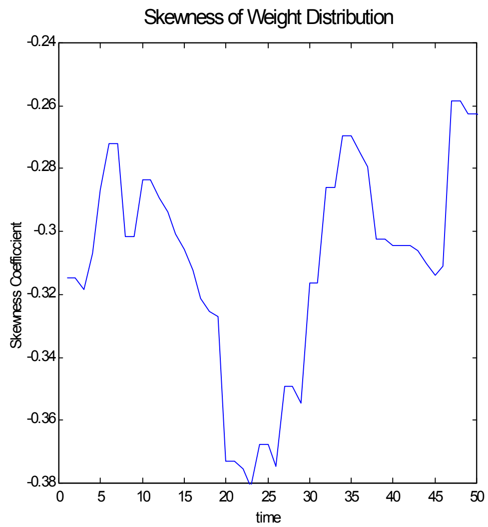

Figure 13:

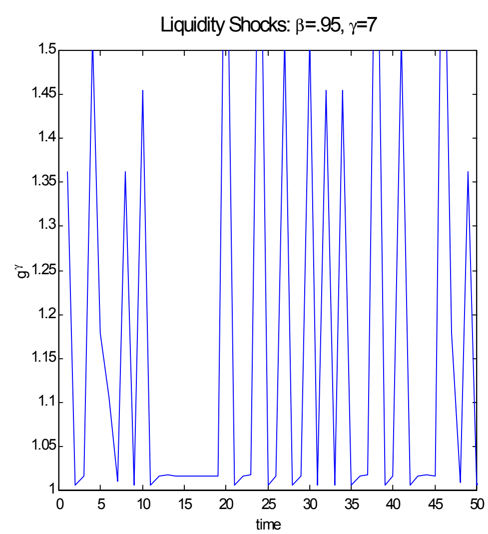

Figure 14:

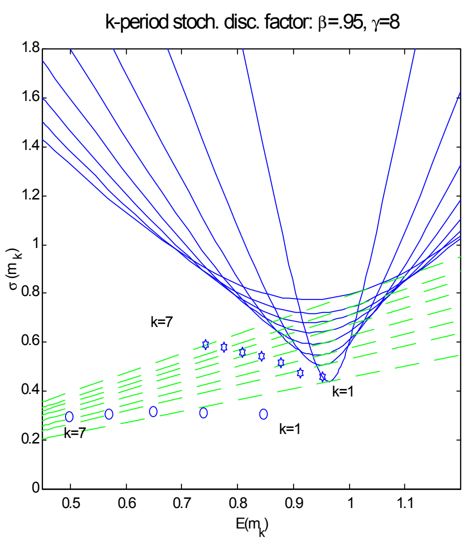
k-period stoch. disc. factor: β=.95, γ=8
σ (m_k)
E(m_k)
k=7
k=1
k=7
k=1
0
0.2
0.4
0.6
0.8
1
1.2
1.4
1.6
1.8
0.5
0.6
0.7
0.8
0.9
1
1.1

Figure 15: