

Evaluating Asset Pricing Models with Limited Commitment using Household Consumption Data¹

Dirk Krueger

University of Pennsylvania, CEPR and NBER

Hanno Lustig

UCLA and NBER

Fabrizio Perri

University of Minnesota, Federal Reserve Bank of Minneapolis, CEPR and NBER

Abstract

We evaluate the asset pricing implications of a class of models in which risk sharing is imperfect because of limited enforcement of intertemporal contracts. Lustig (2004) has shown that in such a model the asset pricing kernel can be written as a simple function of the aggregate consumption growth rate and the growth rate of consumption of the set of households that do not face binding enforcement constraints. These unconstrained households have lower consumption growth rates than all other households in the economy. We use household data on consumption growth from the U.S. Consumer Expenditure Survey to identify unconstrained households, to estimate the pricing kernel implied by these models and evaluate their performance in pricing aggregate risk. We find that for high values of the relative risk aversion coefficient, the limited enforcement pricing kernel generates a market price of risk that is substantially closer to the data than the one obtained using the standard complete markets asset pricing kernel.

Keywords: Limited Commitment, Equity Premium, Stochastic Discount Factor, Household Consumption Growth.

JEL Classification: G12, D53, D52, E44

¹We thank seminar participants at the 2006 EEA Meetings in Vienna and the 2007 AEA Meetings in Chicago for helpful comments and the NSF for financial support.

1 Introduction

Consumption-based asset pricing kernels derived under the complete risk sharing, representative agent (RA) assumption have a hard time explaining the observed large equity premium (see e.g. Mehra and Prescott, 1985). Models in which the sharing of idiosyncratic risk is limited have the potential to solve the puzzle (see for example Constantinides and Duffie, 1996). In these models, the asset pricing kernel, in general, does not only depend on aggregate consumption but also on the entire distribution of consumption across agents. Different models provide different links between the distribution of consumption and asset pricing kernels. An important task is to evaluate whether these models are useful in solving the equity premium puzzle. Recently some studies have done work along this line, either evaluating several types of incomplete risk sharing models (see for example Attanasio, Banks and Tanner, 2002, Brav, Constantinides and Geczy, 2002, Vissing-Jorgensen, 2002 and Kocherlakota and Pistaferri, 2006) or exploring the empirical link between asset prices and higher moments of the consumption growth distribution (Cogley, 2002).

This paper contributes to this research agenda. It evaluates the asset pricing implications of a class of models in which risk sharing is imperfect because of the limited enforcement (henceforth LE) of intertemporal contracts, as in Thomas and Worrall (1988) or Kehoe and Levine (1993). No restrictions are imposed on the menu of traded assets. Alvarez and Jermann (2001) have explored the asset pricing implications of LE in a two agent economy, but they have not evaluated its empirical implications for the cross-sectional distribution of consumption and asset prices. Lustig (2004) has shown that in a version of this model with a continuum of agents the asset pricing kernel can be written as a simple function of the growth rate of consumption of the set of households that do not face binding enforcement constraints in the current state of the world. These unconstrained households have lower consumption growth rates than households that face binding enforcement constraints. This implication of the model allows us to identify unconstrained households as those in the

lower tail of the cross-sectional consumption growth distribution.

We construct the LE pricing kernel using data on household consumption expenditures from the U.S. Consumer Expenditure Survey (CE) and evaluate its performance in pricing aggregate risk. We show that the LE pricing kernel can be estimated using a simple modification of the standard RA pricing kernel. As documented in previous studies, the RA pricing kernel only explains a small part of the equity premium. The power of the LE pricing kernel depends crucially on how we identify unconstrained households but, in general, it explains a larger fraction of the equity premium than the RA pricing kernel.

2 The Model

We consider a pure exchange economy with a continuum of agents that face aggregate and idiosyncratic endowment shocks, trade state-contingent claims to consumption on competitive markets and face solvency constraints that limit the extent to which agents can go short in these consumption claims. In this section we first describe the underlying physical environment and the market structure, then we define a competitive equilibrium and finally we provide a characterization of the asset pricing kernel implied by this model.

2.1 Physical Environment

We denote the current aggregate shock by $z_t \in Z$ and the current idiosyncratic shock by $y_t \in Y$, with Z and Y finite. Let $z^t = (z_0, \dots, z_t)$ and $y^t = (y_0, \dots, y_t)$ denote the history of aggregate and idiosyncratic shocks. We use the notation $s_t = (y_t, z_t)$ and $s^t = (y^t, z^t)$ and let the economy start at initial node z_0 . Conditional on idiosyncratic shock y_0 and thus $s_0 = (y_0, z_0)$, the probability of a history s^t is given by $\pi_t(s^t|s_0)$. Individual endowments are given by $e_t(s^t)$.

At time 0 households are indexed by their idiosyncratic income shock y_0 and their initial asset position a_0 . We denote by $\Theta_0(y_0, a_0)$ the initial distribution of agents over (y_0, a_0) ; this

initial distribution, together with the initial aggregate shock z_0 serves as initial condition for our economy.

Consumers rank stochastic consumption streams $\{c_t(a_0, s^t)\}$ according to

$$U(c)(s_0) = \sum_{t=0}^{\infty} \sum_{s^t \geq s^0} \beta^t \pi(s^t | s_0) \frac{c_t(a_0, s^t)^{1-\gamma}}{1-\gamma} \quad (1)$$

where $\gamma > 0$ is the coefficient of relative risk aversion and $\beta \in (0, 1)$ is the constant time discount factor.

2.2 Market Structure

Households can trade a complete set of contingent consumption claims $\{a_t(a_0, s^t, s_{t+1})\}$ at prices $q_t(s^t, s_{t+1})$. Thus their budget constraints read as

$$c_t(a_0, s^t) + \sum_{s_{t+1}} q_t(s^t, s_{t+1}) a_t(a_0, s^t, s_{t+1}) = e_t(s^t) + a_{t-1}(a_0, s^t) \quad (2)$$

These trades are subject to solvency constraints $\{J(a_0, s^t, s_{t+1})\}$ such that

$$-a_t(a_0, s^t, s_{t+1}) \leq J(a_0, s^t, s_{t+1}) \quad (3)$$

The solvency constraints, precisely spelled out below, are not too tight, in the sense of Alvarez and Jermann (2000): a household that has borrowed exactly up to the constraint (that is $-a_t(a_0, s^t, s_{t+1}) = J(a_0, s^t, s_{t+1})$) is indifferent between defaulting on her debt (and suffering the corresponding consequences spelled out below) and repaying (and thus avoiding these consequences). In the standard complete markets model $J(a_0, s^t, s_{t+1}) = \infty$, since in that model households can fully commit to repay any debt they take on.²

Denote by $V(a, s^t)$ the maximized continuation expected lifetime utility an agent can

²Of course, even in this model an appropriately chosen no-Ponzi constraint is needed to make the household decision problem well-defined.

attain, if she comes into the current period with assets $a_{t-1}(a_0, s^t) = a$ and faces constraints (2) and (3). Furthermore let $V^{Aut}(s^t)$ denote the expected lifetime utility of an agent from consuming the autarkic allocation $c_t(a_0, s^t) = e_t(s^t)$ from node s^t on.³ Finally let c_t^a denote aggregate consumption (equal to the aggregate endowment). The market clearing condition reads as

$$\sum_{s^t} \int c_t(a_0, s^t) \pi(s^t | s_0) d\Theta_0 = \sum_{s^t} \int e_t(s^t) \pi(s^t | s_0) d\Theta_0 \equiv c_t^a(z^t) \text{ for all } z^t. \quad (4)$$

2.3 Equilibrium

We are now ready to define an equilibrium for this economy.

Definition 1 *Given z_0 and an initial distribution $\Theta_0(y_0, a_0)$, an equilibrium with solvency constraints $\{J(a_0, s^t, s_{t+1})\}$ that are not too tight are consumption and asset allocations $\{c_t(a_0, s^t), a_t(a_0, s^t, s_{t+1})\}$ and prices $\{q_t(s^t, s_{t+1})\}$ such that*

1. *Given prices $\{q_t(s^t, s_{t+1})\}$ and constraints $\{J(a_0, s^t, s_{t+1})\}$, for all (y_0, a_0) allocation $\{c_t(a_0, s^t), a_t(a_0, s^t, s_{t+1})\}$ maximizes (1) subject to (2) and (3).*
2. *The solvency constraints are not too tight, that is, satisfy, for all (y_0, a_0) and all s^{t+1} ,*

$$V(J(a_0, s^t, s_{t+1}), s^{t+1}) = V^{Aut}(s^{t+1}).$$

3. *Market clearing: Equation (4) holds.*

2.4 Characterization of Equilibria

Let (y_0, a_0) denote the characteristics of a generic household. In order to characterize equilibrium consumption allocations and the pricing kernel we make use of cumulative

³The specification of the outside option as autarky is done for simplicity. Any other specification of the outside option that is only a function of (a, s^t) gives rise to the same characterization of the asset pricing kernel derived below.

Lagrange multipliers $\{\xi_t(y_0, a_0)\}$, in the spirit of Marcet and Marimon (1998). In period 0 there is a one to one map between Lagrange multipliers ξ_0 and initial wealth a_0 . Thus let the initial distribution of Lagrange multipliers associated with the distribution of initial wealth $\Theta_0(y_0, a_0)$ be denoted by $\Phi_0(y_0, \xi_0)$. Henceforth we will use the notation $\xi_t(y_0, a_0)$ and $\xi_t(y_0, \xi_0)$ interchangeably. Over time these Lagrange multipliers increase whenever the solvency constraint of a household binds, and remains unchanged otherwise. Crucially, this implies that for all agents that are *unconstrained* in a current state, their Lagrange multipliers *all* remain unchanged.

As shown by Lustig (2004) the consumption process of a given household is related to aggregate consumption (endowment) by the risk sharing rule

$$c_t(\xi_0, s^t) = [\xi_t(\xi_0, s^t)]^{1/\gamma} \frac{c_t^a(z^t)}{h_t(z^t)} \quad (5)$$

where

$$h_t(z^t) = \int [\xi_t(\xi_0, s^t)]^{1/\gamma} d\Phi_t$$

and Φ_t is the cross-sectional measure over consumption weights (Lagrange multipliers) $\xi_t(\xi_0, s^t)$ in period t , state z^t .

To rule out arbitrage opportunities, payoffs in state z^{t+1} are priced off the intertemporal marginal rate of substitution (IMRS) of those agents who do not face any binding constraints in transferring resources to and from that state (see Alvarez and Jermann, 2000). Let $UC(s^t, s_{t+1})$ denote the set of these agents. The stochastic discount factor is the IMRS of those agents with labels $\xi_0^* \in UC(s^t, s_{t+1})$, who are unconstrained in state s^t in their sale of securities that deliver goods in state s^{t+1} :

$$m_{t+1}^{LE}(z^{t+1}) = \beta \left(\frac{c_{t+1}(\xi_0^*, s^{t+1})}{c_t(\xi_0^*, s^t)} \right)^{-\gamma}$$

The risk sharing rule in (5) and the fact that for all unconstrained agents the consumption

weights do not change $\xi_{t+1}(\xi_0^*, s^{t+1}) = \xi_t(\xi_0^*, s^t)$ then immediately imply that the pricing kernel is given by:

$$m_{t+1}^{LE}(z^{t+1}) = \beta \left(\frac{c_{t+1}^a(z^{t+1})}{c_t^a(z^t)} \right)^{-\gamma} [g(z^{t+1})]^\gamma \quad (6)$$

where $g(z^{t+1}) = \frac{h_{t+1}(z^{t+1})}{h_t(z^t)}$. Note that in the standard complete markets model the solvency constraints are never binding, thus the distribution of consumption weights (Lagrange multipliers) is never changing, and consequently $h_{t+1}(z^{t+1}) = h_t(z^t)$ and $g(z^{t+1}) = 1$ for all z^{t+1} . Therefore, we recover the well-known stochastic discount factor of the RA model

$$m_{t+1}^{RA}(z^{t+1}) = \beta \left(\frac{c_{t+1}^a(z^{t+1})}{c_t^a(z^t)} \right)^{-\gamma} \quad (7)$$

The only effect of LE on asset prices is the component contributed by the shocks to the cross-sectional distribution of consumption weights $h_{t+1}(z^{t+1})/h_t(z^t)$.

2.5 Implementation

In order to generate an empirical time series for the LE stochastic discount factor in (6) from cross-sectional consumption data we need to estimate the aggregate consumption growth rate (readily available from NIPA) and the growth rate of the consumption weight distribution:

$$g(z^{t+1}) = \frac{h_{t+1}(z^{t+1})}{h_t(z^t)}.$$

From the risk sharing rule in (5) we know that this moment of the consumption weight distribution satisfies:

$$h_t(z^t) = [\xi_t(\xi_0, s^t)]^{1/\gamma} \frac{c_t^a(z^t)}{c_t(\xi_0, s^t)}$$

For all unconstrained households the consumption weight does not change in state s^{t+1} , i.e. $\xi_{t+1}(\xi_0, s^{t+1}) = \xi_t(\xi_0, s^t)$, and hence their consumption growth rate satisfies:

$$g(z^{t+1}) = \frac{h_{t+1}(z^{t+1})}{h_t(z^t)} = \frac{c_{t+1}^a(z^{t+1})}{c_t^a(z^t)} * \frac{c_t(\xi_0, s^t)}{c_{t+1}(\xi_0, s^{t+1})}$$

All unconstrained households have the same growth rate of consumption

$$\begin{aligned} \frac{c_{t+1}(\xi_0, s^{t+1})}{c_t(\xi_0, s^t)} &= \frac{c_{t+1}^a(z^{t+1})/h_{t+1}(z^{t+1})}{c_t^a(z^t)/h_t(z^t)} := g^{UC}(z^{t+1}) \text{ or} \\ g(z^{t+1}) &= \left(\frac{c_{t+1}(\xi_0, s^{t+1})/c_{t+1}^a(z^{t+1})}{c_t(\xi_0, s^t)/c_t^a(z^t)} \right)^{-1} = \frac{g^a(z^{t+1})}{g^{UC}(z^{t+1})}, \end{aligned} \quad (8)$$

where $g^a(z^{t+1})$ is the growth rate of aggregate consumption and $g^{UC}(z^{t+1})$ is the common consumption growth rate of currently unconstrained households. Expression (8) suggests that the term $g(z^{t+1})$ can be measured using the consumption growth of unconstrained households $g^{UC}(z^{t+1})$ relative to the growth rate of aggregate consumption. The key question we deal with in the next section is how to identify unconstrained households in the data.

3 Testing the Empirical Asset Pricing Implications

3.1 Data

The crucial difference between the RA and the LE pricing kernel is that the former can be estimated using aggregate consumption data while for the latter we need to measure the consumption growth rate of a specific group of households (the unconstrained). Therefore household level consumption growth data is needed. The U.S. Consumer Expenditure Survey provides such data since the majority of households sampled in this data set reports consumption expenditures for at least two subsequent quarters.

We use quarterly data from 1980.1 to 2006.1. For each quarter t we select all house-

holds which are complete income respondents and which report positive expenditures on nondurable goods and services for quarters t and $t + 1$. For each household we then compute quarterly growth rates of real (each component is deflated with specific CPI's), per-adult equivalent expenditures on nondurables (this includes nondurable goods and services, excluding services from housing). We have a total of 307778 consumption growth rate observations.⁴

The return data comes from the CRSP (the Center for Research on Securities Prices). As stock returns we use the quarterly value-weighted return on the entire US market (NYSE/AMEX/NASDAQ), deflated by the inflation rate computed from the Consumer Price Index by the Bureau of Labor Statistics. Bond returns are based on the average yield of a 3-month US T-bill, again deflated by the CPI.

3.2 Empirical Construction of the Asset Pricing Kernels

The model predicts that in each period all unconstrained agents share the same consumption growth rate, and that this growth rate is the lowest growth rate across all households. So, interpreting the model literally, one could obtain $\hat{g}(z^{t+1})$, the empirical estimate of $g(z^{t+1})$ over the period 1980.1-2005.4, as simply the ratio between the average consumption growth rate and the minimum (across households) consumption growth rate in a given quarter. Obviously in the data there is significant measurement error in consumption. Also, consumption growth in the data is likely to depend on idiosyncratic events (for example changes in personal health, or educational expenses) which are not explicitly considered in our model. For these reasons this procedure would not be a very robust way of selecting the relative consumption growth of unconstrained households. Instead we identify unconstrained households as all households having consumption growth lower than a certain

⁴Due to a change in the household sample there are no observations in the last quarter of 1985. For more details on the deflation method and on the categories included in nondurable consumption expenditures see appendix A in Krueger and Perri (2006)

threshold $\bar{g}(z^{t+1})$ and then we obtain the following estimate for $g(z^{t+1})$

$$\hat{g}(z^{t+1}) = \frac{\hat{g}^a(z^{t+1})}{\hat{g}^{UC}(z^{t+1})} = \frac{\frac{\sum_i c_i(z^{t+1})}{\sum_i c_i(z^t)}}{\frac{\sum_{i \in U} c_i(z^{t+1})}{\sum_{i \in U} c_i(z^t)}}$$

where $U(z^{t+1}) = \{i : \frac{c_i(z^{t+1})}{c_i(z^t)} \leq \bar{g}(z^{t+1})\}$

An important issue is how to set the threshold $\bar{g}(z^{t+1})$. We start with a natural initial guess which is $\bar{g}(z^{t+1}) = \hat{g}^a(z^{t+1})$ (so that unconstrained households in a given quarter are all households who report consumption growth lower than aggregate growth in that quarter); we also experiment with other possibilities. Once $\hat{g}(z^{t+1})$ has been constructed using CE household data and deseasonalized, it is easy to construct a time series for the LE pricing kernel according to (6). For a given risk aversion γ , time discount factor β and aggregate growth rate of consumption $\frac{c_{t+1}^a(z^{t+1})}{c_t^a(z^t)}$ the LE pricing kernel is given by

$$\hat{m}_{t+1}^{LE}(\gamma, \beta) = \beta \left(\frac{c_{t+1}^a(z^{t+1})}{c_t^a(z^t)} \right)^{-\gamma} (\hat{g}_{t+1})^\gamma \quad (9)$$

where the aggregate growth rate of consumption is constructed using deseasonalized NIPA consumption data.⁵

Below we discuss the asset pricing properties of $\hat{m}_{t+1}^{LE}(\gamma, \beta)$ together with those of the stochastic discount factor of the standard representative agent model, $\hat{m}_{t+1}^{RA}(\gamma, \beta)$, which can be estimated using only NIPA data and is given by

$$\hat{m}_{t+1}^{RA}(\gamma, \beta) = \beta \left(\frac{c_{t+1}^a(z^{t+1})}{c_t^a(z^t)} \right)^{-\gamma}.$$

⁵The definition of NIPA consumption used to compute the aggregate consumption growth rate includes real expenditures on nondurable goods and services (excluding housing). It therefore includes the same categories of consumption expenditures used to estimate the relative growth rate of consumption of unconstrained households.

3.3 Results

In this section we evaluate the performance of the two stochastic discount factors discussed above in explaining the equity premium for different values of the risk aversion parameter γ and under different assumptions for the identification of unconstrained agents. For each specification (including the representative agent stochastic discount factor) we set the time discount factor β so that the sample mean of the estimated stochastic discount factor $E(\hat{m})$ is equal to 1. With this normalization⁶ we can write $fe(\hat{m})$, the fraction of the equity premium that is being explained by the stochastic discount factor \hat{m} , as

$$fe(\hat{m}) = 1 - \frac{E[\hat{m}(R^S - R^B)]}{E[R^S - R^B]} = -Corr(\hat{m}, R^S - R^B)cv(\hat{m})cv(R^S - R^B)$$

where R_{t+1}^S and R_{t+1}^B denote the gross real return on equity and a risk-free bond, and $Corr$ denotes the correlation coefficient, and cv denotes the coefficient of variation. Thus to explain a large fraction of the equity premium we need a stochastic discount factor that is negatively correlated with the equity premium and which has high percentage volatility. The first panel of figure 1 plots the realized excess return on equity $R^S - R^B$ and the representative agent stochastic discount factor $\hat{m}_{t+1}^{RA}(\gamma, \beta)$ (in percentage deviations from its mean) for a value of $\gamma = 1$.

The stochastic discount factor is indeed negatively correlated with the realized equity premium but its percentage volatility is very low and hence the resulting fraction of the equity premium it can explain is low. The second panel plots again the realized excess equity return together with the limited enforcement stochastic discount factor $\hat{m}_{t+1}^{LE}(\gamma, \beta)$ for a value of $\gamma = 1$. Note how $\hat{m}_{t+1}^{LE}(\gamma, \beta)$ is also negatively correlated (albeit less than $\hat{m}_{t+1}^{RA}(\gamma, \beta)$) with the realized equity premium but at the same time it is significantly more volatile than $\hat{m}_{t+1}^{RA}(\gamma, \beta)$. With low risk aversion the higher (relative to $\hat{m}_{t+1}^{RA}(\gamma, \beta)$)

⁶This normalization essentially guarantees that all stochastic discount factors we consider are consistent with the empirically observed risk-free interest rate.




Figure 1: Stochastic discount factors ($\gamma = 1$) and realized equity premium

volatility of $\hat{m}_{t+1}^{LE}(\gamma, \beta)$ is offset by the lower absolute value of its correlation with the equity premium so roughly the 2 discount factors have the same (low) explanatory power. With high risk aversion though the differences in volatilities between $\hat{m}_{t+1}^{LE}(\gamma, \beta)$ and $\hat{m}_{t+1}^{RA}(\gamma, \beta)$ are amplified and the limited enforcement discount factor can explain a significantly larger fraction of the observed equity premium.

Figure 2 plots the equity premium explained by LE stochastic discount factors estimated in three different ways and by the standard RA stochastic discount factor, as we vary the coefficient of risk aversion from 1 to 80. The differences among the three LE curves lie in how we identify unconstrained agents. The curve labeled baseline is constructed defining unconstrained households in a given quarter as those with consumption growth is below average consumption growth in that quarter. The other two LE curves (labeled 95% and 105%) are constructed defining unconstrained households as those whose consumption growth is below 95% of aggregate consumption growth rate or below 105% of the aggregate consumption growth rate. The figure also displays the average equity premium in the data, which for our sample is around 1.8% per quarter. Notice how for low and moderate levels




Figure 2: Quarterly equity premium: observed and predicted by models

of risk aversion all discount factors fail to account for a substantial fraction of the premium. However, for higher risk aversion the LE model can explain a significantly larger fraction of the observed equity premium than the RA model: for example with a relative risk aversion of around 60 the LE model can explain all the premium while the RA model only explains half of it.

4 Conclusion

The consumption-based asset pricing model relying on the representative agent construct has a hard time accounting for the observed equity premium. In this paper we assess how introducing limited enforcement of inter-temporal contracts improves the empirical performance of the consumption-based asset pricing model. In a LE model with a continuum of households the standard stochastic discount factor is augmented by an additional term

which depends on the consumption growth of households in the left tail of the cross-sectional consumption growth distribution. The reason for this is that agents with low consumption growth are unconstrained and thus it is their consumption growth, as opposed to the aggregate consumption growth, which should price equity. We put forth a first attempt of measuring this extra term using household consumption data and find that the LE pricing kernels can account for a larger share of the empirically observed equity premium, relative to the RA model. However, in order to account for the entire premium the LE model still needs a coefficient of risk aversion of around 60, which many would view as implausibly large.

We interpret our results as mixed news for the asset pricing properties of the LE model. On one hand they suggest that when one uses micro consumption data to evaluate the asset pricing properties of the LE model results are not as satisfactory as, say, those obtained by Alvarez and Jermann (2001) using a calibrated version of the model with two agents. On the other hand the fact that the LE model performs better, along certain dimensions, than the standard RA model suggests that making use of individual consumption data, as opposed to just aggregates, might give us a better understanding of how aggregate risk is priced.

Future work is needed to assess how a careful modelling of measurement error in individual consumption growth would affect the empirical estimation and performance of the proposed asset pricing kernel, and to investigate whether the LE model can shed further light on other well-documented asset pricing puzzles (such as the value premium puzzle).

References

- [1] Alvarez, Fernando and Urban Jermann (2000), “Efficiency, Equilibrium, and Asset Pricing with Risk of Default,” *Econometrica*, 68, 775-798
- [2] Alvarez, Fernando and Urban Jermann (2001), “Quantitative Asset Pricing Implications of Endogenous Solvency Constraints,” *Review of Financial Studies*, 14, 1117-1151
- [3] Attanasio, Orazio, James Banks and Sarah Tanner (2002), Assets Holding and Consumption Volatility, *Journal of Political Economy*, 110, 771-792.
- [4] Brav, Alon, George Constantinides and Christopher Geczy (2002), “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy*, 110, 793-824.
- [5] Cogley, Tim (2002), “Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey,” *Journal of Monetary Economics*, 49, 309-334.
- [6] Constantinides, George and Darrell Duffie, “Asset pricing with heterogeneous consumers” (1996), *Journal of Political Economy*, 104, 219-240.
- [7] Kehoe, Tim and David Levine (1993), “Debt Constrained Asset Markets,” *Review of Economic Studies*, 60, 865–888.
- [8] Kocherlakota, Narayana and Luigi Pistaferri (2005), “Asset Pricing Implications of Pareto Optimality with Private Information,” *Working paper*, University of Minnesota.
- [9] Krueger, Dirk and Fabrizio Perri (2006), “Does Income Inequality lead to Consumption Inequality? Evidence and Theory” *Review of Economic Studies*, 73, 163-193.
- [10] Lustig, Hanno (2004), “The Market Price of Aggregate Risk and the Wealth Distribution,” *Working paper*, UCLA.
- [11] Marcet, Albert and Ramon Marimon (1998), “Recursive Contracts,” *Working paper*, Universitat Pompeu Fabra.
- [12] Mehra, Rajnish and Edward Prescott (1985), “The equity premium: A puzzle” *Journal of Monetary Economics* 15, 145-61
- [13] Thomas, Jonathan and Tim Worrall (1988), “Self-Enforcing Wage Contracts,” *Review of Economic Studies*, 55, 541-554.
- [14] Vissing-Jorgensen, Annette (2002), “Limited Asset Market Participation and the Elasticity of Intertemporal Substitution”, *Journal of Political Economy*, 110, 825–53.