Evaluating Asset Pricing Models with Limited Commitment using Household Consumption Data

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 the 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 in that state of the world. These unconstrained households have lower consumption growth rates than constrained households, i.e. they are located in the lower tail of the cross-sectional consumption growth distribution. We use household consumption data from the U.S. Consumer Expenditure Survey to estimate the pricing kernel implied by the model and to evaluate its performance in pricing aggregate risk. We employ the same data to construct aggregate consumption and to derive the standard complete markets pricing kernel. We find that the limited enforcement pricing kernel generates a market price of risk that is substantially larger than the standard complete markets asset pricing kernel.

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

JEL Classification: G12, D53, D52, E44

1 Introduction

Consumption-based asset pricing kernels derived under the complete risk sharing, representative agent (RE) assumption cannot explain the large equity premium, at least not with plausible preference specifications (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 it also depends 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 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 those 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. To be consistent, we use the same CE data to construct aggregate consumption and to compute the standard RE pricing kernel. As documented in previous studies, the RE 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 RE 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 $z^t = (z_0, \ldots, z_t)$ and $y^t = (y_0, \ldots, y_t)$ 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

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

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

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

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

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

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.

Denote by $V(a, s^t)$ the maximized continuation expected lifetime utility an agent can 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.[1] Finally let c_t^a denote aggregateconsumption (equal to the aggregate endowment). The market clearing condition reads as

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

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 to 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 Lagrange multipliers $\{\xi_t(y_0, a_0)\}$, in the spirit of Marcet and Marimon

[1] The 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.

(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

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

where

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

and Φ_t is the cross-sectional measure over consumption weights $\xi_t(\xi_0, s^t)$ 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:

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

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 RE model

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

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

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 and the growth rate of the consumption weight distribution:

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

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

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

For all unconstrained households the consumption weight does not change in state s^{t+1}, $\xi_{t+1}(\xi_0, s^{t+1}) = \xi_t(\xi_0, s^t)$, 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

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

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. Combining (8) and (6), we obtain the following simple representation of the stochastic discount factor:

$$m_{t+1}^{LE}(z^{t+1}) = \beta \left(g^{UC}(z^{t+1}) \right)^{-\gamma} \tag{9}$$

Furthermore it is clear from (5) that currently constrained households have consumption growth rates strictly higher than $g^{UC}(z^{t+1})$ since for those households $\xi_{t+1}(\xi_0, s^{t+1}) > \xi_t(\xi_0, s^t)$. Consequently unconstrained households can be empirically identified as those households in the lower tail of the cross-sectional consumption growth distribution

3 Testing the Empirical Asset Pricing Implications of the Model

3.1 Data

The crucial difference between the RE and the LE pricing kernel is that the former can be estimated using aggregate consumption data while for the latter data on household level consumption growth 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 2004.1. For each quarter t we select all households which are complete income respondents and which report positive expenditures on non durable 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. We have a total of 284675 consumption growth rate observations.[2]

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 main ingredient of the asset pricing kernel of the LE economy is a time series for the growth rate of consumption expenditures of unconstrained households. The model suggests that in each period all unconstrained agents share the same consumption growth rate, and that this growth rate is lower than consumption growth for all constrained households. In the data there is significant measurement error in consumption. Moreover consumption growth in the data is also likely to depend on idiosyncratic events (for example changes in personal

[2]Due 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)

health, or educational expenses) which are not explicitly considered in our model. For these reasons simply selecting the lowest observed consumption growth rate would not be a very robust nor sensible way of selecting the consumption growth of unconstrained households. Instead we select a level of consumption growth which exhibits bunching (i.e. a large fraction of agents have consumption growth close to that level) and which is in the left tail of the cross sectional consumption growth rate distribution. Figure 1 shows the distribution of growth rate of consumption expenditures in the first quarter of 2003 and suggests that empirically plausible levels of consumption growth for the unconstrained lie between the 30^{th} (denoted as P30 in the figure) and the 50^{th} percentile (median) of the consumption growth distribution. In the following section we take the 40^{th} percentile of the consumption growth rate distribution as our benchmark estimate of the the consumption growth rate of the unconstrained $\widehat{g}_{t+1}^{UC}$, but we also experiment with setting it equal to the 30^{th} or the 50^{th} percentile.

FIGURE 1 HERE

Once the growth rate of the unconstrained consumers in each quarter has been estimated, it is easy to construct a time series for the LE pricing kernel according to (9), for a given risk aversion γ and time discount factor β.

$$\hat{m}_{t+1}^{LE}(\gamma,\beta) = \beta\left(\widehat{g}_{t+1}^{UC}\right)^{-\gamma} \tag{10}$$

Below we also report results for the standard RE model computed with our data for which the relevant stochastic discount factor is

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

Here $\hat{g}_{t+1}^{a}$ is simply the growth rate of aggregate consumption in our CE sample between quarters t and $t+1$.

3.3 Results

In this section we evaluate the performance of our pricing kernel 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 RE stochastic discount factor) we set the time discount factor β so that the sample mean of the estimated stochastic discount factor $E_T(\hat{m})$ is equal to 1. With this normalization[3] we can write $fe(\hat{m})$, the

[3] This normalization essentially guarantees that all stochastic discount factors we consider rationalize the empirically observed risk-free interest rate.

fraction of the equity premium that is being explained by the stochastic discount factor $\hat{m}$, as

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

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 a correlation, and cv denotes a coefficient of variation. Thus to explain a large equity premium we need a stochastic discount factor (and thus a consumption growth rate) that is negatively (positively) correlated with the equity premium and very volatile. The main component of the stochastic discount factors under consideration is consumption growth. Therefore in figure 2 we plot the realized excess return on equity $R^S - R^B$ and aggregate consumption growth $\hat{g}^a$, together with the consumption growth rate of unconstrained households $\widehat{g}^{UC}$, identified as the the growth rate of the household at the 40^{th} percentile of the consumption growth distribution.

FIGURE 2 HERE

Note that both aggregate consumption growth (and thus the RE stochastic discount factor) and the consumption growth rate of the unconstrained agents are much less volatile than the equity premium. The key difference between the two growth rates series lies in their correlation with the excess return on equity. This correlation is slightly negative (-0.02) for $\hat{g}^a$ while it is slightly positive for $\hat{g}^{UC}$ (0.06). This difference in the correlation is at the heart of the difference between the equity premium explained by the two models that we report in figure 3.[4] The figure plots the equity premium produced by the RE and LE models; the growth rate of the unconstrained agents in the LE model is estimated as the 30th, 40th and 50th percentile of the consumption growth distribution. We also plot the average equity premium in the data, 1.76%, on a quarterly basis. For moderate levels of risk aversion. the equity premia implied by the LE models, while larger than the one implied by the RE, still fall significantly short of the 1.76 % target. For higher risk aversion some versions of the LE model do significantly better than the RE model. This suggests that this LE model may provide us with a better understanding of how aggregate risk is priced.

FIGURE 3 HERE

[4] As a robustness check we also measured aggregate consumption growth from NIPA data. In our sample aggregate consumption growth from NIPA is slightly more correlated with the equity premium (the correlation is 0.13) but its volatility is only about one third of the volatility of the growth rate computed from CE data. As a consequence, as it is well known, the RE pricing kernel explains only a small fraction of the observed equity premium.

4 Conclusion

The standard representative agent model can only account for a small fraction of the equity premium. In this paper we show that introducing limited enforcement of intertemporal contracts improves the empirical performance of the consumption-based asset pricing model. In a standard LE model with a continuum of households the stochastic discount factor is a function of consumption growth of households in the left tail of the cross-sectional consumption growth distribution. We find that the LE pricing kernel can account for a significantly larger share of the empirically observed equity premium (but not for the entire premium). Future work is needed to assess how 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 it 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] Brav, Alon, George Constantinides and Christopher Geczy (2002), "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, 110, 793-824.

[4] Cogley, Tim (2002), "Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey," *Journal of Monetary Economics*, 49, 309-334.

[5] Constantinides, George and Darrell Duffie, "Asset pricing with heterogeneous consumers", *Journal of Political Economy*, 104 (1996), 219-240.

[6] Kehoe, Tim and David Levine (1993), "Debt Constrained Asset Markets," *Review of Economic Studies,* 60, 865–888.

[7] Kocherlakota, Narayana and Luigi Pistaferri (2005), "Asset Pricing Implications of Pareto Optimality with Private Information," *Working paper*, University of Minnesota.

[8] Krueger, Dirk and Fabrizio Perri (2006), "Does Income Inequality lead to Consumption Inequality? Evidence and Theory" *Review of Economic*

Studies, 73, 163-193.

[9] Lustig, Hanno (2004), "The Market Price of Aggregate Risk and the Wealth Distribution," *Working paper*, UCLA.

[10] Marcet, Albert and Ramon Marimon (1998), "Recursive Contracts," *Working paper*, Universitat Pompeu Fabra.

[11] Mehra, Rajnish and Edward Prescott (1985), "The equity premium: A puzzle" *Journal of Monetary Economics* 15, 145-61

[12] Thomas, Jonathan and Tim Worrall (1988), "Self-Enforcing Wage Contracts," *Review of Economic Studies*, 55, 541-554.

[13] Vissing-Jorgensen, Annette (2002), "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution", *Journal of Political Economy*, 110, 825–53.

Figure 1. Histogram of the consumption growth distribution in the CE, 2003Q1

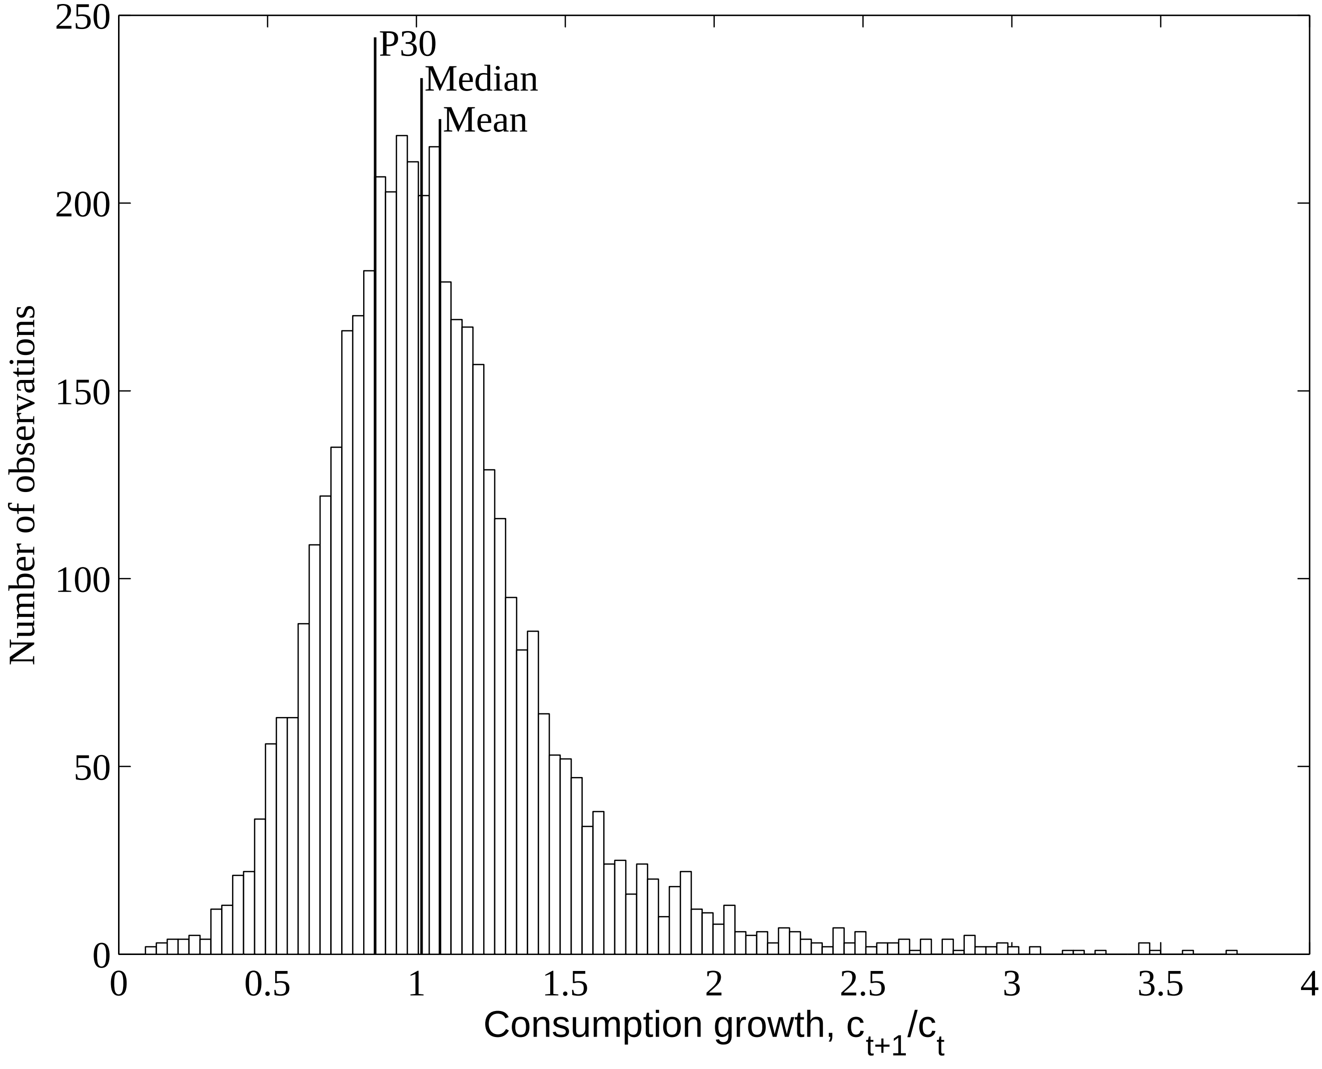

Figure 2. Consumption Growth and Equity Premium

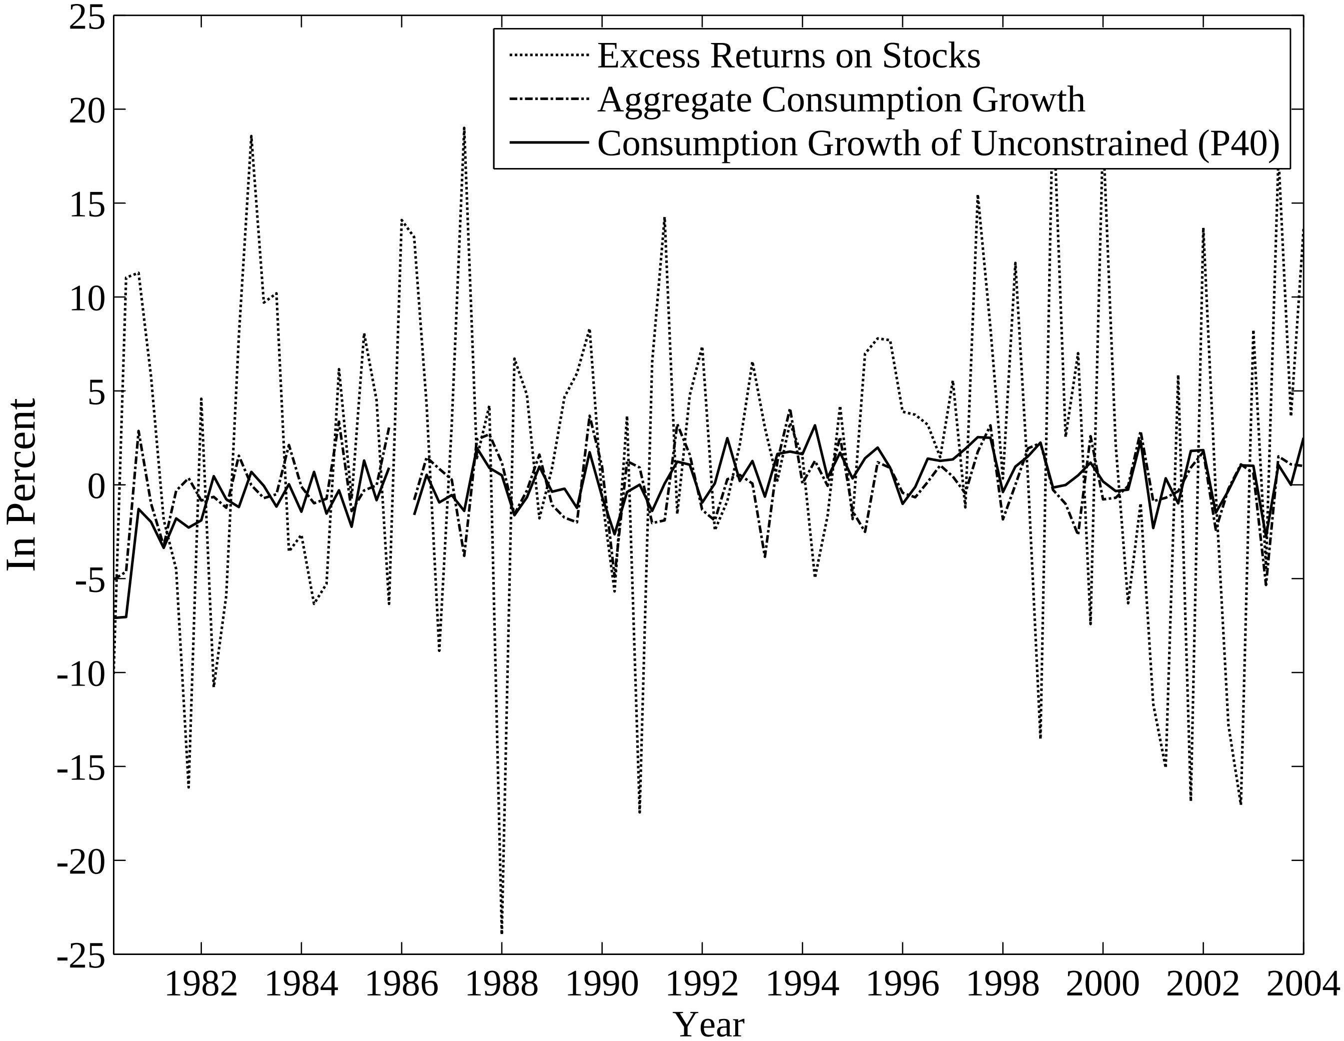

Figure 3. Explained Quarterly Equity Premium