January 2007

Asset Pricing Implications of Pareto Optimality with Private Information*

Narayana Kocherlakota

University of Minnesota, FRB-Minneapolis, and NBER

Luigi Pistaferri

Stanford University and CEPR

ABSTRACT

In this paper, we consider a dynamic economy in which the agents are privately informed about their skills, which evolve stochastically over time in an arbitrary fashion. We consider an asset pricing equilibrium in which equilibrium quantities are constrained Pareto optimal. Under the assumption that agents have constant relative risk aversion, we derive a novel asset pricing kernel for financial asset returns. The kernel equals the reciprocal of the gross growth of the γth moment of the consumption distribution, where γ is the coefficient of relative risk aversion. We use data from the Consumer Expenditure Survey (CEX) to estimate the new stochastic discount factor. We show that its predictions for Treasury bills and the equity premium are arguably better, and certainly no worse, than existing asset pricing models.

*Kocherlakota acknowledges the support of NSF 0305833 and NSF 0606695. The paper has benefited from comments from participants in seminars at London School of Economics, Arizona State University, Stanford University, UCLA, the University of Iowa, Carnegie-Mellon University, the Federal Reserve Bank of Cleveland, the Federal Reserve Bank of Minneapolis, the National Bureau of Economic Research, The University of Chicago, UC-Davis, USC, and Wharton. We thank Sonam Sherpa for research assistance, and Tim Cogley, Peter Hansen, Dirk Krueger, Felix Kubler, Adriano Rampini, and four anonymous referees for their comments. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

1. Introduction

A large theoretical literature has studied the dynamics of allocations in Pareto optima when effort is private information (see Phelan and Townsend (1991) and Atkeson and Lucas (1995), among others). However, there has been little empirical work using these models. One problem is that, because of technical considerations, these models can be solved, even numerically, only when skill shocks follow highly unrealistic laws of motion.

In this paper, we examine the implications of these models for asset pricing. More specifically, we treat individual skills and effort choices as being private information. This informational assumption means that in a Pareto optimum, individual consumption depends on individual-specific shocks, and the representative agent asset pricing model is no longer valid. Instead, we derive a new asset pricing kernel which depends on the moments of the cross-sectional distribution of consumption.[1] An especially attractive feature of our analysis is that this kernel is valid regardless of the data generation process for the underlying skill shocks.

More specifically, our theoretical results and empirical analysis follow directly from three distinct assumptions. First, we assume that the allocation of consumption across agents is Pareto optimal, given that agents are privately informed about their skills. As stated above, we impose *no* restrictions on the time series behavior of the stochastic process governing skills or on the process governing aggregate shocks. Second, we assume that agents have identical preferences that are additively separable over time and between consumption and leisure; as well, agents have power utility functions $\left((1-\gamma)^{-1} c^{1-\gamma}\right)$ over consumption. Third, we assume that the planner's *shadow* stochastic discount factor in the Pareto optimal allocation is in fact a valid *market* stochastic discount factor for asset returns.

Under these assumptions, we find that the following is a valid stochastic discount factor for financial asset returns:

$$\beta C_{\gamma,t}/C_{\gamma,t+1}$$

where $C_{\gamma,t}$ is the γth moment of the cross-sectional distribution of consumption. Here, β is the common discount factor across agents, and γ is their common coefficient of relative risk aversion. (We term this stochastic discount factor the Private Information Pareto Optimal (PIPO) stochastic discount factor.) The key to the

[1] Throughout, when we use the term "moment", we refer to uncentered moments.

construction of this discount factor is the application of a law of large numbers. We assume that the fraction of agents who have a particular history of shocks in the data is the same as the unconditional probability of that history, and thereby convert conclusions about expectations of marginal utility into conclusions about moments of the cross-sectional distribution of consumption.

We go on to show that the estimate of the PIPO stochastic discount factor is robust to measurement error in consumption data. The measurement error must be independent of the true data and be stationary over time, but can be arbitrarily persistent.

We then assess the empirical performance of this new stochastic discount factor. To provide suitable context, we contrast its performance with two alternative stochastic discount factors, derived from two different market structures. The first is an implication of equilibrium in a standard incomplete markets framework without binding borrowing constraints. The discount factor takes the form:

$$\beta C_{-\gamma,t+1}/C_{-\gamma,t}$$

Here, $C_{-\gamma,t}$ is the negative γth moment of the cross-sectional distribution of consumption. We derive this stochastic discount factor by integrating over the intertemporal Euler equations of the consumers in the economy. The second alternative discount factor is an implication of equilibrium when markets are complete. Then, we can use the marginal rate of substitution of the representative agent as the stochastic discount factor:

$$\beta(C_{1,t+1})^{-\gamma}/(C_{1,t})^{-\gamma}$$

In this formula, $C_{1,t}$ is the first moment of the cross-sectional distribution of consumption (i.e., the mean). Both of the alternative stochastic discount factors are also robust to the kind of measurement error described above.

An important feature of all three discount factors is that they can be estimated without longitudinal data on household consumption. Instead, all we need is a time-series of cross-sections of household consumption from which moments of the consumption distribution can be estimated. For each (overlapping) quarter between 1980 and 2004, we construct the three stochastic discount factors using data from the Consumer Expenditure Survey (CEX).

We then apply the Generalized Method of Moments to assess the validity of the three discount factors in terms of their implications for the equity premium and the Treasury bill returns. We focus on three particular restrictions. An arbitrary stochastic discount factor m_t should be consistent with the population restrictions:

$$
\begin{aligned}
&(1) \qquad & E\{(m_t R_t^f - 1)\} &= 0 \\
&(2) \qquad & E\{(m_t R_t^f - 1) R_{t-1}^f\} &= 0 \\
&(3) \qquad & E\{m_t (R_t^{mkt} - R_t^f)\} &= 0
\end{aligned}
$$

where R_t^{mkt} is the gross real return to the stock market and R_t^f is the gross real return to Treasury bills. We chose these restrictions because they are much studied in the macroeconomics literature. The restriction (1) assesses the ability of the discount factor to account for the level of the risk-free rate. The restriction (2) assesses the response of the stochastic discount factor to a key predictor of the Treasury bill return - that is, its own lag. Hall (1988) shows that plausible parameterizations of the standard representative agent model are inconsistent with this last type of restriction. The last restriction (3) assesses the extent to which a candidate discount factor can explain the difference between the stock market and Treasury bill returns. Kocherlakota (1996) argues that (3) is simply a robust re-statement of the equity premium puzzle originally stated by Mehra and Prescott (1985).

Our empirical results are as follows. For all three discount factors, there exist plausible specifications of the preference parameters (β, γ) that zero out the sample analogs of (1) and (2). The resulting estimate of γ for the PIPO stochastic discount factor is about 2, but it is imprecise. The resulting estimates of β are fairly similar for the complete markets and incomplete markets stochastic discount factors but slightly more precise.

Next, we turn to the equity premium error. We find that if we set γ to around 5, the PIPO stochastic discount factor is able to set the sample analog of (3) to zero. The representative agent discount factor is able to do so if γ is approximately 58, but not if γ is smaller. There is no specification of γ (positive or negative) for the incomplete markets case that sets the sample analog of (3) to zero.

Finally, we examine the ability of the discount factors to account simultaneously for the equity premium and the properties of the expected return to the Treasury bill. We use the Hansen and Jagannathan (1997) matrix to weight the three restrictions. As they show, the resulting Generalized Method of Moments objective function can be interpreted as a measure of distance between a given, possibly misspecified, discount factor and the class of true discount factors. We find that the estimated PIPO discount factor is considerably closer to the class of true discount factors than the other two discount factors. However, while the estimate of γ is not overly large (around 5), the estimated value of β is implausibly low at 0.3. Nevertheless, even if we set β to a plausible value of 0.99, we find that the PIPO stochastic discount factor displays a lower Hansen-Jagannathan distance than the other stochastic discount factors throughout the parameter space of γ.

Thus, our new PIPO stochastic discount factor performs distinctly better than the representative agent stochastic discount factor. More surprisingly, it does as well as, if not better than, the incomplete market stochastic discount factor. We interpret these findings as indicating that society-wide incentive costs play a significant role in the determination of asset prices.

2. Prior Literature

There is little prior work that econometrically evaluates the implications of Pareto optimality with private information. An important exception is Ligon (1998), who tests the risk-sharing implications of Pareto optimality with moral hazard. His approach is as follows. He uses consumption data from South Indian villages (the ICRISAT data set). He assumes that there is a risk-neutral banker outside the villages, agents in the village have the same discount rate as the interest rate offered by the outside banker, and all agents have coefficient of relative risk aversion $\gamma > 0$. He asks if the allocation of risk within the village is better described as being Pareto-optimal, given moral hazard, or as the result of risk-free borrowing and lending. He answers this question by estimating the parameter b from the following moment restriction:

(4) $$E_t\{(c_{i,t+1}/c_{i,t})^b\} = 1$$

Under the former hypothesis of constrained Pareto optimality, b equals γ. Under the latter hypothesis of risk-free borrowing and lending, b equals $-\gamma$. Using the Generalized Method of Moments, he estimates b to be positive and interprets this as demonstrating the relative empirical relevance of constrained Pareto optimality.

Our approach bears some similarity to Ligon's. But there are important differences. First, our theoretical analysis is more general than his. We allow for aggregate shocks and do not assume that there is a risk-neutral outsider. Hence, we are able to allow for non-trivial movements in expected asset returns. As well, we do not need to assume that individual productivity shocks are i.i.d. over time (as he does). This assumption of i.i.d. productivity shocks is at odds with the data (Meghir and Pistaferri, 2004). Second, our testable implications are in terms of the cross-sectional consumption distribution, not individual consumptions; we do not need to have panel data on consumption. Finally, our empirical analysis is more robust to measurement error than is his.[2]

Our work is also related to recent papers using data from the CEX to evaluate incomplete markets models of asset pricing. In recent work, Cogley (2002), Brav, Constantinides, and Geczy (BCG) (2002), and Vissing-Jorgensen (2002) use data from the CEX to test the hypothesis that asset prices and household consumption are consistent with an incomplete markets equilibrium. These papers basically proceed as follows. They select all households from the CEX who have two or more periods of observations. They next construct an intertemporal marginal rate of substitution (IMRS) in a given period for each household with observations for that period and the prior one (for BCG, a period is a quarter and for Vissing-Jorgensen, a period is a half-year). Finally, they construct a theoretically valid stochastic discount factor by averaging these IMRS's across households. (Henceforth, we term this the *average IMRS* stochastic discount factor and we use the acronym SDF to refer to stochastic discount factors more generally.)

Note that the average IMRS SDF is not the same as the incomplete markets SDF described in the introduction. The average IMRS SDF used in the prior literature is the average of the ratios of marginal utilities. Our incomplete markets SDF is instead the ratio of averages of marginal utilities. In an incomplete markets economy, with no binding borrowing constraints, both SDFs are valid but they are not necessarily the same.

The findings of this recent work are somewhat mixed. Cogley (2002) argues that the average IMRS SDF does not provide much additional explanatory power over the representative agent SDF in terms of the

[2] Ligon (1996) discusses some ideas for dealing with measurement error, persistence, and relaxing risk-neutrality in Section 6. This working paper is the basis for Ligon (2005).

equity premium. In contrast, BCG (2002) find that the average IMRS SDF does a good job of rationalizing the equity premium. These differences could be explained by differences in the sample period used, sample selection, and the nature of the approximation adopted. Vissing-Jorgensen (2002) considers different samples of households depending on the size of their position in the asset market. She finds that the (log-linearized) average IMRS SDF is a valid SDF for *smaller* values of γ as the average is constructed using samples of agents with *larger* asset positions.[3]

Balduzzi and Yao (2004) also use data from the CEX. Like us, they examine the performance of the incomplete markets SDF, not the average IMRS SDF. They find that if the coefficient of relative risk aversion is around 8 or higher, a log-linearized version of the SDF can account for the equity premium if they restrict attention to households that have $2000 or more in total financial assets.[4]

Our work is novel relative to these other papers because we consider the asset pricing implications of Pareto optimality with private information, as well as the implications of the more traditional incomplete markets formulation. Moreover, our empirical work differs from these papers in two other important respects. First, measurement error in consumption generates a bias in the average IMRS SDF. The bias does not affect the pricing of return differentials (like the equity premium), but it does affect the pricing of returns themselves. Hence, the authors of these other papers are forced to focus only on return differentials. In contrast, all of our SDFs are robust to a wide class of possible measurement error processes. This allows us to explore the ability of the candidate models to account for the Treasury bill return.[5] Second, other than BCG, these other papers rely on Taylor series approximations of the relevant stochastic discount factors. The errors in these approximations may lead to biases in the results. As opposed to dealing with potential outliers in an ad hoc fashion (by discarding data or by using approximations to the theory), we instead deal with them by placing no restriction on the marginal distribution of the measurement errors.

---

[3]As we discuss later, Vissing-Jorgensen (2002) does not look at restrictions involving multiple assets (like the equity premium restriction (1)). Instead, she examines the asset pricing equations for bonds and stocks separately.

[4]Balduzzi and Yao (2004) provide no specific definition of what they mean by "financial assets". When we examine the impact of financial market participation in what follows, we use the definitions of Attanasio, Banks and Tanner (2002) and Vissing-Jorgensen (2002).

[5]This comment about measurement error does not apply to Balduzzi and Yao (2004), because they use the incomplete markets SDF. However, they do not look at the ability of the incomplete markets SDF to account for the level of the Treasury bill return.

3. Environment

In this section, we describe the environment. The description is basically the same as that in Kocherlakota (2005).

The economy lasts for T periods, where T may be infinity, and has a unit measure of agents. We allow for the possibility that the agents can be distinguished from one another by society using an observable but economically irrelevant characteristic. More specifically, suppose each agent is labelled by $s \in S = \{1, 2, ..., N\}$; the measure of agents with label s is equal to π_s. The idea of these labels is to allow for the possibility that in a Pareto optimal allocation, the planner may weight some agents differently from others.

There is a single consumption good that can be produced by labor. The agents have identical preferences. A given agent has von-Neumann-Morgenstern preferences, and ranks deterministic sequences according to the function:

$$(5) \qquad \sum_{t=1}^{T} \beta^{t-1}\{u(c_t) - v(l_t)\}, 1 > \beta > 0$$

where $c_t \in R_+$ is the agent's consumption in period t, and $l_t \in R_+$ is the agent's labor in period t. We assume that $u', -u'', v'$, and v'' all exist and are positive. We also assume that u and v are bounded from above and below. The preferences are restricted to be additively over date and state contingent consumption and labor. This additive separability restriction is essential in the derivation of the new SDF.

There are two kinds of shocks in the economy: public aggregate shocks and private idiosyncratic shocks. The first kind of shocks works as follows. Let Z be a finite set, and let μ_Z be a probability measure over the power set of Z that assigns positive probability to all non-empty subsets of Z. At the beginning of period 1, an element z^T of Z^T is drawn according to μ_Z. The random vector z^T is the sequence of public aggregate shocks; z_t is the realization of the shock in period t.

The idiosyncratic shocks work as follows. Let Θ be a Borel set in R_+, and let μ_Θ be a probability measure over the Borel subsets of Θ^T. At the beginning of period 1, an element of θ^T is drawn for each agent according to the measure μ_Θ. Conditional on z^T, the draws are independent across agents and μ_Θ is the same for all realizations of z^T. We assume that a law of large numbers applies across agents: conditional on any z^T, the measure of agents in the population with type θ^T in Borel set B is given by $\mu_\Theta(B)$. This independence of private idiosyncratic shocks and public aggregate shocks plays a crucial role in what follows.

Any given agent learns the realization of z_t and his own θ_t at the beginning of period t and not before. Thus, at the beginning of period t, the agent knows his own private history $\theta^t = (\theta_1, ..., \theta_t)$ and the history of public shocks $z^t = (z_1, ..., z_t)$. This implies that his choices in period t can only be a function of this history.[6]

The individual-specific and aggregate shocks jointly determine skills. In period t, an agent produces output y_t according to the function:

$$y_t = \phi_t(\theta^T, z^T) l_t \tag{6}$$

$$\phi_t : \Theta^T \times Z^T \to (0, \infty) \tag{7}$$

$$\phi_t \text{ is } (\theta^t, z^t)\text{-measurable} \tag{8}$$

We assume that an agent's output is observable at time t, but his labor input l_t is known only to him. We refer to ϕ_t as an agent's skill in history (θ^t, z^t). Here, we think of l_t as being effort or time actually spent working. Individuals may be required to be in an office or at a job eight hours a day - but it is hard to tell how much of that time they actually spend being productive.

An important element of our analysis is the flexible specification of the stochastic process generating skills. This flexibility takes two forms. First, we are agnostic about the time-series properties of the skill shocks. This generality is crucial, given the current empirical debate about the degree of persistence of individual wages. In particular, we are able to allow for the possibility that individual skills may be at once persistent and stochastic. Both aspects seem to be important empirically.

Second, it has been argued by Storesletten, Telmer, and Yaron (2001) that the cross-sectional variance of wages is higher in recessions than in booms. Thus, the cross-sectional variance of skills varies with aggregate conditions. We can capture this possibility in our setting, because $Var(\phi_t(\theta^t, z^t)|z^t)$ may depend on z_t. The idea here is that the range of ϕ, as a function of θ^t, can be allowed to depend on z_t.[7]

---

[6] The assumptions that Θ is a subset of the real line and that Z is discrete may seem unduly restrictive. It is straightforward to generalize Proposition 1 that follows to allow Θ to be a subset of a Euclidean space. As written, the proof of Proposition 1 in Kocherlakota (2005) relies on the discreteness of Z. However, we believe that the proof could be extended to include the case in which Z is an arbitrary Borel subset of a Euclidean space.

[7] Attanasio and Davis (1996) document that the cross-sectional dispersion of consumption increased in the 1980's in the United States along with the publicly observable change in the cross-sectional dispersion of wages. Sometimes, this finding is interpreted as being evidence that individuals cannot insure themselves against publicly observable shocks. But, as Attanasio and Davis themselves point out, these movements are also consistent with the hypothesis that the increase in the cross-sectional variance of *measured* wages was associated with an increase in the variance of private information about skills. Again, we can specify our function ϕ_t so as to capture this possibility.

The aggregate shocks also affect the aggregate production function as follows. We define an allocation in this society to be (c, y) where:

(9) $\quad c : S \times \Theta^T \times Z^T \rightarrow R_+^T$

(10) $\quad y : S \times \Theta^T \times Z^T \rightarrow R_+^T$

(11) $\quad (c_t, y_t)$ is (s, θ^t, z^t)-measurable

Here, $y_t(s, \theta^T, z^T)$ $(c_t(s, \theta^T, z^T))$ is the amount of effective labor (consumption) assigned in period t to an agent with label s and type θ^T, given that the public aggregate shock sequence is z^T. We define an allocation (c, y) to be *feasible* if for all t, z^T:

$$(12) \quad \sum_{s \in S} \pi_s \int_{\theta^T \in \Theta^T} c_t(s, \theta^T, z^T) d\mu_\Theta \leq \sum_{s \in S} \pi_s \int_{\theta^T \in \Theta^T} y_t(s, \theta^T, z^T) d\mu_\Theta$$

Because θ_t is only privately observable, allocations must respect incentive-compatibility conditions. (The following definitions correspond closely to those in Golosov, Kocherlakota and Tsyvinski (2003).) A *reporting strategy* $\sigma : \Theta^T \times Z^T \rightarrow \Theta^T \times Z^T$, where σ_t is (θ^t, z^t)-measurable and $\sigma(\theta^T, z^T) = (\theta^{T\prime}, z^T)$. Let Σ be the set of all possible reporting strategies, and define:

(13) $\quad W(.; c, y) : S \times \Sigma \rightarrow R$

$$(14) \quad W(s, \sigma; c, y) = \sum_{t=1}^{T} \beta^{t-1} \int_{Z^T} \int_{\Theta^T} \{u(c_t(s, \sigma(\theta^T, z^T))) - v(y_t(s, \sigma(\theta^T, z^T))/\phi_t(\theta^T, z^T))\} d\mu_\Theta d\mu_Z$$

to be the expected utility from reporting strategy σ, given an allocation (c, y). (Note that the integral over Z could also be written as a sum.) Let σ_{TT} be the truth-telling strategy $\sigma_{TT}(\theta^T, z^T) = (\theta^T, z^T)$ for all θ^T, z^T. Then, an allocation (c, y, K) is *incentive-compatible* if:

(15) $\quad W(s, \sigma_{TT}; c, y) \geq W(s, \sigma; c, y)$ for all s in S and all σ in Σ

An allocation which is incentive-compatible and feasible is said to be incentive-feasible.

In this economy, a *Pareto optimal* allocation is an allocation (c, y) that solves the problem of maximizing the utility of agents with label $s = 1$ subject to (c, y) being incentive-feasible, and subject to any agent with label s, $s \neq 1$, receiving ex-ante utility of at least U_s. Note that for any specification of reservation utilities $(U_2, ..., U_S)$ such that the constraint set is non-empty, there is a solution to the planner's maximization problem (the constraint set is compact in the product topology and the objective continuous in the same topology.)

This focus on ex-ante Pareto optima is not restrictive. All of our results are valid for asymmetric interim Pareto optima, in which the planner puts different weights on different agents, and these different weights are allowed to depend on the realization of skills in period 1.

The following proposition is a restatement of Proposition 1 in Kocherlakota (2005)[8] for this environment without capital.

PROPOSITION 1. *Suppose* (c^*, y^*) *is an optimal allocation and that there exists* $t < T$ *and scalars* M_+, M^+ *such that* $M^+ \geq c_t^*, c_{t+1}^* \geq M_+ > 0$ *almost everywhere. Then the random variable:*

$$(16) \quad \lambda_{t+1} \equiv \beta[E\{u'(c_{t+1}^*)^{-1}|s, \theta^t, z^{t+1}\}]^{-1}/u'(c_t^*)$$

is z^{t+1}*-measurable.*

Proof. As in Kocherlakota (2005).

In this proposition, the random variable λ_{t+1} is the ratio of the Lagrange multiplier on the history z^{t+1} resource constraint to the Lagrange multiplier on the history z^t resource constraint (adjusted for the conditional probability of z^{t+1}). For this reason, we later term λ_{t+1} the shadow stochastic discount factor. The proposition shows that λ_{t+1} equals:

$$(17) \quad \beta\{E(u'(c_{t+1}^*)^{-1}|s, \theta^t, z^{t+1})\}^{-1}/u'(c_t^*)$$

The numerator of this ratio may depend on s and θ^t, as may the denominator. However, the ratio itself is independent of s and θ^t.

This result is obviously true without private information, because in that case the optimal c_t^* is such that $c_t^*(s, \theta^t, z^t)$ is independent of θ^t and $c_{t+1}^*(s, \theta^{t+1}, z^{t+1})$. In the presence of private information, it is generally optimal to allow c_t^* to depend on θ^t in order to require high-skilled agents to produce more effective labor. Proposition 1 establishes that in that case, the *harmonic* mean of $\beta u'(c_{t+1}^*)/u'(c_t^*)$, conditional on

[8] See also Golosov, Kocherlakota and Tsyvinski (2003) and Rogerson (1985) for related results.

z^{t+1}, is independent of θ^t. This result is a generalization of a well-known result from the dynamic moral hazard literature (Rogerson (1985)); see Golosov, Kocherlakota and Tsyvinski (2003) for a discussion of the underlying intuition.

4. Asset Pricing

The prior section is based on the analysis in Kocherlakota (2005). In this section, we break new ground. We consider the asset pricing implications of Pareto optimality. We assume that the planner's shadow stochastic discount factor λ is a valid stochastic discount factor for asset returns. We show that for $u(c) = c^{1-\gamma}/(1-\gamma)$, λ is equal to the reciprocal of the (gross) growth of the γth moment of the cross-sectional distribution of consumption. This result remains true even when consumption is mismeasured with possibly biased or persistent measurement errors. We then set forth two other SDFs.

A. Asset Pricing via the Shadow Stochastic Discount Factor

Suppose that in the above environment, agents engage in sequential asset trade: specifically, in each period $t = 1, ..., T-1$, agents can trade (at least) M assets, where the payoff of asset m in period t is a z^t-measurable function of z^T. Let R^m_{t+1} be the equilibrium gross return from period t to period $(t+1)$ of asset m. We assume that the allocation of consumption is Pareto optimal, and the shadow SDF λ_{t+1} is a valid asset pricing kernel for all asset returns. More precisely, we assume that for any asset m:

(18) $\quad 1 = E\{R^m_{t+1}\lambda_{t+1}|z^t\}$ for all t, z^t

Using some algebra, we can use Proposition 1 to express the shadow price λ in terms of moments of the cross-sectional distribution of consumption. Let (c^*, y^*) be an optimal allocation, for $u(c_t) = c^{1-\gamma}/(1-\gamma)$. Define:

(19) $\quad C_{\gamma,t}(z^t) = E\{c_t^{*\gamma}|z^t\}$

to be the γth moment of the cross-sectional distribution of consumption in public history z^t. Proposition 1 implies that:

(20) $\quad \lambda_{t+1}c_t^{*-\gamma} = \beta\{E(c_{t+1}^{*\gamma}|s, \theta^t, z^{t+1})\}^{-1}$

¿From the Law of Iterated Expectations, we get:

$$\lambda_{t+1}^{-1} E(c_t^{*\gamma}|z^{t+1}) = \beta^{-1} E(c_{t+1}^{*\gamma}|z^{t+1})$$

By assumption, (s, θ^t) is independent of z_{t+1}, conditional on z^t. Hence, $E(c_t^{*\gamma}|z^{t+1}) = E(c_t^{*\gamma}|z^t)$ for all t, z^{t+1}, and we can conclude that:

$$\lambda_{t+1}(z^{t+1}) = \beta C_{\gamma,t}(z^t)/C_{\gamma,t+1}(z^{t+1})$$

Thus, the shadow SDF λ is tied to the growth rate of the γth moment of the distribution of consumption. It follows that if equilibrium quantities are Pareto optimal, and λ_{t+1} is a valid *market* SDF, we know that:

$$\text{(APR)} \qquad 1 = \beta E\{C_{\gamma,t} R_{t+1}^m C_{\gamma,t+1}^{-1}|z^t\}$$

where R_{t+1}^m is the equilibrium gross return of asset m. Thus, assets are priced according to a new type of stochastic discount factor which is equal to the growth rate of the γth moment of the cross-sectional distribution of consumption. Henceforth, we use the term Private Information Pareto Optimal (PIPO) stochastic discount factor (SDF) to refer to the expression:

$$\beta C_{\gamma,t}/C_{\gamma,t+1}$$

(Note that this discount factor is the same as the representative agent asset pricing model's discount factor for $\gamma = 1$.)

This result is related to two others in the literature. First, Kocherlakota (1998) derives a similar stochastic discount factor in a two-period setting with moral hazard. Second, this result is in some ways similar to that of Lustig (2002). He shows how in an economy with limited enforcement (but complete information), assets are priced using a stochastic discount factor that depends on the growth rate of a particular moment of the distribution of Pareto-Negishi weights. Relative to Lustig's formulation, the advantage of the above stochastic discount factor is that it is measurable using data from the cross-sectional distribution of consumption.[9]

[9] Pareto optimality also has a host of implications for the behavior of individual-level of consumption. For example, Pareto optimality implies that

$$\frac{E\{u'(c_{t+1})^{-1} | \theta^t \in A, z^{t+1}\}}{E\{u'(c_t)^{-1} | \theta^t \in A, z^t\}}$$

B. Measurement Error in Consumption

One of the difficulties with using cross-sectional data in consumption is that the data are typically measured with error. This measurement error usually creates difficulties when one applies the Generalized Method of Moments to estimate Euler equations of the form:

$$\beta E_t\{(c_{t+1}/c_t)^{-\gamma} R_{t+1}\} = 1$$

Measurement error in the level of consumption can bias the level of measured household consumption growth upward or downward, and so can contaminate the estimates of β and γ in unknown ways.

In our paper, the PIPO SDF is a ratio of moments of the cross-sectional consumption distribution at different dates. Under reasonable assumptions, the impact of measurement error on a particular moment of the consumption distribution is the same at every date and state, because we can aggregate the measurement error across individuals. In this subsection, we prove that this intuition is valid by demonstrating formally that if the asset pricing restriction APR is valid for true consumption, it is also valid for measured consumption, given a relatively weak assumption about the nature of measurement error. (See Chioda (2004) for a similar approach to the measurement error issue.)

In particular, let (c^*, y^*) be a socially optimal allocation, and suppose the shadow SDF λ is a valid market SDF. We allow c^* to be measured with error as follows. Let $(\nu_1, \nu_2, ..., \nu_T)$ be a collection of random variables with joint probability measure μ_ν over the Borel sets in R^T. At the beginning of period 1, after the public shock sequence z^T is drawn, a realization ν^T is drawn according to μ_ν for each agent; conditional on z^T, the draws of ν^T and θ^T are independent from each other and are independent across agents. Note too that ν^T is independent of z^T (because it is drawn from μ_ν for all z^T); however, the measurement error is allowed to have arbitrary serial correlation.

Define $\widehat{c}_t(s, \theta^t, z^t, \nu_t) = \exp(\nu_t) c_t^*(s, \theta^t, z^t)$ to be measured consumption. Define also:

(21) $\widehat{C}_{\gamma,t} = E\{\widehat{c}_t^{\gamma} | z^t\}$

---

is the same for any Borel set A. By varying the Borel set A (say, to be "all employed agents in period t" or "all unemployed agents in period t"), we can generate interesting testable restrictions. We have work in progress in which we assess the validity of such restrictions in the CEX. .

to be the γth moment of cross-sectional measured consumption, in public history z^t. From the definition of measured consumption, we know that:

$$(22) \quad \widehat{C}_{\gamma,t} = E\{c_t^{*\gamma} \exp(\gamma \nu_t) | z^t\}$$

$$(23) \quad = E\{\exp(\gamma \nu_t) | z^t\} E\{c_t^{*\gamma} | z^t\}$$

$$(24) \quad = E\{\exp(\gamma \nu_t)\} C_{\gamma,t}$$

where the penultimate equation comes from the independence of ν_t from θ^t, conditional on z^t.

Now suppose that $E\{\exp(\gamma \nu_t)\} < \infty$ and ν_t is stationary. These assumptions imply that:

$$(25) \quad \beta E\{\widehat{C}_{\gamma,t} R^m_{t+1} / \widehat{C}_{\gamma,t+1} | z^t\} = \beta E\{C_{\gamma,t} R^m_{t+1} / C_{\gamma,t+1} | z^t\}$$

for all (t, z^t). Thus, under these assumptions, the asset pricing restriction APR is also valid for measured consumption, as long as the measurement error is independent across agents, independent from agents' true types, and is stationary over time. These assumptions about the nature of the measurement error are not wholly innocuous. On the other hand, we do not have to make any assumptions at all about the magnitude of the measurement error, beyond assuming the finiteness of a particular moment, or impose any particular restrictions on its autocorrelation structure.[10]

C. Two Other Stochastic Discount Factors

In the prior two subsections, we set forth a new model of a stochastic discount factor for asset pricing. In the empirical work that follows, we contrast its empirical performance with two alternative stochastic discount factors. The first is derived in the same economic environment described in Section 2; it is an implication of equilibrium given that agents trade a possibly limited set of securities, but any borrowing constraints bind with probability zero. The second discount factor is an implication of equilibrium when financial markets are complete and agents' shock histories are publicly observable.

[10] There is no evidence from validation consumption studies that can tell us whether the assumption we make about the nature of the measurement error are truly restrictive. Evidence from validation wage and income studies (Bound and Krueger, 1991) have found that: (a) measurement error appears serially correlated, (b) independent of schooling, and (c) negatively correlated with the true measure. The latter finding will, of course, invalidate our empirical strategy.

The Incomplete Markets SDF

We assume that the economic environment is as described in Section 3. We assume as before that agents engage in sequential asset trade, so that in period $t = 1, ..., T-1$, agents can trade at least M assets, where the payoff of asset m in period t is a z^t-measurable function of z^T. Note that markets are *incomplete,* in the sense that agents cannot directly insure themselves against individual-specific shocks. They may also be incomplete in the different sense that they cannot trade a complete set of z^t-contingent securities.

Let R_{t+1}^m be the equilibrium gross return from period t to period $(t+1)$ of asset m. Let (c^{INC}, y^{INC}) be an equilibrium allocation in this setting such that in equilibrium, agents face no binding borrowing constraints. A necessary condition of individual optimality is:

$$(26) \quad c_t^{INC}(s, \theta^t, z^t)^{-\gamma} = \beta E\{c_{t+1}^{INC}(s, \theta^{t+1}, z^{t+1})^{-\gamma} R_{m,t+1}(z^{t+1}) | s, \theta^t, z^t\}$$

for all t, z^t and almost all θ^t. We can then integrate over s and θ^T on both sides of this equation to get:

$$(27) \quad C_{-\gamma,t}^{INC}(z^t) = \beta E\{C_{-\gamma,t+1}^{INC}(z^{t+1}) R_{m,t+1}(z^{t+1}) | z^t\}$$

and it follows that in this equilibrium, assets are priced according to the following stochastic discount factor:

$$(28) \quad \beta C_{-\gamma,t+1}^{INC}(z^{t+1}) / C_{-\gamma,t}^{INC}(z^t)$$

We will call this the *incomplete markets* SDF.

It is important to distinguish this discount factor from a similar one employed by Brav, Constantinides, and Geczy (BCG) (2002) and Cogley (2002). Those papers make the same assumptions about market structure (incomplete markets with non-binding borrowing constraints) and derive the following SDF:

$$(29) \quad \beta E\{c_{t+1}^{INC}(s, \theta^{t+1}, z^{t+1})^{-\gamma} c_t^{INC}(s, \theta^t, z^t)^{\gamma} | z^{t+1}\}$$

which is the average of the agents' intertemporal marginal rates of substitution. Like the incomplete markets discount factor (28), this *average IMRS* discount factor is also a valid SDF in an incomplete markets equilibrium with non-binding borrowing constraints. Indeed, if agents can trade a complete set of z^t-contingent securities, so that they can completely insure themselves against any publicly observable shock, these two SDFs are identical. Otherwise, though, they may be different.

In this paper, we focus on the incomplete markets SDF (28), because it is more robust to measurement error than the average IMRS SDF (29). To see this, suppose that there is a measurement error process of the kind defined in section 4 and we observe:

$$(30)\quad \widehat{c}_{t+1}^{INC}(s,\theta^{t+1},z^{t+1},\nu_{t+1}) = c_{t+1}^{INC}(s,\theta^{t+1},z^{t+1})\exp(\nu_{t+1})$$

Then, the average IMRS discount factor, calculated using observed consumption, is given by:

$$(31)\qquad \beta E\{c_{t+1}^{INC}(\theta^{t+1},z^{t+1})^{-\gamma}\exp(-\gamma\nu_{t+1})c_t^{INC}(\theta^t,z^t)^{\gamma}\exp(\gamma\nu_t)|z^{t+1}\}$$

$$(32)\qquad = E\{\exp(-\gamma\nu_{t+1})\exp(\gamma\nu_t)\}\beta E\{c_{t+1}^{INC}(\theta^{t+1},z^{t+1})^{-\gamma}c_t^{INC}(\theta^t,z^t)^{\gamma}|z^{t+1}\}$$

which is the true average IMRS discount factor multiplied by a constant. This means that the measured average IMRS discount factor is valid for return *differentials* like the equity premium, but it is not valid for the level of returns.

In contrast, the measured incomplete markets SDF equals:

$$(33)\quad \frac{C_{-\gamma,t+1}^{INC}(z^{t+1})E\{\exp(-\gamma\nu_{t+1})\}}{C_{-\gamma,t}^{INC}(z^t)E\{\exp(-\gamma\nu_t)\}}$$

If ν_t is stationary, and $E\{\exp(-\gamma\nu_t)\} < \infty$, then this measured incomplete markets SDF is equal to the actual incomplete markets SDF. Thus, the incomplete markets SDF defined in this paper is more robust to measurement error than the average IMRS discount factor used by BCG.

Note that in this economic setting, in which agents can freely alter their consumptions by asset exchange, the PIPO SDF will typically *not* be valid. Suppose for example that Z is a singleton, so that there are no aggregate shocks. Then:

$$(34)\quad c_t^{INC}(s,\theta^t)^{\gamma} = \beta^{-1}R_{m,t+1}^{-1}\{E\{c_{t+1}^{INC}(s,\theta^{t+1})^{-\gamma}|s,\theta^t\}\}^{-1}$$

$$(35)\qquad > \beta^{-1}R_{m,t+1}^{-1}E\{c_{t+1}^{INC}(s,\theta^{t+1})^{\gamma}|s,\theta^t\}$$

which implies that:

$$(36)\quad 1 < \beta R_{m,t+1}C_{\gamma,t}/C_{\gamma,t+1}$$

The PIPO SDF is invalid in this incomplete markets setting.

Nonetheless, there are many implementations in which the PIPO SDF is valid in equilibrium (see Kocherlakota (2005) and Golosov and Tsyvinski (2006)). A key feature of any such implementation is that the insurer against productivity shocks must be able to condition insurance payments on the asset holdings of the insuree. This kind of conditioning has empirical counterparts. Two important sources of insurance against large idiosyncratic productivity shocks are the government and informal social networks like families and/or friends. In the United States, the government explicitly conditions social insurance on asset holdings. Thus, an unemployed household with large asset-holdings pays more income taxes and so receives less net payments from the government than does an unemployed household with low asset-holdings. Even more directly, the government does not make welfare payments to households with sufficiently high asset-holdings. For example, to be eligible for Supplemental Security Income, food stamp benefits, welfare benefits under the TANF program, or Medicaid, applicants generally must meet both an income test and an explicit asset test. Specifically, the food stamp asset limit is $2,000 ($3,000 for households with an elderly or disabled member); for SSI, the limits are $2,000 for a single individual and $3,000 for a couple.

In terms of intra-familial insurance, there is some evidence that parents' transfers to their heirs depend on the latter's wealths. McGarry and Schoeni (1995) use microeconomic data from the Health and Retirement Survey (HRS) to show that "equal transfers to all children are the exception rather than the rule" (p. S204). More importantly for our purpose, they find in multivariate controlled regressions that "parents transfer less to wealthy children, and thus less to children with housing wealth than those without" (p. S207).[11] At a more introspective level, it seems intuitive that a family would be more willing to help out a freshly unemployed member if that member were homeless.

More abstractly, the PIPO SDF and the incomplete markets SDF are derived by taking extreme and opposing positions on how insurance against idiosyncratic skill shocks responds to an individual's asset holdings. The incomplete markets SDF assumes that this response is zero. The PIPO SDF assumes that an agent's insurance changes in such a way so as to ensure that his individual consumption remains socially optimal. The world lies somewhere between these two extremes. The goal of our empirical work is to

[11] Light and McGarry (2004) use instead a set of direct questions regarding expectations of *inter vivo* transfers asked to all mothers in the HRS. Mothers who report that they intend to divide their estates unequally are asked to explain why. Of the mothers providing a response, slightly less than 20% "refer to their childrens' financial needs in explaining why some will receive a larger bequest than others (e.g., "the oldest son has more assets than the youngest son")", while another 9% cite asset tests, e.g., "my daughter can't have over $2000 or she will lose her state benefits".

understand which of these two extreme models provides a better approximation to the behavior of asset returns.

Of course, it would be desirable to be able to evaluate less extreme models. The incomplete markets model assumes that agents can freely alter their consumption without the knowledge of insurers. Suppose instead that agents can only secretly alter their consumption by holding some relatively low-return asset (like holding cash, for example). What is the structure of the resultant SDF? Unfortunately, the answer to this important theoretical question is not known.

The Representative Agent SDF

We now consider a different economic environment. We assume that θ_t is public information, instead of only being privately known to the relevant agent. In such an environment, in a Pareto optimal allocation, consumption is independent of θ^t. We assume again that agents engage in sequential trade of at least M assets. Let (c^{RA}, y^{RA}) be an equilibrium allocation in this economy such that agents face no binding short-sales constraints in equilibrium, and assume that this allocation is Pareto optimal; hence, $c_t^{RA}(s, \theta^t, z^t) = \phi_s C_{1,t}^{RA}(z^t)$ for some positive constant ϕ_s.

Then, in equilibrium:

$$(37) \quad c_t^{RA}(s, z^t)^{-\gamma} = E\{c_{t+1}^{RA}(s, z^{t+1})^{-\gamma} R_{t+1}(z^{t+1}) | \theta^t, z^t\}$$

We can therefore construct a valid SDF by using the intertemporal marginal rate of substitution of a representative agent:

$$(38) \quad \beta(C_{1,t+1}^{RA})^{-\gamma} / (C_{1,t}^{RA})^{-\gamma}$$

Note that this representative agent SDF is equivalent to the PIPO SDF when $\gamma = 1$ (the representative agent has log utility).[12]

The representative agent SDF is robust to the same class of measurement error processes described above.

[12] Altug and Miller (1990) also construct a representative agent stochastic discount factor using observations from panel data.

D. The Role of Consumption Inequality in the Three SDFs

The key to our empirical analysis is that the three SDFs differ in how they depend on consumption inequality. To see this, fix a history $\overline{z}^t$, and suppose the cross-sectional distribution of consumption in history $(\overline{z}^t, z')$ is a mean-preserving spread of the cross-sectional distribution of consumption in history $(\overline{z}^t, z)$. The representative agent SDF is the same in states z and z'. Hence, under this SDF, consumption inequality does not affect asset prices.

The other two SDFs depend on consumption inequality, but in different ways. Because $c^{-\gamma}$ is a convex function of c for all $\gamma > 0$, we know that:

$$(39) \quad C_{-\gamma,t+1}(\overline{z}^t, z') > C_{-\gamma,t+1}(\overline{z}^t, z)$$

and so the incomplete markets SDF implies that the state price of consumption is higher in state z'. In contrast, because c^{γ} is a convex function of c for all $\gamma > 1$ (the empirically relevant case), we know that:

$$1/C_{\gamma,t+1}(\overline{z}^t, z') < 1/C_{\gamma,t+1}(\overline{z}^t, z)$$

and so the PIPO SDF implies that the state price of consumption is lower when consumption inequality is higher. When $\gamma > 1$, the PIPO discount factor implies that the state price of consumption is *lower* in state z'. Thus, with the incomplete markets discount factor, assets have high prices and low returns when their payoffs are *positively* correlated with consumption inequality. With the PIPO discount factor when $\gamma > 1$, assets have high prices and low returns when their payoffs are *negatively* correlated with consumption inequality.

Intuitively, there are two forces at work here. First, in a state in which the distribution of consumption is more unequal, individual-specific consumption risk is higher. This higher idiosyncratic risk generates a higher precautionary demand for assets with high payoffs in that state. It is this precautionary effect that drives the positive correlation between inequality and state prices with the incomplete markets discount factor.

The second force is related to incentives. Because of wealth effects on labor supply, it is less costly to provide incentives to people with low consumption. When consumption is less equally distributed, there are more low-consumption people for a given amount of per-capita consumption, and incentive costs are lower.

This additional incentive effect generates the *negative* correlation between inequality and asset prices under the PIPO discount factor when $\gamma > 1$.

This basic intuition can be seen even more clearly by taking a nonparametric approach. Consider a general utility function u. Define its coefficient of relative risk aversion $CRRA(c) = -\frac{u''(c)c}{u'(c)}$ and its coefficient of relative prudence $CRP(c) = -\frac{u'''(c)c}{u''(c)}$. The PIPO SDF is decreasing in consumption inequality if $1/u'$ is a convex function. This is true if:

$$(40) \quad [-CRP(c) + 2 * CRRA(c)] > 0$$

for all c. Thus, the PIPO SDF is decreasing in consumption inequality as long as risk aversion is sufficiently large relative to prudence. Intuitively, higher risk aversion means that incentive costs are lower, while higher prudence means that precautionary demand is higher. The PIPO SDF ends up balancing these two effects against one another. In contrast, and as is well-known, the incomplete markets discount factor is increasing in consumption inequality as long as agents are prudent ($CRP > 0$).

5. Empirical Analysis: Preliminaries

In this section, we describe our data and empirical methodology.

A. The Data: The CEX

The microeconomic data are drawn from the 1980-2004 Consumer Expenditure Survey (CEX). The CEX provides a continuous and comprehensive flow of data on the buying habits of American consumers. The data are collected by the Bureau of Labor Statistics and used primarily for revising the CPI. Consumer units are defined as members of a household related by blood, marriage, adoption, or other legal arrangement, single person living alone or sharing a household with others, or two or more persons living together who are financially dependent. The definition of the head of the household in the CEX is the person or one of the persons who owns or rents the unit.

The CEX is based on two components, the Diary, or record keeping survey, and the Interview survey. The Diary sample interviews households for two consecutive weeks, and it is designed to obtain detailed expenditures data on small and frequently purchased items, such as food, personal care, and household supplies. The Interview sample is in the form of a rotating panel, and it follows survey households for a maximum of 5 quarters, although only inventory and basic sample data are collected in the first quarter (these data are not publicly available.) The data base covers about 95% of all expenditure, with the exclusion of expenditures for housekeeping supplies, personal care products, and non-prescription drugs. Following most previous research, our analysis below uses only the Interview sample.

The CEX collects information on a variety of socio-demographic variables, including characteristics of members, characteristics of housing unit, geographic information, inventory of household appliances, work experience and earnings of members, unearned income, taxes, and other receipts of consumer unit, credit balances, assets and liabilities, occupational expenses and cash contributions of consumer unit. Expenditure is reported in each interview (after the first) and refers to the months of the previous quarter. Thus, a household interviewed in April 1980 reports expenditure for January, February, and March 1980. Income is reported in the second and fifth interview, and it refers to the previous twelve months. Holdings of financial assets are reported only in the fifth interview.

We refer the reader to the Appendix for step-to-step details on sample selection and consumption definition. Our sample selections are aimed at eliminating the most severe reporting errors in consumption.[13] We end up discarding about 25% of observations through our selection procedure.[14] The definition of total non durable consumption is similar to Attanasio and Weber (1995). It includes food (at home and away from home), alcoholic beverages and tobacco, heating fuels and utilities, transports (including gasoline), personal care, clothing and footwear, entertainments, other services (including domestic services). It excludes expenditure on various durables, housing (furniture, appliances, etc.), education and health.

We "deflate" consumption data to account for three phenomena: price differences over time, seasonal differences (i.e., month effects) within a year, and households' demographic differences at a certain point in time. Thus, nondurable consumption is first expressed in real terms using the chained CPI (all items) for Urban Consumers (in 1982-84 dollars, as provided by the BLS). Then, data are de-seasonalized by simple multiplicative regression adjustments. Finally, we convert it into adult-equivalent consumption data.[15] Given

[13]An alternative (or a further sample selection) is to remove observations in the tails of the cross-sectional distribution of consumption. Our sample selection is likely already removing some of these observations. A sample selection of this form drops extreme errors, but also genuine observations (the very rich or the very poor). This is undesirable in the context of the theory we are studying. For similar reasons, we do not use a Taylor series approximation to the various SDFs.

[14]The starting sample has an average of 1911 households in any given (overlapping) quarter. Our final sample has an average of 1412 household per (overlapping) quarter.

[15]The number of adult equivalents is defined as $(A + \alpha K)^{\beta}$ where A is the number of adults (aged 18 or more), K the number

the overlapping panel nature of the CEX, each month a certain number of households enter the panel and an approximately equal number leave it. Monthly consumption data are aggregated to form quarterly consumption data for each household in the sample. Then, we aggregate across households to form moments of the quarterly consumption distribution. Note that households start their second interview (when consumption data are first collected) in different months. Thus, some households' second interview covers the months of January through March, some other households' second interview will have data for the months of February through April, and so forth. By the very design of the CEX, no households contribute multiple observations to adjacent overlapping quarters. In other words, a household that contributes data to January-March 1980 will not contribute data for February-April (or March-May). Its next contribution, if that exists, will be for April-June 1980.

Recently, researchers have noted that for many commodities, the aggregation of CEX data matches poorly National Income and Product Accounts (NIPA) Personal Consumption Expenditure (PCE) data. Some of the discrepancy is undoubtedly due to differences in covered population and definitional issues. But the amount of underestimation of consumer expenditure is sometimes substantial and it raises some important warning flags. Furthermore, there is evidence that the detachment between the CEX aggregate and the NIPA PCE has increased over time.[16] At present, it is not clear why this is so, and whether this is necessarily due to a worsening in the quality of the CEX. For example, Bosworth et al (1991) conclude that most of the discrepancy is explained by the failure of the CEX to sample the super-rich; others have suggested a greater incidence of attrition. According to the BLS, however, the CEX has maintained representativeness of the US population over time, and attrition has not changed much since the redesign of the survey of the early 1980s.

Given these differences between the CEX data and the NIPA data, it is useful to check whether similar results are obtained using the latter. To this purpose, we also estimated the parameters in the representative agent SDF using aggregate NIPA PCE data. We obtain NIPA PCE data from the NIPA Table 2.8.5, which reports Personal Consumption Expenditures by major type of product (durable goods, non durable goods, and services) on a monthly basis.[17] The data are collected by the Bureau of Economic Analysis. Our measure of

of kids, and α and β parameters. We set $\alpha = 0.7$ and $\beta = 0.65$ (following recommendations contained in Citro and Michaels, 1995, which in turn draws from Betson, 1990). Similar results are obtained if we use a more sophisticated Engel approach.

[16] See Attanasio, Battistin, and Ichimura (2004).

[17] All the NIPA tables can be found at http://www.bea.gov/bea/dn/nipaweb/index.asp.

consumption is Personal Consumption Expenditures on nondurable goods (this is comparable to the measure of consumption we construct in the CEX, where services from durables are missing). The data are seasonally adjusted at annual rates, deflated using the same monthly CPI we use to deflate CEX data, and divided by the US population (midperiod estimates). These adjustments mimic those implemented for the micro CEX data as to ensure comparability. The monthly data so obtained are summed to form overlapping quarterly consumption data, the same data construction criterion used in the CEX (thus, consumption in 1980:3 refers to January-March 1980, consumption in 1980:4 to February-April 1980, and so on). However, changing the measure of consumption in this way had little impact on our results for the representative agent case (the results are available on request).

B. The Data: Asset Returns

We use returns data drawn from the Center for Research in Security Prices (CRSP) at the University of Chicago. The construction of the variables of interest (R^{mkt} and R^{f}) is similar to BCG.

The risk free rate R^{f} is obtained in the following way. First, we extract the one-month nominal returns on Treasury bills. Then, we convert it in real terms dividing it by $(1+\pi)$, where π is the monthly inflation rate obtained from the chained CPI-U (in 1982-84 dollars), also used below. Finally, we obtain the quarterly return by compounding the monthly returns.

The market return R^{mkt} is the return on the CRSP value- weighted portfolio. It includes dividends and capital gains. We first take the average one-month nominal return of the pooled sample of stocks listed on the New York Stock Exchange and the American Stock Exchange. We convert it into real terms by dividing it by $(1+\pi)$. Finally, we obtain the quarterly return by compounding the monthly returns. The difference $\left(R_t^{mkt} - R_t^{f}\right)$ is the premium on the value weighted portfolio.

C. Methodology

Our estimation methodology is motivated by that originally described by Hansen and Singleton (1982). Let $\{x_t\}_{t=1}^{T}$ be any stochastic process such that x_t is z^t-measurable, and let $\{R_t^m\}_{t=1}^{T}$ be the gross return process to some financial asset. Then, a valid stochastic discount factor $m_t(\beta,\gamma,z^t)$ satisfies:

$$(41)\quad E\left[\left(m_t(\beta,\gamma,z^t)R_t^m(z^t)-1\right)x_s\right]=0$$

By considering arbitrary instruments x_s's (for $s < t$) and arbitrary returns R_t^m, we can form a large number of such orthogonality conditions. In principle, we can evaluate any of these population restrictions using sample analogs. However, it is important to realize that the small sample properties of the resultant estimators and tests are likely to be poor unless each x has marginal predictive power (over the collection of other x's) for either m, R, or (preferably) both.

In what follows, we focus on three implications that have received a great deal of attention in the macroeconomic literature. The first concerns the ability of the SDF's to rationalize the variation in the expected return to the Treasury bill. The real return to the Treasury bill is highly predictable by its own lag. A valid SDF should eliminate this predictability. To assess this aspect of the SDFs, we investigate the following two restrictions:

$$E\left[\left(m_t(\beta,\gamma,z^t)R_t^f(z^t)-1\right)\right] = 0 \tag{42}$$

$$E\left[\left(m_t(\beta,\gamma,z^t)R_t^f(z^t)-1\right)R_{t-3}^f(z^{t-3})\right] = 0 \tag{43}$$

Next, we investigate the equity premium puzzle of Mehra and Prescott (1985). They point out that, historically, the gap between average stock returns and average Treasury bill returns is very large (on the order of 6% per year) and difficult to rationalize using standard representative agent asset pricing models. As in Kocherlakota (1996), we assess the candidate stochastic discount factors' ability to rationalize the equity premium by considering the restriction that:

$$E\left[m_t(1,\gamma,z^t)\left(R_t^{mkt}(z^t)-R_t^f(z^t)\right)\right]=0 \tag{44}$$

where R_t^{mkt} is the value-weighted return to the stock market and R_t^f is the return to the 90-day Treasury bill (all returns are real). Then, our final comparison of the SDF's is based on their ability to simultaneously rationalize the excess return to the stock market and the two Treasury bill implications.

The CEX provides data of the form $\{\{c_{it}\}_{i=1}^{N_t}\}_{t=-2}^{T}$, where c_{it} is the consumption expenditure of household i for the quarter ending with month t (i.e., covering months $t-2$, $t-1$, and t). We define sample analogs of the various stochastic discount factors using these *cross-sectional* data. In particular, let:

$$(45) \quad \widehat{m}_t^{PIPO}(\beta,\gamma) = \beta\frac{N_{t-3}^{-1}\sum_{i=1}^{N_{t-3}} c_{it-3}^{\gamma}}{N_t^{-1}\sum_{i=1}^{N_t} c_{it}^{\gamma}}$$

$$(46) \quad \widehat{m}_t^{INC}(\beta,\gamma) = \beta\frac{N_t^{-1}\sum_{i=1}^{N_t} c_{it}^{-\gamma}}{N_{t-3}^{-1}\sum_{i=1}^{N_{t-3}} c_{it-3}^{-\gamma}}$$

$$(47) \quad \widehat{m}_t^{RA}(\beta,\gamma) = \beta\frac{\left(N_t^{-1}\sum_{i=1}^{N_t} c_{it}\right)^{-\gamma}}{\left(N_{t-3}^{-1}\sum_{i=1}^{N_{t-3}} c_{it-3}\right)^{-\gamma}}$$

denote the sample analogs of the PIPO, incomplete markets, and representative agent stochastic discount factors. To reiterate, we use overlapping data, so t here indexes the last month of a given quarter. Thus, for example, the first available observation for $\widehat{m}_t^{PIPO}(\beta,\gamma)$ is for 1980:3, and it is constructed as the ratio of the γ-th moment of consumption for 1979:12 (calculated using all households reporting expenditure data for October-December 1979) and the γ-th moment of 1980:3 (calculated using all households reporting expenditure data for January-March 1980). The last observation is for 2004:2. We have overall $T = 288$ observations on $\widehat{m}_t^j(\beta,\gamma)$ ($j = PIPO, INC, RA$).The average N_t is 1412 (the median is 1331).

We then form sample analogs (in the *time series* dimension) of the three restrictions (44)-(43). For example, the sample analog of (44) for the PIPO discount factor is:

$$(48) \quad T^{-1}\sum_{t=1}^{T}\frac{N_{t-3}^{-1}\sum_{i=1}^{N_{t-3}} c_{it-3}^{\gamma}}{N_t^{-1}\sum_{i=1}^{N_t} c_{it}^{\gamma}}(R_t^{mkt}(z^t) - R_t^f(z^t)) = 0$$

We then estimate the unknown parameters and evaluate the three discount factors by applying Generalized Method of Moments (GMM) to these moment conditions. For example, in our discussion of the equity premium, we estimate γ by minimizing the square of the left-hand side of (48) with respect to γ.

Finally, it is worth discussing the non-standard inference problem we face. Note that time-series moment conditions like (48) are functions of cross-sectional non-linear moments of the data. If these cross-sectional moments were known, then we could simply plug in these known moments into the time-series moment conditions and then apply the usual time series GMM formulae to calculate standard errors (perhaps after accounting for serial correlation induced by the use of overlapping data). In reality, we do not know the true cross-sectional moments. In what follows, we will assume that the cross-sectional sample size increases at a rate that is sufficiently fast as to make this source of uncertainty inconsequential and treat sample estimates as if they were the true cross-sectional moments. This is the same approach followed by BCG (2002).[18] We compute our standard errors to account for serial correlation in returns and overlapping data using a correction of the form proposed by Hansen and Hodrick (1980).

6. Empirical Analysis: Results

We provide some simple summary statistics in Table 1. There is a large equity premium contained in Table 1b. The mean return to stocks is about 1.9% per quarter higher than the mean return to Treasury bills. This sample estimate is higher than the 6.2% annual number averaged in the hundred years of data (1889-1978) studied by Mehra and Prescott. The standard deviation of stock returns is about 7.7% per quarter. Importantly for what we do later, the risk-free rate is highly autocorrelated over the sample.

We also plot the PIPO stochastic discount factor in Figures 1 and 2. For large values of γ, the SDF is highly variable. Of course, a valid SDF has to be more than variable: it must covary negatively with stock returns.

A. Treasury Bill Returns

We look first at the ability of the SDF's to rationalize the variation in the expected return to the Treasury bill. Define:

$$(49) \quad \overline{e}^{j}_{b1}(\beta,\gamma) = \frac{1}{T}\sum_{t=1}^{T}\widehat{m}^{j}_{t}(\beta,\gamma)\,R^{f}_{t} - 1$$

$$(50) \quad \overline{e}^{j}_{b2}(\beta,\gamma) = \frac{1}{T}\sum_{t=1}^{T}(\widehat{m}^{j}_{t}(\beta,\gamma)\,R^{f}_{t} - 1)R^{f}_{t-3}$$

We estimate (β,γ) by applying GMM to these pricing errors. Given that the model is exactly identified, our choice of weighting matrix is irrelevant. We find that in all models it is possible to find (β,γ) so as to zero out both pricing errors. We find in Table 2 that the estimate of (β,γ) is about (0.99, 2) for the $PIPO$ discount factor. The estimate for β is slightly lower, considering that it is being estimated over a quarterly frequency. The estimates of β and γ are more plausible and more precise in the other two models.

[18] An alternative would be to assume that the cross-sectional sample size increases as fast as (or even less than) the time-series size. In this case, one could think of computing standard errors by the bootstrap. However, given the complications involved (for instance, the fact that when T goes to infinity so does the number of parameters to estimate), we have decided to leave this as a topic for future research.

Note that there is an extensive amount of dependence between the sample estimates of the two parameters. This will be crucial for understanding the performance of the alternative SDFs in what follows. In particular, the value of β that minimizes the pricing errors $\overline{e}_{b1}^{j}(\beta,\gamma)$ and $\overline{e}_{b2}^{j}(\beta,\gamma)$ shrinks as γ increases. For example, in the PIPO model, $\widehat{\beta}$ equals 0.9922 when $\gamma = 1$, 0.9250 when $\gamma = 3$, and 0.3051 when $\gamma = 5$. The corresponding estimates of β in the IM model are 0.9919, 0.8908, and 0.1706 respectively. The joint confidence interval for (β,γ) is much more concentrated under IM than under PIPO. These confidence intervals are graphed in Figures 3 and 4.

B. The Equity Premium: Results

Next, we look at the ability of the various discount factors to rationalize the large equity premium in the data. Define the sample mean of the equity premium errors to be:

$$(51) \quad \overline{e}_{mkt}^{j}(\gamma) = \frac{1}{T}\sum_{t=1}^{T}\widehat{m}_{t}^{j}(1,\gamma)\,(R_{t}^{mkt} - R_{t}^{f})$$

for $j = PIPO$, INC, and RA. Equation (51) is the empirical analog of (4). A simple way to compare the three models is to compare $\overline{e}_{mkt}^{j}(\gamma)$ for $j = PIPO, INC$, and RA, for different values of γ in an admissible range (we choose the 0-10 range in unit increments). This strategy is similar to Brav et al. (2002) and Kocherlakota (1996).

We report the estimates in Tables 3-5, along with confidence intervals. Our basic finding in Tables 3-5 is that with the PIPO stochastic discount factor, the sample mean of the equity premium error is zeroed out at a value of γ between 5 and 6. In contrast, with the incomplete markets and representative agent discount factors, the sample mean of the equity premium error remains positive for all specifications of γ. The confidence intervals in Table 3 and 4 show that the sample mean of the equity premium error is insignificantly different from zero for $\gamma \geq 4$ using the PIPO stochastic discount factor, and insignificantly different from zero for $\gamma \geq 6$ using the incomplete markets stochastic discount factor. Thus, from an inferential perspective, there is not much evidence to distinguish these two models on the basis of the equity premium. With the representative agent discount factor, the sample mean of the equity premium error is significantly different from zero for all values of $\gamma \leq 10$.

In Table 6, we use a slightly different approach, and formally estimate the coefficient of relative risk aversion that minimizes (51) by applying the Generalized Method of Moments to the equity premium pricing error. We find that the estimate of the coefficient of relative risk aversion is about 5 for the PIPO SDF - which is consistent with our above analysis - and the estimated standard error is about 1. We find that once we loosen the upper bound of 10 for γ, we can eliminate the equity premium pricing error in the representative agent SDF with an estimate of γ equal to approximately 58. (This high estimate of γ, of course, is reminiscent of those obtained in prior work on the equity premium - see Kocherlakota (1996).) The estimate of γ for the incomplete markets SDF is much lower than those found for the other two SDFs. However, this is misleading: the estimated equity premium pricing error for the incomplete markets SDF is around 1.9%. Hence, as our less formal procedure in Tables 3-5 showed, the incomplete markets SDF can explain virtually none of the observed equity premium.[19] Note also that the usual GMM estimate of the standard error is undefined because the (one) moment condition we use is not set equal to zero in sample. Thus, while the PIPO stochastic discount factor seems to provide an improvement over traditional incomplete market stochastic discount factors, the evidence in its favor should be interpreted with caution given the considerable amount of sampling error associated with our estimates.[20]

As pointed out by Mankiw and Zeldes (1991), the incomplete markets SDF may perform better if it is corrected for limitations on participation in financial markets. Attanasio, Banks and Tanner (2002) and Vissing-Jorgensen (2002) use data from the United Kingdom and United States to investigate this hypothesis. They show that limited participation does improve the performance of incomplete markets models to account for the level of returns. However, they do not examine the interaction between limited participation and asset return *differentials* like the equity premium. Can accounting for limited participation make the incomplete markets SDF consistent with the equity premium? To address this issue directly, we repeated our estimation procedure for the incomplete markets case considering only financial market participants.[21] The results,

[19] BCG (2002) restrict attention to households with non-negative financial wealth. When we use this smaller sample, in conjunction with the incomplete markets and complete markets SDFs, the point estimates are similar to what we obtain in Tables 3-5.

[20] We obtain qualitatively similar evidence if we use a market return that includes stocks listed on NASDAQ.

[21] We adopt two alternative classifications of financial market participation at the beginning of period t. The first follows Attanasio, Banks and Tanner (2002) and assumes that it coincides with stockholding. The second follows Vissing-Jorgensen (2002) and equates it with holding of both stocks and bonds. (This is the most appropriate case for the equity premium equation; the equity premium first order condition is only necessarily satisfied for agents who participate both in stock markets and bond markets.) The definition of stockholder and bondholder is as in Vissing-Jorgensen (2002). More details on how we perform this

reported in Table 7, show that accounting for limited financial market participation in this fashion does not substantively change our results for the incomplete markets SDF. Given that only some households are marginal in financial markets, we also changed the definition of participants and include in the group of participants only those with stockholding above a threshold. The results, also reported in Table 7, are qualitatively similar (i.e., we find no zeroing out) for thresholds equal to \$1,000, \$2,000, \$5,000, and \$10,000.

C. Anatomy of the Equity Premium Results

Why is the sample mean of the equity premium error so close to zero when γ is between 5 and 6 for the PIPO stochastic discount factor? It is instructive to look more closely at the data generating this result. Define the time-t error in the PIPO case as $e^{PIPO}_{mkt,t}(1,\gamma) = \widehat{m}^{PIPO}_t(1,\gamma)\left(R^{mkt}_t - R^f_t\right)$ and its time series average as $\overline{e}^{PIPO}_{mkt}(1,\gamma) = \frac{\sum_{t=1}^{T} e^{PIPO}_{mkt,t}(1,\gamma)}{T}$. This is the sample mean of the equity premium error we report in Table 3. The time-t error is negative whenever $\left(R^{mkt}_t - R^f_t\right) < 0$. In particular, if π is the proportion of negative time-t errors, the average error can be rewritten as a weighted average of positive and negative time-t errors,

$$(52)\qquad \overline{e}^{PIPO}_{mkt}(1,\gamma) = \pi \frac{\sum_{+} e^{PIPO}_{mkt,t}(1,\gamma)}{\sum_t \mathbf{1}\left\{e^{PIPO}_{mkt,t}(1,\gamma) < 0\right\}} + (1-\pi)\frac{\sum_{-} e^{PIPO}_{mkt,t}(1,\gamma)}{\sum_t \mathbf{1}\left\{e^{PIPO}_{mkt,t}(1,\gamma) \geq 0\right\}}$$

where $\mathbf{1}\{.\}$ is an indicator function, and $\sum_{+}$ and $\sum_{-}$ are sums over positive and negative errors, respectively. Table 2 shows that $\overline{e}^{PIPO}_{mkt}(1,\gamma) > 0$ for $\gamma < 6$ and $\overline{e}^{PIPO}_{mkt}(1,\gamma) < 0$ for $\gamma \geq 6$. Thus the average of negative time-t errors exceeds (in absolute value) the average of positive time-t errors when $\gamma > 5$. Figure 5 plots the time series of $\overline{e}^{PIPO}_{mkt}(1,\gamma)$ for $\gamma = \{3,4,5,6\}$. Figure 6 plots instead the kernel density estimate of $\overline{e}^{PIPO}_{mkt}(1,\gamma)$ for various values of γ. When γ increases, the distribution shifts to the left and the mean is dominated by spikes exerting larger and larger influence. For $\gamma > 5$ the negative spikes visible from the bottom right panel of Figure 5 get weighted more than the positive one, and $\overline{e}^{PIPO}_{mkt}$ turns negative.

The point where $\overline{e}^{PIPO}_{mkt}(1,\gamma)$ changes sign from positive to negative (if any) clearly depends on the relative weight of realizations of $e^{PIPO}_{mkt,t}(1,\gamma)$ located in the tails. For example, the largest negative value in the distribution of $e^{PIPO}_{mkt,t}(1,6)$ (the far left spike in the bottom right panel of figure 5) occurs in 2001:9. If we exclude it, $\overline{e}^{PIPO}_{mkt}(1,6)$ is positive and we would not get any zeroing-out at $\gamma = 6$ in the PIPO case (but

analysis are in the working paper version.

we would still get it for $\gamma = 7$). However, the counterfactual also works in reverse: If we were to exclude the largest positive value from the distribution of $e^{PIPO}_{mkt,t}(1,5)$ (the far right spike in the bottom middle panel of figure 5), $\overline{e}^{PIPO}_{mkt}(1,5)$ would turn negative, which means that we would get zeroing-out at $\gamma = 5$.[22] It is worth noting that even if we drop the highest possible error realization in the incomplete markets case, the sample equity premium is not eliminated for any value of γ.

The value of $e^{PIPO}_{mkt,t}(1,6)$ for 2001:9 is extremely negative because in that period the premium is negative (-12.89%), $N_t^{-1}\sum_{i=1}^{N_t} c_{it}^{\gamma}$ is small, and $N_{t-3}^{-1}\sum_{i=1}^{N_{t-3}} c_{it-3}^{\gamma}$ is large, relative to other periods (see also Figure 2).[23] There are certainly extreme values of the consumption distribution that are shifting the balance in either direction. Nevertheless, even if we choose to eliminate the four largest and four smallest consumption levels in our sample, we still get zeroing-out in the PIPO case (albeit at $\gamma = 7$) and we still do not get zeroing-out for any value of γ in the incomplete markets or representative agent cases.[24]

Our results for the incomplete markets SDF contrast with the results of BCG (2002) and Semenov (2004) for the average IMRS SDF. They find that the sample equity premium is eliminated when γ is set to a relatively low value (less than 4). Measurement error cannot be the source of the discrepancy; recall that the average IMRS SDF is valid for return differentials like the equity premium under the same class of measurement error processes that we assume in this paper. However, there are two main differences between what we do and what BCG do. First, as we stressed earlier, the incomplete markets SDF and the average IMRS SDF are distinct SDFs. The validity of the latter does not imply the validity of the former, although both should be valid in an incomplete markets equilibrium with no binding borrowing constraints. Second, BCG's sample selection is different than ours: They only keep households who stay in the sample for three or more quarters (because they use the average IMRS SDF) and, to minimize measurement error, discard

[22]Since the results are driven by consumption inequality, we also checked what happens if we remove the months with the four largest and four smallest values of the cross-sectional variance of consumption. The results remain the same, i.e., we get zeroing out for the PIPO discount factor for $\gamma = 6$ and we do not get if for any value of γ for the incomplete market discount factor.

[23]Presumably, both the small fifth moment of the consumption distribution and the low realization of stock returns in 2001:09 were related in some fashion to the terrorist attacks of 9-11.

[24]One could worry that "outliers" are driving our results. We thus experimented by dropping people with a level of (annualized) consumption that is less than 5% (10%) or more than 800% (500%) of combined after-tax household income and financial assets (the sum of amounts held in savings accounts, checking accounts, brokerage accounts, U.S. Savings bonds, stocks, mutual funds, private bonds, government bonds, Treasury notes, and personal loans). Since data on assets are reported only in the 5th interview, while income data are reported only in the 2nd and 5th interview, this robustness check is performed using just 5th interview data. We prefer this "relative" trimming to an "absolute" trimming (which may just be throwing away informative data about the very rich or the very poor; see also Bollinger and Chandra, 2004). We get zeroing-out at a value of γ between 8 and 9 (7 and 8) in the PIPO case, and no-zeroing out in the incomplete markets case.

households who report extremely large increases or decreases in consumption from one quarter to another. Their sample selections end up discarding about 60% of the households in the CEX. Finally, they use the sample period 1982:I-1996:I rather than the sample period 1980:I-2004:I.[25]

D. Joint Restrictions

Finally, we turn to using all three restrictions simultaneously. Here, with two parameters and three moments, the choice of weighting matrix is likely to matter more. Because of the finite sample difficulties documented by Kocherlakota (1990) and others, we are unwilling to use the asymptotically optimal two-step procedure originally used by Hansen and Singleton (1982). Instead, we use the one-step GMM procedure described by Hansen and Jagannathan (1997). They suggest using the weighting matrix given by the inverse of the second moments of the payoffs to be priced. In our setting, that translates into:

$$[E(xx')]^{-1}$$

where x is a column vector given by $\left[(R_t^{mkt} - R_t^f)|R_t^f|R_t^f R_{t-3}^f\right]$. We choose this weighting matrix for two reasons. First, the weighting matrix is the same for all models, which allows for ready comparison across models, even non-nested ones as in our case. Second, Hansen and Jagannathan (1997) show that with this weighting matrix, the minimized GMM objective can be interpreted as a measure of distance between the candidate SDF and the set of true SDFs (the HJ-distance).

Table 8 contains the results. In a statistical sense, the PIPO stochastic discount factor does a better job than the other two SDF's in fitting the data. First, the sample means of the pricing errors associated with the estimated PIPO discount factor are an order of magnitude smaller than those associated with the other two discount factors. Second, the overidentifying restriction is rejected for all models but PIPO. Finally, the estimated Hansen-Jagannathan distance is 0.1039 for the PIPO SDF and more than twice as much for the other two SDF's. To assess whether pricing errors in each model are significantly jointly different from zero, we also calculate the p-value of the test based on the Hansen-Jagannathan distance.[26] We find that in the PIPO model we cannot reject the null that the pricing errors are jointly zero (p-value 1 percent), but we can

[25] We constructed a subsample of the CEX using the selection criteria reported in their paper and were able to replicate most of the results of their paper. In this sample, the sample equity premium is eliminated using the PIPO discount factor when we set γ between 9 and 10.

[26] We use the procedure described in Jagannathan and Wang (1996).

do so for the other two models.

However, from an economic point of view, the PIPO stochastic discount factors gives less convincing evidence. The estimated value of γ is 4.98 and the estimated value of β is 0.3. In our view, the estimate of γ is not implausibly high, but the estimate of β is implausibly low. What generates these estimates? It is necessary to set γ to be approximately 5 in order to set the sample equity premium error to be zero (as in Table 6). Then, given that value of γ, the sample mean of $\widehat{C}_{\gamma,t}/\widehat{C}_{\gamma,t+1}$ is very high (approximately 4); it is then necessary to set β to an implausibly low number to satisfy the other two moment conditions (where β is identified, unlike the first moment condition).

It is also instructive to compare the results in Table 8 with those in Table 2 for the incomplete markets SDF. In Table 2, the incomplete markets SDF prices correctly the two assets with $\widehat{\gamma} = 1.4796$ and $\widehat{\beta} = 0.9903$. In Table 8, however, the performance of this SDF is much worse and the same pricing errors increase dramatically even in the presence of a small change in the estimated value of γ. Why? At $\widehat{\gamma} = 1.4796$ and $\widehat{\beta} = 0.9903$ (the values estimated in Table 2) the estimated pricing error $e_{mkt}(\widehat{\gamma}) = 0.0191167$ and the (implied) HJ-distance is 0.2524. At the values estimated in Table 8 (designed to explicitly minimize the HJ-distance) the estimated pricing error $e_{mkt}(\widehat{\gamma})$ is now 0.0191161 and the HJ-distance is 0.2447. Thus, the HJ-distance trades off (small) reductions in the equity premium pricing error with (large) increases in the Treasury bill return pricing errors. Why are reductions in the equity premium pricing error so heavily weighted? The HJ distance minimizes the maximum pricing error over portfolios, and this (after standardization) happens to be the equity premium error. In fact, as we have seen, it is relatively easy to find SDFs that zero out the Treasury bill return pricing errors, but it is rather hard to find SDFs that zero out the equity premium pricing error. It is impossible to find SDFs in these three models that zero out all pricing errors while also giving plausible parameter estimates.

The estimated β in the PIPO SDF is too low to be plausible. In Figure 7 we plot the HJ-distance for the PIPO and incomplete markets stochastic discount factors setting $\beta = 0.99$, a value that many economists would find plausible for quarterly data (see Gourinchas and Parker, 2000), while letting γ vary. It is evident that the PIPO SDF dominates the incomplete markets SDF throughout the parameter space of γ, in the sense of giving a lower HJ-distance for all values of γ.[27]

In Figure 8 the reverse exercise is considered. Here we set γ to a given value (between 0 and 6) and estimate β by GMM using the three moment conditions (1)-(3). We plot this estimate on the left y-axis. On the right y-axis we also plot the implied HJ-distance. This graph paints the same picture as the one above. The HJ-distance obtained using the PIPO discount factor is always lower than the HJ-distance obtained using the incomplete markets discount factor, and the "degree of implausibility" of the estimate of β is similar. That is, in the unconstrained case of Table 8, the incomplete market stochastic discount factor is throwing out a plausible estimate of β not because it is an economically more reasonable alternative than the PIPO, but because it is estimating a lower value of the coefficient of relative risk aversion γ.

We have not checked if these differences in model performance are statistically significant; we doubt that they are. However, it does seem reasonable to conclude that the PIPO stochastic discount factor is, at the least, no worse than the incomplete markets stochastic discount factor.

7. Conclusions

This paper makes two contributions. The first is theoretical. We consider a Pareto optimal allocation of resources in an economy in which agents are privately informed about their own skills and in which there are publicly observable aggregate shocks. We construct a representation for the shadow stochastic discount factor in terms of moments of the cross-sectional distribution of consumption. The representation is valid regardless of the stochastic process generating the individual-level shocks or the process generating the aggregate shocks. We also show that this representation is robust to a wide class of measurement error processes.

Our second contribution is to document the empirical relevance of this representation. We show that this shadow stochastic discount factor does not fare any worse empirically than leading consumption-based asset pricing models. Our results suggest that society-wide incentive costs play an important factor in the determination of asset prices.

[27] Note that $d^{PIPO}(0.99, \gamma)$ is minimized for a value of γ around 3, while $d^{IM}(0.99, \gamma)$ is minimized for a value of γ around 2.

References

[1] ALTUG, S., and MILLER, R. A. (1990), "Household Choices in Equilibrium," *Econometrica* 58: 543-570.

[2] ATKESON, A., and LUCAS, R. E. (1995), "Efficiency and Equality in a Simple Model of Efficient Unemployment Insurance," *Journal of Economic Theory* 66: 64-88.

[3] ATTANASIO, O., BANKS, J., and TANNER, S. (2002), "Asset Holding and Consumption Volatility," *Journal of Political Economy* 110, 771-792.

[4] ATTANASIO, O., BATTISTIN, E., and ICHIMURA, H. (2004), "What Really Happened to Consumption Inequality in the US?," NBER working paper 10338.

[5] ATTANASIO, O., and DAVIS, S. (1996), "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy* 104, 1227-1262.

[6] ATTANASIO, O., and WEBER, G. (1995), "Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey," *Journal of Political Economy* 103, 1121-1157.

[7] BALDUZZI, P., and YAO, T. (2004), "Testing Heterogenous Agent Models: An Alternative Aggregation Approach," Boston College working paper.

[8] BETSON, D. (1990), "Alternative Estimates of the Cost of Children from the 1980-1986 Consumer Expenditure Survey," US Department of Health and Human Services, Washington: D.C..

[9] BOLLINGER, C.R., and A. CHANDRA (2004), "Iatrogenic Specification Error: A Cautionary Tale of Cleaning Data", NBER Technical Working Paper T0289.

[10] BOSWORTH, B., BURTLESS, B., and SABELHAUS, J. (1991),"The Decline in Saving: Evidence From Household Surveys," *Brooking Papers on Economic Activity*, 183-241.

[11] BRAV, A., CONSTANTINIDES, G., and GECZY, C. (2002), "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy* 110, 793-824.

[12] CAMPBELL, J., and COCHRANE, J. (1999),"By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior," *Journal of Political Economy* 107, 205-251.

[13] CHIODA, L. (200), "Estimating Euler Equations with Measurement Error: a Nonparametric Approach", Princeton University, unpublished manuscript.

[14] CITRO, C., and MICHAELS, R, eds. (1995), *Measuring Poverty: A New Approach (1995),* National Academy Press, Washington: D.C.

[15] COCHRANE, J. (1991), "A Simple Test of Consumption Insurance," *Journal of Political Economy* 96, 893-920.

[16] COCHRANE, J. (2001), *Asset Pricing,* Princeton University Press, Princeton, NJ.

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

[18] EFRON, B., and TIBSHIRANI, R. (1998), *An Introduction to the Bootstrap,* Chapman and Hall, New York City, NY.

[19] GOLOSOV, M., KOCHERLAKOTA, N., and TSYVINSKI, A. (2003), "Optimal Indirect and Capital Taxation," *Review of Economic Studies* 70, 569-587.

[20] HALL, R. (1988), "Intertemporal Substitution in Consumption," *Journal of Political Economy* 96, 339-357.

[21] HANSEN, L., and JAGANNATHAN, R. (1997), "Assessing Specification Errors in Stochastic Discount Factor Models," *Journal of Finance* 52, 557-590.

[22] HANSEN, L., and SINGLETON, K. (1982),"Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models," *Econometrica* 50, 1269-1288.

[23] HOROWITZ, J. (2001),"The Bootstrap in Econometrics," in *Handbook of Econometrics* , Vol. 5, J.J. Heckman and E.E. Leamer, eds., Elsevier Science B.V., Ch. 52, p. 3159-3228.

[24] JAGANNATHAN R., and Z. WANG (1996), "The Conditional CAPM and the Cross-Section of Expected Returns," *Journal of Finance* 51, 3-53.

[25] KOCHERLAKOTA, N. (1990), "On Tests of Representative Consumer Asset Pricing Models," *Journal of Monetary Economics* 26, 285-304.

[26] KOCHERLAKOTA, N. (1996), "The Equity Premium: It's Still A Puzzle," *Journal of Economic Literature* 34*,* 42-71.

[27] KOCHERLAKOTA, N. (1998), "The Effects of Moral Hazard on Asset Prices When Financial Markets Are Complete," *Journal of Monetary Economics* 41, 39-56.

[28] KOCHERLAKOTA, N. (2005), "Zero Expected Wealth Taxes: A Mirrlees Approach to Optimal Dynamic Taxation," *Econometrica* 73, 1587-1622.

[29] LIGON, E. (1996), "Risk-Sharing and Information: Theory and Measurement in Village Economies," unpublished working paper, UC-Berkeley.

[30] LIGON, E. (1998), "Risk-Sharing and Information in Village Economies," *Review of Economic Studies* 65, 847-864.

[31] LIGON, E. (2005), "Formal Markets and Informal Insurance, " forthcoming, *International Review of Law and Economics.*

[32] LIGHT, A., and K. MCGARRY (2004), "Why Parents Play Favorites: Explanations for Unequal Bequests," *American Economic Review* 94(5), 1669-81.

[33] LUSTIG, H. (2002), "The Market Price of Risk and the Wealth Distribution," The University of Chicago working paper.

[34] MANKIW, G., and S. ZELDES (1991), "The Consumption of Stockholders and Non-Stockholdres," *Journal of Financial Economics* 29, 97-112.

[35] MCGARRY, K., and R. SCHOENI (1995), "Transfer Behavior in the Health and Retirement Study: Measurement and the Redistribution of Resources within the Family," *Journal of Human Resources* 30, Special Issue on the Health and Retirement Study: Data Quality and Early Results, S184-S226.

[36] MEGHIR, C., and PISTAFERRI, L., "Income Variance Dynamics and Heterogeneity," *Econometrica* 72, 1-32.

[37] MEHRA, R., and PRESCOTT, E., "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15, 145-161.

[38] MIRRLEES, J. (1971), "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies,* 38, 175-208.

[39] PHELAN, C., (1994), "Incentives, Insurance, and the Variability of Consumption and Leisure," *Journal of Economic Dynamics and Control* 18, 581-599.

[40] PHELAN, C., (1998), "On the Long-Run Implications of Repeated Moral Hazard," *Journal of Economic Theory* 79, 174-191.

[41] PHELAN, C., and R. TOWNSEND (1991), "Computing Multi-Period, Information-Constrained Optima," *Review of Economic Studies* 58, 853-81.

[42] ROGERSON, W. (1985), "Repeated Moral Hazard," *Econometrica,* 53, 69-76.

[43] SEMENOV, A. (2004), "High-Order Consumption Moments and Asset Pricing," York University working paper.

[44] STORESLETTEN, K., TELMER, C., and YARON, A. (2001), "How Important are Idiosyncratic Shocks? Evidence from Labor Supply," *American Economic Review Papers and Proceedings* 91, 413-417.

[45] VISSING-JORGENSEN, A. (2002), "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy* 110, 825-853.

Appendix

1. Sample selection and consumption definition

Each annual tape of the CEX contains four groups of core files, MTAB, ITAB, FMLY, and MEMB. There are also many auxiliary files, which are not used here. The FMLY, MEMB, MTAB, and ITAB files are organized by the calendar quarter of the year in which the data are collected. The FMLY files contain household characteristics, income, and summary level expenditures; the MEMB files contain member characteristics and income data; the MTAB files contain expenditures organized on a monthly basis at the Universal Classification Code (UCC) level; and the ITAB files contain income data converted to a monthly time frame and assigned to UCCs.

There are five quarterly data sets for each of these files. For example, in the 1980 tape, there are files running from the first quarter of 1980 through the first quarter of 1981. This is for the purpose of allowing computation of calendar year statistics. In fact, the MTAB 1980:Q1 file, say, contains expenditure information as reported by households interviewed in the first quarter of 1980. Since households report data for the previous three months, the 1980:Q1 file, say, has expenditure data from October 1979 to February 1980. The 1981:Q1 file is included in the 1980 tape because it contains information that spans the last three months of 1980. With just two exceptions (discussed below), the Y:Q1 file contained in the Y−1 tape is identical to the Y:Q1 file contained in the Y tape. The FMLY file for a given quarter has one record per household. Similarly, the MEMB file for a given quarter has one record per household member.

The CEX has three "household tracking" problems, detailed as follows. In 1980-81 households that re-entered the survey after missing an interview were assigned a new ID. While this is probably a minor proportion of the whole sample, it is for all purposes not a problem in our context, given that we do not focus on the longitudinal aspect of the survey. In 1986 the CEX changed its sample design. The consequence of this is that the core 1986:Q1 files contained in the 1985 tape survey different households than the core 1986:Q1 files contained in the 1986 tape. Indeed, issuing of household IDs starts from scratch beginning with the 1986 tape. Again, this is not a problem for us. We use both "samples", and so end up with a size that is larger than usual. Another sample design change occurs in 1996, but in this case some of the households that are surveyed in the core 1996:Q1 files contained in the 1995 tape appear also in the core 1996:Q1 files of the 1996 tape, although with the same ID. Of course, we eliminate the duplicates.

We use the MTAB files from 1980:Q1 throughout 2004:Q1 to create monthly expenditure records for each household ever surveyed in the CEX. Since households report data for at most four quarters, there are between 3 and 12 observations per household. We merge this information with household characteristics from the FMLY file. The file so compiled contains 1,848,339 observations (where, to reiterate, each observation is a household/month data point).

Our measure of nondurable consumption is as in Attanasio and Weber (1995), and it is the sum of the following items (in parenthesis the UCC codes):[28] food at home (790220, 790230, 190904), food away from home (190901-190903, 790410, 790430, 800700), alcohol (200900, 790310, 790320, 790420), apparel and footwear (360110-420120), clothing services (440110-440140, 440210, 440900), tobacco (630110, 630210), heating (250111-250904), Utilities: gas (260211-260214), Utilities: electricity (260111-260114), Utilities: Water and sewerage (270211-270214, 270411-270414, 270901-270904), public transportation (530110-530902), vehicle expenses (520110-520907), gasoline and oil (470111-470212), vehicle maintenance and repairs (470220-490900), parking fees (220901-220902), newspapers and magazines (590110-590212), books (590220-590230), club membership fees (620110-620115), ticket admissions (620121-620310), miscellaneous entertainment expenses (610900, 620330-620926), home rent (210110-210902, 800710, 350110), home insurance (220111-220122), home maintenance and repairs (230111-230902, 330511, 340914, 790600), telephone and cable (270000-270104, 270310), babysitting (340210-340212), domestic services (340310-340420), other home services (340510-340530, 340906, 340911-340912, 340915), personal care (650110-650900), rentals (340610-340905, 340907-340908, 440150).

Here is a description of our sample selection. As said, we start with 1,848,339 monthly observations. We drop 164,125 observations for which our measure of total nondurable consumption is missing or zero. We further drop 4,258 observations for households that report zero food spending (at home and away from home) during an entire interview (3-month period). Households interviewed in a certain month are supposed to report consumption data only for the previous three months. We drop 4,116 observations reporting data for the same month in which they are interviewed. We next eliminate 398,047 observations corresponding

[28] Erich Battistin at IFS kindly provided assistance in replicating Attanasio and Weber's aggregation procedures.

to households that are classified as ”incomplete income respondents” or report less than three months of data for a given interview. At this point, we aggregate monthly data in (overlapping) quarters, indexed by the last month in the quarter. The resulting sample has 425,931 quarterly observations, corresponding to 147,412 households. We drop 9,635 observations corresponding to households that jump interviews (i.e., exit and re-enter the survey), and 5,498 observations corresponding to households living in college dorms. We end up with a final sample of 410,798 quarterly observations, or 140,364 households. The average number of household per (overlapping) quarter is 1,412. The median is 1,331.

Table 1

Descriptive Statistics

Panel A: Household data from the CEX

	1980	1983	1986	1989	1992	1995	1998	2001	2004
Age	46.17	46.10	46.44	47.25	47.70	48.11	48.50	48.62	49.07
Family size	2.72	2.61	2.62	2.60	2.55	2.58	2.47	2.57	2.54
# of kids	0.79	0.72	0.75	0.72	0.71	0.73	0.67	0.70	0.70
Proportion some college+	0.40	0.44	0.44	0.47	0.49	0.49	0.54	0.56	0.58
Annual before tax income	21599	23738	24165	25926	25023	25017	26607	28109	28437
Stocks	1647	2752	3193	3288	3565	4276	14240	10743	17321
Adult equiv. quarterly cons.	1731	1673	1665	1735	1639	1598	1649	1614	1578
Household quarterly cons.	2884	2705	2707	2803	2609	2555	2569	2586	2528
N	14,426	14,353	21,690	15,801	16,037	14,111	14,551	20,194	4,261

Note: Monetary variables are deflated by the CPI-U (1982-1984=100). The Adult equivalent quarterly consumption is also deaseasonalized as described in the text. Data on stocks (which also include the value of mutual funds, private bonds, government bonds and Treasury notes) refer to the 5th interview only; data on income refer to the 2nd and 5th interviews only.

Panel B: Time series data

	1980-82	1983-85	1986-88	1989-91	1992-94	1995-97	1998-00	2001-04	1980-04
r^f (mean)	0.60	1.23	0.62	0.62	0.14	0.65	0.60	0.05	0.60
r^f (st.dev.)	0.81	0.43	0.61	0.44	0.29	0.33	0.36	0.63	0.64
r^m (mean)	1.56	3.54	3.00	2.76	1.33	5.75	2.39	-0.13	2.51
r^m (st.dev.)	9.95	6.22	10.37	7.45	3.25	4.69	7.65	8.50	7.68
$corr\left(r^f, r^m\right)$	0.0281	0.0532	0.1712	0.7005	0.0676	0.2165	0.1250	0.1217	0.1948
$corr\left(r^f, r^f_{-3}\right)$	0.6602	-0.0250	0.3681	-0.3230	-0.3103	-0.3292	-0.0924	-0.0912	0.4491

Note: data refer to the period 1980:3-2004:2. The table reports (overlapping) quarterly returns.

Table 2

Expected Return to the Treasury Bill

	Pareto-optimal	Incomplete markets	Representative agent
γ	2.0165	1.4796	1.5529
(s.e.)	[3.6762]	[0.6040]	[1.4666]
β	0.9879	0.9903	0.9910
(s.e.)	[0.0454]	[0.0036]	[0.0040]
$\overline{e}_{b1}$	$3.13e-007$	$9.43e-006$	$1.54e-007$
$\overline{e}_{b2}$	$-3.15e-007$	$-9.34e-006$	$-1.25e-007$

Note: In this table, we report the estimates and standard errors associated with estimating β and γ using the restrictions that the pricing errors

$$e_{b1}(\gamma) = (m_t(\beta,\gamma)R_t^f - 1) \qquad e_{b2}(\gamma) = (m_t(\beta,\gamma)R_t^f - 1)R_{t-3}^f$$

have expectation zero. The rows $\overline{e}_{b1}$ and $\overline{e}_{b2}$ report the sample means of the pricing errors at the estimated value of β and γ.

Table 3

The Unexplained Equity Premium: PIPO SDF

γ	Unexplained premium	Time-series 95% C.I. Lower bound	Upper bound
0	0.0192	0.0037	0.0347
1	0.0191	0.0038	0.0346
2	0.0187	0.0033	0.0341
3	0.0178	0.0021	0.0335
4	0.0158	−0.0054	0.0370
5	0.0010	−0.0612	0.0632
6	−0.0858	−0.3443	0.1727
7	−0.4854	−1.5466	0.5758
8	−2.1221	−6.2882	2.0440
9	−8.4229	−24.2229	7.3771
10	−31.8337	−90.3407	26.6733

Note: The unexplained premium is defined as:

$$e_{mkt,t}(\gamma) = T^{-1} \sum_{t=1}^{T} \frac{N_{t-3}^{-1} \sum_{i=1}^{N_{t-3}} c_{it-3}^{\gamma}}{N_t^{-1} \sum_{i=1}^{N_t} c_{it}^{\gamma}} \left(r_t^m - r_t^f \right)$$

and is expressed in percentage form. The time-series C.I. bounds are computed as: $e_{mkt,t}(\gamma) \pm 1.96 \times \sigma(\gamma)$, where $\sigma(\gamma)$ is the standard error of $e_{mkt,t}(\gamma)$ obtained applying a Hansen-Hodrick adjustment.

Table 4

The Unexplained Equity Premium: Incomplete Markets SDF

γ	Unexplained premium	Time-series 95% C.I. Lower bound	Upper bound
0	0.0192	0.0037	0.0347
1	0.0191	0.0037	0.0345
2	0.0192	0.0039	0.0347
3	0.0221	0.0051	0.0391
4	0.0500	0.0089	0.0911
5	0.2158	0.0149	0.4167
6	1.0582	−0.0016	2.1180
7	5.0250	−0.2422	10.2922
8	22.9024	−1.9590	47.7207
9	101.4028	−11.6499	214.4555
10	440.8021	−64.0540	945.6582

Note: The unexplained premium is defined as:

$$e_{mkt,t}(\gamma) = T^{-1} \sum_{t=1}^{T} \frac{N_t^{-1} \sum_{i=1}^{N_t} c_{it}^{-\gamma}}{N_{t-3}^{-1} \sum_{i=1}^{N_{t-3}} c_{it-3}^{-\gamma}} \left(r_t^m - r_t^f \right)$$

and is expressed in percentage form. The time-series C.I. bounds are computed as: $e_{mkt,t}(\gamma) \pm 1.96 \times \sigma(\gamma)$, where $\sigma(\gamma)$ is the standard deviation of $e_{mkt,t}(\gamma)$ obtained applying a Hansen-Hodrick adjustment.

Table 5

The Unexplained Equity Premium: Representative Agent SDF

γ	Unexplained premium	Time-series 95% C.I. Lower bound	Upper bound
0	0.0192	0.0037	0.0347
1	0.0191	0.0037	0.0345
2	0.0190	0.0036	0.0344
3	0.0190	0.0036	0.0344
4	0.0189	0.0034	0.0346
5	0.0189	0.0034	0.0344
6	0.0189	0.0034	0.0344
7	0.0189	0.0033	0.0345
8	0.0189	0.0032	0.0346
9	0.0189	0.0032	0.0346
10	0.0189	0.0031	0.0347

Note: The unexplained premium is defined as:

$$e_{mkt,t}(\gamma) = T^{-1} \sum_{t=1}^{T} \left(\frac{N_t^{-1} \sum_{i=1}^{N_t} c_{it}}{N_{t-3}^{-1} \sum_{i=1}^{N_{t-3}} c_{it-3}} \right)^{-\gamma} \left(r_t^m - r_t^f \right)$$

and is expressed in percentage form. The time-series C.I. bounds are computed as: $e_{mkt,t}(\gamma) \pm 1.96 \times \sigma(\gamma)$, where $\sigma(\gamma)$ is the standard deviation of $e_{mkt,t}(\gamma)$ obtained applying a Hansen-Hodrick adjustment.

Table 6

The Equity Premium

	Pareto-optimal	Incomplete markets	Representative agent
γ	5.0280	1.4104	57.5941
(s.e.)	[0.9408]	$[n.a.]$	[22.9103]
$\overline{e}_{mkt,t}$	$-1.73e-008$	0.0191	$-8.01e-011$

Note: In this table, we report the estimates and standard errors associated with estimating γ using the restriction that $e_{mkt,t}(\gamma)$ has expectation zero, where

$$e_{mkt,t}(\gamma) = m_t(\gamma)\left(r_t^m - r_t^f\right)$$

The row $\overline{e}_{mkt}$ reports the sample mean of the pricing error at the estimated value of γ.

Table 7

The Unexplained Equity Premium for Financial Market Participants

γ	Stockholders	Stockholders and bondholders				
		$A > 0$	$A > 1,000$	$A > 2,000$	$A > 5,000$	$A > 10,000$
0	0.0167	0.0163	0.0167	0.0171	0.0173	0.0168
1	0.0167	0.0168	0.0166	0.0174	0.0179	0.0172
2	0.0174	0.0174	0.0173	0.0188	0.0196	0.0184
3	0.0207	0.0193	0.0192	0.0218	0.0234	0.0218
4	0.0430	0.0234	0.0231	0.0271	0.0312	0.0301
5	0.1903	0.0314	0.0305	0.0363	0.0467	0.0506
6	1.1021	0.0466	0.0444	0.0528	0.0780	0.1058
7	6.3416	0.0753	0.0709	0.0823	0.1439	0.2642
8	34.6687	0.1298	0.1226	0.1391	0.2861	0.7432
9	180.7845	0.2346	0.2248	0.2470	0.5987	2.2413
10	908.0264	0.4373	0.4289	0.4569	1.2946	7.0330

Note: The unexplained premium is defined as:

$$e_{mkt,t}(\gamma) = T^{-1} \sum_{t=1}^{T} \frac{N_t^{-1} \sum_{i=1}^{N_t} c_{it}^{-\gamma}}{N_{t-3}^{-1} \sum_{i=1}^{N_{t-3}} c_{it-3}^{-\gamma}} \left(r_t^m - r_t^f \right)$$

and is expressed in percentage form. A denotes the amount held in stocks.

Table 8

The Equity Premium and the Treasury Bill Return

	Hansen-Jagannathan weighting matrix		
	Pareto optimal	Incomplete markets	Representative agent
γ	4.9844	1.4283	7.9910
(s.e.)	[0.9500]	[1.5131]	[9.2405]
β	0.3074	1.0505	1.0216
(s.e.)	[0.2467]	[0.0238]	[0.0523]
$\overline{e}_{mkt}$	0.0015	0.0191	0.0189
$\overline{e}_{b1}$	-0.0064	0.0606	0.0595
$\overline{e}_{b2}$	-0.0071	0.0609	0.0598
J	0.84	5.96	5.60
(p-value)	[0.3580]	[0.0146]	[0.0180]
HJ-distance	0.1039	0.2447	0.2415
(p-value)	[0.0103]	[0.8763]	[0.8419]

Note: This table contains the results of estimating β and γ using the restrictions that

$$e_{mkt}(\gamma) = m_t(1,\gamma)(R_t^{mkt} - R_t^f) \quad e_{b1}(\gamma) = (m_t(\beta,\gamma)R_t^f - 1) \quad e_{b2}(\gamma) = (m_t(\beta,\gamma)R_t^f - 1)R_{t-3}^f$$

have expectation zero. The rows $\overline{e}_{mkt}$, $\overline{e}_{b1}$ and $\overline{e}_{b2}$ report the sample means of the pricing errors at the estimated value of γ and β. The J-statistic is constructed using the formula in Cochrane (2001, p. 204). The HJ-distance is the square root of the minimized objective. Its p-value is computed using the procedure described in Jagannathan and Wang (1996).

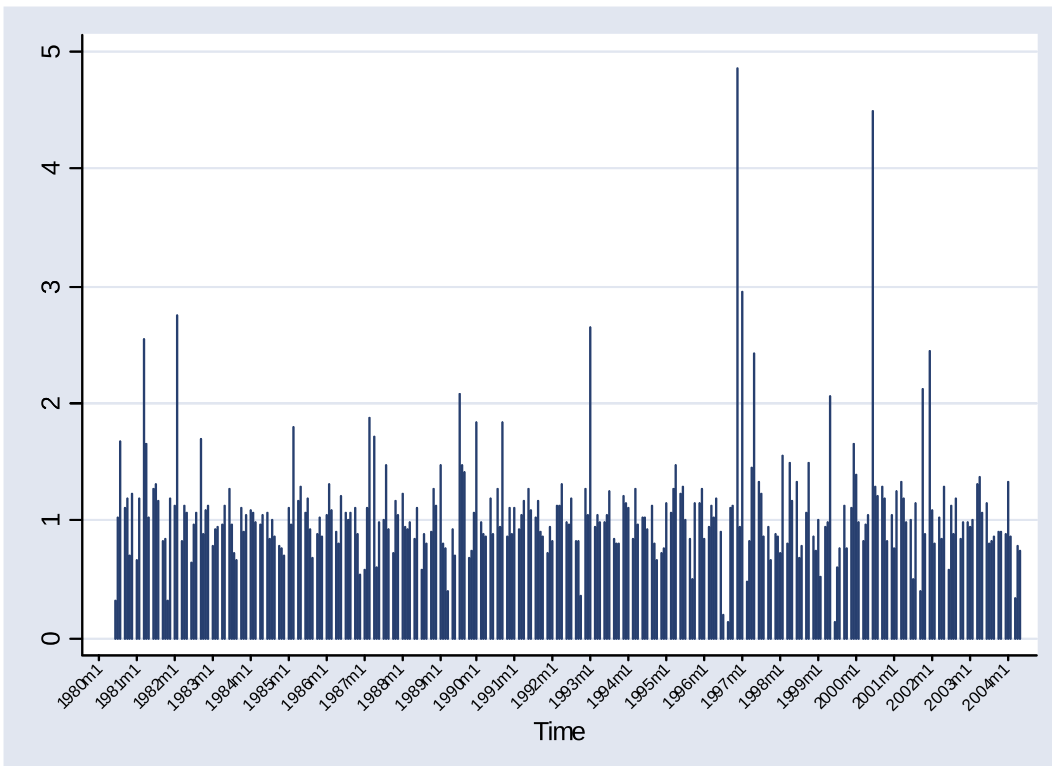

Figure 1: The PIPO stochastic discount factor (with $\beta = 1$ and $\gamma = 2$).

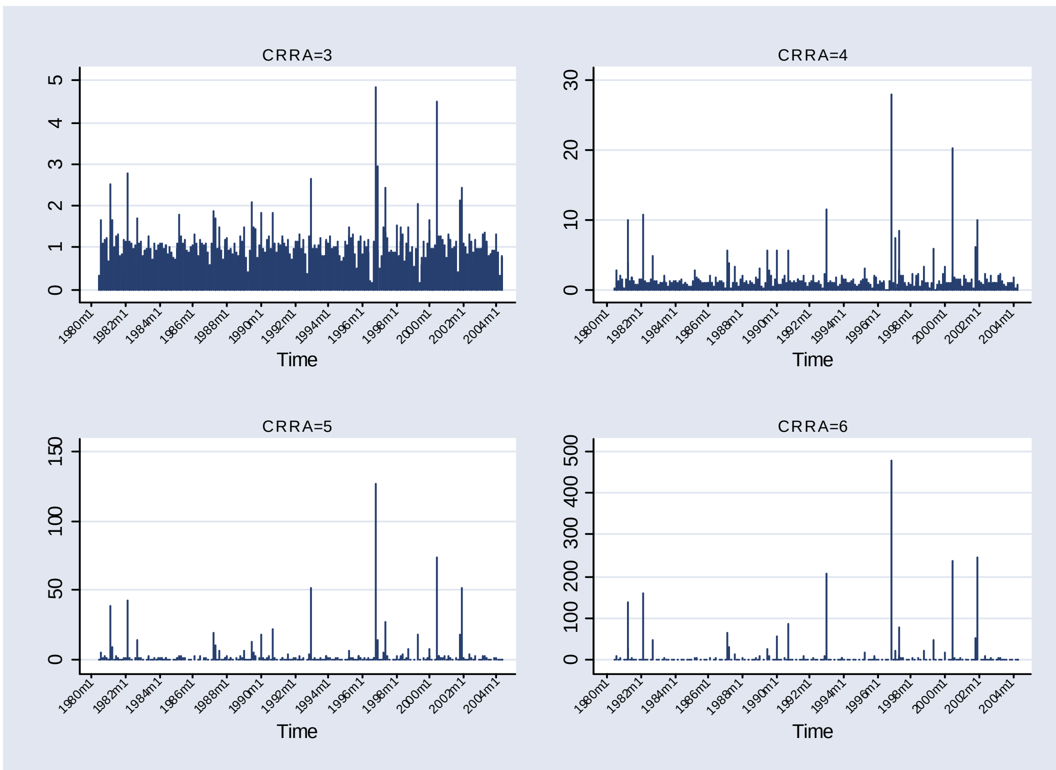

Figure 2: The PIPO stochastic discount factor (with $\beta = 1$ and $\gamma = \{3, 4, 5, 6\}$).

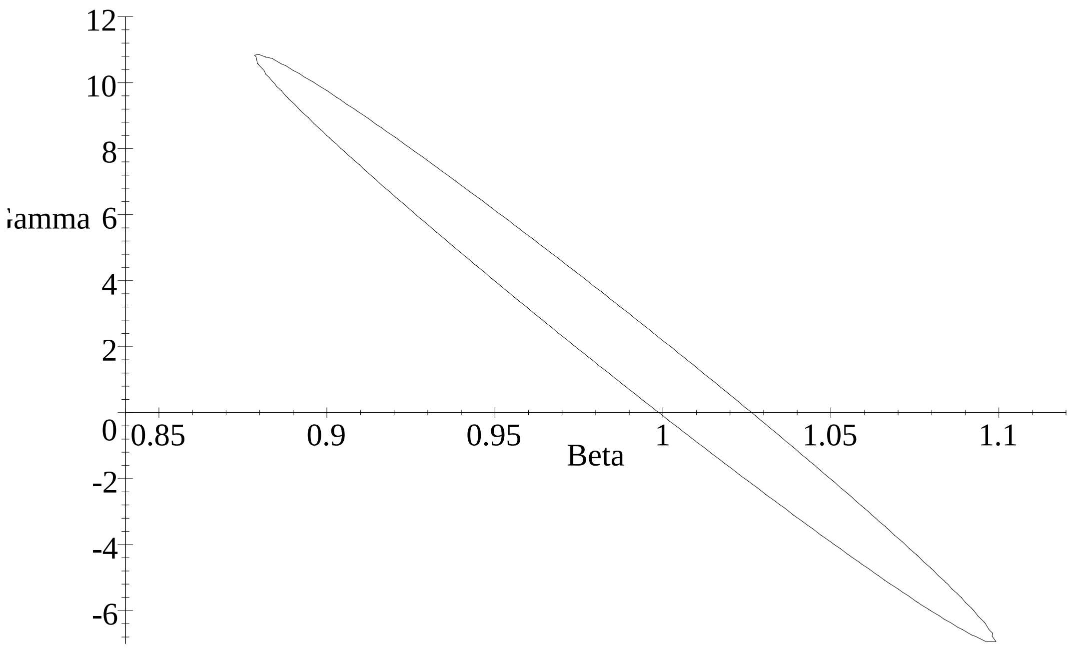

Figure 3: The joint confidence interval (β, γ) in the PIPO model (from Table 2).

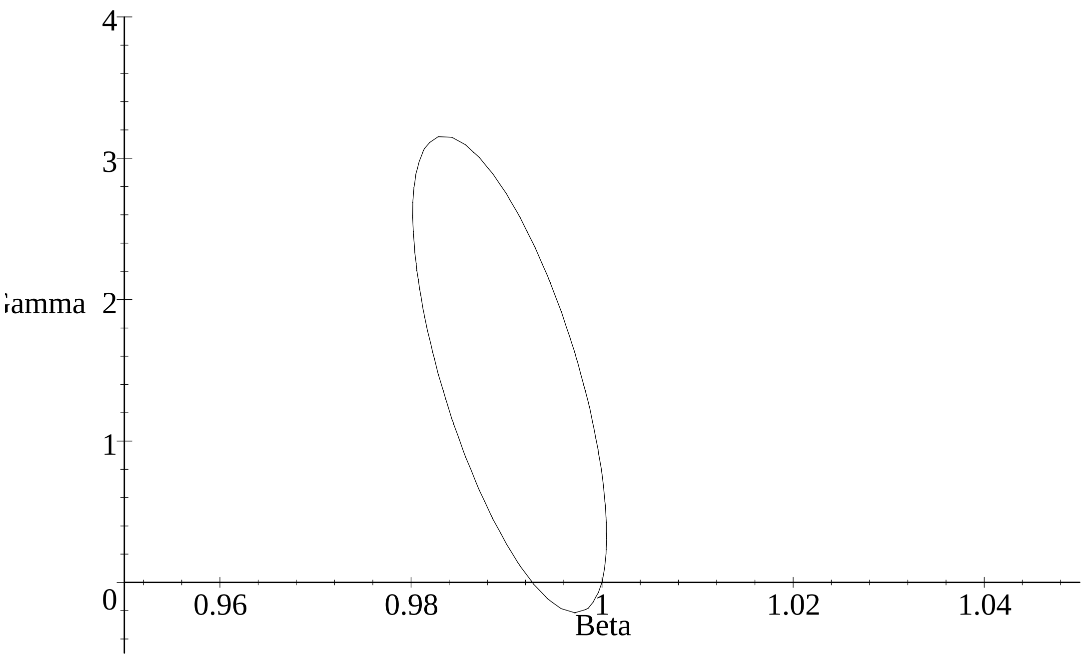

Figure 4: The joint confidence interval (β, γ) for the IM model (from Table 2).

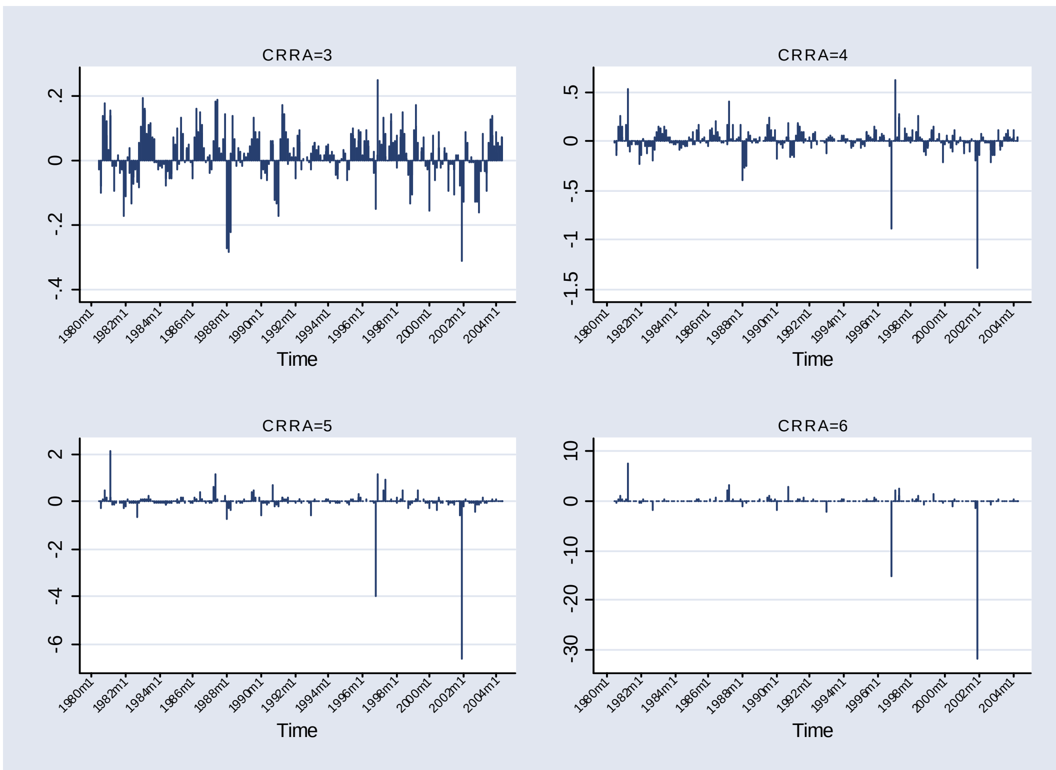

Figure 5: The PIPO unexplained equity premium (with $\beta = 1$ and $\gamma = \{3, 4, 5, 6\}$).

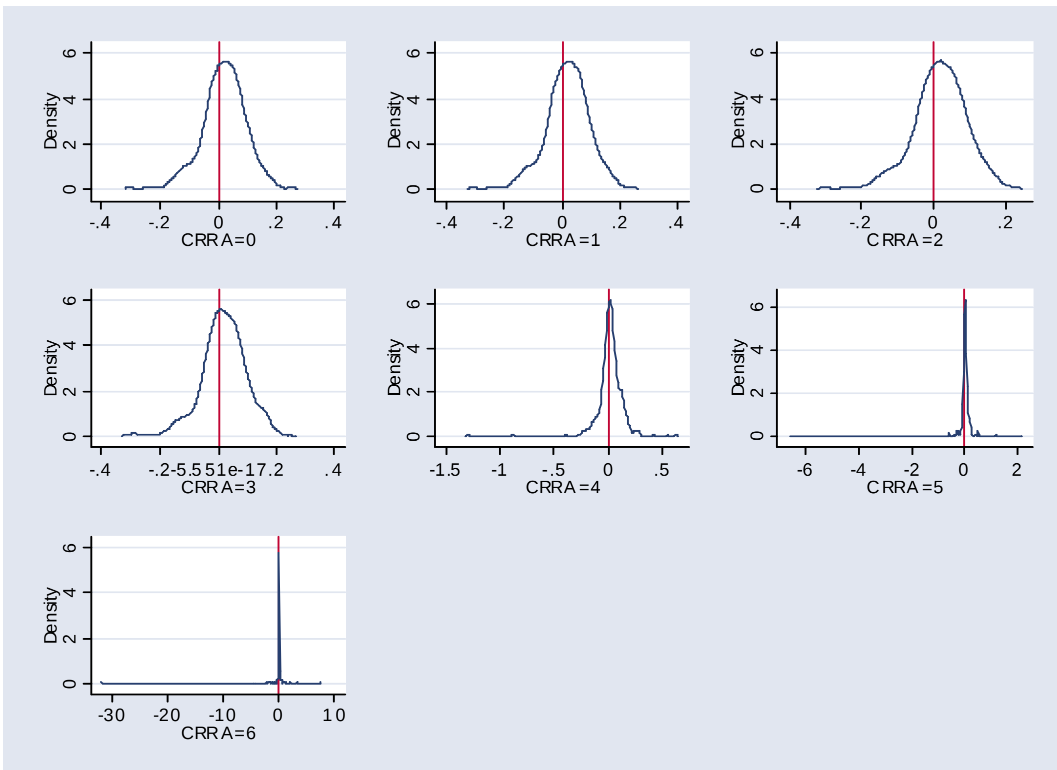

Figure 6: Kernel density estimation of the PIPO unexplained equity premium for various values of γ.

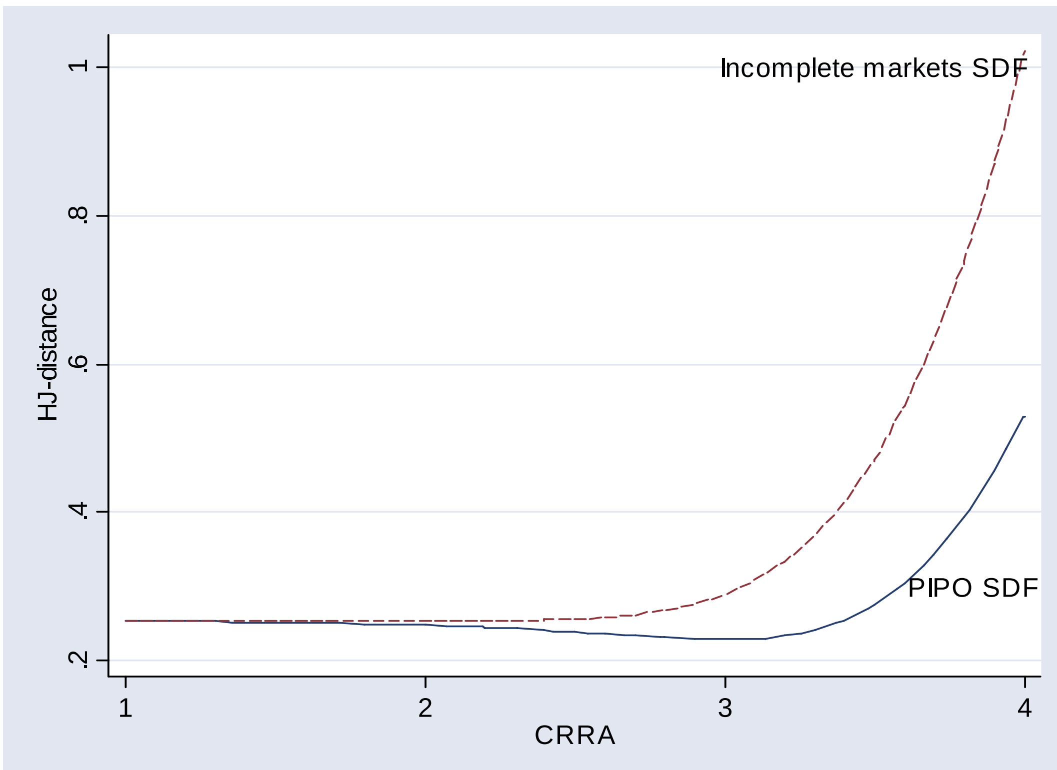

Figure 7: The Hansen-Jagannathan distance setting $\beta = 0.99$.

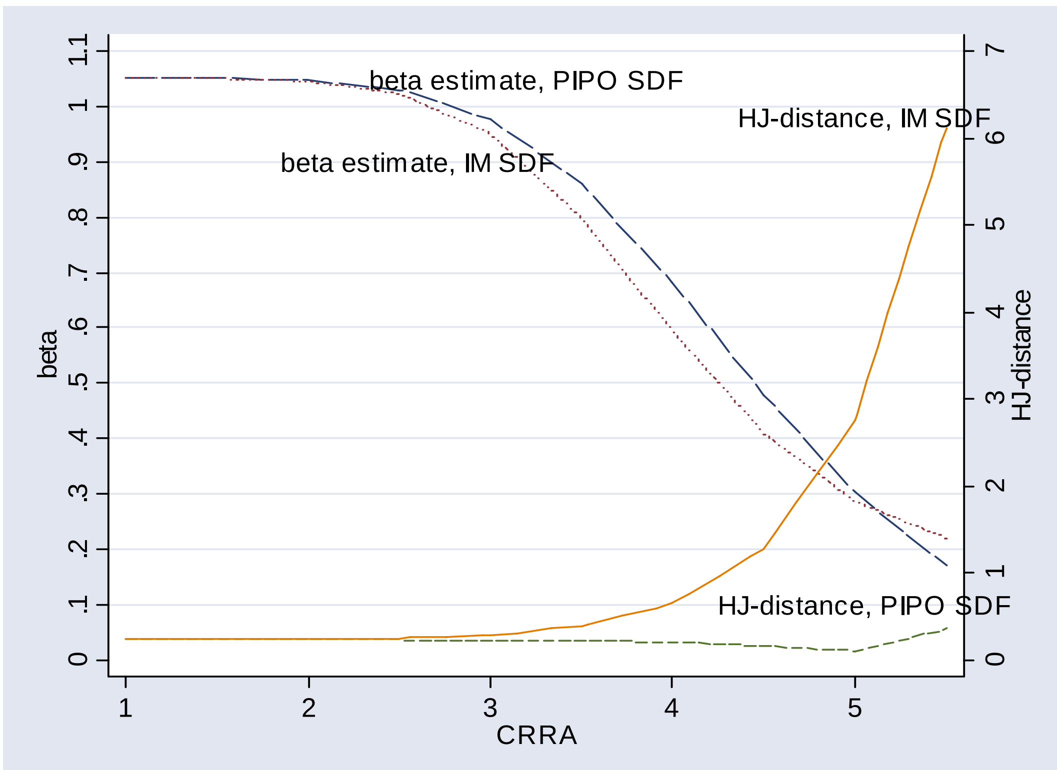

Figure 8: GMM estimates of β and the Hansen-Jagannathan distance for various values of γ.