

September 2004

Asset Pricing Implications of Pareto Optimality with Private Information

Narayana R. Kocherlakota*

Stanford University, Federal Reserve Bank of Minneapolis, and NBER

Luigi Pistaferri

Stanford University

ABSTRACT

In this paper, we consider a dynamic economy in which the agents in the economy 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 γ th moment of the consumption distribution, where γ is the coefficient of relative risk aversion. This implication can be evaluated using a time series of cross sections. We use data from the consumer expenditure survey (CEX) and show that the new stochastic discount factor performs better than existing stochastic discount factors at rationalizing the equity premium. However, its ability to *simultaneously* explain the equity premium and the expected return to the Treasury bill is about the same as existing discount factors.

*Kocherlakota acknowledges the support of NSF 0305833. 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, and the National Bureau of Economic Research. We thank Erich Battistin for providing us with the CEX data, and we thank Adriano Rampini for his 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

The benchmark macroeconomic model of asset pricing assumes that people are fully insured against idiosyncratic shocks. Under this assumption, the marginal investor is a "representative" agent who consumes per-capita quantities. The implications of the representative agent model have been tested in a variety of ways (including by, among many others, Hansen and Singleton (1982) and Mehra and Prescott (1985)). The model has generally not fared well, without adopting somewhat extreme formulations of preferences for the representative agent, such as the high degree of external habit persistence assumed by Campbell and Cochrane (1999).

Of course, there is a great deal of evidence that the allocation of consumption in the United States is such that individuals are *not* fully insured against individual-specific shocks. For example, Cochrane (1991) documents that individual consumption falls as a result of unemployment shocks. This lack of full insurance is not all that surprising. Consider a person who is fully insured against the risk of becoming unemployed. He is unlikely to exert a great deal of effort to avoid becoming unemployed. Nor is he likely to exert a great deal of effort to find a job once unemployed. More generally, imperfect insurance provides incentives to individuals whenever effort choices are hard to monitor or enforce.

In this paper, we present a new model of asset pricing that is based on this incentive consideration. Like the benchmark representative agent model, we assume that the equilibrium allocation of consumption is Pareto optimal. However, 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 moments¹ of the cross-sectional distribution of consumption.

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. We impose *no* restrictions on 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 ($c^{1-\gamma}$) 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 γ th 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 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

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

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.

Using a similar theoretical approach, we construct 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 $-\gamma$ th moment of the cross-sectional distribution of consumption. We derive this stochastic discount factor by integrating over the intertemporal Euler equations of the investors 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:

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

In this formula, C_{1t}^* is the first moment of the cross-sectional distribution of consumption. Hence, as in the two-period model of Kocherlakota (1998), the complete markets discount factor and the PIPO discount factor coincide when $\gamma = 1$. (Of course, all three stochastic

discount factors coincide when $\gamma = 0$.) Both of the alternative stochastic discount factors are also robust to the kind of measurement error described above.

It is important to stress that all three discount factors are valid regardless of the stochastic process generating skills or productivity shocks. Of course, the structure of markets and information imposes a precise mapping between the data generation process for skills and the time-series behavior of the cross-sectional distribution of consumption. But there is a great deal of empirical debate about the time series behavior of wages (see Storesletten, Telmer, and Yaron (2001) and Meghir and Pistaferri (2004)). We regard it as a great strength of our empirical approach that our results are valid regardless of how this empirical debate is eventually resolved.

We then turn to an empirical comparison of the three discount factors. 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. For each month between 1980 and 1998, we construct the three stochastic discount factors using data from the Consumer Expenditure Survey (CEX) from 1980 to 1998. The data is fundamentally overlapping: for each month, we measure household consumption as total nondurable consumption over the preceding quarter. We use these data to construct measures of the quarterly stochastic discount factors.

We then apply the Generalized Method of Moments to assess the three discount factors in terms of two types of implications. The first is the equity premium. An arbitrary stochastic discount factor m_t should be consistent with the population restriction:

$$(1) \quad E\{m_t(R_t^{mkt} - R_t^f)\} = 0$$

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. The second is the intertemporal variation in the Treasury bill return. A stochastic discount factor m_t should be consistent with the two population restrictions:

$$(2) \quad E\{(m_t R_t^f - 1)\} = 0$$

$$(3) \quad E\{(m_t R_t^f - 1)R_{t-1}^f\} = 0$$

Here, it is important to note that the Treasury bill return is highly autocorrelated, so that the two restrictions are both informative. Given our short data set, the predictability of stock returns is too small to be used in a similar fashion.

We chose these restrictions because they are much studied in the macroeconomics literature. The restriction (1) assesses the extent to which a candidate discount factor can explain the difference between the stock market and Treasury bill returns. As Kocherlakota (1996) argues, (1) is simply a robust re-statement of the equity premium puzzle originally stated by Mehra and Prescott (1985). Restrictions (2) and (3) assess 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 these kinds of restrictions.

Our empirical results are as follows. We find that if we set γ near 9, the PIPO stochastic discount factor is able to set the sample analog of (1) to zero. There is no such specification of γ for the other two discount factors. However, once we allow for sampling error, the sample equity premium is consistent with the incomplete markets and PIPO stochastic discount factors even for γ 's as low as 3 or 4.

Next, we turn to the Treasury bill returns data. For all three discount factors, there exist plausible specifications of the preference parameters (β, γ) that zero out the sample analogs of (2) and (3). The resulting estimate of γ for the PIPO SDF is about 3; the bootstrap standard error is around 1.5. The resulting estimates of γ are highly imprecise for the complete markets and incomplete markets stochastic discount factors.

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 find that the Treasury bill returns are highly informative statistically relative to stock returns. Hence, in the joint estimation, the discount factors are estimated in such a way so as to zero out the sample versions of (2) and (3). In all three cases, the resulting estimated stochastic discount factors are unable to explain any of the sample equity premium: the sample estimates of (1) are the same as the mean equity premium. In this sense, the equity premium remains very much a puzzle.

2. Prior Literature

Ligon (1998) tests the risk-sharing implications of Pareto optimality with moral hazard. His approach is as follows. He uses consumption data from South Indian villages. 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 in the following

restriction.

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

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.

Obviously, our approach bears some similarity to Ligon's. But there are important differences. First, our theoretical analysis is much more general than his. We allow for aggregate shocks and do not assume that there is a risk-neutral outsider. Hence, as our empirical work requires, we are able to allow for non-trivial movements in 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. 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 much more robust to measurement error than is his.

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 discard all households from the CEX who have fewer than four quarterly observations (the data is constructed in

such a way that no household has more). They next construct an intertemporal marginal rate of substitution (IMRS) in a given quarter for each household with observations for that quarter and the prior one. Finally, they construct a theoretically valid stochastic discount factor by averaging these IMRS' across households (henceforth, we term this the average IMRS SDF).

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 equity premium. In contrast, BCG (2002) find that a Taylor series approximation to the average IMRS SDF does a good job of rationalizing the equity premium. BCG (implicitly) argue that Cogley's results are flawed because he approximates the average IMRS SDF using moments of the logged consumption distribution. 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.

Obviously, our work is novel 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. However, 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 (like the equity premium). In contrast, we can also explore the ability of the candidate models to account for the Treasury

bill return. Second, these other papers rely on Taylor series approximations of the relevant stochastic discount factors. In our view, the errors in these approximations can potentially lead to large biases in the results.

3. Environment

In this section, we describe the environment. The description is basically the same as that in Kocherlakota (2003).

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 $s \in S = \{1, 2, \dots, N\}$; 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.

The economy is initially endowed with K_1^* units of a capital good. There is a single consumption good that can be produced by capital and labor. The agents have identical preferences. A given agent has von-Neumann-Morgenstern preferences, and ranks deterministic sequences according to the function:

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

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.

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 2^Z that assigns positive probability to all elements 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. We assume that a law of large numbers applies: conditional on any z^T , the measure of agents in the population with type θ^T in Borel set B is given by $\mu_\Theta(B)$.

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 $\theta^t = (\theta_1, \dots, \theta_t)$ and the history of public shocks $z^t = (z_1, \dots, z_t)$. This implies that his choices in period t can only be a function of this history.

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

$$y_t = \phi_t(\theta^t, z^t)l_t$$

$$\phi_t : \Theta^t \times Z^t \rightarrow (0, \infty)$$

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

We assume that an agent's effective labor 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 working.

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 .²

The aggregate shocks also affect the aggregate production function as follows. We

²Attanasio 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.

define an allocation in this society to be (c, y, K) where:

$$K : Z^T \rightarrow R_+^{T+1}$$

$$c : S \times \Theta^T \times Z^T \rightarrow R_+^T$$

$$y : S \times \Theta^T \times Z^T \rightarrow [0, \bar{y}]$$

$$K_{t+1} \text{ is } z^t\text{-measurable}$$

$$(c_t, y_t) \text{ is } (s, \theta^t, z^t)\text{-measurable}$$

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 . K_{t+1} is the amount of the capital good carried over period t into period $(t + 1)$.

As mentioned above, we assume that the initial endowment of capital is K_1^* . We define an allocation (c, y, K) to be *feasible* if for all t, z^T :

$$C_t(z^T) + K_{t+1}(z^T) \leq F_t(K_t(z^T), Y_t(z^T), z^T) + (1 - \delta)K_t(z^T)$$

$$C_t(z^T) = \sum_{s \in S} \pi_s \int_{\theta^T \in \Theta^T} c_t(s, \theta^T, z^T) d\mu_\Theta$$

$$Y_t(z^T) = \sum_{s \in S} \pi_s \int_{\theta^T \in \Theta^T} y_t(s, \theta^T, z^T) d\mu_\Theta$$

$$K_1 \leq K_1^*$$

Here, $F_t : R_+^2 \times Z^T \rightarrow R_+$ is assumed to be strictly increasing, weakly concave, homogeneous of degree one, continuously differentiable with respect to its first two arguments, and z^t -measurable with respect to its last argument. Note that (C_t, Y_t) are z^t -measurable.

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 $\sigma(\theta^T, z^T) = (\theta^T, z^T)$. Let Σ be the set of all possible reporting strategies, and define:

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

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

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:

$$W(s, \sigma_{TT}; c, y) \geq W(s, \sigma; c, y) \text{ for all } s \text{ in } S \text{ and all } \sigma \text{ in } \Sigma$$

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, K) that solves the problem of maximizing the utility of agents with label $s = 1$ subject to (c, y, K) 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, \dots, 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.

4. An Intertemporal Characterization of Optimal Consumption Allocations

In this section, we provide a partial characterization of Pareto optima that is valid for *any* specification of the exogenous elements of the model $(\phi, F, \mu_\Theta, \mu_Z, \pi, u, v, \beta, Z, \Theta)$.

The key proposition is the following. It establishes that any Pareto optimal allocation must satisfy a particular first order condition as long as consumption and capital are uniformly bounded away from zero. (Note that the first order condition is valid even for (s, θ^t, z^t) such that $y_t(s, \theta^t, z^t)$ or $y_{t+1}(s, \theta^{t+1}, z^{t+1})$ are zero.) The first order condition is similar to that derived in Theorem 1 of Golosov, Kocherlakota, and Tsyvinski (2003) and in Rogerson (1985). The proof is equivalent to Proposition 1 in Kocherlakota (2003).

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

λ_{t+1}^* is z^{t+1} -measurable

$$\lambda_{t+1}^* = \beta [E\{u'(c_{t+1}^*)^{-1} | s, \theta^t, z^{t+1}\}]^{-1} / u'(c_t^*) \text{ a.e.}$$

$$E\{\lambda_{t+1}^* (1 - \delta + F_{K_{t+1}}^*) | z^t\} = 1 \text{ a.e.}$$

where $F_{K_{t+1}}^*(z^T) = F_K(K_{t+1}^*(z^T), Y_{t+1}^*(z^T), z^T)$ for all z^T .

Proof. In Kocherlakota (2003).

The content of this proposition is twofold. First, it establishes that:

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

is independent of (s, 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 θ^t , is independent of θ^t .

Second, the theorem establishes that this harmonic conditional mean is equal to the social discount factor (λ) between period t and period $(t + 1)$. The social discount factor can then be used to determine the optimal level of capital accumulation between period t and period $(t + 1)$.

Why does the relationship involve harmonic means, as opposed to arithmetic means? Assume Θ is finite, and assume that all agents are treated identically ex-ante (so that the optimal allocation does not depend on s). Then, think about the marginal benefit to the planner of getting ε extra units of per-capita consumption in history z^t . At first glance, one might think that the marginal benefit is proportional to the arithmetic mean of marginal

utilities:

$$\varepsilon \sum_{\theta^t \in \Theta^t} \mu_{\Theta}(\theta^t) u'(c_t(\theta^t, z^t))$$

(For the purposes of this intuitive argument, we write c_t as a function of (θ^t, z^t) , not (θ^T, z^T) .

This is without loss of generality, because c_t is (θ^t, z^t) -measurable.) But this implicitly assumes that each agent is receiving ε units of consumption regardless of history, which will typically violate incentive constraints.

Instead, the extra resources should be split so that each agent θ^t receives $\eta(\theta^t)$, where $\sum_{\theta^t \in \Theta^t} \eta(\theta^t) \mu_{\Theta}(\theta^t) = \varepsilon$ and for all $\theta^t, \theta^{t'}$:

$$u(c_t(\theta^t, z^t) + \eta(\theta^t)) - u(c_t(\theta^{t'}, z^t) + \eta(\theta^{t'})) = 0$$

or, using a first order approximation:

$$u'(c_t(\theta^t, z^t)) \eta(\theta^t) = u'(c_t(\theta^{t'}, z^t)) \eta(\theta^{t'}) = B$$

for some B . We can solve for B using:

$$\varepsilon = \sum_{\theta^t \in \Theta^t} B \mu_{\Theta}(\theta^t) / u'(c_t(\theta^t, z^t))$$

so that the marginal gain to the planner is given by:

$$\begin{aligned}
& \sum_{\theta^t \in \Theta^t} \mu_{\Theta}(\theta^t) u'(c_t(\theta^t, z^t)) \eta(\theta^t) \\
&= B \\
&= \varepsilon \left[\sum_{\theta^t \in \Theta^t} \mu_{\Theta}(\theta^t) / u'(c_t(\theta^t, z^t)) \right]^{-1}
\end{aligned}$$

The shadow value of resources in a history z^t is given by the harmonic mean of marginal utilities, not the arithmetic mean.³

5. Asset Pricing Implications

The first two sections mimic the analysis in Kocherlakota (2003). 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 γ th moment of the cross-sectional distribution of consumption. This result remains true even when consumption is mismeasured with possibly biased or persistent measurement errors.

A. Asset Pricing via the Shadow Social Discount Factor

Suppose that in the above environment, agents engage in sequential asset trade: specifically, in each period $t = 1, \dots, T - 1$, 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_{t+1}^m be the equilibrium gross return from period t to period $(t + 1)$ of asset m .

³Note that the proposition reduces to Theorem 1 of Golosov, Kocherlakota and Tsyvinski (2003) if Z is a singleton (so there are no aggregate shocks).

There are many ways to implement Pareto optimal allocations with private information. Golosov and Tsyvinski (2004) describe one based on Atkeson and Lucas (1994). In this implementation, agents sign long-lived contracts with intermediaries and then the intermediaries trade assets with one another. Kocherlakota (2003) describes another, in which agents directly trade assets with each other subject to wealth taxes. In both of these implementations, the social discount factor λ_{t+1}^* is in fact a valid asset pricing kernel for the pre-tax asset returns.

We do not take a stand on the nature of the implementation being used by agents. Instead, we simply assume that the allocation of consumption is Pareto optimal, and the social discount factor λ_{t+1}^* is a valid asset pricing kernel for all asset returns. More precisely, we assume that for any asset m :

$$1 = E\{R_{t+1}^m \lambda_{t+1}^* | z^t\} \text{ for all } t, z^t$$

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^*, K^*) be an optimal allocation, for $u(c_t) = c^{1-\gamma}/(1-\gamma)$. Define:

$$C_{\gamma t}^* = E\{c_t^{*\gamma} | z^t\}$$

to be the γ th moment of the cross-sectional distribution of consumption in public history z^t .

Proposition 1 implies that:

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

Taking reciprocals and integrating over (s, θ^t) on both sides, we get:

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

Then, again taking reciprocals we get:

$$\lambda_{t+1}^* = \beta C_{\gamma t}^* / C_{\gamma, t+1}^*$$

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

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

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 γ th moment of the cross-sectional distribution of consumption. Henceforth, we use the term Private

Information Pareto Optimal (PIPO) stochastic discount factor 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.

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 typically 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.

In particular, let (c^*, y^*, K^*) be a socially optimal allocation, and suppose λ^* is a valid stochastic discount factor. We allow c^* to be measured with error as follows. Let (v_1, v_2, \dots, v_T) be a collection of random variables with joint probability measure μ_v over the Borel sets in R_+^T . At the beginning of period 1, after the public shock sequence z^T is drawn, a realization v^T is drawn according to μ_v for each agent; conditional on z^T , the draws of v^T and θ^T are independent from each other and are independent across agents. Note too that v^T is independent of z^T (because it is drawn from μ_v for all z^T); however, the measurement error is allowed to have arbitrary serial correlation.

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

$$\widehat{C}_{\gamma t}^* = E\{\widehat{c}_t^{*\gamma} | z^t\}$$

to be the γ th moment of cross-sectional measured consumption, in public history z^t . From

the definition of measured consumption, we know that:

$$\begin{aligned}
\widehat{C}_{\gamma t}^* &= E\{c_t^{*\gamma} \exp(\gamma \nu_t) | z^t\} \\
&= E\{\exp(\gamma \nu_t) | z^t\} E\{c_t^{*\gamma} | z^t\} \\
&= E\{\exp(\gamma \nu_t)\} C_{\gamma t}^*
\end{aligned}$$

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

Now suppose that:

$$E\{\exp(\gamma \nu_t)\} < \infty$$

ν_t is a stationary process

These assumptions imply that:

$$\begin{aligned}
&\beta E\{\widehat{C}_{\gamma t}^* R_{t+1}^m / \widehat{C}_{\gamma, t+1}^* | z^t\} \\
&= \beta E\left\{ \frac{E\{\exp(\gamma \nu_t)\} C_{\gamma t}^* R_{t+1}^m}{E\{\exp(\gamma \nu_{t+1})\} C_{\gamma, t+1}^*} | z^t \right\} \\
&= \beta E\{C_{\gamma t}^* R_{t+1}^m / C_{\gamma, t+1}^* | z^t\}
\end{aligned}$$

for all (t, z^t) . Thus, under these assumptions, $\beta \widehat{C}_{\gamma t}^* / \widehat{C}_{\gamma, t+1}^*$ is a valid stochastic discount factor for financial asset returns.

This argument implies that 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; in particular, some evidence suggests that there is less measurement error associated with highly skilled people. 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.⁴

6. Two Other Stochastic Discount Factors

In the prior section, 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.

A. The Incomplete Markets SDF

We assume that the economic environment is as described in Section 2. We assume as in Section 4 that agents engage in sequential asset trade, so that in period $t = 1, \dots, T - 1$, 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_{t+1}^m be the equilibrium gross return from period t to period $(t + 1)$ of asset m . Let $(c^{INC}, y^{INC}, K^{INC})$ be an equilibrium allocation in this setting such that in

⁴There 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.

equilibrium, agents face no binding borrowing constraints.

A necessary condition of individual optimality is:

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

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:

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

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

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

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:

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

which is the average of the agents' intertemporal marginal rates of substitution. This *average*

IMRS discount factor is generally different from the incomplete markets SDF. However, both are valid SDF's in an incomplete markets equilibrium with non-binding borrowing constraints (of course, when markets are incomplete, there are many valid SDF's).

In this paper, we focus on the incomplete markets SDF. As stressed in the introduction, the main reason for doing so is measurement error. Suppose that there is a measurement error process of the kind defined in section 4 and we observe:

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

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

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

which is the true average *IMRS* discount factor multiplied by a constant. The measured version of the average *IMRS* discount factor is not valid for arbitrary returns (although it is valid for return differentials like the equity premium).

In contrast, the measured incomplete markets SDF equals:

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

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.

B. 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^{CM}, y^{CM}, K^{CM}) 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 (as would be true, for example, if agents traded a complete set of state-contingent claims).

Then, in equilibrium:

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

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

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

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

7. Empirical Implementation: Preliminaries

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

A. Methodology

Our methodology is similar to 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 $m_{t+1}(\beta, \gamma)$ satisfies:

$$E[\{m_{t+1}(\beta, \gamma, z^{t+1})R_{t+1}^m(z^{t+1}) - 1\}x_t] = 0$$

By considering arbitrary instruments x_t 's and arbitrary returns R_{t+1}^m , we can form an enormous 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_t has marginal predictive power (over the collection of other x_t '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 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:

$$(4) \quad E[m_{t+1}(1, \gamma, z^{t+1})(R_{t+1}^{mkt}(z^{t+1}) - R_{t+1}^f(z^{t+1}))] = 0$$

where R_{t+1}^{mkt} is the value-weighted return to the stock market and R_{t+1}^f is the return to the 90-day Treasury bill (all returns are real).

Next, we investigate 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:

$$E[m_{t+1}(\beta, \gamma, z^{t+1})R_{t+1}^f(z^{t+1})] = 1$$

$$E[\{m_{t+1}(\beta, \gamma, z^{t+1})R_{t+1}^f(z^{t+1}) - 1\}R_t^f(z^t)] = 0$$

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.

We report a chi-squared test of each SDF against any alternative. However, our primary focus instead is on the parameter values (β, γ) that best fit the various restrictions for the model and the relative abilities of the various models to satisfy the various restrictions. We quantify the latter criterion by the sample mean of the error associated with each restriction.

B. The Data

In this section, we describe the data that we use in our empirical analysis.

The CEX

The microeconomic data are drawn from the 1980-1998 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 follows survey households for a maximum of 5 quarters, although only inventory and basic sample data are collected in the first quarter. 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.

Our sample selections are as follows. Our initial 1980-1998 CEX sample includes

1,249,329 monthly observations, corresponding to 141,289 households. We drop observations where expenditure on food and total nondurable goods is missing or reported to be zero. 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 fuel 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 drop duplicate interview months, those who are present for less than three months overall, and those who report less than three months of consumption data in a given interview. We also drop those who miss an interview (i.e., exit and re-enter the survey). Finally, we eliminate incomplete income respondents, i.e., households that do not provide complete information regarding their sources of income. Our sample selections are aimed at eliminating the most severe reporting errors in consumption.⁵ We end up discarding about 2% of households per month through our selection procedure.

We “deflate” consumption data to account for three phenomena: price differences over time, seasonal differences (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). Then, data are de-seasonalized by simple additive regression adjustments (a multiplicative adjustment makes little difference). Finally, we convert it into adult-equivalent consumption data.⁶

⁵An 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.

⁶The number of adult equivalents is defined as $(A + \alpha K)^\beta$ where A is the number of adults (aged 18

Given 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 firstly collected) in different months. Thus, some households' second interview covers the months of January through March, some other households's second interview will have data for the months of February through April, and so forth. By the very design of the CEX, no households contributes multiple observations to adjacent overlapping quarters. In other words, a household that contributes data to January-March 1980 will not contribute to 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 rarely matches 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.⁷ At present, it is not clear why this is so, and whether this is necessarily

or more), K the number 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. This consists of regressing food's budget share on log non-durable expenditure and a set of demographics the "equivalence scale" is assumed to depend on. The baseline household is a childless single. The equivalence scale depends on a dummy of whether children are present, the number of children, and the number of adult members. The equivalence scale is identified by the assumption that, if all households face the same vector of prices, a household i and the baseline household having the same foodshare should be at the same level of welfare.

⁷See Attanasio, Battistin, and Ichimura (2004).

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 estimate the parameters in the complete markets 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.⁸ The data are collected by the Bureau of Economic Analysis. Our measure of 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). We shall see that changing the measure of consumption in this way does not change our results that much.

⁸All the NIPA tables can be found at <http://www.bea.gov/bea/dn/nipaweb/index.asp>.

The returns data

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 then convert it in real terms dividing it by $(1 + \pi)$. Finally, we obtain the quarterly return by compounding the monthly returns. The difference $(R_t^{mkt} - R_t^f)$ is the premium on the value weighted portfolio.

8. Empirical Implementation: Results

The CEX provides data of the form $\{\{c_{it}\}_{i=1}^{N_t}\}_{t=1}^{T+3}$, 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. In particular,

let:

$$\begin{aligned}\widehat{m}_{t+3}^{PIPO}(\beta, \gamma) &= \beta \frac{N_t^{-1} \sum_{i=1}^{N_t} c_{it}^\gamma}{N_{t+3}^{-1} \sum_{i=1}^{N_{t+3}} c_{it+3}^\gamma} \\ \widehat{m}_{t+3}^{INC}(\beta, \gamma) &= \beta \frac{N_{t+3}^{-1} \sum_{i=1}^{N_{t+3}} c_{it+3}^{-\gamma}}{N_t^{-1} \sum_{i=1}^{N_t} c_{it}^{-\gamma}} \\ \widehat{m}_{t+3}^{CM}(\beta, \gamma) &= \beta \frac{N_{t+3}^{-1} (\sum_{i=1}^{N_{t+3}} c_{it+3})^{-\gamma}}{N_t^{-1} (\sum_{i=1}^{N_t} c_{it})^{-\gamma}}\end{aligned}$$

denote the sample analogs of the PIPO, incomplete markets, and complete markets 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 term is the ratio of the γ -th moment of 1980:3 (which is calculated using all households reporting expenditure data for January-March 1980) and the γ -th moment of 1980:6 (which is calculated using all households reporting expenditure data for April-June 1980); the second term is the ratio of the γ -th moment of 1980:4 (which is calculated using all households reporting expenditure data for February-April 1980) and the γ -th moment of 1980:7 (which is calculated using all households reporting expenditure data for May-July 1980); and so forth. The last term of the summation is the ratio of data for 1998:9 to 1998:11 (when our data ends). The average N_t is 1647 (the median is 1642). The maximum value is 3736 (which occurs in 1986, the year where the CEX was redesigned), the minimum 930. It is assumed that we have a time series of $(T + 3)$ observations on $\widehat{m}_t^j(\beta, \gamma)$. In our case, we have data from 1980:3 through 1998:11, and so we have $T = 222$.

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 2.4% per quarter higher than the mean return to Treasury bills. This sample estimate is considerably higher than the 6% annual number averaged in the hundred years of data studied by Mehra and Prescott. The

standard deviation of stock returns is about 7.5% 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 1-2. For all values of γ , the SDF is highly variable. Of course, in order to be a valid SDF, the rate of change has to be more than variable: it must covary negatively with stock returns.

A. The Equity Premium: Results

We look first 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:

$$(5) \quad \bar{e}_{mkt}^j(\gamma) = \frac{1}{T} \sum_{t=1}^T \hat{m}_{t+3}^j(1, \gamma) (R_{t+3}^{mkt} - R_{t+3}^f)$$

for $j = PIPO, INC$ and CM . Equation (5) is the empirical analog of (4). A simple way to compare the three models is to compare $\bar{e}_{mkt}^j(\gamma)$ for $j = PIPO, INC$, and CM , 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).

Throughout the paper, we conduct inference using the block bootstrap. Blocks identify the number of observations per households (from 1 to 4) in the cross-section. Unlike Brav et al.'s calculation of standard errors, our approach takes into account both cross-sectional variability (which influences the "composition" of the γ -th moment of the consumption distribution) and of time series variability (which influences the movements in the premium on the value weighted portfolio and the evolution of the γ -th moments of the consumption

distribution). Note that we calculate block bootstrap standard errors by taking blocks both in the cross-sectional dimension (to account for the fact that individuals may be interviewed multiple times over a 1-year period, which would violate the assumption of independence of errors across individuals) and in the time series dimension (to account for serial correlation in returns, etc.). The optimal block length in the time series is a complicated issue and the literature so far offers little guidance. We choose time-series blocks of length 6. In this way, a block covers two quarters of observations. In contrast to the time series, the length of blocks in the cross-section is a less contentious issue (we know the proportion of people completing 1, 2, 3, and 4 interviews, and the size of blocks is chosen accordingly).

We report the estimates in Tables 2-4, along with standard errors and p-values. Our basic finding in Tables 2-4 is that with the PIPO stochastic discount factor, the estimated equity premium is zeroed out at a value of γ between 8 and 9. In contrast, with the incomplete markets and complete markets discount factors, the estimated equity premium remains positive for all specifications of γ . (Note that Table 4 includes estimates of the complete markets SDF using both CEX and NIPA measures of consumption⁹. The complete markets SDF accounts for little of the equity premium using either measure.) Figure 3 represents these estimates graphically. We conclude that the PIPO stochastic discount factor is somewhat more consistent with the equity premium than are the other two discount factors. We say "somewhat" because 8 or 9 is generally regarded as being very large as a specification of the coefficient of relative risk aversion. As well, both the incomplete markets discount factor and the PIPO discount factor are not statistically significantly different from zero for γ around 3 or 4.

⁹Of course, the bootstrap standard errors for the NIPA data use only the time-series aspect of the data.

In Tables 5 and 6, we re-estimate the equity premium errors for the incomplete markets and complete markets SDFs by using only consumption from households with a positive amount of stockholdings.¹⁰ The idea of this change is to allow for the possibility of incomplete financial market participation (in an admittedly crude way). (There is even less theoretical reason for doing this to the PIPO SDF, and so we do not do so.) The results concerning the equity premium are basically unchanged.

In Table 7, we use a slightly different approach, and formally estimate the coefficient of relative risk aversion 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 around 8.96 for the PIPO SDF - which is consistent with our above analysis - and the standard error is around 3.45. The estimates of γ for the incomplete markets and complete markets SDFs are surprisingly low (at least when we use CEX data for the latter SDF). But these low estimates are misleading. The estimated equity premium pricing error for all of these models is around 0.024. Hence, as our less formal procedure in Tables 2-6 showed, the equity premium pricing error is virtually unaffected by the value of γ for the incomplete markets and complete markets SDFs.

B. Understanding the Equity Premium Results

Why is the sample equity premium zeroed out at $\gamma = 9$ for the PIPO stochastic discount factor? It is perhaps instructive to have a close look at the data generating this result. Define the time- t error in the PIPO case as $e_{mkt,t}^{PIPO}(1, \gamma) = \frac{N_{t-3}^{-1} \sum_{i=1}^{N_{t-3}} c_{it-3}^\gamma}{N_t^{-1} \sum_{i=1}^{N_t} c_{it}^\gamma} (R_t^m - R_t^f)$ and its time series average as $\bar{e}_{mkt}^{PIPO}(1, \gamma) = \frac{\sum_{t=1}^T e_{mkt,t}^{PIPO}(1, \gamma)}{T}$. This is the average unexplained

¹⁰Stockholding is defined as the sum of all stocks, bonds, mutual funds, savings accounts, checking accounts, and U.S. Savings Bonds.

premium we report in Table 2. The time- t error is negative whenever $(R_t^m - R_t^f) < 0$. In particular, if π is the proportion of negative time- t errors, the average unexplained premium can be decomposed as the weighted average of positive and negative time- t errors,

$$\bar{e}_{mkt}^{PIPO}(1, \gamma) = \pi \frac{\sum_{\varepsilon_t(\gamma) < 0} e_{mkt,t}^{PIPO}(1, \gamma)}{\sum \mathbf{1}\{e_{mkt,t}^{PIPO}(1, \gamma) < 0\}} + (1 - \pi) \frac{\sum_{\varepsilon_t(\gamma) > 0} e_{mkt,t}^{PIPO}(1, \gamma)}{\sum \mathbf{1}\{e_{mkt,t}^{PIPO}(1, \gamma) \geq 0\}}$$

where $\mathbf{1}\{\cdot\}$ is an indicator function. Table 2 shows that $\bar{e}_{mkt}^{PIPO}(1, \gamma) > 0$ for $\gamma < 9$ and $\bar{e}_{mkt}^{PIPO}(1, \gamma) < 0$ for $\gamma \geq 9$. Thus the average of negative time- t errors exceeds (in absolute value) the average of positive time- t errors when $\gamma \geq 9$.

The point where $\bar{e}_{mkt}^{PIPO}(1, \gamma)$ turns sign from positive to negative (if any) clearly depends on the relative weight of extreme (positive and negative) realizations of $e_{mkt,t}^{PIPO}(1, \gamma)$. For example, the largest negative value in the distribution of $e_{mkt,t}^{PIPO}(1, 9)$ occurs in 1992:10. If we exclude it, $\bar{e}_{mkt}^{PIPO}(1, 9)$ turns positive and we would not get any zeroing-out at $\gamma = 9$ in the PIPO case. However, the counterfactual also works in reverse: If we exclude the largest positive value from the distribution of $e_{mkt,t}^{PIPO}(1, 8)$, $\bar{e}_{mkt}^{PIPO}(1, 8)$ turns negative, which means that we would get zeroing-out at $\gamma = 8$. 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 (9) in 1992:10 is extremely negative because in that period the premium is negative (-1%), $N_t^{-1} \sum_{i=1}^{N_t} c_{it}^\gamma$ is small, and $N_{t-3}^{-1} \sum_{i=1}^{N_{t-3}} c_{it-3}^\gamma$ is large, relative to other periods (see Figure 2). There are certainly extreme values of the consumption distribution that are shifting the balance in either direction. For example, the 1992:7 CEX cross-section contains

a household reporting quarterly consumption of about \$32,000. Nevertheless, this extreme observation is by no means an outlier (both annual income and assets exceed \$100,000). 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 $\gamma = 10$) and we still do not get zeroing-out for any value of γ in the incomplete market case or complete market case.

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). Of course, 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.

However, BCG (2002) use sample selection criteria that differ from ours in a number of respects. They only keep households who stay in the sample for three or more quarters (because they use the average IMRS SDF). To eliminate outliers, they discard 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. As well, they use the sample period 1982:I-1996:I, not the sample period 1980:I-1998:IV.

We constructed a subsample of the CEX using the selection criteria reported in their paper. (We requested BCG's constructed subsample from Alon Brav, but we have not as yet received these data.) We then recalculated the point estimates in Tables 2-4 using this subsample. Our new results were highly similar to what we report in Tables 2-4. In particular, the sample equity premium is eliminated using the PIPO discount factor when we set $\gamma = 10$. However, just as in Tables 2-4, the sample equity premium is basically unaffected

by the size of γ for the complete markets SDF and it is growing as a function of γ for the incomplete markets SDF (it is about 9 billion for $\gamma = 10$). We also redid BCG (2002)'s Table 2 (see p. 809 of their paper) using our reconstructed version of their sample; like BCG, we find that the equity premium is eliminated using the average IMRS SDF if γ is near 3.¹¹

C. Treasury Bill Returns

We turn next to the two Treasury bill restrictions. Define:

$$\begin{aligned}\bar{e}_{b1}^j(\beta, \gamma) &= \frac{1}{T} \sum_{t=1}^T \hat{m}_{t+3}^j(\beta, \gamma) R_{t+3}^f - 1 \\ \bar{e}_{b2}^j(\beta, \gamma) &= \frac{1}{T} \sum_{t=1}^T (\hat{m}_{t+3}^j(\beta, \gamma) R_{t+3}^f - 1) R_t^f\end{aligned}$$

We estimate (β, γ) by applying *GMM* to these pricing errors. Here, our choice of weighting matrix is irrelevant; in all models, it was possible to find (β, γ) so as to zero out both pricing errors. We find in Table 8 that the estimate of (β, γ) is about (0.93, 3.1) for the *PIPO* discount factor. The estimate for β is somewhat low, considering that it is being estimated over a quarterly frequency. The estimates of β are more plausible in the other two models; the estimates of γ are also plausible in these models but highly imprecise.

Using the NIPA data, rather than the CEX data, for the complete markets SDF does not affect either the estimate of γ or its standard error that greatly. The estimate of β is higher when we use the NIPA data, which is consistent with the fact that the growth rate of consumption is higher in the NIPA data. As well, the standard error for the estimate of β is smaller - but this could be due to the fact that the standard error only accounts for

¹¹The complete results from using BCG's selection criteria are available on request.

time-series variability.

Finally, we turn to using all three restrictions simultaneously. Here, with two possible 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 a one-step GMM procedure. We pick the weighting matrix by using the reciprocal of the variance-covariance matrix of the pricing errors, given that the parameter β is set to 1 and the parameter γ is set to 0. This means that we are using the same variance-covariance matrix for all of the possible discount factors (which is good), and also putting more weight on statistically more informative restrictions (which is also good).

Table 9 contains the results. The basic finding is that the weighting matrix completely downweights the equity premium as being an important source of information; the resulting estimates essentially zero out the Treasury bill pricing errors. However, at the estimated preference parameters, the estimated equity premium error is roughly the same as the equity premium itself (2.4% per quarter). In this sense, even with the PIPO discount factor, the equity premium remains a puzzle.

9. 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 social 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 show too that this representation is robust to a wide class of measurement error processes. We construct similar representations for an asset pricing kernel implied by incomplete markets equilibrium and the unique asset pricing kernel implied by complete markets equilibrium.

The second contribution is empirical. We use data from the CEX to construct sample analogs for the three stochastic discount factors over the period 1980-98. We first compare the stochastic discount factors' ability to explain the size of the equity premium in this period. We show that if the coefficient of relative risk aversion is around 9, the new PIPO discount factor does a good job of rationalizing the equity premium. The other two stochastic discount factors are not able to explain the equity premium for any specification of the coefficient of relative risk aversion. It is worth emphasizing that this empirical result is in sharp contrast to the findings of Brav, Constantinides and Geczy (2002) for their alternative incomplete markets stochastic discount factor.

We then examine the ability of the three discount factors to explain for the level of the risk-free rate and the autocovariance of the risk-free rate over this sample period. All three stochastic discount factors do a good job at matching these two aspects of the data for plausible specifications of the underlying preference parameters. However, none of the discount factors can explain these aspects of the data and also account for the large equity premium.

We draw two conclusions from our empirical analyses. The first is that asset pricing models grounded in fundamentals like private information frictions are better able to account for the large equity premium than are more standard asset pricing models based on incomplete

markets or complete markets models. The second is that simultaneously explaining the equity premium, the level of the risk-free rate, and the autocovariance of the risk-free rate remains challenging for any model. A large amount of empirical research ignores this last aspect of the data (to cite one influential example, Campbell and Cochrane (1999) simply assume that the risk-free rate is constant over time). Yet our analysis shows that, at least statistically, it is more informative than the equity premium. An important challenge for future research is to build asset pricing models that are better able to account for all of these aspects of the data.

References

- [1] ATTANASIO, O., BATTISTIN, E., and ICHIMURA, H. (2004), “What Really Happened to Consumption Inequality in the US?,” NBER working paper 10338.
- [2] ATTANASIO, O., and DAVIS, S. (1996), “Relative Wage Movements and the Distribution of Consumption,” *Journal of Political Economy* 104, 1227-1262.
- [3] 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.
- [4] 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..
- [5] BOSWORTH, B., BURTLESS, B., and SABELHAUS, J. (1991), “The Decline in Saving: Evidence From Household Surveys,” *Brooking Papers on Economic Activity*, 183-241.
- [6] BRAV, A., CONSTANTINIDES, G., and GECZY, C. (2002), “Asset Pricing with Heterogeneous Consumers,” *Journal of Political Economy* 110, 793-824.
- [7] 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.
- [8] CITRO, C., and MICHAELS, R., eds. (1995), *Measuring Poverty: A New Approach (1995)*, National Academy Press, Washington: D.C.

- [9] COCHRANE, J. (1991), “A Simple Test of Consumption Insurance,” *Journal of Political Economy* 96, 893-920.
- [10] COCHRANE, J. (2001), *Asset Pricing*, Princeton University Press, Princeton, NJ.
- [11] COGLEY, T. (2002), “Idiosyncratic Risk and the Equity Premium: Evidence from the Consumer Expenditure Survey,” *Journal of Monetary Economics* 49, 309-334.
- [12] GOLOSOV, M., KOCHERLAKOTA, N., and TSYVINSKI, O. (2003), “Optimal Indirect and Capital Taxation,” *Review of Economic Studies* 70, 569-587.
- [13] HALL, R. (1988), “Intertemporal Substitution in Consumption,” *Journal of Political Economy* 96, 339-357.
- [14] HANSEN, L., and SINGLETON, K. (1982), “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models,” *Econometrica* 50, 1269-1288.
- [15] KOCHERLAKOTA, N. (1990), “On Tests of Representative Consumer Asset Pricing Models,” *Journal of Monetary Economics* 26, 285-304.
- [16] KOCHERLAKOTA, N. (1996), “The Equity Premium: It’s Still A Puzzle,” *Journal of Economic Literature* 34, 42-71.
- [17] KOCHERLAKOTA, N. (1998), “The Effects of Moral Hazard on Asset Prices When Financial Markets Are Complete,” *Journal of Monetary Economics* 41, 39-56.
- [18] KOCHERLAKOTA, N. (2003), “Zero Expected Wealth Taxes: A Mirrlees Approach to Optimal Dynamic Taxation,” Stanford University working paper, <http://www.stanford.edu/~nkocher/mirrlees5.pdf>.

- [19] LIGON, E. (1998), "Risk-Sharing and Information in Village Economies," *Review of Economic Studies* 65, 847-864.
- [20] LUSTIG, H. (2002), "The Market Price of Risk and the Wealth Distribution," The University of Chicago working paper.
- [21] MEGHIR, C., and PISTAFERRI, L., "Income Variance Dynamics and Heterogeneity," *Econometrica* 72, 1-32.
- [22] MEHRA, R., and PRESCOTT, E., "The Equity Premium: A Puzzle," *Journal of Monetary Economics* 15, 145-161.
- [23] MIRRLEES, J. (1971), "An Exploration in the Theory of Optimum Income Taxation," *Review of Economic Studies*, 38, 175-208.
- [24] PHELAN, C., (1994), "Incentives, Insurance, and the Variability of Consumption and Leisure," *Journal of Economic Dynamics and Control* 18, 581-599.
- [25] ROGERSON, W. (1985), "Repeated Moral Hazard," *Econometrica*, 53, 69-76.
- [26] SEMENOV, A. (2004), "High-Order Consumption Moments and Asset Pricing," York University working paper.
- [27] 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.
- [28] VISSING-JORGENSEN, A. (2002), "Limited Asset Market Participation and the Elasticity of Intertemporal Substitution," *Journal of Political Economy* 110, 825-853.

Table 1
Descriptive Statistics

Panel A: Household data from the CEX

	1980	1983	1986	1989	1992	1995	1998
Age	45.37	46.23	46.31	46.92	47.37	47.93	48.56
Family size	2.70	2.63	2.63	2.61	2.57	2.61	2.51
# of kids	0.78	0.73	0.74	0.72	0.73	0.75	0.70
Proportion some college+	0.43	0.44	0.46	0.48	0.50	0.49	0.54
Annual before tax income	21607	23407	24695	26402	25234	25147	26049
Stocks	409	609	735	770	865	1003	3142
Quarterly consumption	1997	1897	1968	2025	1906	1819	1822
Quarterly food exp.	944	826	834	902	831	823	794
N	11,184	14,962	20,028	15,072	15,026	12,849	15,455

Note: Monetary variables are deflated by the CPI-U (1983-1984=100).

Panel B: Time series data

	1980-82	1983-85	1986-88	1989-91	1992-94	1995-98	1980-98
r^f (mean)	1.08	1.23	0.62	0.62	0.14	0.68	0.72
r^f (st.dev.)	0.81	0.43	0.61	0.44	0.29	0.30	0.60
r^m (mean)	2.53	3.54	3.00	2.76	1.33	5.04	3.14
r^m (st.dev.)	9.87	6.22	10.37	7.45	3.25	6.61	7.57
$corr(r^f, r^m)$	-0.1942	0.0531	0.1712	0.7005	0.0675	0.0827	0.1238
$corr(r^f, r_{-3}^f)$	0.6218	-0.0250	0.3681	-0.3230	-0.3103	-0.2598	0.4797

Note: data refer to the period 1980:6-1998:11. The table reports (overlapping) quarterly returns.

Table 2**The Unexplained Equity Premium: PIPO SDF**

γ	Unexplained premium	Bootstrap p -value
0	0.0242	0.0000
1	0.0243	0.0000
2	0.0242	0.0000
3	0.0245	0.0000
4	0.0280	0.1629
5	0.0379	0.7783
6	0.0573	0.8959
7	0.0860	0.7692
8	0.1028	0.7149
9	-0.0110	0.6244
10	-0.7658	0.6063

Note: The unexplained premium is defined as:

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

and is expressed in percentage form. Bootstrap results are based on 200 replications.

Table 3**The Unexplained Equity Premium: Incomplete Markets SDF**

γ	Unexplained premium	Bootstrap p -value
0	0.0242	0.0000
1	0.0242	0.0000
2	0.0249	0.0000
3	0.0419	0.0905
4	0.2843	0.2443
5	2.5039	0.2262
6	19.0415	0.2262
7	134.1667	0.2172
8	917.3239	0.2172
9	6196.249	0.2172
10	41633.55	0.2262

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. Bootstrap results are based on 200 replications.

Table 4**The Unexplained Equity Premium: Complete Markets SDF**

γ	CEX data		NIPA PCE data	
	Unexplained premium	Bootstrap p -value	Unexplained premium	Bootstrap p -value
0	0.0242	0.0000	0.0242	0.0000
1	0.0243	0.0000	0.0242	0.0000
2	0.0243	0.0000	0.0241	0.0000
3	0.0243	0.0000	0.0241	0.0000
4	0.0244	0.0000	0.0240	0.0000
5	0.0244	0.0000	0.0240	0.0000
6	0.0245	0.0000	0.0239	0.0000
7	0.0246	0.0000	0.0239	0.0000
8	0.0248	0.0000	0.0238	0.0000
9	0.0249	0.0000	0.0238	0.0000
10	0.0250	0.0000	0.0237	0.0000

Note: The unexplained premium is defined as:

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

and is expressed in percentage form. Bootstrap results are based on 200 replications.

Table 5**The Unexplained Equity Premium: Incomplete Markets SDF, Positive Assets**

γ	Unexpl. premium	Bootstrap p -value
0	0.0242	0.0000
1	0.0243	0.0000
2	0.0251	0.0000
3	0.0282	0.0181
4	0.0369	0.1267
5	0.0593	0.3077
6	0.1145	0.4525
7	0.2510	0.5430
8	0.5947	0.6154
9	1.4745	0.6787
10	3.7544	0.7240

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. Our sample excludes those with zero assets (defined as the sum of bonds and stocks). Bootstrap results are based on 200 replications.

Table 6**The Unexplained Equity Premium: Complete Markets SDF, Positive Assets**

γ	Unexpl. premium	Bootstrap p -value
0	0.0242	0.0000
1	0.0241	0.0000
2	0.0242	0.0000
3	0.0247	0.0000
4	0.0254	0.0000
5	0.0264	0.0090
6	0.0277	0.1357
7	0.0293	0.4072
8	0.0313	0.6425
9	0.0336	0.8869
10	0.0363	0.9593

Note: The unexplained premium is defined as:

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

and is expressed in percentage form. Our sample excludes those with zero assets (defined as the sum of bonds and stocks). Bootstrap results are based on 200 replications.

Table 7**The Equity Premium**

	<i>Pareto-optimal</i>	<i>Incomplete markets</i>		<i>Complete markets</i>		
		All	Positive	All	Positive	NIPA PCE
		households	assets	households	assets	
γ	8.9599	0.9181	0.4072	0.0793	0.9919	45.6000
	(3.4529)	(1.7804)	(3.0958)	(7.9259)	(4.9064)	(50.1785)
$\bar{e}_{mkt,t}$	7.50e-010	0.0242	0.0242	0.0242	0.0241	0.0228

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 \bar{e}_{mkt} reports the sample mean of the pricing error at the estimated value of γ . For the positive assets columns, our sample excludes those with zero assets (defined as the sum of bonds and stocks). The standard errors are based on 200 block bootstrap replications.

Table 8

Expected Return to the Treasury Bill

	<i>Pareto-optimal</i>	<i>Incomplete markets</i>		<i>Complete markets</i>		
		All	Positive	All	Positive	NIPA PCE
		households	assets	households	assets	
γ	3.0863 (1.5891)	0.5381 (2.4359)	1.7707 (4.4547)	4.9629 (8.8882)	-0.3206 (4.0087)	2.9649 (8.4652)
β	0.9274 (0.3284)	0.9920 (0.3489)	0.9711 (0.3446)	0.9804 (0.0922)	0.9919 (0.2676)	1.0026 (0.0302)
\bar{e}_{b1}	-1.08e-011	-1.25e-009	-1.91e-007	-1.33e-007	-2.20e-009	-2.11e-011
\bar{e}_{b2}	1.53e-009	-5.32e-009	2.76e-005	1.94e-008	1.74e-008	-5.54e-011

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

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

have expectation zero. The rows \bar{e}_{b1} and \bar{e}_{b2} report the sample means of the pricing errors at the estimated value of γ . For the positive assets column, our sample excludes those with zero assets (defined as the sum of bonds and stocks). The standard errors are based on 200 block bootstrap replications.

Table 9

The Equity Premium and the Treasury Bill Return

	<i>Pareto-optimal</i>	<i>Incomplete markets</i>		<i>Complete markets</i>		
		All	Positive	All	Positive	NIPA PCE
		households	assets	households	assets	
γ	3.0926	0.5627	1.5590	1.7715	-0.3395	3.1724
	(2.5680)	(2.5604)	(2.3762)	(7.2411)	(6.4601)	(9.5931)
β	0.9267	0.9920	0.9795	0.9909	0.9919	1.0034
	(0.3655)	(0.3010)	(0.2942)	(0.0671)	(0.3713)	(0.0343)
\bar{e}_{mkt}	0.0247	0.0242	0.0246	0.0243	0.0244	0.0241
\bar{e}_{b1}	9.62e-005	8.61e-005	0.0015	0.0010	9.34e-005	7.86e-005
\bar{e}_{b2}	-3.24e-006	-5.43e-007	3.74e-005	2.40e-005	-3.65e-007	-7.22e-007
J	18.14	21.07	0.30	6.54	20.91	20.75
(p-value)	[0.0409]	[0.0094]	[0.5773]	[0.0500]	[0.0000]	[0.0050]

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

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

have expectation zero. The rows \bar{e}_{mkt} , \bar{e}_{b1} and \bar{e}_{b2} report the sample means of the pricing errors at the estimated value of γ . For the positive assets column, our sample excludes those with zero assets (defined as the sum of bonds and stocks). The J-statistic is constructed using the formula in Cochrane (2001, p. 204). The standard errors and p-values are based on 200 block bootstrap replications.




Figure 1: The PIPO stochastic discount factor (with $\beta = 1$ and $\gamma = 3$).

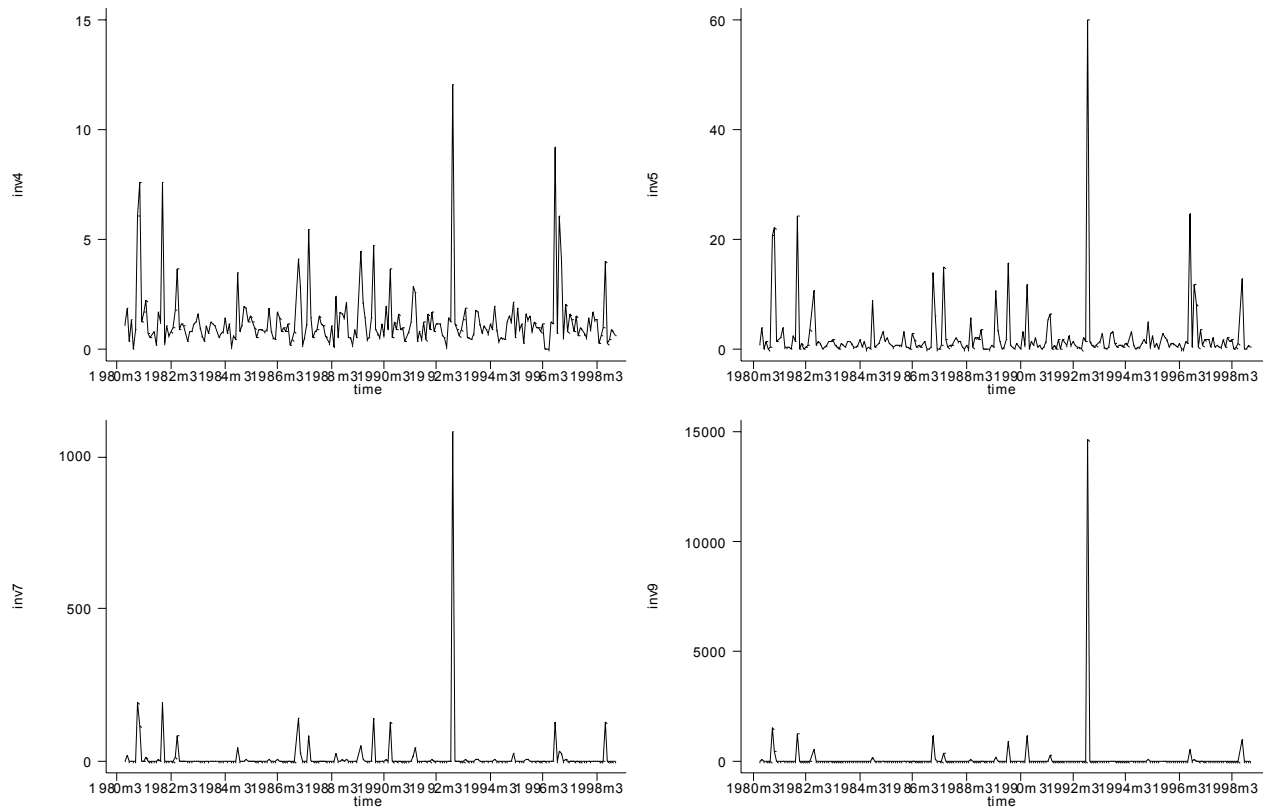


Figure 2: The PIPO stochastic discount factor (when $\beta = 1$ and $\gamma = \{4, 5, 7, 9\}$).




Figure 3: The equity premium pricing error