The representative agent of an economy with external habit-formation and heterogeneous risk-aversion

Costas Xiouros*　　Fernando Zapatero†

*Department of Finance and Business economics, Marshall School of Business, University of Southern California, Los Angeles, CA 90089-0804, e-mail: xiouros@usc.edu

†Department of Finance and Business economics, Marshall School of Business, University of Southern California, Los Angeles, CA 90089-0804, e-mail: fzapatero@marshall.usc.edu

The representative agent of an economy with external habit-formation and heterogeneous risk-aversion

Abstract

For the first time in the literature, we derive an analytic expression for the representative agent of a fairly general class of economies populated by agents with "catching up with the Joneses" preferences, but with heterogeneous risk-aversion. As Chan and Kogan (2002) show numerically, the representative agent has stochastic risk-aversion that moves counter-cyclically with the state variable. However, we show that the heterogeneity of risk-aversion is unlikely to be able to explain the empirical regularities -namely the variability of the Sharpe ratio- that Campbell and Cochrane (1999) explain in a model of a representative agent with stochastic risk-aversion.

The consumption asset pricing literature, starting with Hansen and Singleton (1983), Mehra and Prescott (1985), and others, has tried with limited success to identify the fundamental factors that drive the level, variation, and cyclical movements of asset prices and conditional asset pricing moments. In particular, it seems difficult to reconcile the smoothness of aggregate consumption with the high volatility of asset returns and high average historical returns in excess of the risk free rate. In addition, the historical covariance between aggregate consumption growth and asset returns is very low. These two pieces of evidence imply that the unit price of risk required by investors is on average very high. Further evidence indicates that the unit price of risk also varies significantly across time. In particular, price dividend ratios exhibit high variability compared to dividends, and have some forecasting power in predicting long-run stock returns, as found in Fama and French (1988), Campbell and Shiller (1988a,b) and Campbell (1991), among others. On the same note, the variance decomposition (see Cochrane (1992)) of price dividend ratios reveals that almost all variation can be attributed to varying future excess returns. Additionally, a number of papers have focus on the analysis of stock market volatility, and they have found it to vary over the business cycle. This is shown in, for example, Bollerslev, Chou and Kroner (1992) or Ludvigson and Ng (2007). Several other papers[1] have studied the contemporaneous relation between expected stock returns and conditional return volatilities and found conflicting results. The relation, however, seems to be weak but, more importantly, most of these studies indicate that the Sharpe ratio varies throughout the business cycle. In particular, it appears to increase considerably during recessions, and to fall during expansions.

At the heart of all these stylized facts there seems to be a mean-reverting and counter-cyclically varying risk premium. Furthermore, it seems to vary more than the stock market volatility, giving rise to a counter cyclically varying Sharpe ratio. Campbell and Cochrane (1999) explain many of the mentioned asset pricing features in a model of a representative agent with external habit preferences and counter-cyclical variation in risk aversion. Chan and Kogan (2002) argue that such a variation in the risk aversion of the representative agent can be the result of the endogenous cross-sectional redistribution of wealth in an economy with multiple agents with heterogeneous risk-aversion parameters.

We study a model similar to that of Chan and Kogan (2002), although in discrete time. Based on a detailed study of the mechanics of the stationary equilibrium of the heterogeneous agents economy, we derive explicitly the time-varying risk-aversion of the representative agent. From the resulting expression, we can analyze the properties of the varying risk-aversion parameter of the economy. We find that, although the counter-cyclical pattern of Campbell and Cochrane (1999) is accurate, as shown by Chan and Kogan (2002), many more assumptions are needed

[1] Like Bollerslev, Engle and Wooldridge (1988), Harvey (1989), Whitelaw (1994), Brandt and Kang (2004) and Ludvigson and Ng (2003).

in order to have more than just a marginal effect on asset price dynamics. In particular, we show that enough time-varying risk aversion fails to obtain as the result of aggregation in an economy of rational agents with standard preferences and different risk-attitudes. Even when we inflate the level of heterogeneity and increase the risk in the economy to levels that gives us the ability to predict an average equity premium close to the average excess return of the past 75 years, we are not able to produce enough predictability. For the baseline model for which the level of heterogeneity is calibrated to fit the estimated distribution of Kimball, Sahm and Shapiro (2008), the underlying consumption risk is not enough to even predict an asset pricing behavior that is clearly different, either in patterns or in levels, from a homogeneous agents economy.

Chan and Kogan (2002) assume a highly persistent and slow moving external habit, as well as a particularly high level of heterogeneity in risk-aversion. Our representative agent formulation reveals that if either of these assumptions is missing, a substantially varying Sharpe ratio does not obtain. The empirical results of Kimball, Sahm and Shapiro (2008), as well as our empirical results using the Consumption Expenditure Survey dataset, reveal that it is highly unlikely that the level of heterogeneity assumed by Chan and Kogan (2002) is realistic. In addition, the particular level of persistence of the external habit they assume produces excessive persistence and volatility in the price-dividend ratios. Campbell and Cochrane (1999) also assume a slowly varying state variable, namely the surplus consumption (current aggregate consumption over the external habit), but with the additional feature of a highly time-varying conditional variance, that as we show is related to the risk-aversion of the economy. In this paper we let the persistence in the price-dividend ratio guide us in the selection of the persistence of our habit process. The study of Campbell and Cochrane (1999) is very useful because it identifies the main features that a successful asset pricing model needs to have in order to explain all the aforementioned empirical facts. However, for risk-aversion heterogeneity to have an impact on prices, we would need another source of variability in the wealth distribution across agents. For example, in the overlapping generations model of Gârleanu and Panageas (2008), heterogeneity of risk aversion does have an impact because the re-allocation of wealth is considerably more drastic across time.

The literature on heterogeneity of risk aversion has a long tradition in finance. Dumas (1989) solves numerically a model with two agents with different risk-aversion, one of them with logarithmic utility. Wang (1996) considers also a two-agent economy and concentrates on the dynamics of bond prices. Coen-Pirani (2004) focuses on the dynamics of the wealth distribution among two agents with Epstein-Zin preferences. Kogan, Makarov and Uppal (2007) show that in a two-agent economy with borrowing constraints the Sharpe ratio can be high while the risk-free rate is low. Bhamra and Uppal (2009) show that completing the market in an economy populated with heterogeneous agents might increase the stock price volatility substantially. In this paper we consider an economy with arbitrary number of agents with "catching up with the Joneses" preferences and with markets dynamically complete. Using analytical results, as well as computing the exact equilibrium of several economies, we find that in the absence of any frictions or incompleteness in the market the effect of heterogeneity is in general minimal.

The rest of the paper is structured as follows. In section 1 we describe the heterogeneous agents economy and solve for the competitive equilibrium. In section 2 we consider a representative agent economy that is homeomorphic in its pricing implications with the heterogeneous agents economy of section 1. We derive an expression for the stochastic risk aversion of the representative agent and analyze its properties. In section 3 we parametrize the distribution of agents and fit it to the distribution estimated by Kimball, Sahm and Shapiro (2008). Using the results of sections 1 and 2, we derive analytically the stochastic risk-aversion of the economy. In section 4 we assume a particular process for habit and examine theoretically the asset pricing behavior. In section 5 we consider data from the Consumption Expenditure Survey and calibrate the economy. We carry out an extensive quantitative analysis of the effects of heterogeneity in section 6. We conclude in section 7.

1 The Model

We consider a version in discrete time, but more general, of the infinite horizon endowment economy of Chan and Kogan (2002). We chose discrete time instead of continuous time in order to allow for more general specifications of the uncertainty of the economy that are yet numerically tractable. In our model, uncertainty is driven by an exogenous state that follows a time homogeneous Markov process. The exogenous state is perfectly observable to all agents in the economy. Financial markets are dynamically complete, in the sense that the equilibrium asset structure spans the one period ahead uncertainty at every possible state of nature. There is a single perishable good, and agents exhibit power utility preferences with external habit formation, in the style of the "catching up with the Joneses" preferences of Abel (1990).[2] We present two versions of the model: In the first version the economy is populated by a number of different types of agents with possibly different coefficients of relative risk aversion; in the second version, we replace the heterogeneous agents with a representative agent with a stochastic coefficient of relative risk aversion; we then introduce an expression for the stochastic risk aversion coefficient of the representative agent parameterized by the primitives of the multiple agent economy and derive the rule that makes the two economies equivalent.

As in Chan and Kogan (2002), catching up with the Joneses preferences are not only attractive from an economics point of view (there are some influential papers that assume this type of

[2] Habit formation preferences have been extensively explored in the literature in various forms. Significant contributions include Ryder and Heal (1973), Sundaresan (1989), Constantinides (1990), Detemple and Zapatero (1991), Galí (1994), and Hindy, Huang and Zhu (1997).

preferences, especially Campbell and Cochrane 1999), but they also yield a stationary equilibrium such that the wealth distribution follows a time-homogeneous Markov process. The problem of standard power utility preferences is that, in the limit, wealth is accumulated by the least risk-averse agent.

1.1 Aggregate Uncertainty

The single source of uncertainty in our economy is growth in the aggregate endowment. We denote aggregate endowment by Y, with $y_t = \log(Y_t)$. As it is customary, we model the dynamics of the logarithm of the growth process, which for now we assume to be normal iid,

$$y_t - y_{t-1} = \mu + \sigma \epsilon_t, \qquad t \geq 0, \tag{1}$$

where $\epsilon_t \sim i.i.d.\ N(0,1)$ and y_0 is given. This simple structure will allow us to compare our results to those of Campbell and Cochrane (1999) and Chan and Kogan (2002).

1.2 Financial Markets

We assume that financial markets are dynamically complete, that is, at any point in time the equilibrium asset structure locally spans the one period ahead uncertainty. To keep the model as simple as possible, we assume that there is only one dividend-paying asset, the market security (or simply market), and the risk-free asset. We denote the price of the market by P^m, and the dividend it pays by D_t. We assume for simplicity that the dividend paid by the market is equal to the aggregate endowment,

$$D_t = Y_t, \qquad \forall t \geq 0$$

and that the market asset is in unit net supply. We denote the gross return on the market by R^m. There is also a risk-free security, with its return denoted by R^f, the risk-free rate. R^e is the excess return of the market over the risk-free rate. In addition, we implicitly assume (we do not need a formal characterization for our results) that there is a continuum of derivative securities written on the market asset, so that markets are dynamically complete.

We now introduce additional notation, that we will formalize later on, when we introduce the budget constraint of the optimization problem of an agent. In equilibrium, the process $\delta^t p_t$ represents the price of a unit of consumption good of period t in terms of a unit of consumption of period 0. δ is the common subjective one-period discount factor and p_t would be the normalized marginal utility of consumption in a representative agent economy. Therefore, the equilibrium price of any financial security is the expected value of future dividends priced at $\delta^t p_t$. We implicitly use the result that any financial sequential equilibrium with complete markets is also an Arrow-Debreu equilibrium.

1.3 Heterogeneous Preferences with External Habit

There is a set of infinitely lived agents, Γ. For now we assume that Γ is a compact set of positive values but the results hold for more general assumptions. All agents have the same type of time and state separable preferences,

$$U(c, X|\gamma) = \mathbb{E}_0 \sum_{t \geq 0} \delta^t u(c_t, X_t|\gamma), \tag{2}$$

where $\delta \in (0,1)$ is the common subjective discount factor and c_t is consumption at time t. X_t, with $x_t = \log(X_t)$, is the external habit, common to all agents. The external habit is an indicator of contemporaneous and/or past aggregate consumption. We will discuss its specification later.

The running utility is drawn from the "catching up with the Joneses" literature and is given by,

$$u(c, X|\gamma) = \frac{c^{1-\gamma} X^{\gamma-\rho} - 1}{1-\gamma}, \qquad \forall \gamma \in \Gamma.$$

γ is the coefficient of relative risk aversion of the agent, possibly different across agents, so that different types are characterized by their γ. Let τ be the inverse of γ, i.e., the coefficient of relative risk tolerance. Due to the homotheticity and time-separability properties of preferences, aggregation results hold for agents with the same type γ. Therefore, for our purpose we only need to specify the initial wealth distribution across the different types, which we denote by $\theta_0(\gamma)$. More precisely, $\theta_t(\gamma)$ denotes the proportion of wealth held by agents with type γ at time t, and therefore,

$$\int_\Gamma \theta_t(\gamma) d\gamma = 1, \qquad \forall t \geq 0.$$

Throughout the paper we will be using γ and τ interchangeably, so that when we write $\theta(\gamma)$ or $\theta(\tau)$ we imply the same distribution.

The parameter ρ is common to all agents and determines the relative effect that the external habit has on the marginal utility of each agent. The derivative of the marginal utility of consumption with respect to the external habit is given by,

$$\frac{\partial u_c(c, X|\gamma)}{\partial X} = (\gamma - \rho) c^{-\gamma} X^{\gamma-\rho-1}.$$

Since we would like to have a negative externality for all agents, we impose the restriction that $\rho \leq \min_{\gamma \in \Gamma} \gamma$. With a negative externality, an increase in the level of habit increases the value that each agent places on consumption. We also note that the smaller the habit parameter, the bigger is the effect of the habit on the marginal utility.

As noted by Chan and Kogan (2002), these preferences ensure that the curvature of the value function with respect to wealth is the same as that of the utility function with respect to consumption and the relative risk aversion with respect to wealth is still given by the parameter γ; this is the case because the multiplicative external habit does not affect the curvature of the value function.

1.4 Financial Equilibrium

As it will become clear later in this section, it is convenient to introduce the following variable

$$\omega_t = y_t - x_t, \tag{3}$$

which we will call endowment/habit ratio, for obvious reasons. The dynamics of ω depend on the dynamics of y, given by (1), and the specification -not yet provided- of x. Our results hold for a large class of specifications of x. We just need x to grow on average at the same rate as y and to be a functional of the current and/or past aggregate endowment y_s, $s \leq t$. Therefore ω is a stationary Markov process, and we treat it as our *state variable.* For example, in the continuous time model of Chan and Kogan (2002), x is a weighted average of past aggregate endowment, and the resulting ω is a stationary Markov process. Finally, we denote by $\bar{\omega}$ the unconditional average of the state which, in order to simplify notation, we assume is also the initial state of the economy.

Some of the following derivations are a discrete-time version of the results in Chan and Kogan (2002), and we include them for completeness. The main difference with Chan and Kogan (2002) is that they use as initial condition the weights the social planner assigns to the utility of each type (γ). Ideally, we would want to use as initial condition the wealth distribution across types, which has a clear economic interpretation. Furthermore, at the average state, characterized by $\bar{\omega}$, the equilibrium distributions of wealth and consumption are almost identical.[3] We then choose to use as initial condition the distribution of consumption in the average state as a proxy for the distribution of wealth. As we show later, this allows us to derive the risk aversion of the representative agent in closed form. Also, in our calibration -that we present in section 5- we use the distribution of consumption, instead of the distribution of wealth, for different groups of agents of the Consumer Expenditure Survey dataset.

At any time t, an agent type γ holds a positive proportion of the aggregate wealth $\theta_t(\gamma)$. Since we have complete markets, the budget constraint of each agent can be expressed as a single intertemporal budget constraint. At the initial period the intertemporal budget constraint of

[3] We compute numerically and very efficiently the equilibrium, including the resulting distribution of wealth across types.

an agent of type γ is,

$$\mathbb{E}_0 \sum_{t \geq 0} \delta^t p_t c_t(\gamma) \leq \theta_0(\gamma) p_0 (P_0^m + Y_0), \tag{4}$$

where $(\delta^t p_t, t \geq 0)$ is the equilibrium consumption price process. $\mathbb{E}_t$ is the expectation operator conditional on the consumption growth process up to time t or, alternatively, conditional on the history of the endowment habit ratio since, for a given specification of the external habit process x, we can infer the endowment process from the endowment habit ratio process. Therefore, it suffices to say that the information set is the endowment process and the initial value of the endowment habit ratio. Define also

$$z_t = \log p_t + \rho x_t, \tag{5}$$

which can be interpreted as a normalized and stationary pricing kernel because it is the endogenous process that, along with the exogenous process ω, will determine the equilibrium stochastic discount factor. From now on we will refer to it simply as "the pricing kernel." Given the pricing kernel process, the external habit process, and the initial price of the dividend-paying asset, agents optimally choose their consumption plan in order to maximize their utility. The following proposition characterizes the optimal consumption allocation as a proportion of the aggregate endowment, $\alpha_t(\gamma) = c_t(\gamma)/Y_t$.

Proposition 1. *The optimal consumption allocation of an agent type γ is characterized by,*

$$\alpha_t(\gamma) = \lambda(\gamma) \exp\left[-\frac{z_t}{\gamma} - \omega_t\right], \tag{6}$$

where $\lambda(\gamma)^{1/\gamma}$ is the Lagrange multiplier of the intertemporal budget constraint, and is given by,

$$\lambda(\gamma) = \theta_0(\gamma) \frac{\mathbb{E}_0 \sum_{t \geq 0} \delta^t p_t Y_t}{\mathbb{E}_0 \sum_{t \geq 0} \delta^t p_t Y_t \exp\left[-\frac{z_t}{\gamma} - \omega_t\right]}. \tag{7}$$

The optimality condition (6) is the same as equation (8) in Chan and Kogan (2002). However, in our case, the Lagrange multiplier $\lambda(\gamma)$ is endogenously determined, given the initial wealth distribution. It is well known that there is one-to-one mapping between the equilibria resulting from each set of conditions, but this subtle point has important quantitative implications, as we will see later on, due to the fact that the distribution of wealth across types is a key factor in the determination of the equilibrium. For example, if we increase the wealth held by the most extreme types, then the stochastic risk aversion of the economy is more volatile.

A financial equilibrium is a normalized pricing kernel process, $\{z_t, t \geq 0; z_0 = 0\}$ and a set of consumption allocation processes of ratios of the aggregate endowment, $\{\alpha_t(\gamma), t \geq 0; \gamma \in \Gamma\}$, such that consumption allocations satisfy the optimality conditions of the agents, and the consumption good market clears at all times. We have the following corollary.

Corollary 1. *In equilibrium, the pricing kernel is a function of the endowment/habit ratio and the initial wealth distribution, and is characterized by the following equation,*

$$1 = \int_\Gamma \lambda(\gamma) \exp\left[-\frac{z}{\gamma} - \omega\right] d\gamma, \tag{8}$$

where $\lambda(\gamma)$ is given by (7).

From equation (8) it is not feasible to derive z in closed form. However, it can be computed numerically with very high accuracy after we discretize the distribution of types and the state variable, and solve a large system of equations (a similar method is provided in Judd, Kubler and Schmedders (2003)).

One alternative approach that allows us to derive an expression for z and use it to solve for the risk aversion of the representative agent in closed form, is as follows. Instead of assuming the initial wealth distribution of types we assume that we know the initial (at the average state) consumption distribution. As we have argued before, at the average state is almost identical to the distribution of wealth (as we have verified computationally). In addition, from an empirical point of view, the distribution of consumption across types is as easily observable as the distribution of wealth.

We introduce an auxiliary concept. For a given endowment/habit ratio ω, we define the probability measure $\mathcal{P}_\omega(\tau)$ that assigns to agents of type $\gamma = 1/\tau$ probabilities equal to their equilibrium consumption share. Furthermore, let $\mathcal{E}_\omega$ denote the expectation operator under this probability measure. Then, from Proposition 1, a corollary follows,

Corollary 2. *The probability measure at the average state is given by,*

$$\mathcal{P}_{\bar{\omega}}(\tau) = \theta_0(\tau) \frac{\mathbb{E}_0 \sum_{t \geq 0} \delta^t p_t Y_t}{\mathbb{E}_0 \sum_{t \geq 0} \delta^t p_t Y_t e^{-\tau z_t - (\omega_t - \bar{\omega})}}, \tag{9}$$

and, therefore, the following relation holds,

$$\exp(\omega_t - \bar{\omega}) = \mathcal{E}_{\bar{\omega}} \left[\exp(-\tau z_t)\right], \qquad \forall t \geq 0. \tag{10}$$

We point out that the equilibrium pricing kernel z is a function of the initial distribution of agents -as implied by the expectation operator $\mathcal{E}_{\bar{\omega}}$-, and the state of the economy ω. We also point out that the initial consumption distribution $\mathcal{P}_{\bar{\omega}}$ is very close to the wealth distribution $\theta_0(\tau)$, since the fraction on the right-hand side of (9) is close to one. This is due to the fact that, for any τ, $e^{-\tau z_t-(\omega_t-\bar{\omega})}$ does not vary much in equilibrium and is "centered" around one, its value at the average state.

In addition, the right-hand side of (10) is the moment generating function of τ and hence, for certain distributions it is known and it is straightforward to derive z. However, in general, it is easy to approximate with high precision the pricing kernel by discretizing the distribution and using numerical integration methods. When agents are identical, the pricing kernel is linear in the state, $z(\omega) = -\gamma(\omega - \bar{\omega})$.

In the following lemma, we can analyze the properties of the pricing kernel with the help of the probability measure we just introduced.

Lemma 1. *The pricing kernel is a continuous function of the state with a negative first derivative and positive second derivative,*

$$z'(\omega) = -\frac{1}{\mathcal{E}_\omega(\tau)}, \tag{z2}$$

$$z''(\omega) = \frac{\mathcal{E}_\omega(\tau^2) - \mathcal{E}_\omega(\tau)^2}{\mathcal{E}_\omega(\tau)^3}. \tag{z3}$$

Furthermore,

$$\lim_{\omega\to\pm\infty} z(\omega) = \mp\infty, \qquad \lim_{\omega\to\pm\infty} z''(\omega) = 0,$$

$$\lim_{\omega\to+\infty} z'(\omega) = -\min_{\gamma\in\Gamma}\gamma, \qquad \lim_{\omega\to-\infty} z'(\omega) = -\max_{\gamma\in\Gamma}\gamma.$$

As a natural extension of what happens in an economy populated by a single agent, equation ($z2$) shows that the slope of the pricing kernel is negative and equal to the inverse of the weighted average risk tolerance in the economy. If time were continuous $|\sigma z'(\cdot)|$ would correspond to the price of risk where σ is the volatility of consumption growth. Bhamra and Uppal (2007), as well as Gârleanu and Panageas (2008) also show that the price of risk is determined by the weighted harmonic average of the risk-aversion of the agents. Furthermore, from ($z3$), the curvature of the pricing kernel depends on the dispersion of the risk tolerance types, which from ($z2$) implies that pricing kernel and average risk tolerance in the economy are more volatile when the dispersion of types is higher. This is due to the fact that the more variability there is in types, the more extreme the investment positions of the agents are, and this leads to bigger changes in the cross-sectional wealth distribution.

2 The Representative Agent Equivalent Economy

In this section we construct an economy with a representative agent with state-dependent risk-aversion parameter equivalent to the heterogeneous agents economy. We say that two economies are equivalent when they have the same aggregate endowment process, financial market structure and pricing kernel process. We will derive an expression for the stochastic risk-aversion parameter of the representative agent and study its properties. The risk-aversion of the representative agent provides a natural measure of risk-aversion in the heterogeneous agents economy.

2.1 Preferences and Equilibrium

In this economy, there is a representative agent with risk-aversion coefficient that depends on the state ω. The representative agent has the following utility,

$$U_r(c,x) = \mathbb{E}_0 \sum_{t \geq 0} \delta^t \frac{c_t^{1-\gamma(\omega_t)} (X_t^r)^{\gamma(\omega_t)-\rho} - 1}{1-\gamma(\omega_t)},$$

where the external habit is $X_t^r = X_t e^{\bar{\omega}}$, for all $t \geq 0$. This change in the definition of the habit process guarantees that the stochastic risk-tolerance of the representative agent is equal to the average (using the probability measure $\mathcal{P}_\omega$) risk-tolerance of the economy at the average state. For the standard definition, this would be true at the state $\omega = 0$. Since we would like the risk-aversion of the representative agent to be a good representation of the risk-aversion in the economy, this choice seems more appropriate. As in the heterogeneous agents economy, we require a negative externality from the habit, and for that we need the habit parameter to be always less than the risk-aversion parameter, i.e. $\rho \leq \min_\omega \gamma(\omega)$.

For a process $(p_t, t \geq 0)$, the representative agent maximizes the previous utility subject to the intertemporal budget constraint. Financial equilibrium obtains when there is a pricing kernel process $(z_t^r, t \geq 0; z_0^r = 0)$ such that the consumption process $(Y_t, t \geq 0)$ is optimal for the representative agent. The following proposition states the result.

Proposition 2. *The equilibrium pricing kernel of the representative agent economy is a function of the state, and is characterized by the following equation,*

$$1 = \exp\left[z_t^r + \gamma(\omega_t)(\omega_t - \bar{\omega})\right], \qquad \forall t \geq 0. \tag{11}$$

For the representative agent economy to be equivalent to the heterogeneous agent economy we need to have the same equilibrium pricing kernel processes in the two economies. Hence, we require that $z(\omega) = z^r(\omega)$ for all ω. This gives rise to the following corollary.

Corollary 3. *Let $z(\omega)$ be the equilibrium pricing kernel of the heterogeneous agent economy. Then the representative agent economy is an equivalent economy if the stochastic risk aversion of the representative agent is the following continuous function of the state,*

$$\gamma(\omega) = \begin{cases} -\dfrac{z(\omega)}{\omega - \bar{\omega}}, & \omega \neq \bar{\omega} \\ -z'(\bar{\omega}) = \dfrac{1}{\mathcal{E}_{\bar{\omega}}(\tau)}, & \omega = \bar{\omega} \end{cases} \tag{12}$$

We point out that $\gamma(\bar{\omega})$ can be set to any value,[4] but the choice in (12) makes the stochastic risk aversion function continuous at $\bar{\omega}$. From the properties of the pricing kernel function we can derive certain properties of the risk-aversion function. We summarize them in the following corollary.

Corollary 4. *The stochastic risk aversion function $\gamma(\omega)$, characterized in (12), has the following first and second derivatives,*

$$\gamma'(\omega) = -\frac{1}{\omega - \bar{\omega}} \left[\gamma(\omega) - \frac{1}{\mathcal{E}_{\omega}(\tau)} \right] < 0, \qquad (\gamma 1)$$

$$\gamma''(\omega) = \frac{1}{\omega - \bar{\omega}} \left[-2\gamma'(\omega) - z''(\omega) \right]. \qquad (\gamma 2)$$

Furthermore,

$$\lim_{g \to \pm\infty} \gamma'(\omega) = 0, \qquad \lim_{g \to \pm\infty} \gamma''(\omega) = 0$$

$$\lim_{g \to +\infty} \gamma(\omega) = \min_{\gamma \in \Gamma} \gamma, \qquad \lim_{g \to -\infty} \gamma(\omega) = \max_{\gamma \in \Gamma} \gamma.$$

The negative first derivative of the risk-aversion function implies the counter-cyclicality of risk aversion in the economy. The risk aversion of the economy moves from the highest level to the lowest as the endowment/habit ratio goes from minus infinity to plus infinity. Intuitively, more risk-tolerant agents invest more in the stock market and, therefore, end-up wealthier in good states (high endowment/habit ratio) and poorer in bad states.

3 Agent Distribution and the Variation of $\gamma(\cdot)$

As we have explained before, it is straightforward to find the equilibrium pricing kernel numerically. However, for our analysis, it is convenient to make the following parametric assumption about $\mathcal{P}_{\bar{\omega}}(\tau)$: We assume that at the average state $\bar{\omega}$ the risk-tolerance is gamma distributed,

[4] In equation (11) we have the term $\gamma(\omega_t)(\omega_t - \bar{\omega})$, therefore it does not matter what $\gamma(\bar{\omega})$ is, since it is always multiplied by zero.

with parameters κ and ϑ,

$$\tau(\bar{\omega}) \sim Gamma(\kappa, \vartheta),$$

so that $\mathcal{E}_{\bar{\omega}}(\tau) = \kappa\vartheta$ and $\mathcal{V}_{\bar{\omega}}(\tau) = \kappa\vartheta^2$, where $\mathcal{V}$ denotes the variance. With this assumption, we can derive the pricing kernel, and hence the risk-aversion function, in closed form. In order to study and analyze the predictions of the model, we need estimates of the parameters of the distribution. For that purpose, we use the empirical results of Kimball, Sahm and Shapiro (2008) (henceforth KSS), who through survey responses on hypothetical income gambles, construct an empirical distribution of risk-tolerance, and then fit it to a log-normal distribution. However, the empirical distribution of KSS is in terms of number of people, whereas our assumed distribution is in terms of consumption weights, $\alpha(\gamma)$. Therefore, we make the implicit assumption that at the average state the distribution of consumption is independent of the distribution (in terms of mass of agents) of the risk tolerance. Formally, let $\mathcal{I} = [0, 1]$ be the mass of agents and let $\tilde{\alpha}(i)$ the consumption distribution over the set $\mathcal{I}$ at the average state $\bar{\omega}$. Then we assume that

$$\int_{\mathcal{I}} \mathbf{1}(\tau_i = \tau)di = \int_{\mathcal{I}} \tilde{\alpha}(i)\mathbf{1}(\tau_i = \tau)di,$$

where $\mathbf{1}(\cdot)$ is the indicator function. This assumption, even though necessary, seems natural because it applies to the average state of the economy. If we were to actually measure the consumption shares of the people in the survey we would also have to determine the state of the economy because, unlike in the average state, we would expect to see a positive (negative) correlation of risk-tolerance with consumption shares at a good (bad) state of the economy.

We do not use the log-normal distribution for the model economy because it does not have a moment-generating function. Thus, we choose the closely related gamma distribution.[5] We fit the gamma distribution by minimizing the overall distance between the distribution of risk-aversion implied by the log-normal distribution and that implied by the gamma distribution over a discretized set from 0 to 200. Figure 1 plots the cumulative and density functions of the log-normal distribution of KSS, and our fitted gamma distribution. We observe that they are very close to each other, but the gamma distribution exhibits a fatter tail. This implies slightly higher mean and standard deviation of risk aversion. Let $\bar{\gamma}$ and $\bar{\nu}$ denote, respectively, the inverse of the average risk-tolerance and the standard deviation of risk-tolerance, both at the average state. The estimated parameters imply that $\bar{\gamma} = 5.17$ and $\bar{\nu} = 0.13$.

Given our parametric assumption about the cross-sectional consumption distribution of types we have the following corollary.

Corollary 5. *Let us define* $\eta = (\bar{\gamma}\bar{\nu})^2$. *When the cross-sectional dispersion of types (risk-tolerance) at the average state is gamma distributed with mean* $1/\bar{\gamma}$ *and standard deviation* $\bar{\nu}$*,*

[5] It is straightforward to verify numerically that the results are insignificantly different when we assign weights to a discretized distribution of risk-aversion between 0 and 200 using the log-normal distribution instead.

and weight for each type given by the consumption share, the equilibrium pricing kernel is

$$z(\omega) = \bar{\gamma}\frac{\exp\left[-(\omega-\bar{\omega})\eta\right]-1}{\eta}.$$

We point out that as the level of heterogeneity in the economy $\bar{\nu}$ tends to zero, the pricing kernel tends to $-\bar{\gamma}(\omega-\bar{\omega})$, as noted in the discussion of Corollary 2. We now can study the variability of the risk-aversion of the representative agent. First, we recall from (12) that the risk-aversion of the representative agent is equal to $\bar{\gamma}$. In addition, we define the following function,

$$h(\omega) = \frac{\gamma(\omega)}{\bar{\gamma}},$$

that is, the the ratio of the risk-aversion coefficient of the representative agent in state ω, given in (12), to the risk-aversion coefficient in the average state $\bar{\omega}$. It represents the coefficient of variation (we will call it "multiplier") of risk-aversion in a given state with respect to the average state. Using the result of corollary 5, we can express it in closed form:

$$h(\omega) = \frac{1-\exp\left[-\eta(\omega-\bar{\omega})\right]}{\eta(\omega-\bar{\omega})}.$$

Figure 4 plots h for three different values of $\bar{\nu}$, the value empirically estimated using the distribution of KSS, twice this value, and one half this value; the value of $\bar{\gamma}$ is the one estimated given the distribution in the same paper. The function is plotted for deviations of the average state ranging between -1 to $+1$. As we will argue later, this kind of range is unrealistically large. It is more natural to expect that the deviations will be less than 0.5 in absolute value: a deviation of 0.5 implies that the aggregate consumption surplus ratio is higher than the average by around 65%. From figure 4 we observe that the possible variation in the coefficient of risk-aversion for the representative agent is very small. Unless we assume a level of heterogeneity twice as much as the estimated level, the risk-aversion of the representative agent is not expected to deviate from the average risk aversion by more than 20% at any time. Even for the most extreme case we consider, the risk-aversion of the economy doubles only when the economy is in very bad times. This plot is a first indication that the effect of risk-aversion heterogeneity on asset prices, and particularly on risk-premia, is potentially small.

In an economy with rational investors the risk-attitude of the representative agent can be time-varying due to two reasons. The first one, which drives the risk-aversion coefficient in this model, is the evolution of the cross-sectional wealth distribution. Unless the variation in the state vector that determines the cross-sectional wealth distribution is very high, the wealth reallocation across time cannot have a substantial effect on asset prices, as we have already seen. A second possible reason for time-variation in the risk-aversion coefficient of the representative agent is time-variation in the individual risk-aversion coefficients. In order to have a high variation in the risk-preferences of the representative agent we would also need individual preferences to be moving together, and in the same direction. In particular the entire distribution needs to be moving up and down with the state variable.

4 Habit Process and Asset Prices

In order to solve for equilibrium asset prices, we need to specify a process for the external habit. Following Chan and Kogan (2002) and a good part of the literature that uses "catching up with the Joneses" preferences, we assume that the current value of the habit is a weighted average of the previous habit and aggregate endowment levels,

$$x_{t+1} = \lambda y_t + (1 - \lambda) x_t. \tag{13}$$

From (13) the endowment/habit ratio, the state variable, follows a mean-reverting process,

$$\omega_{t+1} - \omega_t = -\lambda(\omega_t - \bar{\omega}) + \sigma \epsilon_{t+1}. \tag{14}$$

The unconditional mean of the state variable is $\bar{\omega} = \mu/\lambda$, and the unconditional volatility is $\sigma_\omega = \sigma/\sqrt{\lambda(2-\lambda)}$. The reversion rate parameter λ is particularly important in this model because it determines the likely "range" in which the state variable moves through the dependence of σ_ω on λ. The smaller the value of λ, the bigger the unconditional variance. Figure 4 shows that the risk-aversion of the representative agent varies more the larger the "range" of the state variable. Therefore, as λ decreases, the potential effect of agent heterogeneity on asset prices becomes larger. When the state of the economy moves further away from its average state, the wealth allocation tends to be concentrated on the more or less risk-averse agents, depending on whether the deviation is negative or positive, respectively. Hence, with a smaller rate of reversion (equivalent to higher persistence in the state variable) larger reallocations of wealth and, therefore, variations in the risk-aversion of the economy, are possible. However, these large swings in risk-aversion require a long time. For example, Campbell and Cochrane (1999) use a reversion rate of 0.13 which implies a half-life of around 5 years,[6] while Chan and Kogan (2002) use values that imply half-lives of 12 years for their heterogeneous agents economy, and around 17 years for their single agent economy. Persistence in the state variable translates into persistence in the price-dividend ratio. It is natural therefore to select this parameter in order to match the price-dividend ratio persistence implied by the model to the persistence observed in the data.

[6] The half-life is the time required for the deterministic version of the process to cover half of the distance to the unconditional mean. It is given by $-\log(2)/\log(1-\lambda)$

4.1 The Stochastic Discount Factor

The fundamental price of an asset is the expected value of discounted future dividends. Let M_{t+1} ($m = \log(M)$) denote the one-period stochastic discount factor between periods t and $t+1$. Since p_t denotes the price of a unit of consumption in period t, the stochastic discount factor is,

$$M_{t+1} = \delta \frac{p_{t+1}}{p_t} = \delta e^{z_{t+1} - z_t - \rho(x_{t+1} - x_t)}, \qquad \forall t > 0.$$

Using equation (5) and our other assumptions, we have the following corollary:

Corollary 6. *Assume that the habit process is as in (13) and the cross-sectional distribution of types (risk-tolerance) with respect to their consumption share at the average state is gamma distributed, with mean $1/\bar{\gamma}$ and standard deviation $\bar{\nu}$. The equilibrium one period log stochastic discount factor is conditionally log-normally distributed,*

$$m_{t+1} = \log(\delta) - \rho\lambda\omega_t - \frac{\bar{\gamma}}{\eta} e^{-\eta(\omega_t - \bar{\omega})} \left(1 - e^{\lambda\eta(\omega_t - \bar{\omega}) - \sigma\eta\epsilon_{t+1}}\right).$$

When agents are homogeneous, i.e. $\bar{\nu} = 0$,

$$m_{t+1} = \log(\delta) - \rho\lambda\omega_t - \bar{\gamma}(\omega_{t+1} - \omega_t).$$

We observe that the stochastic discount factor of both the standard Lucas tree and the model of Campbell and Cochrane (1999) are particular cases of the stochastic discount factor of corollary 6. Campbell and Cochrane (1999) assume homogenous agents, but their state variable, the consumption surplus ($-x$, using our notation) satisfies some specific dynamics that, Chan and Kogan (2002) argue, might be obtained as the result of simpler dynamics and agents with heterogeneous risk-aversion.[7] One of our objectives is to study this point further. For example, Campbell and Cochrane (1999) assume that the conditional volatility of their state variable which is the main factor of their model risk-premium is time-varying and counter-cyclical. The time-varying conditional volatility has been interpreted as the economy's risk-aversion. In our model (as in Chan and Kogan 2002), the equilibrium risk-aversion of the representative agent

[7] In the representative agent case of Campbell and Cochrane (1999), $s = -x$ is the state-variable and not ω, since it is Markov stationary, following the process,

$$s_{t+1} - s_t = -(1-\phi)(s_t - \bar{s}) + \lambda(s_t)\sigma\epsilon_{t+1},$$

where ϵ is the aggregate endowment growth innovation and $\lambda(x)$ is some given function. Hence we have,

$$\omega_{t+1} - \omega_t = \mu - (1-\phi)(s_t - \bar{s}) - [1 + \lambda(s_t)]\,\sigma\epsilon_{t+1}.$$

The persistence parameter $(1-\phi)$ plays a role similar to our parameter λ. Finally, since $\rho = 0$ we obtain the same stochastic discount factor.

turns out to be negatively related to the state of the economy ($\gamma'(\omega) < 0$). However, we have also showed that the possible variation of the risk-aversion of the representative agent (figure 4) is relatively modest for realistic parameter values, and it seems difficult to argue that it can explain the variation in the risk-premium observed in the data. We next elaborate further on this point.

The main driving forces of Campbell and Cochrane (1999) are: (i) the persistence in the consumption surplus ratio, which produces the persistence in price dividend ratios and the variability of stock expected returns; (ii) the counter-cyclical conditional volatility of the state variable (their consumption surplus ratio or in our case of x). In Campbell and Cochrane (1999), the varying conditional volatility is chosen so that it fixes the risk-free rate at a certain level, and the entire variation in expected returns translates into variation in risk premia. In our model the persistence of the state variable is the result of assuming persistence in habit, but the varying conditional volatility of the stochastic discount factor is endogenous and related to the variation in risk-aversion. To make our exposition simple we will use a first order approximation to $z(\omega_{t+1}) - z(\omega_t) \approx z'(\omega_t)(\omega_{t+1} - \omega_t)$ to obtain an approximate stochastic discount factor,

$$\tilde{m}_{t+1} = \log(\delta) - |z'(\omega_t)|\mu + \left(|z'(\omega_t)| - \rho\right)\lambda\omega_t - |z'(\omega_t)|\sigma\epsilon_{t+1} \tag{15}$$

where

$$-z'(\omega) = \bar{\gamma}e^{-\eta(\omega-\bar{\omega})}. \tag{16}$$

The above expression of $\tilde{m}$ is the true stochastic discount factor if time was continuous or if agents were homogeneous. $\tilde{m}$ is conditionally normally distributed with an endogenously varying conditional volatility equal to $|z'(\omega)|$ which determines the price of risk. Therefore, $\tilde{m}$ has a similar form as the stochastic discount factor of Campbell and Cochrane (1999). The main difference is that Campbell and Cochrane (1999) choose an exogenously given price of risk that varies considerably more than ours, as we show next, in order to explain the variability of the price-dividend ratio with a constant risk-free rate.

In figure 5 we plot the function $|z'(\omega)|$ for three different levels of agent heterogeneity. The label of each line is the number that multiplies the value of the risk-tolerance standard deviation $\bar{\nu}$ estimated by KSS. The price of risk in Campbell and Cochrane (1999) ranges in value from around 0 to 100. Clearly, the level of variation that can be generated endogenously in our economy is substantially smaller. First we see that for the estimated parameters (graph labeled "1") the variation in the shown region is from about 3 to about 8. Furthermore, for a realistic set of parameters the state variable is not expected to move more than 20% away from the mean. All these imply that our economy will be able to predict variation in risk-premia of about an order of magnitude less than what is required by Campbell and Cochrane (1999). In order to get significant variation we will need at least a doubling of the level of heterogeneity, as well as an increase in the variation of the state variable, $\sigma_\omega = \sigma/\sqrt{\lambda(2-\lambda)}$. This could be achieved by a substantial increase in consumption risk or by a high level of persistence. The one-lag autocorrelation of the price-dividend ratio in annual data is about 0.88, and for this reason λ should be set to 0.12, a parameter value also selected by Campbell and Cochrane (1999).

Figures 4 and 5 also serve to understand the results of Chan and Kogan (2002), which can be summarized by their Sharpe ratio. In their model, the Sharpe ratio has an average of about 0.32 and it varies from around 0.5 down to about 0.2 over the relevant region of the state variable. In order to get an average of 0.32 they need assume that the standard deviation of consumption growth is 4% while the risk-aversion of the economy at the average state (which is equal to $z'(\bar{\omega})$) is around 8. From the results of KSS we use an average of around 5 while we estimate the annual standard deviation of consumption growth to be slightly higher than 2%. Furthermore, they assume the persistence for the external habit to be around 0.94, which along with the 4% consumption risk yields a standard deviation of the state variable a bit higher than 0.1. Hence the relevant region is about a 40% deviation around the average state. The distribution of risk-aversion they assume is close to the distribution that we get if we double the standard deviation of risk-tolerance in our fitted distribution. Therefore, the relevant graphs in figures 4 and 5 are those with the label 2. Over the relevant region (deviations of 40%) the coefficient of variation of the risk-aversion of the economy is from about 1.5 down to about 0.7. Therefore the variation of the risk-aversion is from about 12 to about 5 which along with the assumed standard deviation of consumption growth gives the indicated variation in the Sharpe ratio.

This shows unambiguously that both a high level of heterogeneity and a high unconditional variance of the state variable are required in order to explain the high variation in the Sharpe ratio. The empirical results of KSS as well as our empirical results with respect to the cross-sectional variation in the consumption growth exposure to market returns -to be shown later-indicate that doubling the risk-tolerance standard deviation yields an unrealistically high level of heterogeneity. Also, Chan and Kogan (2002) obtain their high variability in the state variable in part through their choice of the level of persistence. The trade-off is that with such high persistence a full financial cycle would require many more years than what we observe, and the price-dividend ratio would be excessively persistent and volatile. In addition, the predictability of future excess returns by the price-dividend ratio will be over much longer periods than we detect in the data.

4.2 Asset Prices

The assets of interest are the risk-free bond, that pays a unit of consumption next period, and the infinitely lived market security, that pays the aggregate endowment. The price of the risk-free bond is

$$P^f(\omega_t) = \mathbb{E}_t\left[e^{m(\omega_t,\omega_{t+1})}\right].$$

The price of the market security is increasing in the level of endowment, but the price-endowment (or price-dividend) ratio, that we denote PD, is stationary,

$$PD(\omega_t) = \mathbb{E}_t\left[e^{m(\omega_t,\omega_{t+1})+y_{t+1}-y_t}\left(PD(\omega_{t+1})+1\right)\right].$$

In the appendix we describe the method we follow to compute the prices numerically.

The continuously compounded risk-free rate is the negative log of the bond price. Using (15) we can derive a good approximation,

$$\tilde{r}_t^f = -\log(\delta) + \mu|z'(\omega_t)| - \left(|z'(\omega_t)| - \rho\right)\lambda\omega_t - \frac{1}{2}\left(|z'(\omega_t)|\sigma\right)^2. \tag{17}$$

The model of Campbell and Cochrane (1999) explains the observed low volatility of the interest rate by assuming that the precautionary savings term is inversely proportional to the habit term.[8] The "habit term" -which in equation (17) is $-\left(|z'(\omega_t)| - \rho\right)\lambda\omega_t$- measures the incentive to postpone consumption and therefore save more when consumption is high today with respect to habit. The "precautionary savings term" -which in equation (17) is $-\frac{1}{2}\left(|z'(\omega_t)|\sigma\right)^2$- measures the incentive to save less when the real risk or risk premium in the economy is low. In our model these two terms are inversely proportional to each other since the risk-aversion of the economy, and hence the price of risk, is counter-cyclical in the state variable. However, the habit term dominates the precautionary savings term unless the level of heterogeneity in the economy is very high, and the fundamental risk of the economy is significantly higher than in reality. We show this next.

In (17) we have an additional term, $\mu|z'(\omega_t)|$, that causes variation which is related to the varying elasticity of intertemporal substitution but, compared to the habit term, its effect is negligible. In a homogeneous agent economy or in an economy with a low level of heterogeneity, almost all variation in interest rates is caused by the habit term since the variation of the precautionary term is very small. The level of heterogeneity required for the risk-premium to offset the variation in the risk-free rate due to the habit term is unreasonable. A back of the envelope calculation based on the derived price of risk reveals that the standard deviation of risk-tolerance in the average state required to keep the risk-free rate close to constant is at

[8] In fact the conditional volatility of the pricing kernel of their model is derived by making the risk-free rate constant.

least 1.37, about 10 times bigger than the estimated value.[9] The standard deviation of the risk-tolerance estimated by KSS is about 0.17, while our fitted gamma distribution gives a risk-tolerance standard deviation of about 0.13. The number computed is based on the approximate expression,

$$\nu \approx \frac{\sqrt{2\lambda}}{\gamma\sigma\sqrt{\bar{\gamma}}},$$

after imposing the condition of a constant risk-free rate, and using a value of 10 for the average risk-aversion, a value of 0.12 for the persistence parameter, and consumption risk of 2.18%. The derived expression is very useful because it provides information about the relative effect of the different factors. First we observe the intuitive relation that a higher level of risk (σ) requires a lower level of heterogeneity. We also note that a higher level of persistence (lower λ) also requires a lower level of heterogeneity, since the state variable becomes more volatile.

5 Calibrating the Economy

The focus of this paper is to assess quantitatively up to what extent risk-aversion heterogeneity can explain financial variables observed in the economy. The main economic implications of the model come from two key elements. First, the persistent habit, which is able to produce persistence and high variability in the price dividend ratio due to persistence and high variability in discount rates. Second, the endogenously generated counter cyclical conditional volatility of the stochastic discount factor due to the heterogeneity in risk-preferences. More heterogeneity induces agents to take more extreme positions, and this leads to higher variability in cross-sectional wealth and consumption distributions, and hence higher variability in the risk-aversion of the representative agent of the economy. Essentially, the main question we try to address with this model is whether risk-preference heterogeneity can produce enough variation in the price of risk so as to be able to explain the low variability in the risk-free rate, the variation in equity premia and the long-run predictability of excess returns.

First and foremost, we need to calibrate our model to the data. For this purpose, we use two sources: First, as described before, we parameterize the distribution of types using the empirical findings of KSS who use surveys of hypothetical income gambles and assume power utility

[9] We impose the condition $r^f(\omega) = r^f(\bar{\omega})$ and then rearrange to get the following expression,

$$\exp\left[-\eta(\omega-\bar{\omega})\right] = 1 - \frac{2\lambda}{\bar{\gamma}\sigma^2}(\omega-\bar{\omega})\left(1+\frac{\bar{\gamma}-\rho}{|z'(\omega)|-\bar{\gamma}}\right).$$

We ignore the fraction in the second parenthesis in the right hand side, which would require an even higher level of heterogeneity, and a first order approximation yields the following condition,

$$\eta \approx \frac{2\lambda}{\gamma\sigma^2}.$$

preferences;[10] their study provides plausible levels of risk-aversion and probably a good estimate of the cross-sectional variation. Second, we use individual consumption data as obtained by the Consumption Expenditure Survey (CEX) conducted by the Bureau of Labor Statistics from 1980 to 2005; the cross-sectional variation of individual consumption growth responses to market returns or changes in the level of the market gives us an estimate of the cross-sectional variation of risk-tolerance. We discussed before the first calibration. In the remainder of this section we analyze the second. We first calibrate the cross-sectional distribution of risk-tolerance, and then calibrate the rest of the parameters using market data.

5.1 Cross-Sectional Distribution

We can infer the heterogeneity in risk-aversion from the cross-sectional variation of the covariance between the consumption of the different individuals and stock returns. In particular, according to our model, the individual consumption growth processes is given by:

$$g_{t+1}(\gamma) = \mu + \lambda(\omega_t - \bar{\omega}) - \frac{1}{\gamma}\Delta z_{t+1},$$

where $g_{t+1}(\gamma)$ is the one period log-consumption growth of an individual with risk-aversion γ and $\Delta z_{t+1} = (z_{t+1} - z_t)$. The change in the pricing kernel Δz_{t+1} is not observable but it is highly correlated with the change in the price-dividend ratio or the return on the stock market. Therefore, we wish to run the following regression,

$$g^i_{t+1} = b^i_0 + b^i_1 \Delta p_{t+1}$$

across different agents, and compare the coefficients b^i_1. The ratio of the coefficients b of two different consumers is an estimate of the ratio of their risk-aversion coefficients.

As a measure of the price-dividend ratio we chose to use the price-earnings ratio from Robert Shiller's website. We also present results of regressions on the CRSP value weighted index return. Individual consumption processes were obtained from the Consumer Expenditure Survey (CEX) data.

CEX data go from the start of 1980 to the first quarter of 2005. About 5,000 households are interviewed every quarter prior to 1999 and the sample increases to about 7,500 households afterwards. Every household is interviewed 5 times, three months apart from each other, but the first interview is not included in the data. Interviews are spread out over every quarter and therefore individual quarterly consumption growth data is available every month. Financial information is gathered at the last interview. About 40% of the households have one or more

[10] Figure 1 provides the empirically estimated log-normal distribution of relative risk-tolerance of KSS and our fitted gamma distribution.

interviews missing and are excluded from the dataset. Also, we drop data before 1982, since the definition of food consumption changed, causing an artificial change in consumption values. There is missing data at the beginning of 1986 and 1996 due to changes in the household identification numbers. We use other standard filters, similar to those of Malloy, Moskowitz and Vissing-Jorgensen (2008) . We exclude non-urban households and households that reside in student housing. We also exclude cases were quarterly log consumption growth was either very low (< -1.6) or veryy high (> 1.6). We also exclude households that report consumption for more or less than three months or report zero or negative consumption for one or more quarters.

The total consumption of each household every quarter includes expenditure on non-durable goods and services that were aggregated from the various reported categories in order to match as closely as possible the consumption definitions of non-durables and services of the National Income and Product Accounts. The log real quarterly consumption growth for each household is then computed after adjusting for inflation with the Consumer Price Index for urban households constructed by the Bureau of Labor Statistics.

Since there are only three observations of consumption growth for each household, it is necessary to construct groups of households that we expect to have similar risk-aversion coefficients and create long series of consumption growth processes. The members of a group are chosen according to their holdings of "stocks, bonds, mutual funds and other such securities" reported at the final interview compared to their total consumption over their four interview quarters. We exclude households with zero reported investments because we consider them market non-participants. Given the limitations of the data, we ignore that differences in individual income risk or differences in beliefs can also be determining factors of household holdings of securities.

For every month of the data sample we collect all the households that report investment holdings at the end of the particular month and compute the ratio of holdings over their total reported consumption,

$$ratio^i_m = \frac{investment^i_m - 1}{\sum_{j=1}^{12} consumption^i_{m-j}}$$

where m is the month of the final interview. Using the holdings/consumption ratio we allocate all the households of our dataset into eight groups. For this purpose, we fit a log-normal distribution to the ratio of holdings to consumption every month of our data period, and allocate households using quantiles.

From the final dataset we excluded households that report their investments in December, since differences in holdings across households at the end of every year might be due to tax selling. The reported amounts of "stocks, bonds, mutual funds and other securities" were top-coded in the dataset. This means that holdings above a certain amount were replaced by the maximum amount allowable for reporting. This means that the computed ratios for top-coded households are lower than their true ratios. For this reason we exclude top-coded households that did not fall in the highest ratio group. Fortunately, there were not many such cases and the results are not significantly affected by this adjustment.

In the next step, for every month, we compute the quarterly consumption growth of each group in three ways: (a) First, by averaging the consumption growth across all households within each group. (b) Second, by taking a weighted average, using as weights the consumption shares of the initial quarter for which the consumption growth was reported. (c) Third, by summing up the consumption of all households the previous quarter and the current quarter and then computing the consumption growth of the whole group. The timing convention we use is that consumption growth g_{t+1} represents the quarterly consumption growth over the previous period, which includes consumption from months $t-4$ to $t-2$, and the current consumption is over the months $t-1$ to $t+1$. The price change or return over the same period is from month $t-2$ to month $t+1$.

Table 1 displays the consumption growth sensitivities to changes in the price-earnings ratio of all the CEX groups for the three different consumption growth series, where (a) refers to simple average, (b) to weighted average and (c) to group consumption growth, as described before. We first observe that there is some variation in the estimated coefficients over the different groups, and an indication of increasing sensitivity with the group number. The exception is the 5^{th} group, but its coefficients are statistically insignificant. In table 2 we show the results of using returns of the CRSP valued-weighted index. We observe higher variation over the coefficients but with lower statistical significance. Overall, these results give us some confidence that the constructed groups capture significant cross-sectional variation in sensitivities to market changes, possibly coming from different risk-aversion coefficients.

To finalize the derivation of the cross-sectional distribution of risk-aversion, we use two additional pieces of information. First, we need to assign a consumption share to each of our eight groups at the average state. Since we do not know when the average state takes place, we use the average consumption shares over the whole period. Second, we know the ratios of the risk-aversion coefficients across the different groups, but we need to pin down the distribution in one point. For that purpose, we assume that the average risk-tolerance is 1/5.17, which is the one that results from the gamma distribution that was fitted to the empirical distribution of KSS.

In figures 2 and 3 we show the resulting distributions for regressions using the price-earnings ratio and the market returns, respectively, along with the fitted gamma distribution to the results of KSS, for comparison purposes. The distribution implied by the regression results with the price-earnings ratio shows overall smaller variation than the fitted gamma distribution. The regression results of the market returns exhibit almost the same variation as the fitted gamma distribution. In light of these results, it seems clear that the distribution estimated by KSS is a good proxy for cross-sectional variation of risk-aversion. Additionally, it exhibits higher variation than the distributions implied by the regression results. For that reason, we will use it for our analysis in the next section, so that we get an upper bound of how much variability in the data can be explained by heterogeneity of risk-aversion.

5.2 Choice of Parameter Values

The second channel through which agent heterogeneity affects prices is the unconditional volatility of the state variable. If the state of the economy is very volatile, then there is more wealth re-distribution over time and, therefore, more volatility in the risk-aversion of the representative agent. The unconditional volatility of our state variable depends on the persistence parameter λ and the volatility of consumption growth. To estimate the mean and standard deviation of aggregate consumption we use NIPA data on real consumption growth between 1930 and 2005. The persistence parameter is not directly observable, but can be selected so as to match the persistence it induces in the price-dividend ratio: We estimate λ by fitting the model implied price-dividend ratio autocorrelation function (up to lag 7) to the autocorrelation parameter found in the data. We use the annual series of Boudoukh et al. (2007), which includes common share repurchases from cash flow statements.

Both the persistence parameter λ and the habit parameter ρ have a strong effect on the unconditional volatilities of the price-dividend ratio and the risk-free rate. In our calibration exercise we include both volatilities. We have already listed our source for the price-dividend ratio. For the (real) risk-free rate we take the yield of the 3-month Treasury bill after subtracting the realized inflation provided in NIPA. The risk-free rate volatility for the full sample is slightly over 3%. The post-war period, during which inflation was more predictable, the risk-free rate volatility drops to a bit less than 2%. For this reason, we put a smaller weight on the risk-free rate volatility in the calibration exercise.

The last parameter we need to calibrate is the subjective discount factor δ. This parameter affects the average level of the price-dividend ratio, as well as the average risk-free rate. Since we are unable to fit both of them at the same time, following the emphasis of the literature, we exclude the average risk-free rate and focus on the average price-dividend ratio. This underscores the inability of the model to explain the average excess return found in the data. As we will see, in order to produce an excess return high enough, we need to assume a volatility of consumption growth significantly larger than the one estimated from the data. As we have explained, such an additional real risk amplifies the effect of risk-aversion heterogeneity. For that reason, in our calibration exercise we include the average excess return, but with a small weight.

The model solves for the price-dividend ratio, pd, and risk-free rate, r^f, endogenously. Then, for a given set of parameters, we take the observed consumption growth series and the estimated value for λ and generate a time series for ω. We use this time-series to compute the *model-implied time-series* for pd and r^f. From the price-dividend time-series and the dividend growth time-series we derive the market returns.

We calibrate our model for two different sets of parameters that we henceforth call models (or parameterizations) 1 and 2. These models differ in the parameter values of consumption growth volatility, as well as the parameters that characterize the distribution of agents. Model 1 assumes the distribution of agents estimated using the distribution of KSS, with $\bar{\gamma} = 5.17$ and $\bar{\nu} = 0.13$. The objective of the second model is to produce a sizable equity premium -similar to that observed in the economy-, as well as high variability in the Sharpe ratio -that may produce predictability-; we choose the parameters of this model with this objective in mind: we assume that volatility of consumption growth is 2.5 times the estimated level, and that the standard deviation of risk-tolerance is twice as high as that estimated using KSS. Chan and Kogan (2002) produce high variability in the Sharpe ratio through high persistence in the state variable and a high level of heterogeneity. However, this does not produce predictability, while it generates excessive variability in all market series as well as excessive autocorrelation. The problem, as they observe, is that most of the equity premium is in fact term premium. For this reason we amplify the aggregate risk and the level of heterogeneity, so that we produce a sizeable and variable risk premium. The second parametrization gives us the opportunity to discuss whether heterogeneity in preferences may be able to explain market behavior.

The parameters values of λ, ρ and δ minimize the following objective,

$$\begin{aligned} J_j(\lambda,\rho,\delta) &= [\mu_j(PD)-\mu(PD)]^2 + [\sigma_j(PD)-\sigma(PD)]^2 + \sum_{i=1}^{7}\left[acf_j^i(PD)-acf^i(PD)\right]^2 \\ &+ w_e\,[\mu_j(R^e)-\mu(R^e)]^2 + w_f\left[\sigma_j(R^f)-\sigma(R^f)\right]^2, \end{aligned} \quad (18)$$

for each model j. $\mu_j(x)$ and $\sigma_j(x)$ denote the model-implied (for the observed consumption growth), average and standard deviation, respectively, of the time series of variable x. acf^i denotes the autocorrelation for lag i. μ, σ and acf denote the values estimated from the data. In table 3 we collect the values of the parameters for each model both those assumed and those computed according to (18). We estimate the persistence parameter λ as explained before, and we find it to be 0.12 for the first model, similar to the value in Campbell and Cochrane (1999). The persistence of the price-dividend ratio in the model-implied time-series (not shown), is in fact slightly higher than the persistence in the data. In model 2 the resulting value for λ is smaller (more persistence), while δ is higher than for model 1.

6 Quantifying The Effect of Agent-Heterogeneity

The main purpose of this section is to study whether time variation in the risk-aversion of the economy coming from agent heterogeneity can produce the asset pricing facts that Campbell and Cochrane (1999) focus on. Namely, the high and counter-cyclical equity premium, the high variability and persistence of the price-dividend ratio, the low volatility of the risk-free rate and the ability of the price-dividend ratio to predict long-run excess returns. First we show that under the baseline calibration of the model the impact of risk-aversion heterogeneity is negligible, from the comparison of the predictions of the calibrated heterogeneous agent economy with a homogeneous-agent -but otherwise identical- economy. The reason is twofold: First the consumption risk is small, which even with a persistent state variable, yields very low wealth reallocation. Second, the level of risk-aversion heterogeneity as estimated by KSS -which is higher than that estimated from CEX data- is also small, and the wealth reallocation does not lead to significant changes in the risk-aversion of the representative agent.

We next consider substantially higher volatility of consumption and heterogeneity of risk-aversion than those implied by the data and estimated as explained before. We are able to generate a sizeable equity premium and a Sharpe ratio that exhibits high variability, in fact higher variability than in Chan and Kogan (2002). However, the risk-free rate is still high on average and very volatile and, as a result, the predictability of long-run excess returns does not come close to that found in the data.

In our analysis, we first study price and return unconditional statistics (table 4). To derive the statistics of each parametrization we generate two sets of results: (i) *model* statistics, that correspond to the long-run unconditional values as time tends to infinity, and we compute by running multiple simulations; (ii) *historical simulation* statistics, based on the actual consumption growth. We emphasize that the model-implied time-series is generated using the model solutions of the *pd*-ratio and r^f-rate and the historical data on consumption and dividend growth. These time-series are shown for each model in figures 6 and 7 respectively.

We then consider autocorrelations of the *pd*-ratio and the excess return, as well as the *pd*-ratio predictability regressions (tables 5, 6 and 7). These statistics are the result of simulating the economy 1000 times. For each simulation we generate 75 years of annual artificial data, which corresponds to the frequency and length of our real data. The reported statistics are the averages of the estimated values of each simulation. Finally we analyze the model implied conditional price and return statistics in figures 8 and 9, corresponding to the two parameterizations.

6.1 The baseline model

The top panel of table 4 compares unconditional price and return statistics of the data and model 1. In the heterogeneous economy, the variability, as well as the level of the price-dividend ratio, is fitted relatively well. The average market return is predicted at slightly more than 1% higher than what is found in the data. The model completely misses the average risk free rate, predicted at around 7.22%; the historical simulation average (explained before) is around 6% while the data average is less than 1%. This is a result of the inability of the model to generate conditional volatility for the stochastic discount factor high enough to explain both the high equity premium along with the low risk-free rate. The historical simulation captures the return variability but the model predicts a number around 3% lower than the data. The risk-free rate, on the other hand, is more volatile in the model by more than 1%. From the previous results, it is not surprising that the model predicts a very small average Sharpe ratio of around 0.11, while the data average is more than three times higher. The results of the homogeneous economy with the first parametrization are almost identical and, therefore, the two economies cannot be separated based on such data. The endogenous variation in the volatility of the stochastic discount factor (function $|z'(\omega)|$ and figure 5) generated by the variation in the risk-aversion of the economy is so small that it cannot be detected either in the risk-free rate, the risk premium, or return volatility.

Figure 6 shows the model 1-implied time-series of the *pd*-ratio and r^f-rate for both the heterogeneous agents and the homogeneous agents economies. From the theoretical and empirical results we have obtained so far it is no surprise that the two time-series are indistinguishable. The two economies are also indistinguishable when it comes to the autocorrelation coefficients shown in the top panels of tables 5 and 6. There is a slight difference in the predictability regressions in the top panel of table 7. Predictability in the data starts with a modest value of around 7%, increases monotonically with the horizon and reaches a level of 37%. For the homogeneous economy the predictability starts at 0.13% and goes up to 5.81%, while for the heterogeneous economy these numbers range from 0.15% to 5.87%.

In figure 8 we plot some conditional statistics. The state variable is shown within four standard deviations from the unconditional mean. The risk-aversion of the representative agent varies very little, as it ranges from slightly less than 5 to only 5.4. The risk-aversion coefficient for the homogeneous economy is 5.17. For this reason, the curvature in the pricing kernel is almost unnoticeable. From lemma 1 we know that the slope of the pricing kernel is equal to the consumption-weighted harmonic average risk-aversion in the economy. Similarly, the price-dividend ratio, as well as the risk-free rate of the homogeneous and heterogeneous economies of model 1, are again almost indistinguishable.

The curvature of the price-dividend ratio function is closely related to the conditional return volatility. When the price-dividend ratio is flatter around the expected future state, then the return volatility is smaller. Since heterogeneity in the economy makes the *pd*-ratio slightly flatter during expansions and slightly steeper during bad times, the conditional volatility of returns becomes slightly more counter-cyclical with heterogeneity. However, this is an effect of our choice of preferences and the fact that, along with the variation of the risk-aversion of the heterogeneous economy, we also have variation in the elasticity of intertemporal substitution (EIS). The higher the EIS, i.e. the lower the risk-aversion of the economy, the lower the variability of the discount rates and hence of the price-dividend ratio. Therefore since the risk-aversion of the heterogeneous economy is counter-cyclical, the EIS is pro-cyclical and the conditional volatility of returns becomes more counter-cyclical for the heterogeneous economy.

Finally we look at the variation in the Sharpe ratio. For the homogeneous agent economy the variation is only due to the conditional correlation between market returns and consumption growth. The maximum Sharpe ratio in an economy where consumption growth is *iid* normal is constant. For model 1, the Sharpe ratio of the heterogeneous economy becomes more counter-cyclical, but only marginally, since it decreases from around 0.12 to 0.105, reflecting again the small variability in the risk-aversion of the economy.

6.2 Abnormally high risk and level of heterogeneity

In order to show how far our model can go to produce the positive results of Campbell and Cochrane (1999) we multiply the consumption risk estimated from the data by two and a half times, and we double the level of heterogeneity found by KSS (we call this parametrization 2). We present the results in figure 9.

Even though the new model statistics are significantly improved with respect to the data, and we now observe a noticeable difference between the homogeneous and the heterogeneous economy, the predictability results of the heterogeneous economy are still far from approximating the data. The particularly volatile Sharpe ratio of the heterogeneous economy does not seem to be able to produce enough predictability which seems to be due to the fact that the risk-free rate is still very volatile. R^2 for the 7-year horizon is just less than 10% for the heterogeneous-agents economy, while it is almost 6.5% for the homogeneous-agents economy. In addition, as we observe in the lower panel of table 6, the heterogeneity does indeed have a noticeable impact on the counter-cyclicality of the excess returns, but still does not yield the level of mean-reversion shown in the data.

7 Conclusion

Campbell and Cochrane (1999) consider a stochastic discount factor that can explain a number of well documented properties of asset prices. It is central to their model the assumption that risk-aversion is stochastic and counter-cyclical. However, for their model to be successful, they need risk-aversion to have a very high variation in a very broad range. Chan and Kogan (2002) show that a model with multiple agents with heterogeneous risk-aversion can produce the counter-cyclical pattern of the stochastic discount factor of Campbell and Cochrane (1999), as well as a varying Sharpe ratio.

In this paper we consider a model similar to that in Chan and Kogan (2002), we derive explicitly the risk-aversion coefficient of the representative agent and find that the variation required by the stochastic discount factor of Campbell and Cochrane (1999) is unlikely to be produced by such a model with reasonable (as observed in the economy) parameter values.

From our analysis, it appears that the assumed risk-premium of Campbell and Cochrane (1999) has to be driven by something other than the risk-aversion of the economy. A promising alternative would be to examine asset prices as they are formed on expectations or beliefs of investors about the underlying risks in the economy. The way they update these beliefs or the uncertainty they have about the conditional distribution might be able to explain why investors require such high compensation for every unit of ex-post risk they take, and why the price of risk appears to vary so radically across time.

References

Abel, A. B. (1990). Asset prices under habit formation and catching up with the joneses. *American Economic Review Papers and Proceedings*, *80*, 38-42.

Abel, A. B. (1999). Risk premia and term premia in general equilibrium. *Journal of Monetary Economics*, *43*, 3-33.

Bhamra, H. S., & Uppal, R. (2009). The effect of introducing a non-redundant derivative on the volatility of stock-market returns. *Review of Financial Studies*, *22*, 2303-2330.

Bollerslev, T., Chou, R. Y., & Kroner, K. F. (1992). Arch modeling in finance: A review of the theory and empirical evidence. *Journal of Econometrics*, *52*(1-2), 5-59.

Bollerslev, T., Engle, R. F., & Wooldridge, J. M. (1988, February). A capital asset pricing model with time-varying covariances. *Journal of Political Economy*, *96*(1), 116-31.

Boudoukh, J., Michaely, R., Richardson, M., & Michael, R. R. (2007). On the importance of measuring payout yield: Implications for empirical asset pricing. *Journal of Finance*, *62*, 877-915.

Brandt, M. W., & Kang, Q. (2004). On the relation between the conditional mean and volatility of stock returns: a latent var approach. *Journal of Financial Economics*, *72*, 217257.

Campbell, J. Y. (1991). A variance decomposition of stock returns. *Economic Journal*, *101*(157-179).

Campbell, J. Y., & Cochrane, J. H. (1999). By force of habit: A consumption-based explanation of aggregate stock market behavior. *Journal of Political Economy*, *107*, 205-251.

Campbell, J. Y., & Shiller, R. J. (1988a). The dividend-price ratio and expectations of future dividends and discount factors. *Review of Financial Studies*, *1*(3), 195-228.

Campbell, J. Y., & Shiller, R. J. (1988b). Stock prices, earnings, and expected dividends. *Journal of Finance*, *43*(3), 661-676.

Chan, Y. L., & Kogan, L. (2002). Catching up with the joneses: Heterogeneous preferences and the dynamics of asset prices. *Journal of Political Economy*, *110*, 1255-1285.

Cochrane, J. H. (1992). Explaining the variance of price-dividend ratios. *Review of Financial Studies*, *5*(2), 243-280.

Coen-Pirani, D. (2004, November). Effects of differences in risk aversion on the distribution of wealth. *Macroeconomic Dynamics*, *8*, 617-632.

Constantinides, G. (1990). Habit formation: A resolution of the equity premium puzzle. *Journal of Political Economy*, *98(3)*, 519-543.

Detemple, J. B., & Zapatero, F. (1991, November). Asset prices in an exchange economy with habit formation. *Econometrica*, *59*(6), 1633-57.

Dumas, B. (1989). Two person dynamic equilibrium in the capital market. *Review of Financial Studies*, *2*(2), 157-188.

Famma, E. F., & French, K. R. (1988). Dividend yields and expected stock returns. *Journal of Financial Economics*, *22*, 3-27.

Gali, J. (1994). Keeping up with joneses: Consumption externalities, portfolio choice, and asset prices. *Journal of Money, Credit and Banking*, *26*(1), 1-8.

Gârleanu, N., & Panageas, S. (2008, April). *Young, old, conservative and bold: The implications of heterogeneity and finite lives for asset pricing.* (Working Paper)

Hansen, L. P., & Singleton, K. J. (1983, April). Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy*, *91*(2), 249-65.

Harvey, C. R. (1989). Time-varying conditional covariances in tests of asset pricing models. *Journal of Financial Economics*, *24*(2), 289-317.

Hindy, A., Huang, C., & Zhu, S. H. (1997). Optimal consumption and portfolio rules with durability and habit formation. *Journal of Economic Dynamics and Control*, *21*(2-3), 525-550.

Judd, K. L., Kubler, F., & Schmedders, K. (2003, October). Asset trading volume with dynamically complete markets and heterogeneous agents. *Journal of Finance*, *58*(5), 2203-2217.

Kimball, M. S., Sahm, C. R., & Shapiro, M. D. (2008, September). Imputing risk tolerance from survey responses. *Journal of the American Statistical Association*, *103*(483), 1028-1038.

Kogan, L., Makarov, I., & Uppal, R. (2007, April). The equity risk premium and the riskfree rate in an economy with borrowing constraints. *Mathematics and Financial Economics*, *1*(1), 1-19.

Ludvigson, S. S., & Ng, S. (2007). The empirical risk-return relation: A factor analysis approach. *Journal of Financial Economics*, *83*, 171-222.

Malloy, C., Moskowitz, T., & Vissing-Jorgensen, A. (2008). Long-run stockholder consumption risk and asset returns. *Journal of Finance*. (forthcoming)

Mehra, R., & Prescott, E. (1985). The equity premium. *Journal of Monetary Economics*, *15*, 145-161.

Ryder, H. E., & Heal, G. M. (1973). Optimal growth and intertemporally dependent preferences. *The Review of Economic Studies*, *40*, 1-31.

Sundaresan, S. M. (1989). Intertemporally dependent preferences and the volatility of consumption and wealth. *Review of Financial Studies*, *2*, 73-88.

Wang, J. (1996). The term structure of interest rates in a pure exchange economy with heterogeneous investors. *Journal of Financial Economics*, *41*, 75-110.

Whitelaw, R. F. (1994, June). Time variations and covariations in the expectation and volatility of stock market returns. *Journal of Finance*, *49*(2), 515-41.

A Proofs

Below we provide the proofs of the propositions, lemmas, and corollaries found in the text. The preference assumption is that agents have time and state separable power preferences with external habit, and the running utility is given by

$$u(c, X|\gamma) = \frac{c^{1-\gamma} X^{\gamma-\rho} - 1}{1-\gamma}.$$

where c denotes consumption. We denote by α the consumption proportion c/Y. Agents' types are characterized by the initial distribution of wealth $\theta(\gamma)$, which is a density function over the set of types Γ. The state variable of the economy is the aggregate endowment habit ratio $\omega = y - x$, where y and x are the natural logs of Y and X respectively. We also define the pricing kernel $z_t = \log(p_t) + \rho x_t$.

Proof of Proposition 1. Given the habit process $(X_t, t \geq 0)$, the consumption price process $(p_t, t \geq 0)$, and the initial total wealth W_0, the optimization problem of an agent of type γ with initial wealth allocation $\theta(\gamma)$ is given by,

$$\max_{(c_t(\gamma), t\geq 0)} \mathcal{L}(\gamma) = \mathbb{E}_0 \sum_{t\geq 0} \delta^t u(c_t, X_t|\gamma) - \lambda(\gamma)^{-\gamma} \left[\mathbb{E}_0 \sum_{t\geq 0} \delta^t p_t c_t(\gamma) - \theta_0(\gamma) p_0 W_0 \right]$$

Due to our preference assumptions, the optimal solution is interior, and is characterized by the first order condition,

$$\left(\frac{c_t(\gamma)}{X_t} \right)^{-\gamma} = \lambda(\gamma)^{-\gamma} p_t X_t^{\rho}, \qquad \forall t \geq 0.$$

and the equality of the inter-temporal budget constraint. Using the definition of z_t and the consumption proportion α, we arrive at the optimality condition (6). We then substitute in the budget constraint, satisfied with inequality, the optimal consumption share from (6), to arrive at the equilibrium value for $\lambda(\gamma)$. □

Proof of Corollary 1. Market equilibrium is obtained when the optimal consumption share in each period as function of the type γ integrates to one. Given the optimality condition (6) it is evident that the pricing kernel is a function of the endowment habit ratio, $z_t = z(\omega_t)$. □

Proof of Corollary 2. The pricing kernel is normalized to have a value of zero at the average state, $z(\bar{\omega}) = 0$. The initial state is assumed to be the average state and hence expression (9) is easily derived from (6) at the initial state and (7). Since the equilibrium pricing kernel is a function of ω, so is the optimal consumption share as given by (6). At the average state it is given by

$$\alpha(\bar{\omega}, \gamma) = \lambda(\gamma) e^{-\bar{\omega}}.$$

Using the previous result, we express $\lambda(\gamma)$ in terms of $\alpha(\bar{\omega},\gamma) = \mathcal{P}_{\bar{\omega}}(1/\gamma)$ and substitute it in (8) to arrive at (10). We have replaced the coefficient of risk-aversion with the coefficient of risk-tolerance, $\tau = 1/\gamma$. □

Proof of Lemma 1. Differentiate expression (8) with respect to ω to get

$$0 = \int_\Gamma \left(-\frac{z'(\omega)}{\gamma} - 1\right) \lambda(\gamma) \exp\left(-\frac{z(\omega)}{\gamma} - \omega\right) d\gamma$$

and note that $\mathcal{P}_\omega(\tau) = \lambda(1/\tau)\exp(-\tau z(\omega) - \omega)$ to derive ($z2$). Differentiate again with respect to ω to get,

$$\begin{aligned} 0 =& \int_\Gamma \left(-\frac{z''(\omega)}{\gamma} + \frac{z'(\omega)^2}{\gamma^2} + 2\frac{z'(\omega)}{\gamma} + 1\right) \mathcal{P}_\omega(1/\gamma) d\gamma \\ =& - z''(\omega)\mathcal{E}_\omega(\tau) + z'(\omega)\mathcal{E}_\omega(\tau^2) - 2z'(\omega)\mathcal{E}_\omega(\tau) + 1. \end{aligned}$$

Using the result for $z'(\omega)$ and rearranging we get the second derivative of z. Since the numerator is a variance term then the second derivative is positive.

For the rest of the results we assume explicitly that the set Γ has only positive values and is bounded above and below. Now, since the term $\mathcal{E}_\omega\left[\exp(-\tau z(\omega))\right]$ is strictly positive and finite for all finite and positive values of $z(\omega)$, from (10) it has to be that z tends to minus and plus infinity, as ω tends to plus and minus infinity, respectively. Let γ_{max} denote the maximum γ of the set Γ, and consider the following ratio of optimal consumption shares,

$$\frac{\alpha(\omega,\gamma_{max})}{\alpha(\omega,\gamma)} = \frac{\lambda(\gamma_{max})}{\lambda(\gamma)} \exp\left[-z(\omega)\left(\frac{1}{\gamma_{max}} - \frac{1}{\gamma}\right)\right]$$

Now since $\lim_{\omega\to-\infty} z(\omega) = +\infty$ then from the previous ratio, we have to conclude that

$$\lim_{\omega\to-\infty} \alpha(\omega,\gamma) = 0, \qquad \forall\gamma < \gamma_{max}.$$

Hence, the consumption distribution collapses to a spike at the maximum γ. Similarly, when the state tends to plus infinity, the distribution collapses to the minimum γ. Finally, since in either case the variance tends to zero, then the second derivative tends to zero as well. □

Proof of Proposition 2. Given the habit process $(X^r_t = X_t e^{\bar{\omega}}, t \geq 0)$, the consumption price process $(p^r_t, t \geq 0)$ and the initial total wealth W^r_0, the optimization problem of the representative agent is given by,

$$\max_{(c_t, t\geq 0)} \mathcal{L}_r = \mathbb{E}_0 \sum_{t\geq 0} \delta^t u_r(c_t, X^r_t) - \lambda_r \left[\mathbb{E}_0 \sum_{t\geq 0} \delta^t p^r_t c_t - p^r_0 W^r_0\right]$$

where

$$u_r(c_t, X_t^r) = \frac{c_t^{1-\gamma(\omega_t)}(X_t^r)^{\gamma(\omega_t)-\rho}}{1-\gamma(\omega_t)},$$

$\gamma(\omega)$ is the stochastic risk-aversion of the representative agent, and ω is as before. Then, the representative agent optimally consumes the aggregate endowment,, and the first order condition is given by,

$$\left(\frac{Y_t}{X_t}\right)^{-\gamma(\omega_t)} e^{\bar{\omega}[\gamma(\omega_t)-\rho]} = \lambda_r e^{z_t^r},$$

where $z_t^r = \log(p_t^r) + \rho x_t$. To normalize z^r to be zero at the average state the Lagrange multiplier has to equal $e^{\rho\bar{\omega}}$. □

Proof of Corollary 3. From the definitions of z^r and z, the two pricing kernels are equal in all states iff the consumption processes p^r and p are equal. Therefore, the asset price processes of the two economies are equal if $z^r \equiv z$. Hence, from Proposition 2 we must have that,

$$\gamma(\omega) = -\frac{z(\omega)}{\omega - \bar{\omega}}, \qquad \forall \omega \neq \bar{\omega}.$$

For continuity, we use l'Hôpital's rule to define $\gamma(\omega)$ at the average state. □

Proof of Corollary 4. Since

$$\gamma(\omega) = -\frac{z(\omega)}{\omega - \bar{\omega}},$$

the first and second derivatives are obtained using the first two derivatives of z from Lemma 1. In order to prove that $\gamma'(\omega)$ is negative, we need to show that the expression in the square brackets $\gamma(\omega) + z'(\omega) \gtreqless 0$ when $\omega \gtreqless \bar{\omega}$. From the convexity of function z and the fact that $z(\bar{\omega}) = 0$, we have that

$$-z(\omega) + z'(\omega)(\omega - \bar{\omega}) > 0$$

The result then obtains after dividing by $(\omega - \bar{\omega})$, and noting that the inequality switches when $\omega < \bar{\omega}$. The limits of $\gamma(\omega)$, $\gamma'(\omega)$ and $\gamma''(\omega)$ are straightforward applications of the limit properties of z. □

Proof of Corollary 5. If the risk-tolerance at the average state is gamma distributed with parameters (κ, ϑ), from the moment generating function we know that,

$$\mathcal{E}_{\bar{\omega}}\left[\exp(-\tau z(\omega))\right] = \left[1 + \vartheta z(\omega)\right]^{-\kappa}.$$

Using the equilibrium condition (10) and rearranging, we get,

$$z(\omega) = \frac{\exp\left(-\frac{\omega - \bar{\omega}}{\kappa}\right) - 1}{\vartheta}$$

$\bar{\gamma}$ denotes the inverse of the average risk-tolerance, and therefore is equal to $(\kappa\vartheta)^{-1}$,and $\bar{\nu}^2$ denotes the risk-tolerance variance, given by $\kappa\vartheta^2$. Furthermore, we use $\eta = (\bar{\gamma}\bar{\nu})^2$ to obtain the result. □

Proof of Corollary 6. From the definition of z we have that,

$$m(\omega, \omega') = log(\delta) + z(\omega') - z(\omega) - \rho(x' - x).$$

The habit process is given by,

$$x' = \lambda y + (1 - \lambda)x$$

and therefore $x' - x = \lambda\omega$. The result is then obtained once we substitute in the expression for z derived in Corollary 5. When $\bar{\nu} = 0$ then it is obvious from (10) that $z(\omega) = -\bar{\gamma}(\omega - \bar{\omega})$. □

B Computational Approach

We outline below the numerical method used to compute the price-dividend ratio and the risk-free rate as functions of the state. The price-dividend ratio of the market security that pays a dividend that grows linearly in the aggregate endowment growth is given by

$$PD(\omega) = \delta\mathbb{E}_\omega \left[\exp\left(z(\omega') - z(\omega) - \rho\lambda\omega + \mu_d + \sigma_d\varrho\epsilon + \sigma_d\sqrt{1 - \varrho^2}\epsilon^d\right)\left(PD(\omega') + 1\right)\right]$$

where ϵ and ϵ^d are independent and normally distributed variables. The distributional assumptions made here are not required, in general. The method can be adapted to any arbitrary distribution. Let the law of motion for ω be some known function,

$$\omega' = L(\omega, \epsilon).$$

Let $\tilde{\delta} = \delta \exp\left(\mu_d + \frac{1}{2}\sigma_d^2(1 - \varrho^2)\right)$, and also let,

$$N(\omega, \epsilon) = \exp\left[z(L(\omega, \epsilon)) - z(\omega) - \rho\lambda\omega + \sigma_d\varrho\epsilon\right].$$

Then the price-dividend equation can be rewritten as follows:

$$PD(\omega) = \tilde{\delta}\int_{-\infty}^{+\infty} \varphi(\epsilon)N(\omega, \epsilon)\left[PD(L(\omega, \epsilon)) + 1\right]d\epsilon,$$

where $\varphi(\epsilon)$ is the standard normal density. We then discretize the state variable in a set of n values Ω and compute the $n \times n$ transition matrix $\Pi = [\pi(\omega, \omega')]$ based on L and φ. Let $\mathbf{PD}$ be the vector of price-dividend ratios and let $\mathbf{N}$ be the $n \times n$ matrix computed from $N(\omega, \epsilon)$. The price-dividend ratio equation can now be written simultaneously for all states in Ω in the following matrix form:

$$\mathbf{PD} = \tilde{\delta}\left[\Pi \circ \mathbf{N}\right]\left(\mathbf{PD} + \mathbf{1}_n\right),$$

where $\circ$ denotes the element-wise matrix multiplication and $\mathbf{1}_n$ is the $n-$dimentional unit vector. After a matrix inversion the price-dividend ratio vector is computed according to,

$$\mathbf{PD} = \tilde{\delta}\left[\mathbf{I}_n - (\Pi \circ \mathbf{N})\right]^{-1}\mathbf{1}_n,$$

where $\mathbf{I}_n$ is the $n-$dimentional identity matrix.

The price of the risk-free bond that pays a unit of consumption the next period is computed in a similar fashion. Let

$$M(\omega, \epsilon) = \delta \exp\left(z(L(\omega, \epsilon)) - z(\omega) - \rho\lambda\omega\right)$$

and let $\mathbf{M}$ denote the square matrix for all the combination of values in the set Ω. Then the vector of bond price values is computed by

$$\mathbf{P}^f = (\Pi \circ \mathbf{M})\mathbf{1}_n.$$

The risk-free rate is then computed as $R^f(\omega) = 1/PD(\omega) - 1$.

For the simulation part of our results we also approximate the price-dividend ratio function, as well as the risk-free rate function, with cubic-splines. The data used to estimate the piece-wise polynomial coefficients were the computed vectors $\mathbf{PD}$ and $\mathbf{P}^f$ and the vector of arguments Ω. For our computational results we use a specific assumption for the state variable that gives the following law of motion,

$$\omega' - \omega = -\lambda(\omega - \bar{\omega}) + \sigma\epsilon,$$

where σ is the volatility of consumption growth. Hence the state variable is conditionally normally distributed. The set Ω used was an equidistant grid of 251 values where the minimum and the maximum values are 8 standard deviations, left and right from the mean, respectively.

Table 1: Exposure of consumption growth to changes in the price-earnings ratio

	(a)		(b)		(c)	
	Δpe_{t+1}	R^2_{adj}	Δpe_{t+1}	R^2_{adj}	Δpe_{t+1}	R^2_{adj}
g^1_{t+1}	0.087	0.196	0.113	0.201	0.097	0.240
	(1.465)		(1.504)		(1.423)	
g^2_{t+1}	0.178	0.056	0.212	0.035	0.229	0.050
	(1.877)		(1.658)		(2.055)	
g^3_{t+1}	0.160	0.131	0.206	0.133	0.179	0.129
	(1.209)		(1.148)		(1.092)	
g^4_{t+1}	0.180	0.061	0.224	0.058	0.212	0.088
	(2.282)		(2.206)		(2.137)	
g^5_{t+1}	0.063	0.201	0.067	0.163	0.043	0.184
	(1.192)		(0.857)		(0.640)	
g^6_{t+1}	0.169	0.231	0.209	0.186	0.239	0.210
	(2.020)		(1.720)		(1.949)	
g^7_{t+1}	0.138	0.072	0.152	0.067	0.171	0.094
	(1.711)		(1.606)		(1.891)	
g^8_{t+1}	0.279	0.060	0.258	0.048	0.294	0.046
	(2.353)		(1.636)		(1.988)	

In parenthesis the t-statistics were estimated using Newey-West. Log-consumption growths and changes in the price-earnings ratio were over quarterly periods. The CEX groups (1-8) were constructed based on increasing ratio of investment over consumption and using a fitted log-normal distribution.

Table 2: Exposure of consumption growth to market returns

	(a)		(b)		(c)	
	r_{t+1}	R^2_{adj}	r_{t+1}	R^2_{adj}	r_{t+1}	R^2_{adj}
g^1_{t+1}	0.031	0.192	0.077	0.199	0.037	0.238
	(0.512)		(1.030)		(0.528)	
g^2_{t+1}	0.127	0.054	0.135	0.032	0.161	0.047
	(1.504)		(1.175)		(1.596)	
g^3_{t+1}	0.157	0.133	0.155	0.131	0.149	0.129
	(1.318)		(0.986)		(1.005)	
g^4_{t+1}	0.114	0.054	0.138	0.049	0.148	0.082
	(1.659)		(1.624)		(1.814)	
g^5_{t+1}	0.089	0.204	0.120	0.168	0.090	0.188
	(1.785)		(1.726)		(1.383)	
g^6_{t+1}	0.112	0.225	0.121	0.179	0.164	0.203
	(1.360)		(1.095)		(1.524)	
g^7_{t+1}	0.076	0.068	0.103	0.064	0.104	0.089
	(0.868)		(1.171)		(1.179)	
g^8_{t+1}	0.177	0.050	0.223	0.047	0.231	0.043
	(1.698)		(1.748)		(1.907)	

In parenthesis the t-statistics were estimated using Newey-West. Log-consumption growths and market returns were over quarterly periods. The CEX groups (1-8) were constructed based on increasing ratio of investment over consumption and using a fitted log-normal distribution.

Table 3: Model parameterizations

	Data	1	2
μ	3.1197	3.1197	3.1197
σ	2.1834	2.1834	5.4584
λ^*		0.1203	0.0960
$\bar{\gamma}$	5.1739	5.1739	5.1739
$\bar{\nu}$	0.1289	0.1289	0.2579
ρ^*		-3.1654	1.4488
δ^*		0.8404	0.9642
$\bar{\omega}$		0.2593	0.3250
σ_ω		0.0459	0.1277
$half-life$		5.4077	6.8683

The parameters with * were chosen by minimizing an objective function that measures the squared errors in certain quantities between the data and those of the model implied time-series. The quantities were the mean, standard deviation and autocorrelation function of the price-dividend ratio, the standard deviation of the risk-free rate and the mean equity premium. The model implied time-series of the price-dividend ratio and the risk-free rate were generated by using the true consumption and dividend growth series.

Table 4: Price and return statistics

		Parametrization 1			
		Heterogeneous ec.		Homogeneous ec.	
	Data	Model	Hist.sim.	Model	Hist.sim.
$\mathbb{E}(R^f)$	0.5667	7.2238	6.0343	7.2261	6.0234
$\mathbb{E}(R^e)$	7.6821	1.9119	2.5181	1.9116	2.4963
$\sigma(R^e)$	20.0214	17.2009	20.3598	17.2248	20.2110
$\sigma(R^f)$	3.8288	4.9264	5.1531	4.9418	5.0990
$\mathbb{E}(SR)$	0.3849	0.1112	0.1154	0.1110	0.1154
$\log(\mathbb{E}(PD))$	3.1928	3.2118	3.2765	3.2120	3.2763
$\log(\sigma(PD))$	1.9715	2.0222	1.9062	2.0299	1.9090
$\varrho(y_{t+1}-y_t, r^m_{t+1})$	0.0839	0.9548	0.5903	0.9552	0.5889
		Parametrization 2			
$\mathbb{E}(R^f)$		3.7433	3.0993	4.3756	3.0271
$\mathbb{E}(R^e)$		6.1415	6.3984	5.7333	5.6413
$\sigma(R^e)$		20.9119	24.7746	21.1446	20.0496
$\sigma(R^f)$		3.4194	4.5910	4.7691	4.7444
$\mathbb{E}(SR)$		0.2937	0.2395	0.2711	0.2642
$\log(\mathbb{E}(PD))$		3.1970	3.2818	3.1606	3.2417
$\log(\sigma(PD))$		1.9573	1.8995	2.1081	1.9781
$\varrho(y_{t+1}-y_t, r^m_{t+1})$		0.9293	0.5901	0.9695	0.5949

The historical simulations were generated by using the true consumption and dividend growth series. The homogeneous economy of each parametrization sets the level of heterogeneity to zero by setting $\bar{\nu} = 0$.

Table 5: Autocorrelations of price-dividend ratio

	Lag				
	1	2	3	5	7
Data	0.88	0.75	0.65	0.46	0.28
Heter.ec. 1	0.83	0.68	0.56	0.36	0.22
Homog.ec. 1	0.83	0.68	0.56	0.36	0.22
Heter.ec. 2	0.86	0.74	0.64	0.46	0.32
Homog.ec. 2	0.85	0.72	0.61	0.43	0.30

The autocorrelation function for each model was estimated by averaging the sample estimates of 1000 simulations of the same length as the data.

Table 6: Autocorrelations of excess returns

	Lag				
	1	2	3	5	7
Data	-0.06	-0.20	0.03	-0.05	0.02
Heter.ec. 1	-0.01	-0.02	-0.01	-0.02	-0.01
Homog.ec. 1	-0.01	-0.02	-0.01	-0.02	-0.01
Heter.ec. 2	-0.04	-0.03	-0.03	-0.02	-0.01
Homog.ec. 2	-0.01	-0.01	-0.02	-0.01	-0.01

The autocorrelation function for each model was estimated by averaging the sample estimates of 1000 simulations of the same length as the data.

Table 7: Long-run predictability regressions

				Parametrization 1					
	Data			Heterogeneous ec.			Homogeneous ec.		
Horizon	coef.	std.er.	R^2	coef.	std.er.	R^2	coef.	std.er.	R^2
1	-0.18	(0.05)	6.86	-0.03	(0.05)	0.15	-0.03	(0.05)	0.13
2	-0.30	(0.10)	11.29	-0.02	(0.10)	1.34	-0.02	(0.10)	1.32
3	-0.41	(0.12)	18.66	-0.01	(0.13)	2.48	-0.01	(0.13)	2.44
5	-0.61	(0.16)	29.07	-0.01	(0.20)	4.35	-0.01	(0.20)	4.30
7	-0.82	(0.21)	36.82	-0.03	(0.24)	5.87	-0.03	(0.24)	5.81
				Parametrization 2					
1				-0.08	(0.05)	2.23	-0.03	(0.06)	0.12
2				-0.10	(0.09)	3.50	-0.02	(0.10)	1.42
3				-0.12	(0.13)	4.98	-0.01	(0.15)	2.73
5				-0.16	(0.18)	7.53	-0.01	(0.22)	4.79
7				-0.20	(0.22)	9.67	-0.01	(0.27)	6.40

Regressions were run with log price-dividend ratio as the predictive variable and j-period realized excess log returns on the left hand side,

$$r^e_{t+j} = \beta_0 + \beta_1 pd_t + \varepsilon_{t+j}.$$

The regression estimates for each model are the averages of the sample estimates of 1000 simulations of the same length as the data. The standard errors were corrected using a GMM procedure and the Newey-West weighting scheme.

Figure 1: Fitted cross-sectional distribution of agents at $\bar{\omega}$

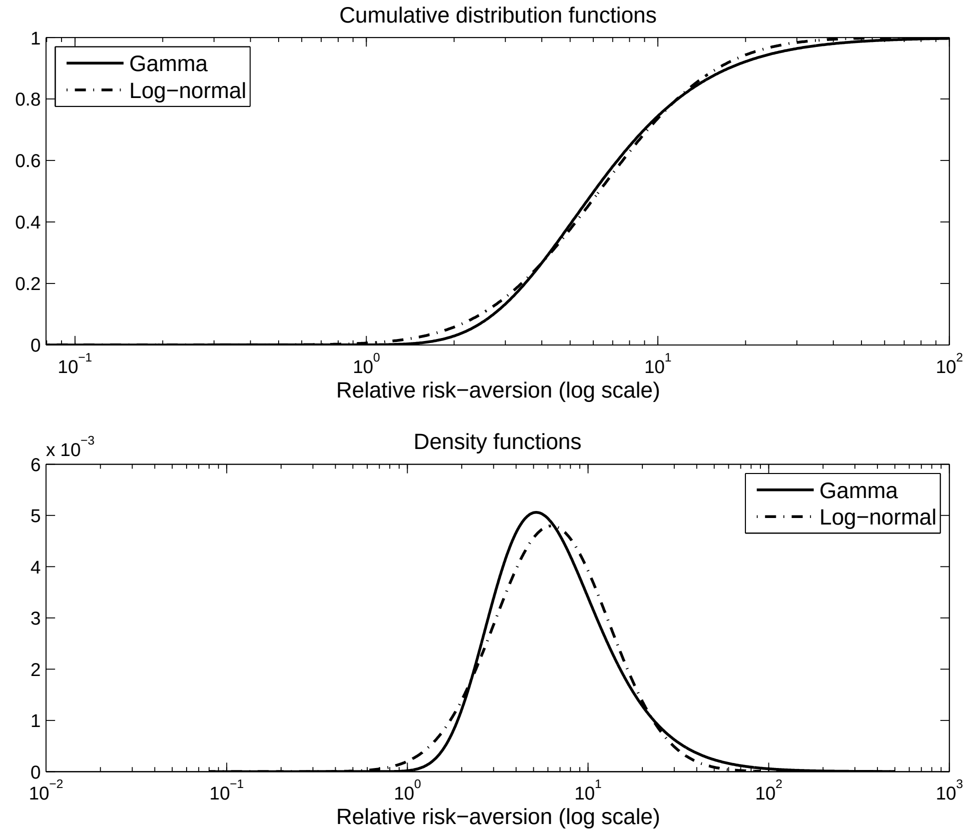

Given the estimated risk-tolerance distribution $\log(\tau) \sim N(-1.84, 0.73)$, by Kimball, Sahm and Shapiro (2008) we fit the gamma distribution by minimizing the distance between the two distribution functions,

$$\min_{\kappa,\theta} \int_0^{+\infty} \left[F_1(\tau) - F_2(\tau)\right]^2 d\tau$$

where F_1 and F_2 are the log-normal and gamma distribution functions. The integral was approximated numerically. The estimated parameters where $\kappa = 2.2474$ and $\theta = 0.0860$.

Figure 2: CEX group risk-aversion estimates using exposure to price-earnings ratio changes

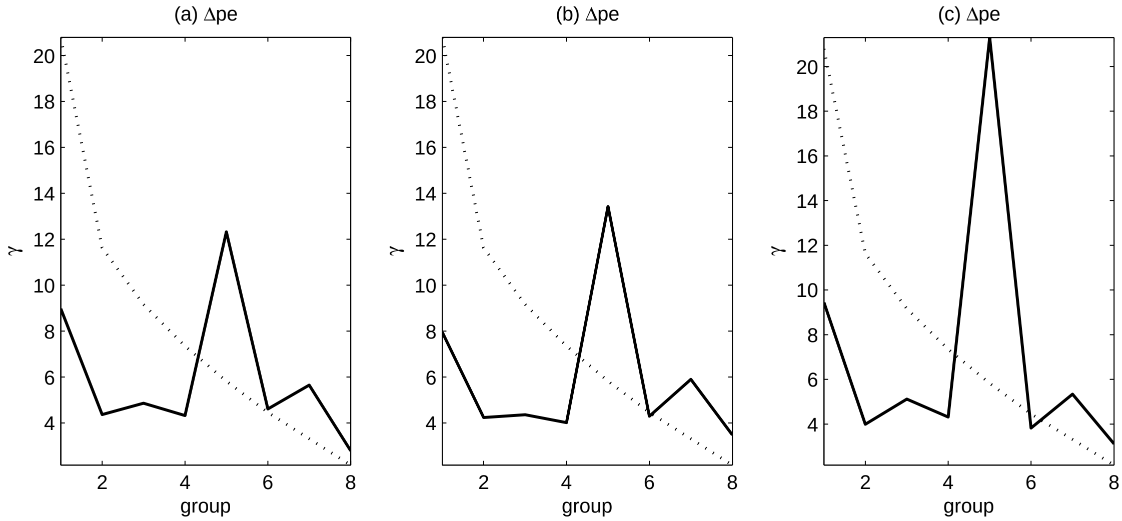

The straight line is computed based on the regression estimates and the average risk-tolerance of the fitted distribution. The dotted line is the fitted distribution using the average consumption shares of the CEX groups. The CEX groups (1-8) were constructed based on increasing ratio of investment over consumption and using a fitted log-normal distribution.

Figure 3: CEX group risk-aversion estimates using exposure to market returns

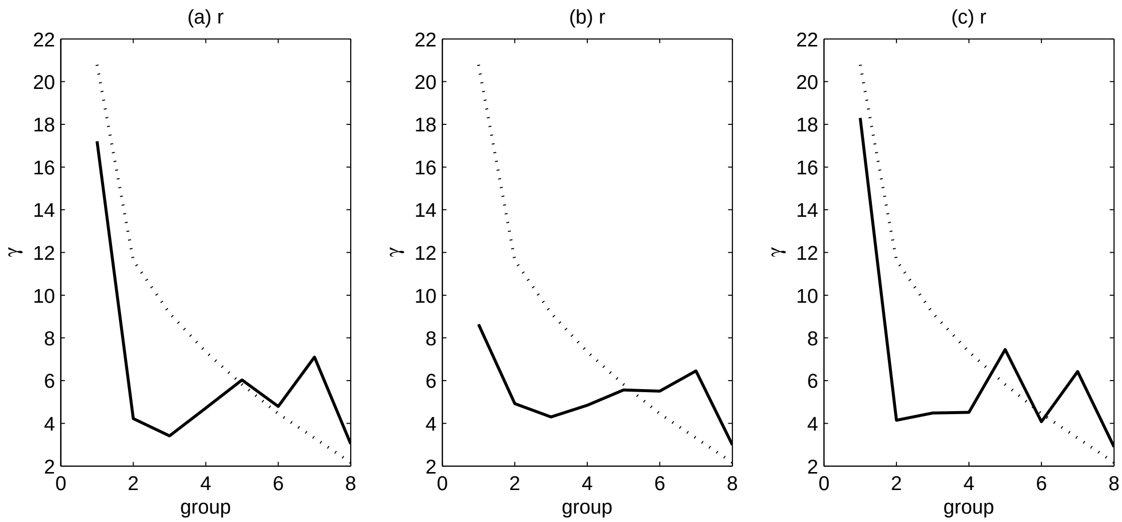

The straight line is computed based on the regression estimates and the average risk-tolerance of the fitted distribution. The dotted line is the fitted distribution using the average consumption shares of the CEX groups. The CEX groups (1-8) were constructed based on increasing ratio of investment over consumption and using a fitted log-normal distribution.

Figure 4: Variation in the representative agent risk aversion.

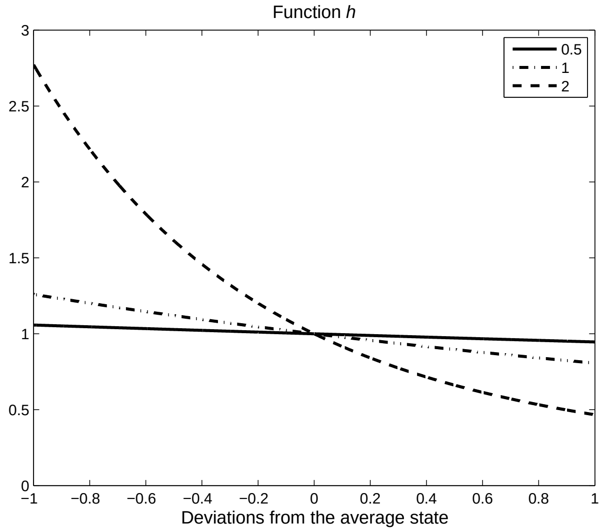

Function h is the multiplier of the risk-aversion of the representative agent as the economy deviates from the average state. For the three plots we multiply the estimated standard deviation of risk-tolerance with the corresponding number while the average risk-tolerance is the one estimated.

Figure 5: Variation in the conditional volatility of $\tilde{m}$

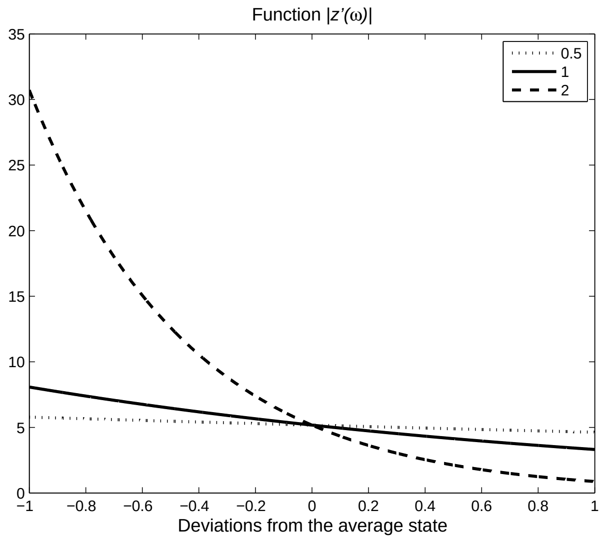

The conditional volatility of the approximate pricing kernel is,

$$|z'(\omega)|\sigma.$$

For the three plots we multiply the estimated standard deviation of risk-tolerance with the corresponding number while the average risk-tolerance is the one estimated. Parameter λ was set to 0.13 for these plots.

Figure 6: Data and model 1 implied time series of price-dividend ratio and risk-free rate

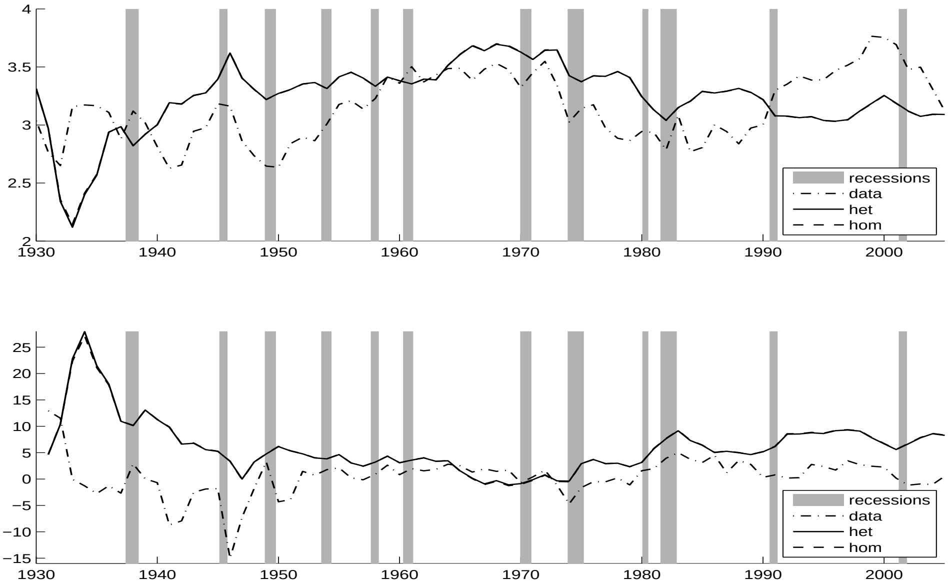

The model implied time-series of the price-dividend ratio and the risk-free rate were generated by using the true consumption and dividend growth series.

Figure 7: Data and model 2 implied time series of price-dividend ratio and risk-free rate

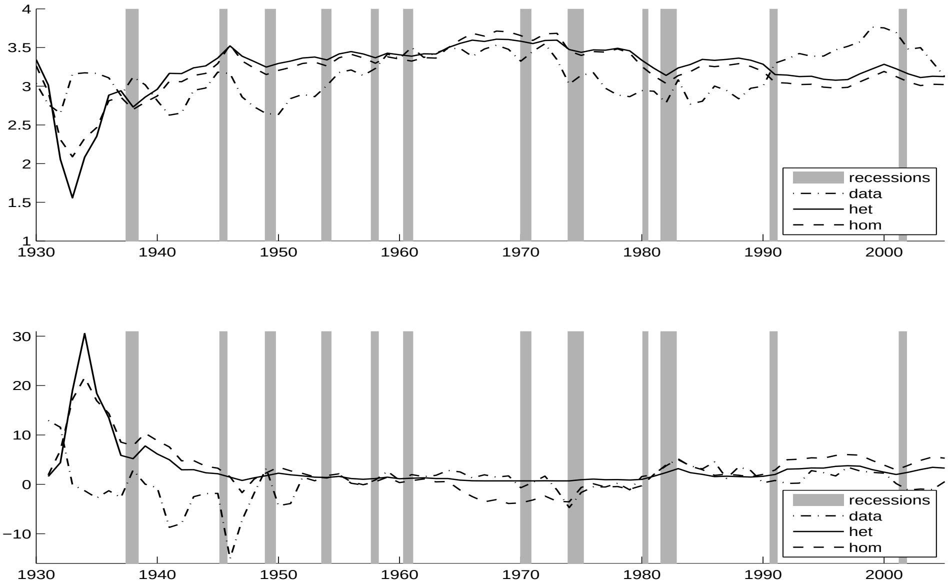

The model implied time-series of the price-dividend ratio and the risk-free rate were generated by using the true consumption and dividend growth series.

Figure 8: Model parametrization 1

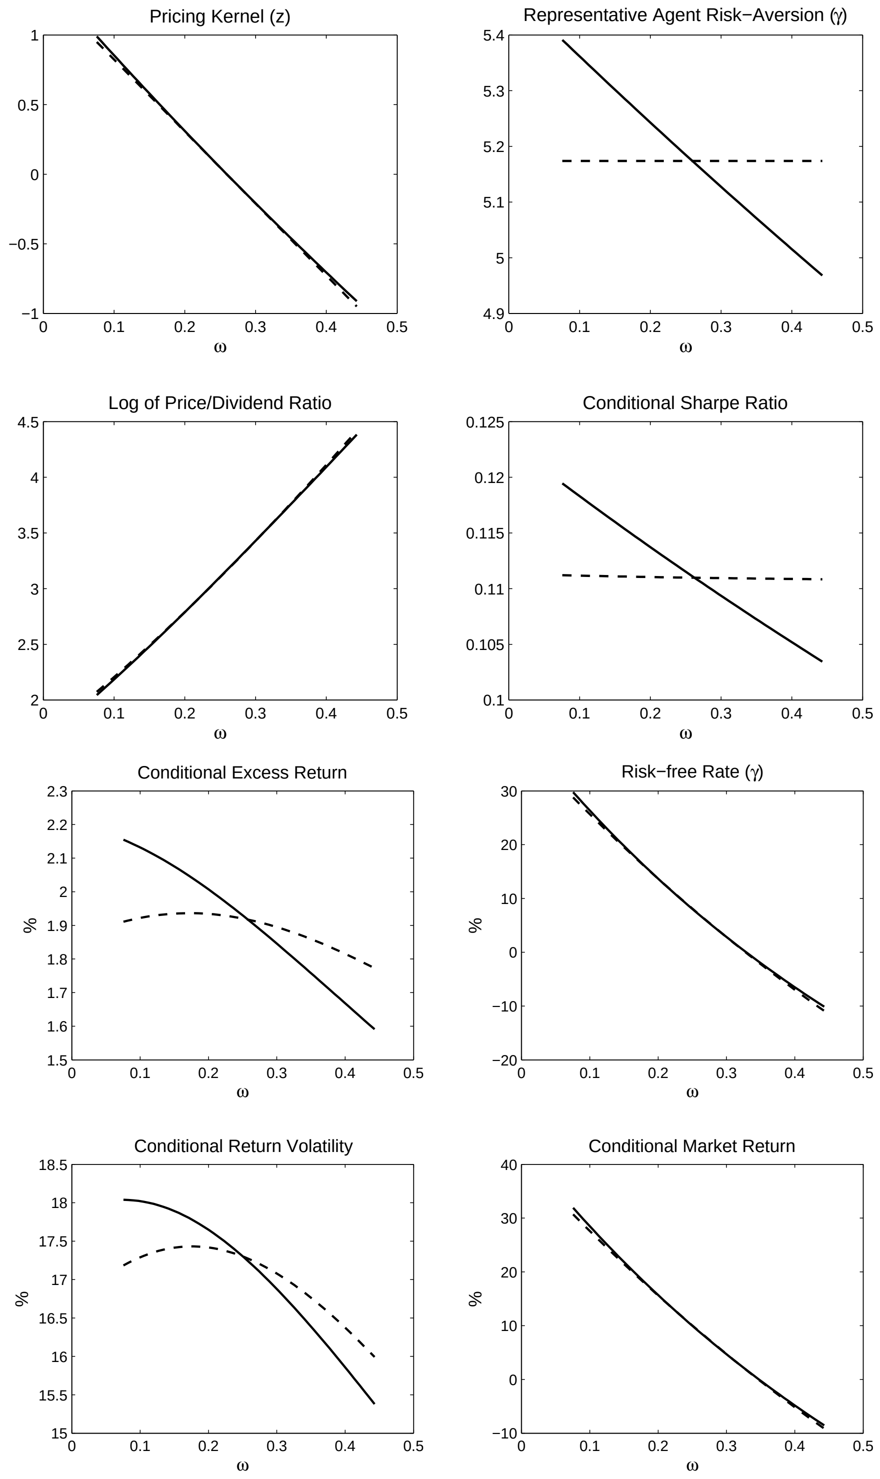

The range of the state variable is four standard deviations around its mean.

Figure 9: Model parametrization 2

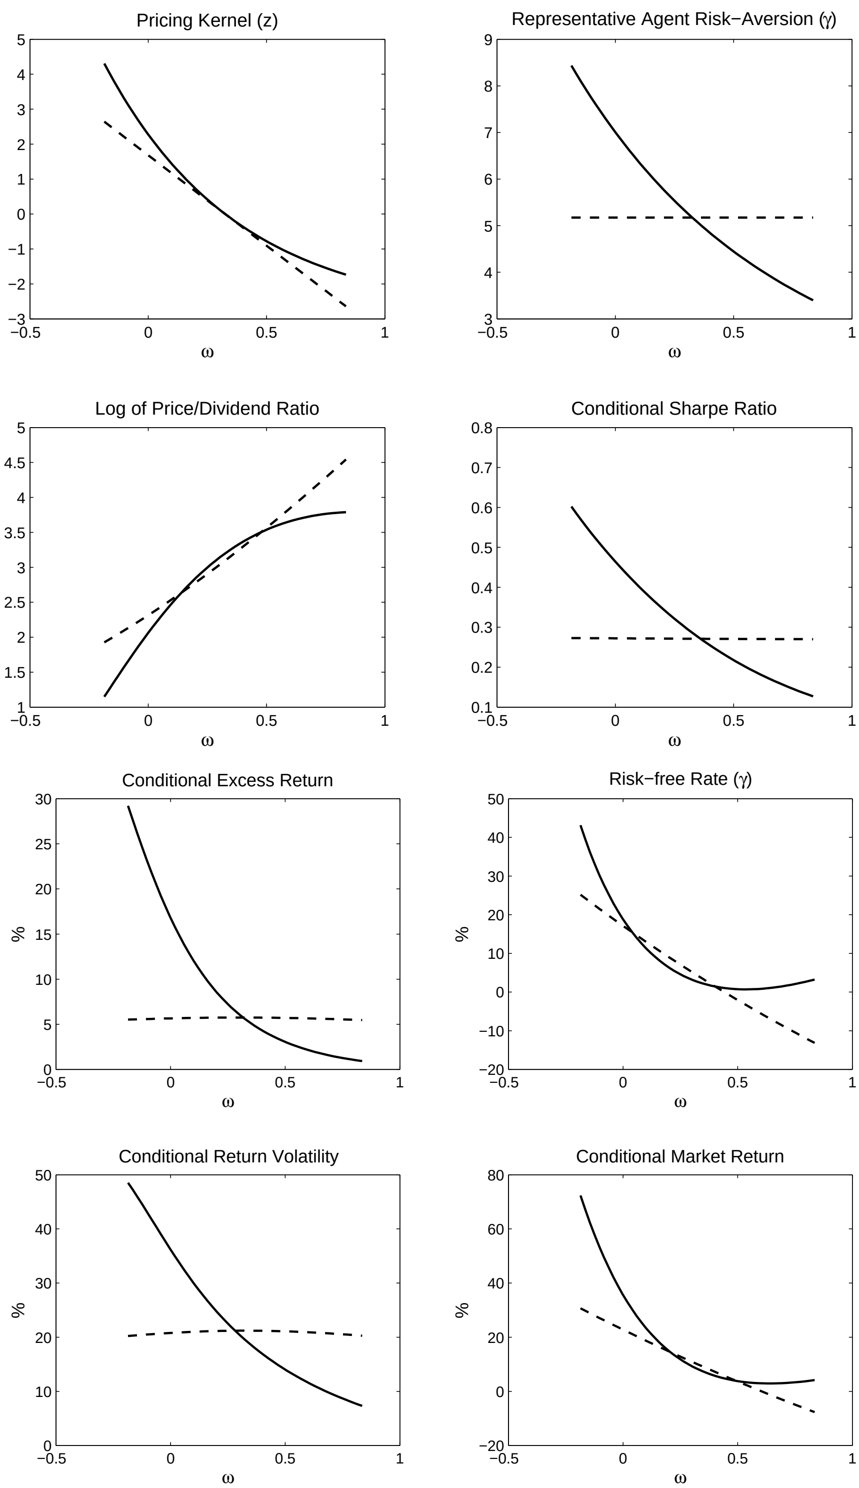

The range of the state variable is four standard deviations around its mean.