Limited participation in the joint behavior of asset prices and individual consumptions

Veronika Czellar, René Garcia and François Le Grand*

First version: August 2014
This version: August 22, 2017

*Czellar: Department of Accounting, Law, Finance and Economics, EDHEC Business School, 24 avenue Gustave Delory, 59057 Roubaix, France, veronika.czellar@edhec.edu. Garcia: Department of Economics, University of Montréal, Pavillon Lionel-Groulx, 3150 rue Jean-Brillant, Montréal QC H3T 1N8, Canada, rene.garcia@umontreal.ca, Toulouse School of Economics. Le Grand: Department of Finance, Economics and Control, EMLYON Business School, 23 avenue Guy de Collongue, 69134 Ecully, France, legrand@em-lyon.com, ETH Zurich (Chair of Integrative Risk Management and Economics). We received helpful comments from Torben Andersen, Abraham Lioui, Kevin Sheppard, George Tauchen, Raman Uppal and participants of the EMLYON Research Workshop in May 2015, the Workshop on Methodological Advances in Statistics Related to Big Data in June 2015, the EDHEC finance seminar in November 2015, the Computational and Financial Econometrics Conference in December 2015, the Toulouse Financial Econometrics Conference in May 2016, and the 10$^{\text{th}}$ Annual SOFIE meeting in June 2017.

Abstract

We propose an asset pricing model featuring both limited participation and heterogeneity, in which agents randomly participate in the bond and stock markets according to a probability that depends on their non-financial income. We develop an indirect inference method to estimate our model on individual US consumption (CEX) and financial data. Our estimated model performs very well at jointly replicating the equity premium and the unequal distribution of individual consumptions. As an external validity check, our model accurately predicts the estimated stock market participation cost and its decline over the period 1980-2004, as well as observed financial market participation.

Keywords: limited participation, heterogeneity, indirect inference, individual consumption distribution.

1 Introduction

The seminal paper of Mehra and Prescott (1985) makes it clear that aggregate consumption growth is not fluctuating enough to explain the equity premium in a model with time-additive utility and frictionless complete markets, unless agents are extremely risk averse. While the main stream of the literature has focused on elaborating economically-reasonable preferences for the representative agent that would make the consumption-based stochastic discount factor more variable[1], another fruitful research effort has relaxed the assumption of full insurance against individual income risk and introduced consumer heterogeneity. Under

[1] The two benchmark models feature habit formation (Campbell and Cochrane, 1999) and recursive preferences à la Epstein and Zin (1989) for the representative agent. Recent contributions have combined recursive utility with long-run risk to capture several stylized facts regarding the equity premium, its volatility and the predictability of asset returns (Bansal and Yaron, 2004; Bonomo et al., 2011).

certain conditions on the individual income processes relative to aggregate income, Constantinides and Duffie (1996) show that asset prices can be supported by an exchange economy equilibrium[2]. Limited participation in financial markets is another important economic reason for agent heterogeneity to influence asset prices. Indeed, not all US households have significant amounts of savings in financial markets, and among those who do, not all trade stocks[3].

In this paper, we construct an asset pricing model featuring time-varying limited participation in both bond and stock markets and agent heterogeneity. Our model considers an economy populated by a large number of agents who participate in the financial markets with a certain probability that depends on their individual income. Non-participants, participants in the bond market, and bond and stock holders are respectively endowed with an individual log-consumption growth process, which is consistent with their behavior. More precisely, if an agent participates in the bond market, the Euler condition linking its consumption growth to bond return will hold with equality. For a non-participant, an inequality will prevail. For a stockholder, who also holds bonds in our model, the Euler equation will hold with equality and include both bond and stock returns. These Euler conditions allow us to infer the endogenous presence of limited participation and to assess the magnitude of financial market participation costs.

We propose an indirect inference (Gouriéroux, Monfort, and Renault, 1993; Smith, 1993) method to estimate the parameters of our model. Indirect inference is a simulation-based estimation method which is increasingly used in the financial economics literature[4], as well

[2]An essential feature is that labor income must be a unit-root process with innovations which become more volatile during aggregate downturns. Mankiw (1986) introduces the idea that aggregate shocks and the volatility of idiosyncratic shocks are negatively related. Storesletten et al. (2007) add life-cycle effects and capital accumulation to the model of Constantinides and Duffie (1996).

[3]See supporting references in Guvenen (2009). Favilukis (2013) provides percentages of stockholders over time.

[4]See for instance Calvet and Czellar (2015), Calzolari et al. (2004), Czellar et al. (2007), Garcia et al. (2011), Sentana et al. (2008) among many others.

as in the macroeconomics literature with agent heterogeneity[5]. This estimation method is particularly well suited to problems where the structural model – our heterogenous-agent model with limited participation in financial markets – is hard to estimate with usual methods such as maximum likelihood but easy to simulate. The idea is then to formulate a reduced-form model (so-called auxiliary model) that is easy to estimate but reflects only partially the structural model. The last step is to minimize the distance between the auxiliary estimators obtained with simulated data generated by the structural model and observed empirical data.

In our case, we simulate the structural model involving the individual consumption growth processes and the individual income-based probabilities of participating in the bond and stock markets, as well as the equity premium stochastic process, and estimate a reduced-form regression model linking aggregate consumption to the time series of the risk-free rate, the equity returns, and aggregate income. We build aggregate consumption growth from the weighted average of both simulated and observed individual consumption growths, and aggregate income from the weighted average of observed individual incomes. We validate our estimation procedure with a thorough Monte Carlo study of the resulting indirect inference estimator.

We estimate the model with individual consumption and income data provided in the Consumer Expenditure Survey (CEX) conducted in the United States together with the returns on the 3-month Treasury and the S&P 500 equity index. The CEX is a rotating panel of households that collects consumption data over a maximum of four consecutive quarters for each household, and we consider in each quarter all available data on consumption and income[6]. Participation in financial markets is determined endogenously through the estimation of the model[7].

[5]See in particular Guvenen and Smith (2014) and Berger and Vavra (2015).

[6]Even though the survey includes partial information on financial asset holdings of the household we do not use this information to elicit participation in financial markets.

[7]In order to keep as many data points as possible, we only discard households that do not report two

The paper delivers three main results. First, our estimated model jointly replicates the asset price properties together with the tremendous heterogeneity in individual consumption observed in the data. This parsimonious asset pricing model relying on limited participation is sufficient to account for most of the dispersion in individual consumptions, while being consistent with bond and stock returns, as well as their standard deviations. Second, we show that limited market participation is necessary for accurate replication of asset prices and individual consumption. The unlimited participation variant of the model leads to much less accurate replications of empirical individual consumptions. Furthermore, our estimation enables us to discuss endogenous limited participation and compute financial market participation costs. On the one hand, we find that there is no significant cost for participating in the bond market and that bond-market non-participation is endogenous. On the other hand, the stock market participation is consistent with the presence of fixed transaction costs. We find that there exists a level of stock market participation cost such that richer agents are willing to trade stocks, while other agents prefer not to do so and only trade bonds. These results confirm previous findings in the literature, such as Vissing-Jorgensen (2002b) for instance. More precisely, we obtain an average yearly cost amounting to 330 USD (in 1980 value) and a decreasing per period cost over 1980-2004. Finally, the estimated model generates a financial market participation that is consistent with empirical data provided in the Survey of Consumer Finances[8]. The model accurately estimates that 40 to 45% of households participated in stock markets in 2004 and that this share has been steadily increasing from 1980 on. These last two findings are strong external validity checks for our model and our estimation strategy, since they are obtained while using only consumption and asset prices data – and no data on financial market participation or financial

consecutive periods of consumption and those who report strictly negative revenues. Indeed, since we estimate log-consumption growth processes, we must have consumption level data for at least two consecutive periods. In addition, probabilities of participating in financial markets directly depend on household revenues and hence require positive revenues. Beyond this criterion, we do not discard further data.

[8]The CEX does not allow to measure participation rate.

wealth – for our estimation.

Our paper is related to four main strands of literature. First, it provides a new and important contribution to the equity premium literature. We show that a model where agents have the same power utility with reasonable discount rate and elasticity of intertemporal substitution (EIS) can be consistent with the observed equity premium level and volatility. Apart from using individual income and consumption data, the sole source of heterogeneity comes from the fact that some agents do not participe in financial markets, others only participate in the bond market, yet others hold bonds and stocks. This is different from the model in Guvenen (2009) where there are two categories of agents with different EIS but also with recursive utility preferences[9]. Vissing-Jorgensen (2002) reports higher estimates of the EIS for non-stockholders but emphasizes that differences in EIS estimates between asset holders and non-asset holders should not be interpreted as evidence of heterogeneity in the EIS across households[10]. Similarly, Brav et al. (2002) show that asset prices are consistent with the assumptions of limited-participation and market participant heterogeneity.

Second, we relate to papers linking individual consumption to income. Heathcote et al. (2010) document a continuous and sizable increase in wage inequality over the 1980-2004 period but find that access to financial markets has curbed the level and growth of consumption inequality. Primiceri and van Rens (2009) use the CEX data on consumption and income to decompose idiosyncratic changes in income into predictable life-cycle changes, transitory and permanent shocks and estimate the contribution of each to total inequality of income and consumption. Guvenen and Smith (2014) use the joint dynamics of individuals' labor earnings and consumption-choice decisions to quantify individual income risk and their access to informal insurance against this risk. Gourinchas and Parker (2002) use income and

[9]However, as mentioned in footnote 3 of Guvenen (2009), the disentangling of the risk aversion from the EIS is not essential for the main substantive results of the paper.

[10]The main reason is that EIS estimates are not consistent estimates for non-asset holders since their Euler condition does not hold with equality if they do not hold assets. This is precisely what distinguishes our model where Euler conditions of non-asset holders are not related to asset returns.

consumption data in sequence to explain the hump-shaped consumption profile in a life-cycle consumption-savings model.

Third, our estimation strategy allows us to obtain an estimate of the participation cost per period in the stock market. Several figures have been selected through estimation or calibration in the literature. Vissing-Jorgensen (2002b) finds evidence of a positive relation between nonfinancial income and the probability of stock market participation. This is the central ingredient in determining participation in our model. Furthermore, she estimates that a per period stock market participation cost of just 350 USD is sufficient to explain the choices of 75% of stock market nonparticipants. Our estimation is of the same order of magnitude. In an incomplete markets, overlapping generations model, Favilukis (2013) shows that lower participation costs can explain the substantial increase in stock market participation and the smaller increase in consumption inequality observed in the last 30 years. The model also tracks the decline in interest rates and the expected equity premium.

Finally, our indirect inference estimation method has been used recently in two macroeconomic models with heterogeneity. Guvenen and Smith (2014) estimates a structural consumption-savings model where they study the joint dynamics of individuals' labor earnings and consumption-choice decisions. Their auxiliary model is based on approximations of the true structural equations. Berger and Vavra (2015) use indirect inference together with calibration to estimate a heterogeneous agent incomplete markets model with fixed costs of durable adjustment to study durable expenditures dynamics during recessions.

The remainder of the paper is organized as follows. Section 2 presents our theoretical framework. Section 3 describes the data that we use. Section 4 provides an estimation method for the model introduced in Section 2. We discuss estimation results in Section 5 and run a robustness check in Section 6. Section 7 concludes. Proofs and calculation details can be found in the Appendix.

2 Theoretical Setup

Our setup is a very standard consumption-based asset pricing model, with a main twist: agents can randomly enter and quit financial markets, with a probability that depends on their income. We specify consumption and asset return processes that jointly satisfy Euler conditions for stockholders, bondholders and non-financial market participants.

We consider a discrete-time economy, where time is denoted by $t = 1, \ldots, T$. The economy is populated by N finite-life agents who enter the economy at various dates and live for at most $T' \leq T$ consecutive periods.

2.1 Financial Markets

Agents consume a single good and can intertemporally smooth out their consumption using financial securities. Two securities can be traded: a riskless bond and a risky stock. Financial security returns are real and expressed in consumption units. First, investing one consumption unit in bonds at date t pays off R^f_{t+1} consumption units at date $t+1$. This return is assumed to be deterministic from date t point of view. Similarly, the investment of one unit of consumption in stocks at date t pays off R^s_{t+1} consumption units at date $t+1$. The stock return is stochastic and we assume that it has the following dynamics:

$$\log R^s_t = \log R^f_t + (\mu - 1/2)\sigma^2_u + \sigma_u u_t, \quad t = 1, \ldots, T\,, \tag{2.1}$$

where $\{u_t\}_{t=1,\ldots,T}$ are independent, identically distributed standard normal variables. Equation (2.1) determines the dynamics of the equity premium, $\log R^s_t - \log R^f_t$, governed by two parameters μ and σ_u[11]. The parameter σ_u is the conditional variance of the equity premium,

[11]We ignore voluntarily the well-documented heteroscedastic behavior of the stock returns. Since we aim at finding a simple consumption-based model with heterogeneity and limited participation, we focus on estimating the mean and variance parameters of the stock returns process consistently with reasonable preference parameters. A time-varying variance could change the proportions of agents participating in the bond and stock markets but the main determinant of these proportions remains the income distributions

while the conditional mean $(\mu - 1/2)\sigma_u$ is affected by the parameter μ.

2.2 Participation in Financial Markets

Agents participation in financial markets is stochastic. They face a two-stage financial market participation risk. A first shock makes agents enter or quit financial markets. When out of financial markets, agents cannot save and are hand-to-mouth consumers. When participating in financial markets, agents trade bonds. These participating agents face a second shock driving their stock market participation. Formally, we denote by $\{\tilde{h}_t^i\}_{t\geq 1}$ the participation process of agent i in the bond market. If $\tilde{h}_t^i = 0$, agent i is excluded from the bond market at date t. The bond market participations $\{\tilde{h}_t^i\}_{t\geq 1}$ are independent processes defined on $\{1, 0\}$ in the following way. Denote by $\tilde{p}_t^i$ the probability of $\tilde{h}_t^i = 1$. We suppose that:

$$\tilde{p}_t^i = \min(\tilde{a}\zeta_t^i, 1)\,, \tag{2.2}$$

where $\tilde{a} > 0$ and ζ_t^i is the individual income of agent i at date t (in 10,000 USD) divided by the household size[12]. Similarly, $\{h_t^i\}_{t\geq 1}$ characterizes agent i's willingness to participate in the stock market. The $\{h_t^i\}_{t\geq 1}$ processes are defined in the same way as bond market participations but with probability p_t^i of $h_t^i = 1$ defined as:

$$p_t^i = \min(a\zeta_t^i, 1)\,, \tag{2.3}$$

where $a > 0$. However, willingness to participate in the stock market is not a sufficient condition to actually participate in it. In order to participate in the stock market, an agent willing to participate in the stock market needs to be a bond market participant first. Hence, the stock market participations are given by $\{\tilde{h}_t^i h_t^i\}_{t\geq 1}$. We suppose that $\tilde{h}_t^i$ and h_t^i

across agents over time.

[12]The literature on consumption and saving establishes clearly that saving decreases with family size (see the survey by Browning and Lusardi (1996)).

are independent which implies that the probability for agent i to participate in the stock market is:

$$\tilde{p}_t^i p_t^i = \min(\tilde{a}\zeta_t^i, 1)\min(a\zeta_t^i, 1)\,. \tag{2.4}$$

We can also interpret the $\{h_t^i\}_{t\geq 1}$ process as a conditional stock market participation. Indeed, conditional on participating in bond markets ($\tilde{h}_t^i = 1$), the participation of agent i in the stock market will be driven by h_t^i. More precisely, if $\tilde{h}_t^i h_t^i = 1$, the agent i will trade stocks. Our choice of h_t^i and $\tilde{h}_t^i$ enables us to model conditional stock market participation by using two independent processes[13].

The above-mentioned idiosyncratic shocks may cover several aspects of individual risk, such as unemployment risk, income risk and health risk among many others. For several reasons, these shocks are to a large extent non-insurable, or at least very imperfectly insurable. For instance, a full employment insurance coverage would be very likely to fail due to moral hazard. Being non-insurable, these shocks generate a precautionary behavior among agents. It notably implies that agents save so as to self-insure against these idiosyncratic risks. This is the so-called precautionary saving motive which has been shown to have a strong impact on asset prices and on riskless returns (Carroll, 1997; Huggett, 1993). These individual shocks may also generate endogenous limited participation. A succession of bad shock realizations – such as being unemployed for a number of consecutive periods – may lead agents to sell off their savings and to quit financial markets[14].

Instead of detailing the idiosyncratic risk machinery, we jointly model asset prices and

[13] The literature on the determinants of stock ownership reports that the probability of holding stocks is increasing in wealth and education. So income should be a good proxy given its high correlation with wealth and education. Moreover, financial wealth data in the CEX is only reported in the last quarter in which the household participates in the survey and does not allow for a perfect separation between stockholders and non-stockholders (see Vissing-Jorgensen, 2002). Finally, there is obviously no unicity for the functional form in (2.2) or (2.3). For the indirect inference estimation to remain tractable, we have chosen a simple functional form featuring an increasing and concave relationship between income and participation probability.

[14] The combination of limited market participation with idiosyncratic risks is very fruitful to jointly explain consumption inequalities and asset prices (equity premium and riskless return) – see LeGrand and Ragot (2015).

income and consumption processes and combine idiosyncratic risks with financial markets via time-varying financial market participation statuses.

2.3 Agents' Programs

Agents are endowed with time-separable preferences. Their constant discount factor is denoted by β. Agents consume a single consumption good and their elasticity of intertemporal substitution (EIS, henceforth) is assumed to be constant and common to all agents. We denote the inverse of EIS by γ. In consequence, the utility function is of CRRA type and for a given consumption level c, can be expressed as follows:

$$u(c) = \begin{cases} \frac{c^{1-\gamma}-1}{1-\gamma} \text{ if } \gamma \neq 1, \\ \ln(c) \text{ otherwise.} \end{cases} \tag{2.5}$$

Note that the case with $\gamma = 1$ in (2.5) is a continuous extension of the case with $\gamma \neq 1$.

We assume that agents are expected-utility maximizers. At any date t, a given agent i chooses her consumption level c_t^i, her bond holdings b_t^i and her stock holdings x_t^i, so as to maximize her intertemporal utility function, subject to budget and financial market participation constraints. The program of an agent i can be expressed as follows

$$\max_{(c_t^i, b_t^i, x_t^i)_{t=1,\ldots,T}} \mathbb{E} \sum_{t=1}^{T} \beta^t u(c_t^i) \,, \tag{2.6}$$

$$\text{s.t. } c_t^i + b_t^i + x_t^i + F\, 1_{x_t^i > 0} \leq \zeta_t^i + R_t^f b_{t-1}^i + R_t^s x_{t-1}^i \,, \tag{2.7}$$

$$b_t^i = 0 \quad \text{if} \quad \tilde{h}_t^i = 0 \,, \tag{2.8}$$

$$x_t^i = 0 \quad \text{if} \quad \tilde{h}_t^i h_t^i = 0 \,, \tag{2.9}$$

$$c_t^i > 0 \,, \tag{2.10}$$

where $1_{x_t^i > 0}$ is an indicator function equal to 1 when $x_t^i > 0$ and 0 otherwise. The $\mathbb{E}$ operator is the expectation with respect to individual and aggregate shocks in the consumption $\{c_t^i\}$ and return $\{R_t^s\}$ processes. The budget constraint equation (2.7) states that the resources of agent i at date t, consisting of an income revenue ζ_t^i, bond and stock payoffs $R_t^f b_{t-1}^i + R_t^s x_{t-1}^i$ should cover the agent's consumption c_t^i, and her investment in bonds, b_t^i, and in stocks, x_t^i, and the stock market participation cost F in case she eventually invests in stocks ($x_t^i \neq 0$). Consistently with our description of limited participation in Section 2.2, the agent's financial market participation is constrained by equations (2.8)-(2.9), which state that the agent can trade neither bonds nor stocks if $\tilde{h}_t^i = 0$. In addition, equation (2.9) also states that the agent cannot trade stocks if $h_t^i = 0$.

2.4 Consumption Processes

We now specify the dynamics of individual consumption processes. We denote by $\Delta \log c_{t+1}^i = \log(c_{t+1}^i / c_t^i)$ the log-consumption growth process of household i between dates t and $t+1$. For any date $t = 1, \dots, T-1$, and any household i in the sample at date t, we assume that the log-consumption process satisfies the following dynamics:

$$\begin{aligned} \Delta \log c_{t+1}^i = \frac{1}{\gamma} \Bigg\{ & \log \beta + \log(R_{t+1}^f) + \frac{\mu^2 \sigma_u^2}{2} \lambda^{2 - 2 h_t^i \tilde{h}_t^i} \\ & + (1 - \tilde{h}_t^i) \frac{\sigma_e^2}{2} + (1 - \tilde{h}_t^i) \sigma_e \varepsilon_{t+1}^i + \tilde{h}_t^i \mu \lambda^{1 - h_t^i \tilde{h}_t^i} \sigma_u u_{t+1} \Bigg\}, \end{aligned} \tag{2.11}$$

where $\{\varepsilon_t^i\}$ are independent, identically distributed standard normal variables and $\{u_t\}$ are the stock return shocks in (2.1). The variables $\{\varepsilon_t^i\}$ and $\{u_t\}$ are mutually independent. Note that the model (2.11) can be rewritten as

$$\begin{aligned} \Delta \log c_{t+1}^i = \frac{1}{\gamma} \Bigg\{ & \log \beta + \log(R_{t+1}^f) + \frac{\mu^2 \sigma_u^2}{2} [(1 - \lambda^2) h_t^i \tilde{h}_t^i + \lambda^2] + (1 - \tilde{h}_t^i) \frac{\sigma_e^2}{2} \\ & + (1 - \tilde{h}_t^i) \sigma_e \varepsilon_{t+1}^i + \mu [(1 - \lambda) h_t^i \tilde{h}_t^i + \lambda \tilde{h}_t^i] \sigma_u u_{t+1} \Bigg\} . \end{aligned} \tag{2.12}$$

The log-consumption growth of a household i between dates t and $t+1$ in equation (2.12) depends on the statuses $\tilde{h}_t^i$ and h_t^i. When the agent i faces $\tilde{h}_t^i = 0$, the value of variable h_t^i has no impact on the log-consumption growth: this is consistent with the fact that the variable $\tilde{h}_t^i$ characterizes participation in financial markets ($\tilde{h}_t^i = 0$ implies that agent i does not participate in either bond or stock markets), while h_t^i only characterizes stock market participation ($h_t^i = 0$ implies that agent i does not participate in the stock market but she can participate in bond markets).

The log-consumption growth is affected by two shocks. The first one, denoted ε_t^i, is an individual shock for agents who are not participating in financial markets, while the second one, denoted u_t, is an aggregate shock. The variance of the shock for non-participants is equal to σ_e^2/γ^2. The aggregate shock affects all agents participating in financial markets ($\tilde{h}_t^i = 1$), but has no effect for agents out of financial markets ($\tilde{h}_t^i = 0$). We recall that this shock also impacts the stock return, as visible in equation (2.1). The variance of the aggregate shock is driven by the stock return variance σ_u^2 multiplied by μ^2/γ^2. This variance is also affected by the parameter λ^2 for agents who do not participate in the stock market (but participate in the bond market).

Our model assumes that non-asset holders support the full brunt of the idiosyncratic shocks, having no way to partially insure themselves, but none of the aggregate shocks. At the other end of the spectrum, stockholders are fully insured against the idiosyncratic shocks but face the macroeconomic uncertainty associated with the stock market. Bondholders insure themselves against idiosyncratic shocks by holding the risk-free asset but are exposed to the aggregate shock. This is certainly a simplified characterization of the risk positions of the three groups but it should serve to identify better the volatility and exposure (λ) parameters.

2.5 Euler Conditions

We now derive the conditions under which individual consumption processes (2.11) are consistent with Euler equations derived from the agent's maximization program (2.6)–(2.9)[15]. In particular, we check that if $\tilde{h}_t^i = 1$, agent i prices bonds at date t and similarly that if $\tilde{h}_t^i h_t^i = 1$, agent i prices stocks at date t.

More precisely, for the bond market, the following proposition holds.

Proposition 1 (Euler conditions for bonds) *The bond Euler equation holds with equality for agent i at date t if $\tilde{h}_t^i = 1$:*

$$\mathbb{E}\left[\exp\{-\gamma \Delta \log c_{t+1}^i\} \middle| \, \tilde{h}_t^i = 1\right] = \frac{1}{\beta R_{t+1}^f}, \tag{2.13}$$

while it holds with strict inequality for agent i at date t if $\tilde{h}_t^i = 0$:

$$\mathbb{E}\left[\exp\{-\gamma \Delta \log c_{t+1}^i\} \middle| \, \tilde{h}_t^i = 0\right] < \frac{1}{\beta R_{t+1}^f}. \tag{2.14}$$

Proposition 1 states that the variable $\tilde{h}_t^i$ is consistent with the behavior of the agent i on the bond (riskless saving) market at date t. If $\tilde{h}_t^i = 1$, agent i prices the bond at date t and equality (2.13) holds. If $\tilde{h}_t^i = 0$, riskless savings pay off a return which is too small compared with agent i's own valuation at date t. The intuition is as follows. Agents with $\tilde{h}_t^i = 0$ have a log-consumption growth which is not affected by the aggregate shock u_t. In consequence, they are less willing to save and are excluded from bond markets. Bond savings are too expensive – or equivalently riskless savings pay off a return which is too small – for them to accept investing in riskless savings.

The next proposition states a similar result for stocks.

[15] The fact that we use log linearized Euler conditions should reduce measurement errors compared with nonlinear Euler equations (see the Monte Carlo study of Vissing-Jorgensen, 2002).

Proposition 2 (Euler conditions for stocks) *The stock Euler condition holds with equality for agent i at date t if $\tilde{h}_t^i = \tilde{h}_t^i h_t^i = 1$:*

$$\mathbb{E}\left[R_{t+1}^s \exp\{-\gamma \Delta \log c_{t+1}^i\} \middle| \tilde{h}_t^i = 1, \tilde{h}_t^i h_t^i = 1 \right] = \frac{1}{\beta}. \tag{2.15}$$

The stock Euler condition holds with strict inequality for agent i at date t

$$\mathbb{E}\left[R_{t+1}^s \exp\{-\gamma \Delta \log c_{t+1}^i\} \middle| \tilde{h}_t^i = 0, \tilde{h}_t^i h_t^i = 0 \right] < \frac{1}{\beta}, \tag{2.16}$$

if $\tilde{h}_t^i = 0$ and $\mu(1 - \frac{\mu\lambda^2}{2}) < 0$ and

$$\mathbb{E}\left[R_{t+1}^s \exp\{-\gamma \Delta \log c_{t+1}^i\} \middle| \tilde{h}_t^i = 1, \tilde{h}_t^i h_t^i = 0 \right] < \frac{1}{\beta}, \tag{2.17}$$

if $\tilde{h}_t^i = 1$, $h_t^i = 0$ and $\mu(1 - \lambda) < 0$.

The h_t^i variable is consistent with the behavior of agents on stock markets. An agent i with $\tilde{h}_t^i = h_t^i = 1$ prices the stocks at date t. Alternatively, if $\tilde{h}_t^i = 0$, agent i consistently finds stocks too expensive to be traded at date t whenever the condition $\mu(1 - \lambda) < 0$ holds. An agent facing $\tilde{h}_t^i = 0$ refuses therefore to trade both stocks and bonds.

3 Data

Our data stem from the Consumer Expenditure Survey (CEX) conducted between the first quarter of 1980 and the first quarter of 2004. The CEX contains quarterly data about US households collected via regular surveys. Each selected household is interviewed for a maximum of 4 consecutive quarters. In each interview, the household reports a detailed account of consumption, income and household composition. We use the dataset of Krueger and Perri (2005).

We are mainly interested in consumption data and in particular in the total consumption of households expressed in real terms. Since our model relies on log-consumption growth, we only retain in our dataset households that report consumption data for at least two consecutive periods. Contrarily to what is commonly done in the literature, we do not drop further households out of the sample. This way, we have $N = 117,857$ households that are occasionally observed over $T = 97$ periods, providing the consumption matrix:

$$\begin{pmatrix} c_1^1 & c_2^1 & \cdots & c_T^1 \\ c_1^2 & c_2^2 & \cdots & c_T^2 \\ \vdots & \vdots & \ddots & \vdots \\ c_1^N & c_2^N & \cdots & c_T^N \end{pmatrix}. \tag{3.1}$$

Note, however, that the data matrix is sparse and embeds many zeros since every household is observed for at most $T' = 4$ consecutive periods.

Our choice for consumption data deserves some further comments. We rely on a measure of total consumption expenditures. These expenditures include spending in nondurable goods (such as food), in equipment, in entertainment, as well as in lodging and vehicles. The first elements, that can be classified as nondurable or small durable goods and services, are flow expenditures. The two last elements, lodging and vehicle expenditures, are durable goods. However, lodging expenditures mainly include the rent of primary residence, which is either the actual rent or an estimate for home owners (both values being provided in the CEX). Lodging expenditures are therefore mainly accounted as a flow[16]. Vehicle expenses include purchases of vehicles – which is therefore the only actual durable good expenditure not accounted as a flow – and other vehicle expenses. All these expenditures data come directly from the CEX. Similarly to our choice to retain as many households as possible in our sample, we have decided not to transform data for consumption expenditures and to rely

[16]Note that these expenditures also include owned dwelling and other lodging expenditures. They can be considered as small durable goods and are therefore close to a flow expenditure.

on raw data from the CEX, even though it involves including vehicle purchases, which are not accounted as flows.

Krueger and Perri (2005) propose a solution to make the conversion, from vehicle purchases into equivalent consumption flows. The vehicle flows are defined as a fixed fraction (1/32) of the stock value of vehicles of the household. Since the value of this stock is not observed in the CEX, it must be estimated from vehicles that are currently purchased and then imputed to all households (procedure of Cutler and Katz, 1991). In consequence, transforming vehicle expenditures into flows implies some arbitrariness. To avoid this arbitrariness, we decided to use raw data from the CEX, without any transformation. However, to check that our results are not driven by our choice concerning vehicle data, we run a robustness check in Section 6, in which we estimate our model using the pure flow consumption expenditures. It turns out that it has little impact on our results.

We also use the reported income of each household i at time t to estimate the parameters of the probability functions determining the participation in financial markets as modeled in equations (2.2) and (2.3). The measure of income includes wages, salaries, and government transfers (such as unemployment insurance benefits) for all household members, net of taxes and of social contributions.

Besides total consumption and income, we also use the weights of households provided in the CEX dataset, which represent their relative share in the US population. If either the date t or date $t+1$ consumption is missing in the dataset, we set $\omega_t^i = 0$. Similarly, if the revenue of agent i at date t is negative, we set $\omega_t^i = 0$. Finally, we define normalized weights $w_t^i = \omega_t^i / \sum_{i=1}^N \omega_t^i$ which guarantee that $\sum_{i=1}^N w_t^i = 1$ for each $t = 1, \ldots, T$.

Financial returns have been computed on a quarterly basis using Shiller's dataset[17]. The riskless return is the 3-month US T-Bill rate, while the stock return is the S&P 500 return (dividends included). Financial returns cover the same time period as the CEX data.

[17] We obtained the data from the Online Data Robert Shiller's website http://www.econ.yale.edu/ shiller/data.htm.

4 Estimation

We now turn to the estimation of the structural model defined in Section 2. A direct estimation of the model is challenging due to the fact that participations in financial markets are unobserved and individual consumptions are only observed for a handful of periods among the T dates. However the model is clearly easy to simulate. We therefore introduce an indirect inference (Gouriéroux, Monfort, and Renault, 1993; Smith, 1993) method which is a two-step estimation procedure. In a first step, an auxiliary model different from the structural model is chosen such that it is easy to estimate. Note that the auxiliary model is misspecified with respect to the structural model and is not unique[18]. Its simplicity also offers a way to obtain estimates of the structural model that are robust to model misspecification. Even though our structural model is hard to estimate it does not completely capture the optimal consumption and investment decisions of the household and is therefore itself misspecified. The second step of the estimation consists in generating by simulation pseudo-data under the structural model and compute auxiliary estimates based on these simulated data. The structural parameter estimates are obtained when they produce simulated auxiliary estimates that are close to the empirical auxiliary estimates.

We denote the structural parameter vector by

$$\theta = (\beta, \gamma, \lambda, \mu, \sigma_e, \sigma_u, \tilde{a}, a)' \, . \tag{4.1}$$

4.1 An Auxiliary Model

We consider an auxiliary model in which there exists only one representative agent whose consumption c_t and revenue ζ_t are observed at each date $t = 1, \ldots, T$. We create such an

[18]Indirect inference leaves considerable flexibility in choosing the auxiliary model. An important requirement is that the relationship that binds the structural and the auxiliary parameters must identify the respective effects of each structural parameter on at least one auxiliary parameter. Formally, the mapping must be injective in the vicinity of the true structural parameters. In our case we will have as many auxiliary parameters as structural parameters and the relationship will be obviously invertible.

artificial single agent by aggregating individual consumption and revenue as follows:

$$\Delta \log c_{t+1} = \sum_{i=1}^{N} w_{t+1}^{i} \Delta \log c_{t+1}^{i} \, , \tag{4.2}$$

$$\zeta_t = \sum_{i=1}^{N} w_{t+1}^{i} \zeta_t^{i} \, , \tag{4.3}$$

$$\zeta_t^2 = \sum_{i=1}^{N} w_{t+1}^{i} (\zeta_t^{i})^2 \, . \tag{4.4}$$

We assume that this auxiliary model is governed by the dynamics in (2.1)-(2.3) and a modified version of (2.12) in which the nonobservable $\tilde{h}_t$ is replaced by $\tilde{a}\zeta_t$ (and h_t by $a\zeta_t$):

$$\begin{aligned} \Delta \log c_{t+1} = \frac{1}{\gamma} \Big\{ & \log \beta + \log(R_{t+1}^{f}) + \frac{\mu^2 \sigma_u^2}{2} [(1-\lambda^2)\tilde{a}a\zeta_t^2 + \lambda^2] + (1-\tilde{a}\zeta_t)\frac{\sigma_e^2}{2} \\ & + (1-\tilde{a}\zeta_t)\sigma_e \varepsilon_{t+1} + \mu[(1-\lambda)\tilde{a}a\zeta_t^2 + \lambda\tilde{a}\zeta_t]\sigma_u u_{t+1} \Big\} \, . \end{aligned} \tag{4.5}$$

We now proceed to calculating the auxiliary estimator

$$\hat{\eta} = (\hat{\eta}_1, \ldots, \hat{\eta}_8)' \, , \tag{4.6}$$

from which the estimators of the parameters in θ can be successfully recovered.

We start with $\hat{\eta}_1$ and $\hat{\eta}_2$, which only depend on stock returns. More precisely, in the stock return equation (2.1), we denote the risk premium's mean and variance by $\eta_1 = (\mu - 1/2)\sigma_u^2$ and $\eta_2 = \sigma_u^2$, respectively. These parameters can be estimated by

$$\hat{\eta}_1 = T^{-1} \sum_{t=1}^{T} \left(\log R_t^s - \log R_t^f \right) \, , \quad \hat{\eta}_2 = (T-1)^{-1} \sum_{t=1}^{T} \left(\log R_t^s - \log R_t^f - \hat{\eta}_1 \right)^2 \, . \tag{4.7}$$

We now turn to the other components of $\hat{\eta}$. From the stock return equation (2.1), we deduce that $\sigma_u u_{t+1} = \log(R_{t+1}^s) - \log(R_{t+1}^f) - \eta_1$. The equation (4.5) driving the log-consumption dynamics of the aggregate agent becomes:

$$\begin{aligned}\Delta \log c_{t+1} = \eta_3 + \eta_4 \log(R^f_{t+1}) + \eta_5 \zeta_t + \eta_6 \zeta_t^2 + \eta_7[\log(R^s_{t+1}) - \log(R^f_{t+1}) - \eta_1]\zeta_t + \\ + \eta_8[\log(R^s_{t+1}) - \log(R^f_{t+1}) - \eta_1]\zeta_t^2 + (1 - \tilde{a}\zeta_t)\sigma_e \varepsilon_{t+1}\,,\end{aligned} \tag{4.8}$$

where the components of η are defined as follows: $\eta_3 = (\log \beta + \frac{\mu^2 \sigma_u^2 \lambda^2}{2} + \frac{\sigma_e^2}{2})/\gamma$, $\eta_4 = 1/\gamma$, $\eta_5 = -\tilde{a}\sigma_e^2/(2\gamma)$, $\eta_6 = [\mu^2 \sigma_u^2 (1 - \lambda^2)\tilde{a}a]/(2\gamma)$, $\eta_7 = \mu\lambda\tilde{a}/\gamma$ and $\eta_8 = \mu(1 - \lambda)\tilde{a}a/\gamma$.

The above equation (4.8) for the log-consumption growth of the aggregate agent can simply be rewritten using matrix notation as

$$\Delta \log c_{t+1} = (\eta_3, \ldots, \eta_8)x_t' + (1 - \tilde{a}\zeta_t)\sigma_e \varepsilon_{t+1}\,, \tag{4.9}$$

where for any $t \in \{1, \ldots, T\}$, the vector $x_t \in \mathbb{R}^6$ is defined as

$$x_t = \left(1, \log(R^f_{t+1}), \zeta_t, \zeta_t^2, [\log(R^s_{t+1}) - \log(R^f_{t+1}) - \hat{\eta}_1]\zeta_t, [\log(R^s_{t+1}) - \log(R^f_{t+1}) - \hat{\eta}_1]\zeta_t^2\right)\,.$$

Equation (4.9) makes it clear that $(\eta_3, \ldots, \eta_8)'$ can simply be estimated using a least squares technique. More formally, denoting by $X = (x_1, \ldots, x_{T-1})'$, the estimators can be expressed as:

$$(\hat{\eta}_3, \ldots, \hat{\eta}_8)' = (X'X)^{-1}X'\Delta \log c_{2:T}\,, \tag{4.10}$$

where $\Delta \log c_{2:T}$ is the vector of log-consumption growth from dates 2 to T.

4.2 Indirect Inference

The method to estimate the structural parameter θ requires that for any given parameter vector θ, pseudo-data samples can be generated from the structural model provided in Sections 2.1, 2.2 and 2.4. These pseudo-data are used to compute the components of the auxiliary estimator $\hat{\eta}$ defined in (4.7) and (4.10). The method finally selects the parameter vector $\hat{\theta}_{II}$ that minimizes the distance between the auxiliary estimates calculated with simulated and empirical data.

We now describe the method more formally. Let us now consider $M \geq 1$ pseudo-data samples denoted by:

$$[\Delta \log c_t^{i,m}(\theta), \log R_t^{s,m}(\theta)]_{t=1,\dots,T}^{i=1,\dots,N} \ , \quad m = 1, \dots, M \, .$$

The mth simulated sample $[\Delta \log c_t^{i,m}(\theta), \log R_t^{s,m}(\theta)]$ consists of individual consumption data for the whole population and of stock market returns. The dependence in the vector of parameters θ is made explicit. For each sample m, we can calculate the auxiliary estimate $\hat{\eta}^m(\theta)$ using equations (4.7) and (4.10). Using these auxiliary estimates, we can finally define the simulated auxiliary statistic: $\hat{\eta}(\theta) = M^{-1} \sum_{m=1}^{M} \hat{\eta}^m(\theta)$. An indirect inference estimator is then defined as

$$\hat{\theta}_{II} = \arg\min_{\theta} [\hat{\eta}(\theta) - \hat{\eta}]' \Omega [\hat{\eta}(\theta) - \hat{\eta}] \, , \tag{4.11}$$

where Ω is a symmetric, positive definite matrix. Using the optimal weighting matrix Ω^* defined in Gouriéroux, Monfort, and Renault (1993) and Gouriéroux, and Monfort (1996), and under regularity conditions given therein, an indirect inference estimator is consistent of the structural parameter θ and asymptotically normally distributed.

In addition, this indirect inference environment allows us to perform tests on parameters. More precisely, we can test whether a subset of parameter components is equal to specific values. To do so, we partition the structural parameter as: $\theta = (\theta_1, \theta_2)'$ with $\theta_1 \in \mathbb{R}^{p_1}$ and $\theta_2 \in \mathbb{R}^{8-p_1}$ for $p_1 \in \{1, \dots, 8\}$. We define the null hypothesis as $H_0 : \theta_1 = \theta_{10}$ where θ_{10} contains fixed parameter values. We denote by $\hat{\theta}_0$ the constrained indirect inference estimator satisfying the null hypothesis, i.e.

$$\hat{\theta}_0 = \arg\min_{\theta_2} \left[\hat{\eta} \begin{pmatrix} \theta_{10} \\ \theta_2 \end{pmatrix} - \hat{\eta} \right]' \Omega^* \left[\hat{\eta} \begin{pmatrix} \theta_{10} \\ \theta_2 \end{pmatrix} - \hat{\eta} \right] \, .$$

The test statistic

$$\xi = \frac{TM}{M+1} \left\{ [\hat{\eta}(\hat{\theta}_0) - \hat{\eta}]' \Omega^* [\hat{\eta}(\hat{\theta}_0) - \hat{\eta}] - [\hat{\eta}(\hat{\theta}_{II}) - \hat{\eta}]' \Omega^* [\hat{\eta}(\hat{\theta}_{II}) - \hat{\eta}] \right\}$$

is asymptotically χ^2-distributed with p_1 degrees of freedom.

In empirical and Monte Carlo applications, we use $M = 10$ pseudo-data samples and Ω^* estimated by the inverse of the covariance matrix of 200 auxiliary estimates obtained from pseudo-samples simulated from the structural model with the empirical auxiliary parameter estimates.

4.3 Monte Carlo Results

Before discussing our empirical results, we assess the accuracy of our estimation method by Monte Carlo simulation. We generate 100 samples from the structural model defined in Sections 2.1, 2.2 and 2.4 with the following parameter values: $\beta = 0.735$, $\gamma = 1.196$, $\lambda = -1.188$, $\sigma_e = 0.558$, $\mu = 10.024$, $\sigma_u = 0.062$, $\tilde{a} = 1.019$, $a = 0.330$. These parameter choices are motivated by the empirical estimates reported in Table 1. For each sample, we calculate the auxiliary and indirect inference estimates and summarize the 100 estimates for each method in boxplots reported in Figure 1. Auxiliary estimates are reported on the left, while indirect inference estimates are on the right of each panel. True parameter values are reported by horizontal lines in each panel.

The boxplots show that the auxiliary estimators are severely biased for the parameters λ, σ_e^2, $\tilde{a}$ and a. Indirect inference estimators capture well the heterogeneity and limited participation features of the model and greatly outperform the auxiliary estimators by correcting the biases[19]. It can be noted that the bias and the correction are particularly important for parameters affecting households individually, such as the variance of individual shocks

[19] Gouriéroux, and Monfort (1996) have shown that in the just-identified case, indirect inference corrects finite-sample bias.

σ_e^2, and parameters a and $\tilde{a}$ driving individual financial market participation. Even though our theoretical model can be seen as a simple twist of a standard asset pricing model, these biases show that the heterogeneity in the data and the related empirical properties can only be captured thanks to the indirect inference procedure.

5 Empirical Results

In this section, we first report the parameter estimates of the limited participation model. We then look at implied quantities such as consumption inequality and its dynamics over time. We conclude our analysis by computing the proportions of agents participating in the bond and stock markets and their costs of participation and their evolution over time in our sample.

To highlight the role of the time-varying financial market participation, we compare our estimates to those of an unlimited participation model, in which all agents participate at all dates. The unlimited participation model is thus nested in our general model and corresponds to a constrained specification where we impose $\tilde{h}_t^i = h_t^i = 1$ for all households i at all dates t. The constraint on processes h and $\tilde{h}$ is equivalent to setting the probability parameters a and $\tilde{a}$ to infinity: $\tilde{a} = a = \infty$[20]. As a direct consequence of this specification, the parameters λ and σ_e cannot be estimated in the unlimited participation model, since they are not relevant and correspond to households not participating in some financial markets. Finally, from an estimation perspective, since households all share the same preference parameters, the unlimited participation model aggregates to the standard consumption-based asset pricing model with a representative agent built from the distribution of consumption instead of aggregate consumption.

[20] If $\tilde{a} = a = \infty$, the participation probabilities verify $\tilde{p}_t^i = p_t^i = 1$ for all t and i.

5.1 Parameter Estimates

Results of our empirical estimation are gathered in Table 1. Finite standard errors using the Monte Carlo replicates of Section 4.3 are reported in parentheses[21]. All parameters in the limited participation model are significant, while only σ_u is significant in the unlimited participation model. The first column reports the indirect inference estimates of our unrestricted model with limited participation and the second column summarizes the constrained indirect inference estimates of the restricted model with unlimited participation ($\tilde{h}_t^i = h_t^i = 1$ for all i and t). We also report the value of the ξ statistic when testing the unlimited participation ($H_0 : \tilde{a} = a = \infty$) model. The associated p-value is in parentheses below the ξ statistic value.

We start by discussing the discount factor. In the unlimited participation model, we estimate a value of 0.226 which is nonsignificant. We estimate a value of 0.735 with a standard error of 0.083 under the limited participation model. This value is consistent with the calibrated value recently used in a paper by Favilukis (2013). His model features incomplete markets and overlapping generations. It aims at explaining several stylized facts related to wealth inequality, stock market participation and the equity premium. He calibrates the discount factor at 0.78 and argues that it is a value that, combined with a strong bequest motive for a small portion of the population, explains the high level of wealth inequality in the data. We do not have wealth in our model but we do use the income data, which mirrors the strong inequality in wealth.

The estimated value is also consistent with several experimental studies. For example, Andreoni and Sprenger (2012) estimate an aggregate rate of time discounting ($-\ln(\beta)$, equal to 30% in our estimation) between 25% and 40%, with a standard error between 6% and 17%. When allowing for individual heterogeneous discount rates, they find a median value of approximately 40%, with a very large heterogeneity among agents. The dispersion in the

[21]Finite standard errors are calculated using the 100 Monte Carlo estimates in Section 4.3.

value of the rate of time discounting is confirmed in the review of Andersen et al. (2014, Table 3, p. 28) of experimental studies on the value of the rate of time discounting.

Our results are also in line with several empirical estimates of the discount rate. Warner and Pleeter (2001) estimate the discount rate between 0% and 30% using the US military drawdown program of the early 90's. Carroll and Samwick (1997) find values between 7.5 % and 20%. Using SCF data, Samwick (1998) reports a value between 6% and 11%, depending on the agent's planning horizon. Lawrance (1991) using PSID data estimate a discount rate between 12% and 19%[22].

For empirical estimates in models examining consumption and income dynamics, both Gourinchas and Parker (2002) and Guvenen and Smith (2014) estimate a β close to 0.96. We note that the former paper estimates a model with synthetic cohorts (which means representative agents at different ages) that tends to push the β towards one and that the coefficient of relative risk aversion is set outside the estimation at a value of one. In the latter paper, the coefficient of relative risk aversion is not estimated either and is set to a value of 2. The authors stress that identification of the risk aversion parameter and the time discount factor in consumption-savings without a stochastic interest is tenuous at best.

We can conclude from these various comparisons that our estimate of the time discount factor is not out of line with estimates obtained from studies based on individuals and that the large proportion of non-participants in the stock market tends to lower the estimated β.

For the second parameter γ, the risk aversion parameter in a standard expected-utility model, we find a value of 1.196 in the limited participation model. The unlimited participation model provides a much higher value of 187.960. A very high risk aversion estimate is often associated with aggregate consumption-based asset pricing models. Given low consumption growth volatility[23] and covariance with the stock returns, a very high risk aversion

[22]Note that a couple of studies based on surveys have reached opposite conclusions and report very low and even negative discount rates, such as Loewenstein and Prelec (1991) or Barsky et al. (1994) among others.

[23]In Table 2, the estimated value of log-consumption growth is close to zero under the unlimited partici-

is needed to match the equity premium. The risk aversion value estimated under the limited participation appears much more reasonable and manages to match the equity premium given a good estimate for the implied aggregate growth of log-consumption. We need to remember that it is based on the whole cross-section of individual consumptions. This is in the spirit of Constantinides and Duffie (1996) whereby the variance of the cross-sectional distribution of individual consumptions is added to the Euler condition to better match the premium even under the unlimited participation model. Both limited and unlimited participation models match well the mean and volatility of the equity premium due to agent heterogeneity.

The literature on limited asset market participation has focused instead on the EIS, the inverse of γ in a standard expected power utility model, which in our limited participation model is 0.836. The seminal paper of Vissing-Jorgensen (2002) provides estimates of the EIS based on log linearized Euler conditions of asset holders and non-asset holders. She finds estimates of 0.3–0.4 for stockholders, 0.8–1.0 for bondholders, and small and insignificantly different from zero for non-stockholders and non-bondholders. She particularly emphasizes that these different values should not be interpreted as evidence of heterogeneity in the EIS across households[24]. Our model assumes that all households have the same elasticity of substitution and the overall estimate of 0.836 is not unreasonable given the higher number of bondholders compared to stockholders.

The magnitudes of the volatilities of shocks, σ_e and σ_u, appear reasonable. The latter matches very well the volatility of the equity premium (see Table 2). Given the estimated value of -1.188 for λ, we can discuss the estimation of idiosyncratic and aggregate shocks for agents as a function of their market participation statuses. The estimated volatility of aggregate shock for bondholders (only, i.e. $\tilde{h}_t^i = 1$, $h_t^i = 0$) is $|\mu\lambda\sigma_u/\gamma| = 61.733\%$, while the one for stockholders ($\tilde{h}_t^i = h_t^i = 1$) is equal to $\mu\sigma_u/\gamma = 51.964\%$. The estimation

pation version of the model.

[24]See her discussion on this point and other studies on page 827.

implies therefore that the aggregate shock volatility is slightly higher for non-stockholders than for stockholders. Finally, the estimated volatility of the shock for non-bondholders is $\sigma_e/\gamma = 46.656\%$.

The remaining parameters $\tilde{a}$ and a are at the core of the model, determining the probabilities of participation in the bond and stock markets. We examine in detail in Section 5.4 the stock market participation and the participation cost over time. To gauge the reasonableness of the parameters $\tilde{a}$ and a entering the probability functions, we plot in Figure 2 the probabilities of $\tilde{h}_t^i = 1$ and $\tilde{h}_t^i h_t^i = 1$ as functions of per capita household income (in $10,000$ USD). The graph shows that the probability of being a bondholder becomes one at a per capita income just below $10,000$ USD, while the threshold is close to $30,000$ USD for stockholders. These figures appear intuitively sensible, and we will see what they imply in terms of percentages of participation over time.

5.2 Consumption Inequality Dynamics

We now compare the abilities of our limited and unlimited participation models to replicate the cross-sectional heterogeneity in consumption data. Table 2 reports characteristics of aggregate and individual consumption under the two model specifications based on the estimated parameters in Table 1. The moments and quantiles are calculated by replicating 100 panels of data of the same size as the empirical dataset, generated from the estimated unlimited and limited participation models, and averaging the statistics obtained in the 100 replicated panels. In each simulated panel, individual log-consumption growth values $\{\widetilde{\Delta \log c_t^i}\}$ are generated from both models and individual consumption levels $\{\tilde{c}_t^i\}$ are calculated by using $\tilde{c}_t^i = \text{median}\{c_{t_0^i}^j\}_{j\geq 1} \exp(\sum_{t'=t_0^i+1}^{t} \widetilde{\Delta \log c_{t'}^i})$ where t_0^i is the first date at which agent i participates in the CEX survey and $\{c_{t_0^i}^j\}_{j\geq 1}$ are the empirical consumption levels.

The aggregate log-consumption growth moments are accurately estimated by the limited participation model, while the unlimited participation version completely misses the variance of log-consumption growth. Limited participation is also providing much more accurate replications of individual log-consumption growth quantiles and consumption quantiles. We conclude that limited participation is necessary for reliable estimation of the individual consumption distribution.

To assess the ability of the limited participation model to capture the inequality in consumption, we also plot in Figure 3 the Gini coefficients for consumption (top left panel), the 10th (P10), 50th (P50) and 90th (P90) percentiles of the consumption levels (top right panel), the P50/P10 ratios (bottom left panel) and P90/P50 ratios (bottom right panel) quarter per quarter during 1980-2004. Empirical series are plotted with continuous lines and simulated ones with dotted lines. Most of the time, the simulated series follow closely the movements in the observed ones. There are two periods associated with turmoil in financial markets where our model reacts much more than in the data. The market crash in 1987, as well as the Asian and Russian episodes in 1996-1998, make all our consumption measures feature a marked high-low pattern associated with the boom and bust movement in the stock market. This is especially visible in the measures associated with the highest consumption quantiles where the households participate more in the stock market. Our model does not embed any mechanism that accounts for crisis and it should not come as a surprise that it falls a bit short during episodes of financial market turmoil. However, except during these two crises, the model does a very good job in replicating the tremendous consumption heterogeneity observed in the data. It should also be noted that the inequality measures are lower at the beginning of the sample where the high interest rates favor relatively more the bondholders with a negative equity premium.

Overall, our results strongly support the limited participation model. In the following sections, we will explore further the implications of the estimated limited participation model for measuring the evolution of bond and stock ownership over time and of the cost of participating in the stock market.

5.3 Financial Market Participation

In our model, limited participation is not exogenously set but is an endogenous outcome of the estimation on consumption data and asset prices. The fraction of agents participating in financial markets is therefore not fixed over time but time-varying. Figure 4 plots the temporal evolution of participation in both bond and stock markets across our sample. The main lesson is that our model features an increasing participation in stock markets over our sample, between 1980 and 2004. The participation rate gains almost 10 percentage points and varies from approximately 32% to slightly more than 40%. Our model therefore predicts a sizable increase in the stock market participation rate. Interestingly, the participation rate values generated by the model are close to their empirical counterparts. Even though the CEX does not allow for measuring the participation rate, a partial answer is provided by the Survey of Consumer Finances (SCF), which is "a triennial cross-sectional survey of U.S. families" conducted by the Federal Reserve Board, starting in 1983. Indeed, as reported in Ameriks and Zeldes (2001) for instance, there is a strong evidence of an increase in stock market participation during the 1990s. More precisely, the stock market participation rate in the SCF increases from 30.6% in 1983 to 43.9% in 2004 (see Favilukis, 2013)[25]. The participation rate in our model is therefore quite consistent, both for absolute values and for the temporal evolution, with empirical data. The combination of a parsimonious limited participation asset pricing model with the indirect inference estimation enables us to reproduce a reasonable pattern for the stock market participation rate.

The other lesson of Figure 4 is that the participation rate in bond markets does not exhibit any particular trend and is approximately equal to 65%. Conversely, this implies that the fraction of hand-to-mouth consumers amounts to 35%. At first glance, this number

[25]These participation rate figures reflect direct and indirect stock holdings. Indirect stock holding takes into account households with a positive wealth in Individual Retirement Accounts (IRAs). But, IRAs, as well as mutual funds, may include financial assets other than stocks but the SCF does not provide information about the composition of portfolios. Even in the SCF, the participation rate is difficult to properly and exactly measure.

may seem to be a bit large. Indeed, in the SCF, as reported in Cagetti and De Nardi (2008), the share of households with a negative net worth lies in the range 7%–8%, and considering households with a low net worth, below 1000 USD, implies a share of 12 to 15%. These two values lie well below our proportion of 35% of hand-to-mouth consumers. However, looking at liquid wealth, which actually matters for buffer-stock savings and thus for consumption smoothing, the picture is quite different. Focusing on non-home wealth distribution in the SCF, which is a proxy for liquid wealth, Wolff (2010) reports that the bottom 40% has a negative non-home wealth for the period 1983–2007. These agents, with a small liquid wealth are very likely to be hand-to-mouth consumers. Kaplan et al. (2014) directly estimate the proportion of hand-to-mouth consumers in the US and find it to vary between 30% and 35% for the period 1989-2010. The two latter measures are thus overall consistent with our 35% fraction of hand-to-mouth consumers generated by our model.

Overall, our model generates a limited participation in financial markets that is strongly supported by empirical data. Importantly, limited participation is generated as an endogenous outcome of the model estimated on consumption data and asset prices. We thus see this participation figure as being a compelling validation check of our model.

5.4 Stock Market Participation Cost

We now analyze the implications on Propositions 1 and 2 of the estimation results in Table 1 for the limited participation model. We first show that our model is compatible with a stock market participation cost and then provide an estimation of this cost.

Our model is consistent with a stock market participation cost. For the bond market participation, Proposition 1 guarantees that agents excluded from financial markets – with $\tilde{h}_t^i = 0$ – are never willing to trade bonds and this result holds for all model parameter values.

The stock market participation in Proposition 2 is subject to two conditions (2.16) and (2.17), one for agents with $\tilde{h}_t^i = h_t^i = 0$ and the other one for agents with $\tilde{h}_t^i = 1, h_t^i = 0$. The first row of Table 3 reports the empirical tests associated with condition (2.16), while the second row corresponds to condition (2.17). Finite standard errors are reported in parentheses.

Only condition (2.16) holds and it holds significantly at the 5% level with our parameter estimates. Agents excluded from financial markets – with $\tilde{h}_t^i = h_t^i = 0$ – do not wish to trade stocks. Their behavior is consistent with market prices: stock market prices are too high – or equivalently, stock market returns are too low – for them to hold stocks.

The second condition (2.17) does not hold with our estimates and this result is significant at the 5% level. This result means that given a market price, agents with $\tilde{h}_t^i = 1, h_t^i = 0$ would like to hold stocks but do not do so. A possible explanation for this behavior is the presence of stock market participation costs that agents need to pay for participating in stock markets. This participation cost can cover a number of various costs, either monetary or not. For instance, on the monetary side, this cost may include trading fees or financial intermediation fees. On the non-monetary side, the cost may reflect acquiring and maintaining financial literacy and knowledge about stock markets, as well as monitoring financial news on a regular basis. Participation costs have already been used in a number of papers, such as for instance Allen and Gale (1990, 1994), Heaton and Lucas (1996), Orosel (1998), Campbell et al. (2001), Calvet et al. (2004), Polkovnichenko (2004), Gomes and Michaelides (2008), Chien et al. (2011), and Favilukis (2013). Participation costs can be of two types. First, agents can pay a once-in-a-lifetime participation cost, which would then imply that the main hurdle to participate in financial markets would be informational and the cost would correspond to acquiring financial literacy. Second, the cost can be recurring and corresponds to trading-related fees. Our estimation is compatible with the presence of one or both of these costs[26].

[26] An alternative explanation to the presence of stock market participation cost would be to state that

Computing the stock market participation cost. We propose a method that enables to evaluate the stock market participation cost in consumption equivalents. We focus here on a per period cost, consistently with our model presented in Section 2.3.

To do so, we consider an agent, participating only in bond markets. The stock market participation cost will be computed as the transfer, in consumption units, that would make this agent indifferent between participating in stock markets or not. More formally, we can prove (see equation (A.4) in Appendix A.3) that the individual participation cost is always positive if $|\lambda| > 1$ and is an increasing function of $|\lambda|$. The parameter λ can therefore be interpreted as a driver of the stock market participation. We then deduce that the average individual cost at date t, denoted $\overline{\tau}_t$, can be expressed as follows:

$$\overline{\tau}_t = \frac{1}{\#\{j : \tilde{h}_t^j = 1, h_t^j = 0\}} \sum_{i \in \{j : \tilde{h}_t^j = 1, h_t^j = 0\}} c_t^i \left(\frac{\left(1 + \beta^{\frac{1}{\gamma}} \left(R_{t+1}^f\right)^{\frac{1-\gamma}{\gamma}} e^{\frac{1-\gamma}{\gamma^2} \frac{\mu^2 \sigma_u^2}{2} \lambda^2}\right)^{\frac{1}{1-\gamma}}}{\left(1 + \beta^{\frac{1}{\gamma}} \left(R_{t+1}^f\right)^{\frac{1-\gamma}{\gamma}} e^{\frac{1-\gamma}{\gamma^2} \frac{\mu^2 \sigma_u^2}{2}}\right)^{\frac{1}{1-\gamma}}} - 1 \right),$$

where $\#\{j : \tilde{h}_t^j = 1, h_t^j = 0\}$ denotes the number of agents j with $\tilde{h}_t^j = 1$ and $h_t^j = 0$. We deduce then a per-period stock market participation cost $\tilde{\tau}_{c,t}$, expressed as a percentage of the average total consumption:

$$\tau_{c,t} = \frac{\overline{\tau}_t}{\frac{1}{\#\{j : c_t^j > 0\}} \sum_{i \in \{j : c_t^j > 0\}} c_t^i}, \tag{5.1}$$

where $\#\{j : c_t^j > 0\}$ represents the number of households participating in the CEX at date t.

Figure 5 plots the temporal evolution of the participation cost $\tau_{c,t}$ where $\overline{\tau}_t$ is estimated using 100 Monte Carlo replicates of participation panels of the same size as the empirical

stock holdings are forbidden for some agents. This is for instance the case in Basak and Cuoco (1998), or Guvenen (2009).

dataset. First, on the whole period, the cost approximately amounts to 2.3% of the average quarterly consumption, or approximately 330 USD (in 1980 value) per year. This empirical estimate is in line with the literature. Vissing-Jorgensen (2002b) estimates an annual participation cost varying between 150 and 350 USD. Our value is also comparable to the values chosen by Gomes and Michaelides (2008) and Favilukis (2013) in their respective calibrations.

The second feature of Figure 5 is that our estimation of the per-period cost exhibits a decreasing temporal pattern. Indeed, it diminishes from 2.5% in 1980 to 2.2% in 2004, which means a relative decrease (in real terms) greater than 10%. This decrease is consistent with the technological evolution over the period. First, financial innovation, through the development of mutual funds, new securities such as ETFs or of other intermediation vehicles with low transaction fees, can contribute to foster stock market participation. Using Swedish data, Calvet et al. (2016b) report a positive impact of financial innovation on stock market participation. Second, the development of Internet, lowering both transaction costs and informational frictions, may also contribute to boost stock market participation. Barber and Odean (2002) for instance report a decrease in financial transaction costs due to the development of online trading. Such an effect is also documented in Bogan (2008). In conclusion, our model endogenously features a decreasing stock market participation cost, similarly to what can be found in the data. This is an additional external validation check for our model.

6 Robustness Check

We now conduct a robustness check to verify that our results do not rely on our choice of consumption expenditure data. Instead of using non-transformed consumption expenditures, directly measured from the CEX dataset, we estimate our model using the flow consumption expenditures, described in Section 3 and that mainly differ from raw consumption data through vehicle expenditures. More precisely, we perform exactly the same exercise – with the same model and the same estimation technique – as in Section 4, except that we use flow consumption expenditures. Similarly to Table 1, Table 4 gathers the results of the estimation on the flow consumption data (finite standard errors, computed using 100 Monte Carlo replicates are in parentheses). Note that due to the change in consumption data, the dataset does not contain exactly the same households as in the initial dataset, and the number of data points is much smaller. This number is down from $117,857$ to $99,204$, implying a loss of $18,653$ households.

The comparison with Table 1 is unambiguous: there is no significant difference in parameter estimates between raw and flow datasets.

In Table 5, we report the characteristics of aggregate and individual consumption for the model estimated on flow consumption data. We can compare these results to those of the main estimation, that can be found in Table 2. First, the equity premium statistics are well matched. Second, aggregate log-consumption growth moments are also accurately replicated, even though the standard deviation of the model is somehow lower than the empirical one. Finally, individual log-consumption growth and consumption level quartiles are relatively well replicated by the model. However, the top and bottom quantiles deviate from their empirical counterparts. Overall, the model generates slightly less heterogeneity than the data.

This robustness check validates our estimation method. Indeed, even though the two datasets differ in size and consumption accounting, the estimation provides similar results.

7 Conclusion

This paper proposes an asset pricing model based on limited participation and heterogeneity which is able to accurately replicate asset pricing properties and the distribution of individual consumption. Euler conditions for bonds and stocks are the basis of our model and are not only used as moment conditions for parameter estimation. We develop an indirect inference method to estimate the parameters of the model. We compare limited and unlimited participation models and show that the latter one fails to replicate both aggregate and individual consumption behaviors. Finally, our model permits us to test whether Euler conditions are satisfied empirically. Application to CEX data shows evidence of endogenous bond market participation and of non-zero stock market participation cost.

A Appendix

A.1 Proof of Proposition 1

$$
\begin{aligned}
&\mathbb{E}[\exp\{-\gamma\Delta\log c^i_{t+1}\}] = \\
&\quad = \frac{1}{\beta R^f_{t+1}}\mathbb{E}\left[\exp\{-\frac{\mu^2\sigma_u^2}{2}\lambda^{2-2h^i_t\tilde{h}^i_t} - (1-\tilde{h}^i_t)\frac{\sigma_e^2}{2} - (1-\tilde{h}^i_t)\sigma_e\varepsilon^i_{t+1} - \mu\tilde{h}^i_t\sigma_u\lambda^{1-h^i_t\tilde{h}^i_t}u_{t+1}\}\right] \\
&\quad = \frac{1}{\beta R^f_{t+1}}\mathbb{E}\left[\exp\{-\frac{\mu^2}{2}\sigma_u^2\lambda^{2-2h^i_t\tilde{h}^i_t} + \frac{\mu^2}{2}\sigma_u^2\tilde{h}^i_t\lambda^{2-2h^i_t\tilde{h}^i_t}\}\right] \\
&\quad \begin{cases} = \frac{1}{\beta R^f_{t+1}} & \text{if } \tilde{h}^i_t = 1 \\ = \frac{1}{\beta R^f_{t+1}}\mathbb{E}\left[\exp\{-\frac{\mu^2\sigma_u^2\lambda^2}{2}\}\right] & \text{if } \tilde{h}^i_t = 0 \end{cases} \\
&\quad \begin{cases} = \frac{1}{\beta R^f_{t+1}} & \text{if } \tilde{h}^i_t = 1 \\ < \frac{1}{\beta R^f_{t+1}} & \text{if } \tilde{h}^i_t = 0 \end{cases}
\end{aligned}
$$

A.2 Proof of Proposition 2

$$
\begin{aligned}
&\mathbb{E}[\exp\{\log R^s_{t+1} - \gamma\Delta\log c^i_{t+1}\}] = \\
&\quad = \frac{1}{\beta}\mathbb{E}\left[\exp\{(\mu - \frac{1}{2})\sigma_u^2 + \sigma_u u_{t+1} - \frac{\mu^2\sigma_u^2}{2}\lambda^{2-2h^i_t\tilde{h}^i_t} - (1-\tilde{h}^i_t)\frac{\sigma_e^2}{2} - (1-\tilde{h}^i_t)\sigma_e\varepsilon^i_{t+1}\right. \\
&\qquad \left. -\mu\tilde{h}^i_t\sigma_u\sqrt{\lambda}^{1-h^i_t\tilde{h}^i_t}u_{t+1}\}\right] \\
&\quad = \frac{1}{\beta}\mathbb{E}\left[\exp\{(\mu - \frac{1}{2})\sigma_u^2 - \frac{\mu^2\sigma_u^2}{2}\lambda^{2-2h^i_t\tilde{h}^i_t} + \sigma_u(1 - \mu\tilde{h}^i_t\lambda^{1-h^i_t\tilde{h}^i_t})u_{t+1}\}\right] \\
&\quad = \frac{1}{\beta}\mathbb{E}\left[\exp\{\mu\sigma_u^2 - \frac{\mu^2\sigma_u^2}{2}\lambda^{2-2h^i_t\tilde{h}^i_t} + \frac{\sigma_u^2\tilde{h}^i_t}{2}(-2\mu\lambda^{1-h^i_t\tilde{h}^i_t} + \mu^2\lambda^{2-2h^i_t\tilde{h}^i_t})\}\right].
\end{aligned}
$$

If $h_t^i = 1$:

$$\mathbb{E}[\exp\{\log R^s_{t+1} - \gamma \Delta \log c^i_{t+1}\}] =$$
$$= \frac{1}{\beta}\mathbb{E}\left[\exp\{\mu\sigma_u^2 - \frac{\mu^2\sigma_u^2}{2}\lambda^{2-2\tilde{h}_t^i} + \frac{\sigma_u^2\tilde{h}_t^i}{2}(-2\mu\lambda^{1-\tilde{h}_t^i} + \mu^2\lambda^{2-2\tilde{h}_t^i})\}\right]$$
$$\begin{cases} = \frac{1}{\beta} & \text{if } \tilde{h}_t^i = 1 \\ = \frac{1}{\beta}\mathbb{E}\left[\exp\{\sigma_u^2\mu(1-\frac{\mu\lambda^2}{2})\}\right] & \text{if } \tilde{h}_t^i = 0 \end{cases}$$
$$\begin{cases} = \frac{1}{\beta} & \text{if } \tilde{h}_t^i = 1 \\ < \frac{1}{\beta} & \text{if } \tilde{h}_t^i = 0 \text{ and } \mu(1-\frac{\mu\lambda^2}{2}) < 0 \end{cases}.$$

If $h_t^i = 0$:

$$\mathbb{E}[\exp\{\log R^s_{t+1} - \gamma \Delta \log c^i_{t+1}\}] =$$
$$= \frac{1}{\beta}\mathbb{E}\left[\exp\{\mu\sigma_u^2 - \frac{\mu^2\sigma_u^2}{2}\lambda^{2} + \frac{\sigma_u^2\tilde{h}_t^i}{2}(-2\mu\lambda + \mu^2\lambda^{2})\}\right]$$
$$\begin{cases} = \frac{1}{\beta}\mathbb{E}\left[\exp\{\mu\sigma_u^2(1-\lambda)\}\right] & \text{if } \tilde{h}_t^i = 1 \\ = \frac{1}{\beta}\mathbb{E}\left[\exp\{\mu\sigma_u^2(1-\frac{\mu\lambda^2}{2})\}\right] & \text{if } \tilde{h}_t^i = 0 \end{cases}$$
$$\begin{cases} < \frac{1}{\beta} & \text{if } \tilde{h}_t^i = 1 \text{ and } \mu(1-\lambda) < 0 \\ < \frac{1}{\beta} & \text{if } \tilde{h}_t^i = 0 \text{ and } \mu(1-\frac{\mu\lambda^2}{2}) < 0 \end{cases}.$$

A.3 Deriving the Participation Cost Expression (A.4)

Consider at date t an agent i with $\tilde{h}_t^i = 1$ and $h_t^i = 0$. Her intertemporal utility V_t^i is

$$V_t^i = \frac{\left(c_t^i\right)^{1-\gamma}}{1-\gamma} + \beta \mathbb{E}_t \left[\frac{\left(c_{t+1}^i\right)^{1-\gamma}}{1-\gamma} \right] + \beta^2 \mathbb{E}_t \left[V_{t+2}^i \right],$$

$$= \frac{\left(c_t^i\right)^{1-\gamma}}{1-\gamma} \left(1 + \beta \mathbb{E}_t e^{(1-\gamma)\Delta \log c_{t+1}^i} \right) + \beta^2 \mathbb{E}_t \left[V_{t+2}^i \right]. \tag{A.1}$$

We consider the following thought experiment. The agent i is constrained to participate in the stock market at date t. We denote by the superscript c the corresponding variables. The constrained agent therefore faces $h_t^{i,c} = 1$ (and $\tilde{h}_t^{i,c} = 1$). To compensate her for this constrained stock market participation, she receives a flat amount τ_t^i at date t. We assume that this amount is fully consumed at date t. The intertemporal utility $V_t^{i,c}$ becomes then

$$V_t^{i,c} = \frac{\left(c_t^i + \tau_t^i\right)^{1-\gamma}}{1-\gamma} \left(1 + \beta \mathbb{E}_t e^{(1-\gamma)\Delta \log c_{t+1}^{i,c}} \right) + \beta^2 \mathbb{E}_t \left[V_{t+2}^{i,c} \right]. \tag{A.2}$$

The compensation τ_t^i exactly offsets the forced participation if $V_t^{i,c} = V_t^i$. If we assume that the constrained participation has no effect after date $t+2$, the equality $V_t^{i,c} = V_t^i$ can be simplified using equations $(A.1)$ and $(A.2)$ as follows

$$\tau_t^i = c_t^i \left(\frac{\left(1 + \beta \mathbb{E}_t e^{(1-\gamma)\Delta \log c_{t+1}^i} \right)^{\frac{1}{1-\gamma}}}{\left(1 + \beta \mathbb{E}_t e^{(1-\gamma)\Delta \log c_{t+1}^{i,c}} \right)^{\frac{1}{1-\gamma}}} - 1 \right). \tag{A.3}$$

Using the dynamics of the log-consumption growth, equality (A.3) defining τ_t^i becomes

$$\tau_t^i = c_t^i \left(\frac{\left(1 + \beta^{\frac{1}{\gamma}} \left(R_{t+1}^f \right)^{\frac{1-\gamma}{\gamma}} e^{\frac{1-\gamma}{\gamma^2} \frac{\mu^2 \sigma_u^2}{2} \lambda^2} \right)^{\frac{1}{1-\gamma}}}{\left(1 + \beta^{\frac{1}{\gamma}} \left(R_{t+1}^f \right)^{\frac{1-\gamma}{\gamma}} e^{\frac{1-\gamma}{\gamma^2} \frac{\mu^2 \sigma_u^2}{2}} \right)^{\frac{1}{1-\gamma}}} - 1 \right). \tag{A.4}$$

References

Allen, F., and D. Gale (1990). Incomplete Markets, and Incentives to Set Up an Options Exchange. *Geneva Papers on Risk, and Insurance Theory* **15**, 17–46.

Allen, F., and D. Gale (1994). Limited Market Participation, and Volatility of Asset Prices. *American Economic Review* **84**, 933–955.

Ameriks, J., and S. Zeldes (2001). How do Household Portfolio Shares Vary with Age? *Working Paper, Columbia University.*

Andersen, S., G. Harrison, I. Morten, and E. Rutström (2014). Discounting Behavior: A Reconsideration. *European Economic Review* **71**, 15–33.

Andreoni, J., and C. Sprenger (2012). Estimating Time Preferences from Convex Budgets. *American Economic Review* **102**, 3333–3356.

Attanasio, O., and G. Weber (1995). Is Consumption Growth Consistent with Intertemporal Optimization? Evidence from the Consumer Expenditure Survey. *Journal of Political Economy* **103**, 1121–1157.

Bansal, R., and A. Yaron (2004). Risks for the Long Run: A Potential Resolution of Asset Pricing Puzzles. *Journal of Finance* **59**, 1481–1509.

Barber, B., and T. Odean (2002). Online Investors: Do the Slow Die First? *Review of Financial Studies* **15**, 455–487.

Barsky, B., F. Juster, M. Kimball, and M. Shapiro (1997). Preference Parameters, and Behavioral Heterogeneity: An Experimental Approach in the Health, and Retirement Study. *Quarterly Journal of Economics* **112**, 537–579.

Basak, S., and D. Cuoco (1998). An Equilibrium Model with Restricted Stock Market Participation. *Review of Financial Studies* **11**, 309–341.

Berger, D., and J. Vavra (2015). Consumption Dynamics During Recessions. *Econometrica* **83**, 101–154.

Blundell R., M. Browning, and C. Meghir (1994). Consumer Demand, and the Life-Cycle Allocation of Household Expenditures. *Review of Economic Studies* **61**, 57–80.

Bogan, V. (2008). Stock Market Participation, and the Internet. *Journal of Financial, and Quantitative Analysis* **43**, 191–211.

Bonomo, M., R. Garcia, N. Meddahi, and R. Tedongap (2011). Generalized Disappointment Aversion, Long-Run Volatility Risk, and Asset Prices. *Review of Financial Studies* **24**, 82–122.

Browning, M., and A. Lusardi (1996). Household Saving: Micro Theories, and Micro Facts. *Journal of Economics Literature* **34**, 1797–1855.

Brav, A., G. Constantinides, and C. Geczy (2002). Asset Pricing with Heterogeneous Consumers, and Limited Participation: Empirical Evidence. *Journal of Political Economy* **110**, 793–824.

Cagetti M., and C. De Nardi (2008). Wealth Inequality: Data, and Models. *Macroeconomic Dynamics* **12**, special supplement S2 on Inequality edited by R. Townsend, 285–313.

Calvet L., and V. Czellar (2015). Through the Looking Glass: Indirect Inference Via Simple Equilibria. *Journal of Econometrics* **185**, 343–358.

Calvet L., C. C élérier, P. Sodini, and B. Vallée (2016). Financial Innovation, and Stock Market Participation. *IFSID working paper* **WP–16-01**.

Calvet L., M. Gonzalez-Eiras, and P. Sodini (2004). Financial Innovation, Market Participation, and Asset Prices. *Journal of Financial, and Quantitative Analysis* **39**, 431–459.

Calzolari G., G. Fiorentini, and E. Sentana (2004). Constrained Indirect Estimation. *Review of Economic Studies* **71**, 945–973.

Campbell, J., J. Cocco, F. Gomes, and P. Maenhout (2001). Investing Retirement Wealth: A Life-Cycle Model. In *Risk Aspects of Investment-Based Social Security Reform*, J. Campbell, and M. Feldstein, eds. University of Chicago Press (2001), Chapter 11. (Also reproduced in Campbell, and Viceira, 2002).

Campbell, J., and J. Cochrane (1999). By Force of Habit: A Consumption-Based Explanation of Aggregate Stock Market Behavior. *Journal of Political Economy* **107**, 205–251.

Carroll, C. (1997). Buffer-Stock Saving, and the Life Cycle/Permanent Income Hypothesis. *Quarterly Journal of Economics* **112**, 1–55.

Carroll, C., and A. Samwick (1997). The Nature of Precautionary Wealth. *Journal of Monetary Economics* **40**, 41–71.

Challe, E., F. LeGrand, and X. Ragot (2013). Incomplete Markets, Liquidation Risk, and the Term Structure of Interest Rates. *Journal of Economic Theory* **148**, 2483–2519.

Chien, Y., H. Cole, and H. Lustig (2011). A Multiplier Approach to Understanding the Macro Implications of Household Finance. *Review of Economic Studies* **78**, 199–234.

Constantinides, G., and D. Duffie (1996). Asset Pricing with Heterogeneous Consumers. *Journal of Political Economy* **104**, 219–240.

Cutler, D., and J. Katz (1991). Macroeconomic Performance and the Disadvantaged. *Brookings Papers on Economic Activity* **2**, 1–74.

Czellar, V., G. Karolyi, and E. Ronchetti (2007). Indirect Robust Estimation of the Short-Term Interest Process. *Journal of Empirical Finance* **14**, 546–563.

Epstein, L., and S. Zin (1989). Substitution, Risk Aversion, and the Temporal Behavior of Consumption, and Asset Returns: A Theoretical Framework. *Econometrica* **57**, 937–969.

Favilukis, J. (2013). Wealth Inequality, Stock Market Participation, and the Equity Premium. *Journal of Financial Economics* **107**, 740–759.

Gallant, A., and G. Tauchen (1996). Which Moments to Match? *Econometric Theory* **12**, 657–681.

Garcia, R., E. Renault, and D. Veredas (2011). Estimation of Stable Distributions by Indirect Inference. *Journal of Econometrics* **161**, 325–337.

Gomes, F., and A. Michaelides (2008). Asset Pricing with Limited Risk Sharing, and Heterogeneous Agents. *Review of Financial Studies* **21**, 415–448.

Gouriéroux, C., and A. Monfort (1996). *Simulation-Based Econometric Methods.* Oxford University Press.

Gouriéroux, C., A. Monfort, and E. Renault (1993). Indirect Inference. *Journal of Applied Econometrics* **8**, S85–S118.

Gourinchas, PO., and J. Parker (2002). Consumption over the Life Cycle. *Econometrica* **70**, 47–89.

Guvenen, F. (2006). Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective. *Journal of Monetary Economics* **53**, 1451–1472.

Guvenen, F. (2009). A Parsimonious Macroeconomic Model for Asset Pricing. *Econometrica* **77**, 1711–1750.

Guvenen, F., and A. Smith (2014). Inferring Labor Income Risk, and Partial Insurance from Economic Choices. *Econometrica* **82**, 2085–2129.

Heathcote, J., F. Perri, and G. Violante (2010). Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 1967-2006. *Review of Economic Dynamics* **13**, 519–548.

Heaton, J., and D. Lucas (1996). Evaluating the Effects of Incomplete Markets on Risk Sharing, and Asset Pricing. *Journal of Political Economy* **104**, 443–487.

Huggett, M. (1993). The Risk Free Rate in Heterogeneous-Agent Incomplete-Insurance Economies. *Journal of Economic Dynamics, and Control* **17**, 953–969.

Kaplan, G., G. Violante, and J. Weidner (2014). The Wealthy Hand-to-Mouth. *Brookings Papers on Economic Activity* **Spring**, 77–153.

Krueger, D., and F. Perri (2005). Does Income Inequality Lead to Consumption Inequality? Evidence and Theory. *Review of Economic Studies* **73**, 163–193.

Krussel, P., T. Mukoyama, and A. Smith (2011). Asset Prices in a Huggett Economy. *Journal of Economic Theory* **146**, 812–844.

Lawrance, E. (1991). Poverty, and the Rate of Time Preference: Evidence from Panel Data. *Journal of Political Economy* **99**, 54–77.

LeGrand F., and X. Ragot (2015). Asset Returns, Idiosyncratic, and Aggregate Risk Exposure. *Working Paper SciencesPo.*

Loewenstein G., and D. Prelec (1991). Negative Time Preference. *American Economic Review, Papers, and Proceedings* **81**, 347–352.

Mankiw, N. (1986). The Equity Premium, and the Concentration of Aggregate Shocks. *Journal of Financial Economics* **17**, 211–219.

Mehra, R., and E. Prescott (1985). The Equity Premium: A Puzzle. *Journal of Monetary Economics* **15**, 145–61.

Merton, R. (1971). Optimum Consumption, and Portfolio Rules in a Continuous-Time Model. *Journal of Economic Theory* **3**, 373–413.

Orosel, G. (1998). Participation Costs, Trend Chasing, and Volatility of Stock Prices. *Review of Financial Studies* **11**, 521–557.

Parker, J., and A. Vissing-Jorgensen (2009). Who Bears Aggregate Fluctuations, and How? *American Economic Review* **99**, 399–405.

Polkovnichenko, V. (2004). Limited Stock Market Participation, and the Equity Premium. *Finance Research Letters* **1**, 24–34.

Primiceri, G., and T. van Rens (2009). Heterogeneous Life-Cycle Profiles, Income Risk, and Consumption Inequality. *Journal of Monetary Economics* **56**, 20–39.

Samwick, A. (1998). Discount Rate Heterogeneity, and Social Security Reform. *Journal of Development Economics* **57**, 117–146.

Sentana, E., G. Calzolari, and G. Fiorentini (2008). Indirect Estimation of Large conditionally Heteroskedastic Factor Models, with an Application to the Dow 30 Stocks. *Journal of Econometrics* **146**, 10–25.

Smith, A. (1993). Estimating Nonlinear Time Series Models Using Simulated Vector Autoregressions. *Journal of Applied Econometrics* **8**, S63–S84.

Storesletten K., C. Telmer, and A. Yaron 2007. Asset Pricing with Idiosyncratic Risk, and Overlapping Generations. *Review of Economic Dynamics* **10**, 519–48.

Vissing-Jorgensen, A. (2002). Limited Asset Market Participation, and the Elasticity of Intertemporal Substitution. *Journal of Political Economy* **110**, 825–853.

Vissing-Jorgensen, A. (2002). Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income, and Participation Cost Structures. *NBER Working Paper* **No 8884**.

Wolff, E. (2010). Recent Trends in Household Wealth in the United States: Rising Debt, and the Middle-Class Squeeze – an Update to 2007. *Levy Economics Institute Working Papers Series* **No 159**.

Warner, J., and S. Pleeter (2001). The Personal Discount Rate: Evidence from Military Downsizing Programs. *B.E. Journal of Macroeconomics (Contributions)* **10**, 33–53.

Parameter	Symbol	Limited participation	Unlimited participation
Time discount factor	β	0.735 (0.083)	0.226 (0.200)
Inverse of EIS	γ	1.196 (0.354)	187.960 (138.646)
Multiplier of agg. volatility	λ	-1.188 (0.301)	–
Individual volatility	σ_e	0.558 (0.279)	–
Constant in mean equity premium	μ	10.024 (2.310)	10.658 (7.601)
Aggregate Volatility	σ_u	0.062 (0.005)	0.062 (0.028)
Income multiplier in $P(\tilde{h}_t = 1)$	$\tilde{a}$	1.019 (0.422)	–
Income multiplier in $P(h_t = 1)$	a	0.330 (0.043)	–
Chi-square statistic	ξ		1.085 (0.581)

Table 1: Empirical Estimates. Parameter estimates of the limited participation model are reported in the first column and those of the unlimited participation model are in the second column. Finite standard errors are reported in parentheses. The ξ statistic when testing the unlimited participation model is reported in the last row of the table. The associated p-value is in parentheses below the ξ statistic value.

	Statistics	Data	Limited participation	Unlimited participation
Equity premium	$\mathbb{E}[\log(R_t^s/R_t^f)]$	3.7961%	3.6729%	3.9238%
	$\sigma[\log(R_t^s/R_t^f)]$	6.1104%	6.1551%	6.1602%
Aggregate log consumption growth	$\mathbb{E}(\Delta\log c^{tot})$	−0.7019%	−0.7329%	−0.6359%
	$\sigma(\Delta\log c^{tot})$	3.1439%	2.8369%	0.3013%
Individual log-consumption growth quantiles	$q_{0.01}^{\Delta\log c}$	−1.5152	−1.3049	−0.0151
	$q_{0.05}^{\Delta\log c}$	−0.8454	−0.8983	−0.0125
	$q_{0.10}^{\Delta\log c}$	−0.5519	−0.6973	−0.0111
	$q_{0.25}^{\Delta\log c}$	−0.2454	−0.3701	−0.0089
	$q_{0.50}^{\Delta\log c}$	−0.0080	−0.0046	−0.0064
	$q_{0.75}^{\Delta\log c}$	0.2283	0.3519	−0.0040
	$q_{0.90}^{\Delta\log c}$	0.5349	0.6725	−0.0015
	$q_{0.95}^{\Delta\log c}$	0.8303	0.8770	0.0001
	$q_{0.99}^{\Delta\log c}$	1.5285	1.2575	0.0037
Individual consumption quantiles	$q_{0.01}^{c}$	628.1	460.6	3368.4
	$q_{0.05}^{c}$	1089.0	921.9	3410.2
	$q_{0.10}^{c}$	1447.0	1300.2	3438.4
	$q_{0.25}^{c}$	2274.0	2165.3	3504.6
	$q_{0.50}^{c}$	3619.0	3579.5	3581.7
	$q_{0.75}^{c}$	5662.0	5896.1	3648.4
	$q_{0.90}^{c}$	8744.0	9579.1	3710.6
	$q_{0.95}^{c}$	11957.0	13197.8	3745.3
	$q_{0.99}^{c}$	21094.2	25672.2	3872.8

Table 2: Equity premium and consumption features under estimated models. The moments and quantiles are calculated by replicating 100 panels of data of the same size as the empirical dataset, generated from the estimated limited and unlimited participation models, and averaging the statistics obtained in the 100 replicated panels. In each simulated panel, individual log-consumption growth values $\{\widetilde{\Delta\log c_t^i}\}$ are generated from both models and individual consumption levels $\{\tilde{c}_t^i\}$ are calculated by using $\tilde{c}_t^i = \text{median}\{c_{t_0^i}^j\}_{j\geq 1}\exp(\sum_{t'=t_0^i+1}^{t}\widetilde{\Delta\log c_{t'}^i})$ where t_0^i is the first date at which agent i participates in the CEX survey and $\{c_{t_0^i}^j\}_{j\geq 1}$ are the empirical consumption levels.

Market participation	Euler condition	Empirical estimate	Conclusion
Nonparticipants in financial markets ($\tilde{h}_t^i = h_t^i = 0$)	$\mu(1 - \frac{\mu\lambda^2}{2}) < 0$	-60.865 (36.232)	Endogenous nonparticipation in bond and stock markets
Bondparticipants only ($\tilde{h}_t^i = 1, h_t^i = 0$)	$\mu(1 - \lambda) < 0$	21.931 (4.617)	Stock-market participation cost

Table 3: Testing the Euler conditions. The first row reports the empirical tests associated with condition (2.16), while the second row corresponds to condition (2.17). Finite standard errors are reported in parentheses.

Parameter	Symbol	Limited participation
Time discount factor	β	0.772 (0.048)
Inverse of EIS	γ	2.693 (0.566)
Multiplier of agg. volatility	λ	-1.034 (0.087)
Individual volatility	σ_e	0.494 (0.108)
Constant in mean equity premium	μ	10.054 (2.397)
Aggregate Volatility	σ_u	0.062 (0.005)
Income multiplier in $P(\tilde{h}_t = 1)$	$\tilde{a}$	0.748 (0.215)
Income multiplier in $P(h_t = 1)$	a	0.276 (0.017)

Table 4: Empirical estimates using flow consumption data. The table reports indirect inference estimates of the limited participation model. Finite standard errors are reported in parentheses.

	Statistics	Data	Limited participation
Equity premium	$\mathbb{E}[\log(R^s_t/R^f_t)]$	3.7961%	3.7165%
	$\sigma[\log(R^s_t/R^f_t)]$	6.1104%	6.1820%
Aggregate log-consumption growth	$\mathbb{E}(\Delta \log c^{tot})$	−0.3593%	−0.3050%
	$\sigma(\Delta \log c^{tot})$	3.1045%	1.2052%
Individual log-consumption growth quantiles	$q_{0.01}^{\Delta \log c}$	−0.8188	−0.5379
	$q_{0.05}^{\Delta \log c}$	−0.4926	−0.3708
	$q_{0.10}^{\Delta \log c}$	−0.3596	−0.2867
	$q_{0.25}^{\Delta \log c}$	−0.1764	−0.1510
	$q_{0.50}^{\Delta \log c}$	−0.0036	−0.0002
	$q_{0.75}^{\Delta \log c}$	0.1667	0.1455
	$q_{0.90}^{\Delta \log c}$	0.3509	0.2763
	$q_{0.95}^{\Delta \log c}$	0.4829	0.3562
	$q_{0.99}^{\Delta \log c}$	0.8138	0.5123
Individual consumption quantiles	$q_{0.01}^{c}$	918.0	1844.0
	$q_{0.05}^{c}$	1576.0	2453.8
	$q_{0.10}^{c}$	2034.0	2827.4
	$q_{0.25}^{c}$	2956.0	3483.4
	$q_{0.50}^{c}$	4292.0	4281.4
	$q_{0.75}^{c}$	6072.0	5249.4
	$q_{0.90}^{c}$	8305.0	6394.5
	$q_{0.95}^{c}$	10131.1	7285.0
	$q_{0.99}^{c}$	15585.0	9507.0

Table 5: Equity premium and consumption features for the flow consumption data. The moments and quantiles are calculated by replicating 100 panels of data of the same size as the empirical dataset, generated from the estimated limited participation model, and averaging the statistics obtained in the 100 replicated panels. In each simulated panel, individual log-consumption growth values $\{\widetilde{\Delta \log \tilde{c}^i_t}\}$ are generated from both models and individual consumption levels $\{\tilde{c}^i_t\}$ are calculated by using $\tilde{c}^i_t = \text{median}\{c^j_{t^i_0}\}_{j\geq 1} \exp(\sum_{t'=t^i_0+1}^{t} \widetilde{\Delta \log c^i_{t'}})$ where t^i_0 is the first date at which agent i participates in the CEX survey and $\{c^j_{t^i_0}\}_{j\geq 1}$ are the empirical consumption levels.

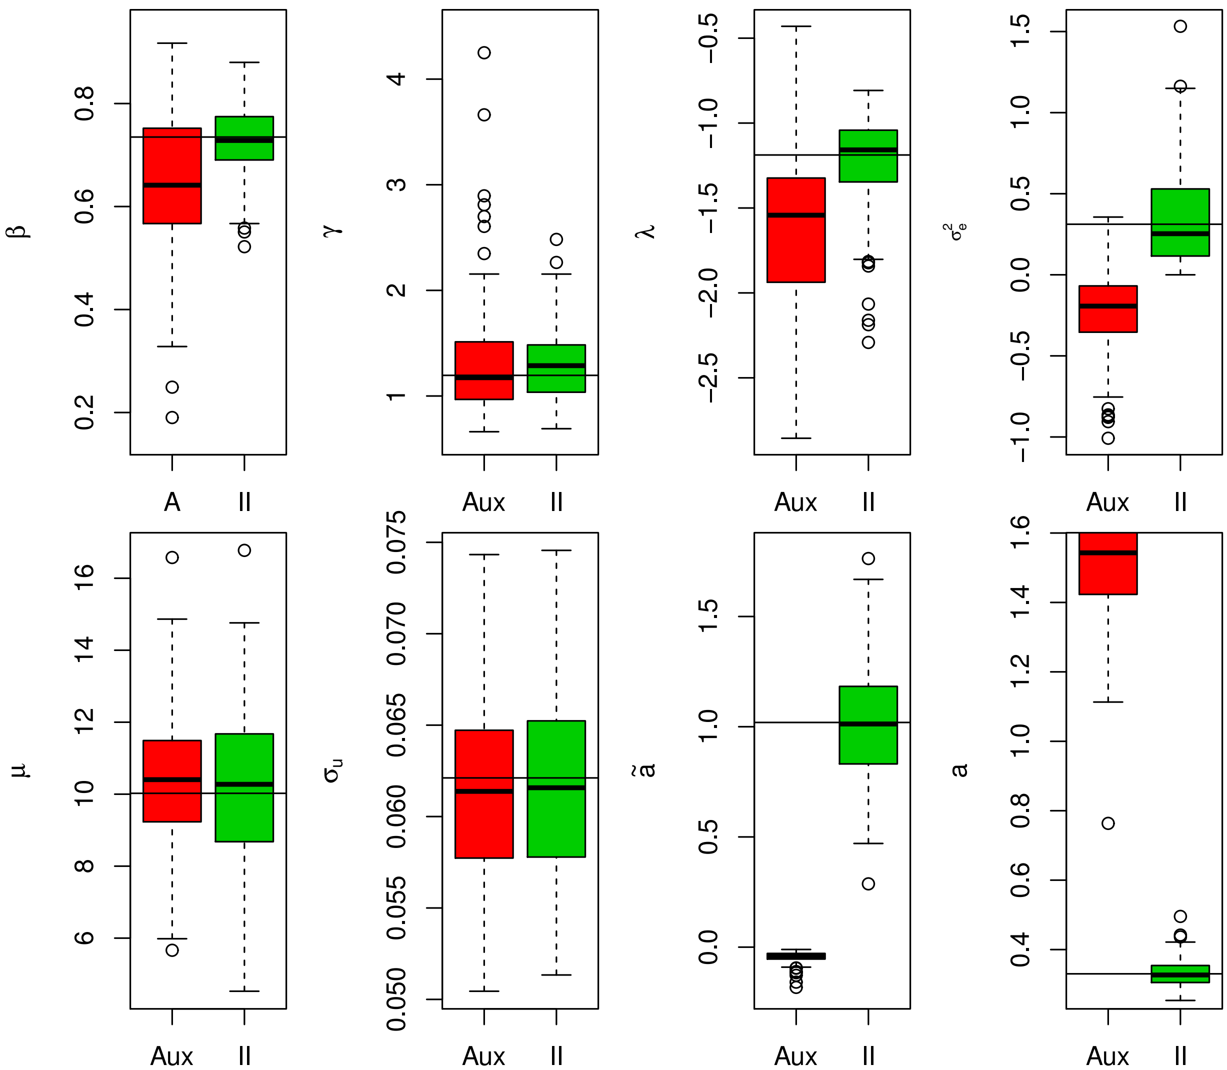

Figure 1: Boxplots of auxiliary and indirect inference parameter estimates. 100 auxiliary estimates are reported on the left and 100 indirect inference estimates are on the right of each panel. True parameter values are reported by horizontal lines in each panel and correspond to $\beta = 0.735$, $\gamma = 1.196$, $\lambda = -1.188$, $\sigma_e = 0.558$, $\mu = 10.024$, $\sigma_u = 0.062$, $\tilde{a} = 1.019$, $a = 0.330$.

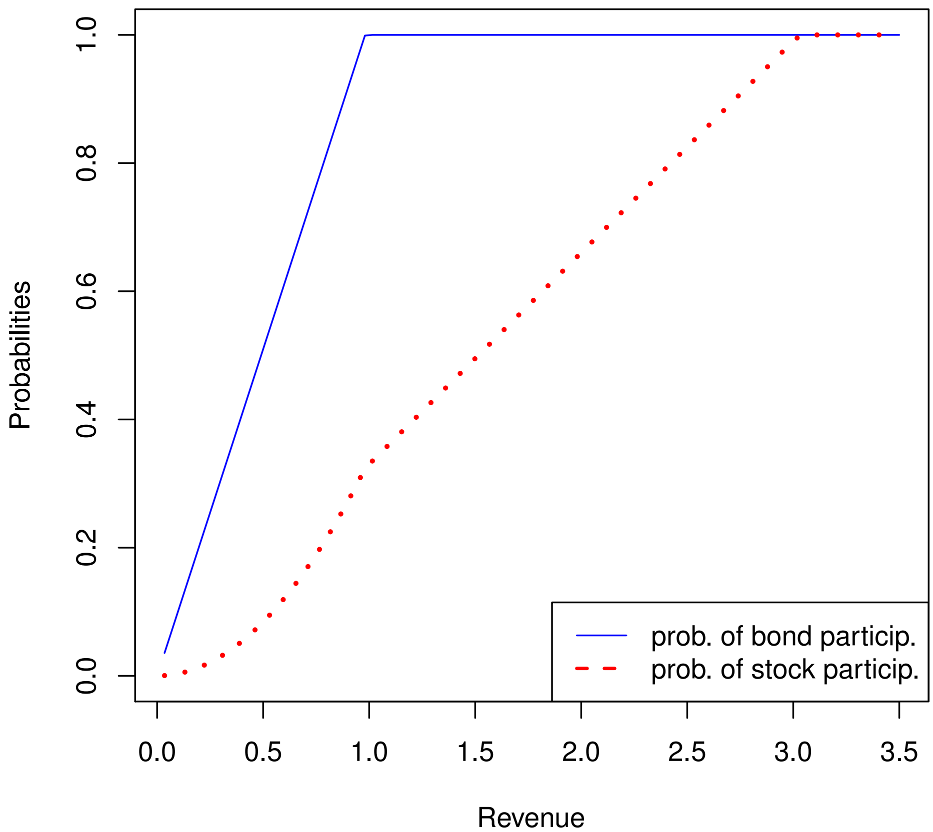

Figure 2: Probabilities of $\tilde{h}_t^i = 1$ (bond) and $\tilde{h}_t^i h_t^i = 1$ (stock) as functions of the household per capita income (in $10,000$ USD).

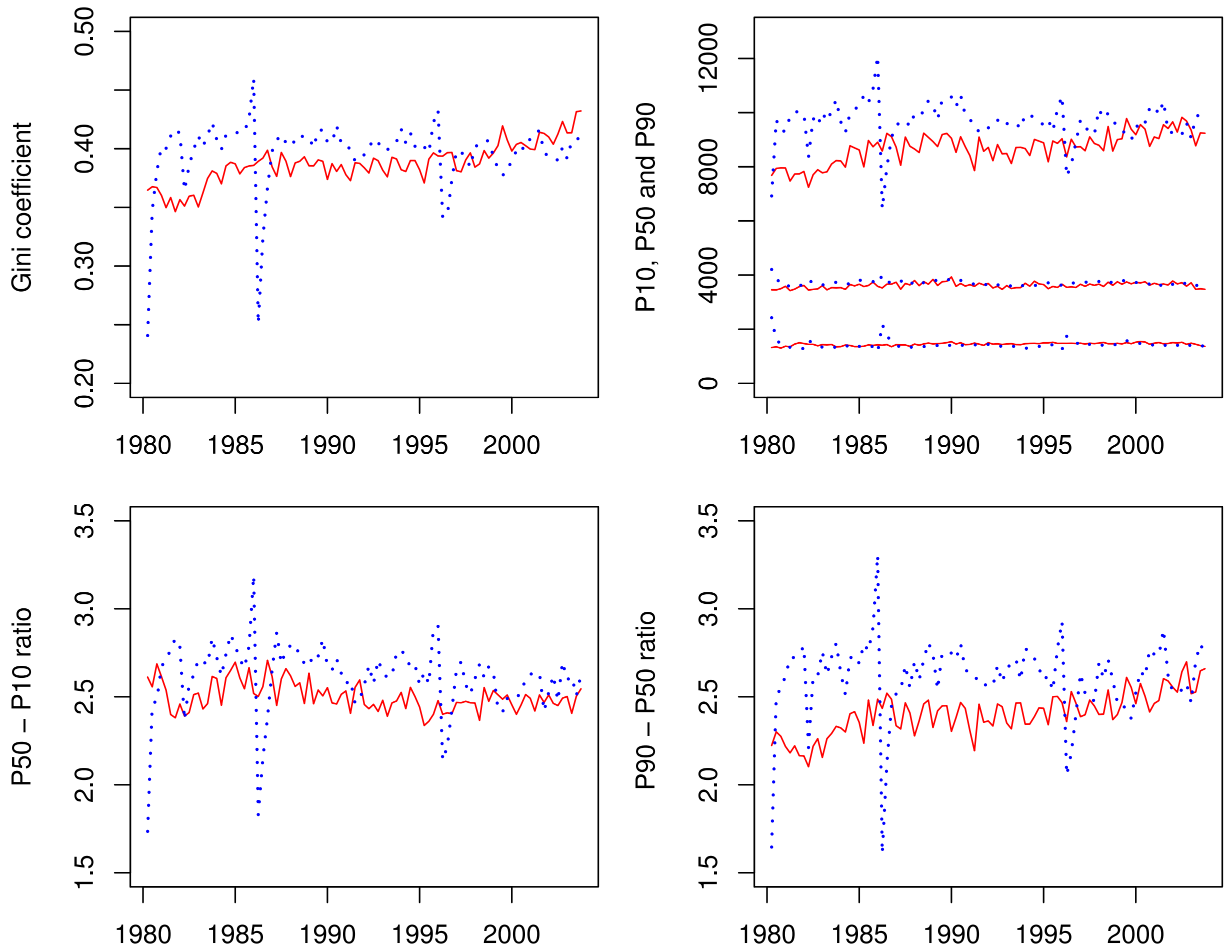

Figure 3: Observed and simulated consumption quantiles and Gini coefficients. Gini coefficients (top left panel), the 10th (P10), 50th (P50) and 90th (P90) percentiles of the consumption levels (top right panel), the P50/P10 ratios (bottom left panel) and P90/P50 ratios (bottom right panel) quarter per quarter during 1980-2004. Empirical series are plotted with continuous lines and simulated ones with dotted lines.

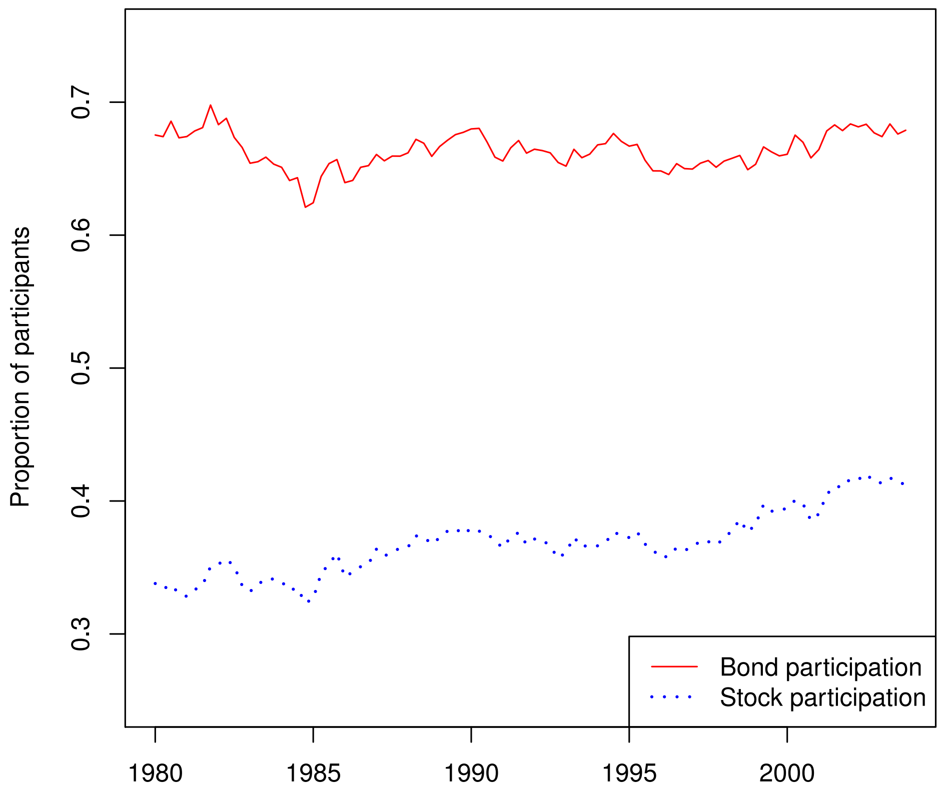

Figure 4: Temporal evolution of the estimated participation in financial markets. The estimated proportion of households participating in the bond markets is reported with a continuous line and the one for stock markets is reported with a dotted line.

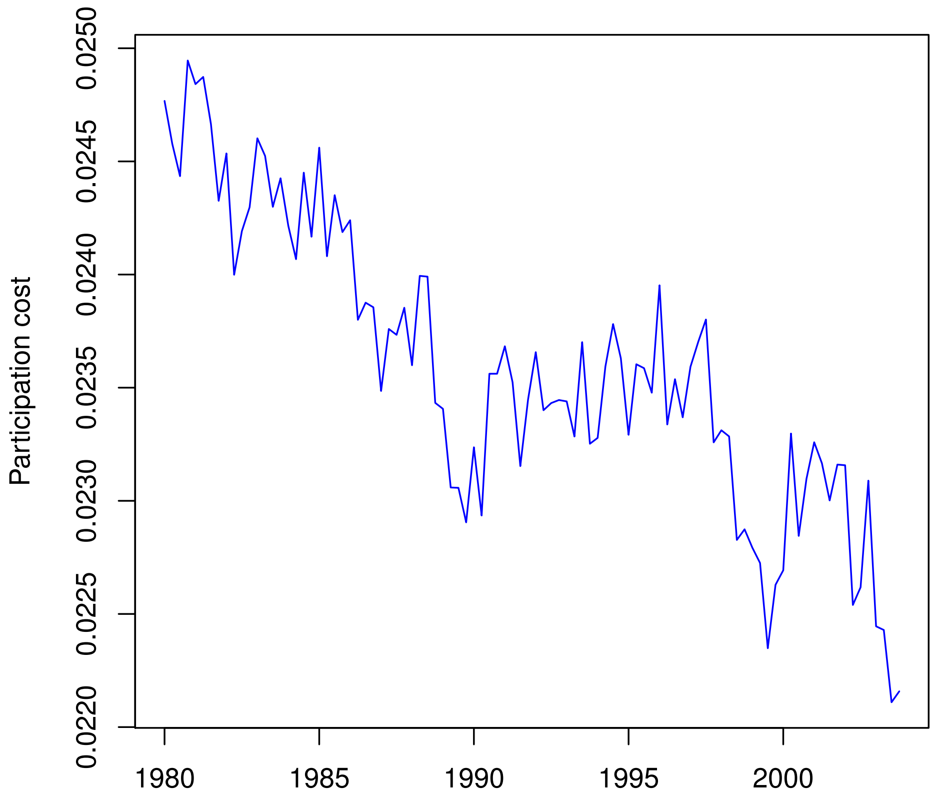

Figure 5: Temporal evolution of the participation cost. This figure plots the participation cost $\tau_{c,t}$ where $\overline{\tau}_t$ is estimated using 100 Monte Carlo replicates of participation panels of the same size as the empirical dataset.