



# Reconciling conflicting evidence on the elasticity of intertemporal substitution: A macroeconomic perspective<sup>☆</sup>

Fatih Guvenen\*

*Department of Economics, University of Rochester, 216 Harkness Hall, Rochester, NY 14627, USA*

Received 21 March 2005; received in revised form 19 May 2005; accepted 24 June 2005

Available online 26 May 2006

---

## Abstract

In this paper we reconcile two opposing views about the elasticity of intertemporal substitution (EIS). Empirical studies using aggregate consumption data typically find the EIS to be close to zero, whereas calibrated models designed to match growth and fluctuations facts typically require it to be close to one. This contradiction is resolved when two kinds of heterogeneity are acknowledged: one, the majority of households do not participate in stock markets; and two, the EIS increases with wealth. We introduce these two features into a standard real business cycle model. First, limited participation creates substantial wealth inequality as in the U.S. data. Consequently, the properties of aggregates directly linked to wealth (e.g., investment and output) are mainly determined by the

---

<sup>☆</sup> Special thanks to Tony Smith and Per Krusell for their advice and encouragement. For helpful conversations and comments, I thank Daron Acemoglu, Aydogan Altı, Olivier Blanchard, Jeremy Greenwood, Robert Hall, Lars Hansen, Robert King (the Editor), Torsten Persson, David Laibson, Nick Souleles, Stan Zin, an anonymous referee, as well as the seminar participants at Carnegie-Mellon, Chicago GSB, IIES at Stockholm University, Koc University, MIT, Penn State, Queen's, Richmond FED, Rochester, Sabanci University, Upenn, UCSD, the NBER Economic Fluctuations and Growth Meeting, the Minnesota Macroeconomics Workshop, and the Stanford SITE conference. An earlier version of this paper was circulated under the title "Mismeasurement of the Elasticity of Intertemporal Substitution: The Role of Limited Stock Market Participation."

\*Corresponding author at: Department of Economics, University of Texas at Austin, 1 University Station, C3100, BRB 3.118, Austin, TX 78712, USA.

*E-mail address:* [guvenen@eco.utexas.edu](mailto:guvenen@eco.utexas.edu).

(high-EIS) stockholders. Since consumption is much more evenly distributed than is wealth, estimation from aggregate consumption uncovers the low EIS of the majority (i.e., the poor).

© 2006 Elsevier B.V. All rights reserved.

*JEL classification:* E32; E44; E62

*Keywords:* The elasticity of intertemporal substitution; Limited stock market participation; Business cycle fluctuations; Incomplete markets; Wealth inequality

## 1. Introduction

A rational agent will seize intertemporal trade opportunities revealed in asset prices by adjusting her consumption growth, by an amount inversely related to her (counteracting) desire for a smooth consumption profile. The degree of this consumption growth response is called the elasticity of intertemporal substitution in consumption (EIS). Research in many fields of macroeconomics has established the EIS as crucial for many questions ranging from the effects of monetary and fiscal policies to the determinants of long-run economic growth (see, for example, Summers, 1981; King and Rebelo, 1990; Jones et al., 1999). The goal of this paper is to use macroeconomic analysis to reconcile seemingly contradictory evidence about the value of this parameter.

On the one hand, macroeconomists generally use a value close to 1, reflecting the view that a high degree of intertemporal substitution is more consistent with aggregate data viewed through the lens of dynamic macroeconomic models. To illustrate some of the arguments, let us rearrange the consumption Euler equation to obtain

$$R_t^f = \eta + (1/EIS) \times \log(C_{t+1}/C_t), \quad (1)$$

where  $\eta$  is the time preference rate. For convenience we abstracted from uncertainty.<sup>1</sup> Given that the annual per-capita consumption growth in the U.S. is about 2%, an elasticity of 0.1 as estimated by Hall (1988) together with  $\eta > 0$  implies a lower bound of 20% for the real interest rate! This observation is well known as Weil's (1989) risk-free rate puzzle. Alternatively, substituting a realistic average interest rate of 3%, and a consumption growth of 2%, requires the EIS to be at least 0.66 if  $\eta$  is to be positive. In fact, a similar observation has led Lucas (1990) to rule out an elasticity below 0.5 as implausible (in his notation  $\sigma \equiv 1/EIS$ ):

If two countries have consumption growth rates differing by one percentage point, their interest rates must differ by  $\sigma$  percentage points (assuming similar time discount rates). A value of  $\sigma$  as high as 4 would thus produce cross-country interest differentials much higher than anything we observe, and from this viewpoint even  $\sigma = 2$  seems high.

Similarly, Jones et al. (2000) have examined the volatilities of the macroeconomic time-series generated by a real business cycle model and concluded that they match the U.S.

<sup>1</sup>The following argument is robust to the introduction of uncertainty and generalizing preferences to the Epstein and Zin (1989) utility function, because an equation very similar to (1) can still be derived under these conditions; see, for example, Attanasio and Weber (1989).

data the best when the EIS is calibrated to be between 0.8 and 1. For many macroeconomists reasoning like these constitute convincing evidence that the EIS is quite high, probably close to 1.

On the other hand, an alternative approach is to use the *conditional* consumption Euler equation, and estimate the EIS from the co-movement of aggregate consumption with asset returns (Hansen and Singleton, 1983; Hall, 1988). This line of research, however, has reached a completely different conclusion. In an influential paper Hall (1988) has argued that consumption growth is completely insensitive to changes in interest rates and, hence, the EIS is very close to zero. The subsequent empirical macro literature has provided further support (see Campbell and Mankiw, 1989; Browning et al., 1999, and the references therein).<sup>2</sup> Thus, there is an apparent contradiction between the dynamic macroeconomics literature and econometric studies which both use aggregate data in different ways.

We show that this apparent inconsistency is largely a consequence of the “representative agent” perspective widely adopted in both literatures. To this end, we study an otherwise standard real business cycle model and introduce two key sources of heterogeneity: limited participation in the stock market and heterogeneity in the elasticity of intertemporal substitution. A large body of empirical evidence will be presented in Section 3 documenting these facts. Specifically, we consider an economy with neoclassical production and competitive markets. There are two types of agents. The majority of households (second type) do not participate in the stock market where one-period risky claims to the aggregate capital stock are traded. However, a risk-free bond is available to all households, so non-stockholders can also accumulate wealth and smooth consumption intertemporally. Finally, consistent with the empirical evidence reviewed in Section 3, stockholders are assumed to have a higher EIS (around 1.0) than non-stockholders (around 0.1).

The main result of the paper can be explained as follows. In the model, limited participation creates substantial wealth inequality matching the extreme skewness observed in the U.S. data. Consequently, the properties of aggregate variables directly linked to wealth, such as savings, investment, and output, are almost entirely determined by the (high-elasticity) stockholders who own virtually all the wealth (capital) in the economy. On the other hand, consumption turns out to be much more evenly distributed across households, again as in the U.S. data, so aggregate consumption—and hence Euler equation estimations—mainly reveal the low EIS of the majority, that is, the poor. As a result, the model delivers several business cycle statistics that appear to be generated from a representative agent RBC model with an EIS of 1 (Section 4), while at the same time a Hall-type Euler equation estimation uncovers an EIS of around 0.25 (Section 6).

This explanation clearly relies on the premise that the *average investor* is very different than the *average consumer*. Section 4.1 documents the remarkably different concentrations of wealth and consumption in the U.S. data. For example, the richest 20% own 83% of net worth and 95% of financial assets, but account for only about 30% of aggregate consumption (Fig. 1). Moreover, it is clear that the ability of the *model* to generate this substantial wealth inequality is crucial for our results, otherwise the average consumer and investor would be similar in the model. Further, since similar models with a dynastic structure typically yield too little wealth inequality, Section 5 explains how the introduction of limited participation into such a framework generates substantial

---

<sup>2</sup>We discuss some econometric studies which obtain higher estimates of the EIS below.

heterogeneity. Finally, an appendix shows—in the context of a capital income taxation problem—that policy conclusions drawn from a representative-agent model calibrated to the EIS obtained from Euler equation estimations can be seriously misleading.

### 1.1. Related literature

This paper is related to two strands of literature. The first one, starting with an influential paper by Mankiw and Zeldes (1991), investigates the role of limited stock market participation in resolving asset pricing puzzles (Basak and Cuoco, 1998; Attanasio et al., 2002; Guvenen, 2004). However, despite the increasing number of studies in this field, the main focus so far has been on asset prices and less attention has been devoted to the macroeconomic implications of this phenomenon which remain largely unexplored. In this paper we attempt to close this gap by studying the role of limited participation in generating inequality in wealth and consumption, and consequently, in shaping the determination of macroeconomic aggregates. A second literature estimates the EIS using different versions of the log-linearized Euler equation. Although some early studies (Summers, 1981; Hansen and Singleton, 1983) obtained estimates of the EIS around 1.0, Hall (1988) argued that these estimates were biased upward because of the time aggregation in consumption data. More recently, Ogaki and Reinhart (1998) have argued that non-separability between durables and non-durables could bias the estimates of EIS, such as Hall's (1988) in the opposite direction—downward—if not accounted for. Similarly, Basu and Kimball (2000) have shown that non-separability between consumption and leisure could create a similar downward bias. Both papers obtained estimates of the EIS around 0.35. Compared to these studies, this paper stresses a different economic point: even when the Euler equation is estimated without any problems, the resulting estimate of the EIS (which is necessarily an “average” of individual elasticities) is not the appropriate measure of elasticity for many questions, especially those related to savings, investment, and output.

## 2. The model

For transparency of results, our modeling goal is to stay as close to the standard real business cycle framework as possible and only introduce the two key features discussed above.

### 2.1. The firm

There is a single perishable consumption–investment good in this economy. The single aggregate firm converts capital ( $K_t$ ) and labor ( $L_t$ ) inputs into output according to a Cobb–Douglas technology:  $Y_t = Z_t K_t^\theta L_t^{1-\theta}$ , where  $\theta \in (0, 1)$  is the factor share parameter. The stochastic technology level  $Z_t$  follows a first-order Markov process with a strictly positive support. In the absence of any intertemporal links in the firm's problem (such as capital adjustment costs, etc.) the firm's decision can be simply expressed as a series of one-period profit maximization problems:

$$\text{Max}_{K_t, L_t} [Z_t K_t^\theta L_t^{1-\theta} - (R_t^s + \delta)K_t - W_t L_t],$$

where  $R_t^s$  and  $W_t$  are the market return on capital and the wage rate, respectively, and  $\delta$  is the depreciation rate of capital.

Finally, capital and labor markets are competitive implying that the factors are paid their respective marginal products after production takes place:

$$\begin{aligned} R_t^s &= \theta Z_t (K_t/L_t)^{\theta-1} - \delta, \\ W_t &= (1 - \theta) Z_t (K_t/L_t)^\theta. \end{aligned} \quad (2)$$

## 2.2. Households

We consider an economy populated by two types of agents who live forever. The population is constant and is normalized to unity. Let  $\lambda$  ( $0 < \lambda < 1$ ) denote the measure of the first type of agents (who will be called “stockholders” later) in the total population. Both agents have one unit of time endowment in each period, which they supply inelastically to the firm.

### 2.2.1. Preferences

Agents value temporal consumption lotteries according to the following Epstein–Zin (1989) recursive utility function:

$$U_t^i = [(1 - \beta)(C_t)^{\varphi^i} + \beta(E_t(U_{t+1}^i)^{1-\alpha})^{\varphi^i/(1-\alpha)}]^{1/\varphi^i}, \quad 0 \neq (1 - \alpha), \quad \varphi^i < 1 \quad (3)$$

for  $i = h, n$ . Throughout this paper the superscripts  $h$  and  $n$  denote stockholders and non-stockholders, respectively. The subjective time discount factor under certainty is given by  $\beta \equiv 1/(1 + \eta)$ , and  $\alpha$  is the risk aversion parameter for wealth gambles, common to both types. The focus of attention in this paper is the EIS parameter which is denoted by  $\rho^i \equiv (1 - \varphi^i)^{-1}$  and as the superscript  $i$  indicates, types may differ in their intertemporal elasticities.

Before proceeding further, it is important to stress that our choice of recursive preferences is for clarity: by disentangling risk aversion and the elasticity of intertemporal substitution, this specification allows us to introduce heterogeneity in the EIS without generating corresponding differences in risk aversion. It is worth noting though that risk aversion plays little role in the model: even assuming CRRA utility delivers practically the same results as long as the calibration of the EIS remains the same as here. Finally, note that these preferences nest expected utility as a special case: when  $\varphi = 1 - \alpha$ , this specification reduces to the familiar CRRA expected utility function.

### 2.2.2. Stock market participation

Besides the productive capital asset there is also a one-period risk-less household bond (in zero net supply) that is traded in this economy. The crucial difference between the two groups is in their investment opportunity sets: the “non-stockholders” can freely trade the bond, but they are restricted from participating in the capital market. The “stockholders,” on the other hand, have access to both markets and hence are the sole capital owners in the economy.<sup>3</sup>

<sup>3</sup>It is possible to think of the participation structure assumed here as the *endogenous* outcome of a model where every period agents have the option of paying a one-time fixed cost of entering the stock market (if they had not already done so in previous periods). With a cost of appropriate magnitude, the group of agents with high EIS will

Finally, we impose portfolio constraints as a convenient way to prevent Ponzi schemes. As we discuss later, for the main results of the paper these constraints need not bind; they can be as loose as possible.

### 2.3. Households' dynamic problem and the equilibrium

To state the individual's problem recursively, the aggregate state space for this economy needs to be specified. The Markov characteristic of the exogenous driving force naturally suggests concentrating on equilibria that are dynamically simple. That is, we assume that the portfolio holdings of each group together with the exogenous technology shock constitute a sufficient state space which summarizes all the relevant information for the equilibrium functions.

In a given period, the portfolios of each group can be expressed as functions of the *beginning-of-period* capital stock,  $K$ , the aggregate bond holdings of the non-stockholders *after* production,  $B$ , and the technology level,  $Z$ . Let us denote the aggregate state by  $\mathbf{Y} \equiv (K, B, Z)$ , and the financial wealth of an agent by  $\omega$  where superscripts are suppressed for clarity of notation. Given the recursivity of the utility and the stationarity of the environment, maximization of (3) for the stockholders can be expressed as the solution to the following dynamic programming problem:

$$\begin{aligned}
 V(\omega; \mathbf{Y}) = \max_{C, b', s'} & \quad ((1 - \beta)(C)^\varphi + \beta(E[V(\omega'; \mathbf{Y}')^{1-\alpha} | Z])^{\varphi/(1-\alpha)})^{1/\varphi} \\
 \text{s.t.} & \\
 C + q(\mathbf{Y})b' + s' & \leq \omega + W(K, Z), \\
 \omega' & = b' + s'(1 + R^s(K', Z')), \\
 K' & = \Gamma_K(\mathbf{Y}), \\
 B' & = \Gamma_B(\mathbf{Y}), \\
 b' & \geq \underline{b}^h,
 \end{aligned}$$

where  $b'$  and  $s'$  denote bond and stock (capital) choice of the agent, respectively. The endogenous functions  $\Gamma_K$  and  $\Gamma_B$  denote the laws of motion for aggregate wealth distribution which are determined in equilibrium, and  $q$  is the equilibrium bond pricing function. Note that each agent is facing a constraint on bond holdings with possibly different (and negative) lower bounds. The problem of the non-stockholder can be written as above with  $s' \equiv 0$ .

A *recursive competitive equilibrium* for this economy is given by a pair of value functions  $V^i(\omega^i; \mathbf{Y})$  ( $i = h, n$ ), consumption and bondholding decision rules for each agent,  $C^i(\omega^i; \mathbf{Y})$  and  $b^i(\omega^i; \mathbf{Y})$ , stockholding decision for the stockholder,  $s(\omega^h; \mathbf{Y})$ , a bond pricing function,

(footnote continued)

enter the stock market whereas the other group will stay out. This is because, loosely speaking, agents with a low EIS are reluctant to accumulate wealth quickly and hence their benefits from participation is lower than agents with a high EIS. See the working paper version of this article for further discussion. In Guvenen (2004) we quantify the magnitude of this participation cost in a version of this model with capital adjustment costs in production. For the purposes of this paper, the present setup seems a reasonable simplification.

$q(\mathbf{Y})$ , competitive factor prices,  $R^s(K, Z)$ ,  $W(K, Z)$ , and laws of motion for aggregate capital and aggregate bond holdings of non-stockholders,  $\Gamma_K(\mathbf{Y})$ ,  $\Gamma_B(\mathbf{Y})$ , such that:

(1) Given the pricing functions and the laws of motion, the value functions and decision rules of each agent solve that agent's dynamic problem.

(2) Factors are paid their respective marginal products (Eq. (2) is satisfied).

(3) The bond market clears:  $\lambda b^h(\varpi^h; \mathbf{Y}) + (1 - \lambda)b^n(\varpi^n; \mathbf{Y}) = 0$ , where  $\varpi^i$  denotes the aggregate wealth of a given group; and the labor market clears:  $L = \lambda \times 1 + (1 - \lambda) \times 1 = 1$ .

(4) Aggregates result from individual behavior:

$$K' = \lambda s(\varpi^h, \mathbf{Y}), \quad (4)$$

$$B' = (1 - \lambda)b^n(\varpi^n, \mathbf{Y}). \quad (5)$$

### 3. Numerical solution and calibration

The model is solved using numerical methods since an analytical solution is not available. As usual the existence of aggregate shocks introduces an additional layer of complexity. Instead of approximating the equilibrium functions around the stochastic steady state (for example, as in Krusell and Smith, 1998) we solve for all the functions globally over the state space. Although this generality has additional computational costs, it also allows us to conduct policy experiments involving transitions (such as the capital income taxation problem studied in the Appendix) very easily. Moreover, to our knowledge this is the first attempt at numerically solving a dynamic programming problem with general recursive utility. A computational appendix available from the author's website contains the algorithm as well as a discussion of the accuracy of the solution.

#### 3.1. Baseline parameterization

The model parameters are chosen to replicate the long-run empirical facts of the U.S. economy. The time period in the model corresponds to one year of calendar time. Following Cooley and Prescott (1995) the capital share of output is set equal to 0.4. The technology shock  $Z$  is assumed to follow a first-order, two-state Markov process with transition probabilities  $\pi_{ij} = P(Z_{t+1} = j | Z_t = i)$  chosen such that business cycles are symmetric and last for 6 years on average. This condition implies  $\pi_{11} = \pi_{22} = 2/3$ . The mean of the technology shock is a scaling parameter and is normalized to 1. The standard deviation of the technology shock,  $\sigma(Z)$ , is set equal to 3.1% which is the implied annual volatility assuming that the quarterly Solow residuals follow an AR(1) process with a persistence of 0.95 and a coefficient of variation of 1%. We also investigate the sensitivity of the results to alternative calibrations.

##### 3.1.1. Participation rates

Our model assumes a constant participation rate in the stock market, which appears to be a reasonable approximation for the period before the 1990s when the participation rate was either stable or increasing gradually (Poterba and Samwick, 1995). In contrast, during the 1990s participation increased substantially: from 1989 to 2002 the number of households who owned stocks increased by 74%, and by 2002 half of the U.S. households

had become stock owners (Investment Company Institute, 2002). Modeling the participation boom in this later period would require going beyond the stationary structure of our model, so we leave it for future work. In this paper, we calibrate the participation rate excluding this later period (1990 onward). We set the measure of stockholders,  $\lambda$ , equal to 30%, which roughly corresponds to the stock market participation rate during the 1980s, when participation is defined as holding any positive amount of stocks. A significant fraction of stockholders during this period owned no more than a few thousand dollars worth of stocks (see again Poterba and Samwick, 1995), so this percentage should probably be interpreted as an upper bound on the percentage of households actively participating in the stock market. In the next section we also report the results by assuming a participation rate of 15%.

### 3.1.2. Heterogeneity in the EIS

There is a fairly large and active literature documenting heterogeneity in the EIS. A first group of papers consider flexible preference specifications allowing for non-homotheticity and estimate their parameters from the consumption Euler equation. Using this approach, Blundell et al. (1994) find that the EIS is monotonically increasing in income, and in some specifications it is more than three times larger for the highest decile compared to the lowest decile. They also investigate heterogeneity in the EIS due to other factors but conclude that “most of the variation in the EIS across the population is due to differences in consumption (which can loosely be thought of as a proxy for lifetime wealth) and not to differences in demographics and labor supply variables” (p. 73). Similarly, Attanasio and Browning (1995) confirm this finding by employing an even more general functional form, although they do not report point estimates for the EIS.<sup>4</sup> Since stockholders are on average much wealthier than the rest of the population, the finding that elasticity is increasing in wealth also provides evidence of heterogeneity between the EIS of stockholders and others.

A second group of papers focus *directly* on stockholders and non-stockholders and estimate a separate elasticity parameter for each group (assuming homothetic preferences within each group). For example, Attanasio et al. (2002) use the Family Expenditure Survey data set from the U.K. and obtain elasticity values around 1 for stockholders, and between 0.1 and 0.2 for non-stockholders. Vissing-Jørgensen (1998) obtains very similar estimates from the Consumer Expenditure Survey data on U.S. households.

With Epstein–Zin preferences, the risk aversion and the EIS can be calibrated independently of each other. For transparency of results, we abstract from heterogeneity in risk aversion and set  $\alpha^h = \alpha^n = 3$ , which is within the range viewed as plausible by many economists. This allows us to focus purely on the effects of heterogeneity in the elasticity of substitution. In light of the empirical evidence reviewed above, the EIS of the non-stockholders is set equal to 0.1, and the EIS of the stockholders to 1.0.

<sup>4</sup>It is possible to give some theoretical justifications for why the EIS may increase with an individual's wealth level. For example, if individuals consume a variety of goods with different income elasticities, the intertemporal elasticity of the *total* consumption bundle will be higher for wealthier individuals. This is because the share of necessity goods in total consumption is large when the individual is poor, making her less willing to substitute these necessities across time. In contrast, wealthy individuals will have a larger fraction of luxury goods in their bundle, which are more easily substituted across time (see Browning and Crossley, 2000, for a proof). Similarly, non-homotheticity in preferences due to subsistence requirements, habit formation and so on also typically implies a higher EIS for the wealthy.

Finally, the subjective discount factor,  $\beta$ , is set equal to 0.96 in order to match the U.S. capital–output ratio of 3.3 reported by Cooley and Prescott (1995).

### 3.1.3. Borrowing constraints

These bounds are chosen to reflect the fact that the stockholders can potentially accumulate capital which can then be used as collateral for borrowing in the risk-free asset, whereas the non-stockholders have to pay all their debt through future wages. In the baseline case, the stockholders' borrowing constraint is set equal to four years of expected labor income ( $\underline{B}^h = 4 \times E(W)$ ). The borrowing limit of the non-stockholders is set to 30% of one year's expected income, which is the average credit limit typically imposed by short-term creditors, such as credit card companies. Table 1 summarizes the baseline parameterization.

## 4. Macroeconomic results

This section *quantitatively* examines if limited participation and heterogeneity in the EIS can explain the conflicting evidence discussed in the Introduction. Our main argument is that the majority of the population with low elasticity has quantitatively little effect on aggregates directly linked to wealth, such as investment and output, which are determined by the high elasticity of the wealthy stockholders. On the other hand, aggregate consumption mainly reveals the preferences of the poor, who contribute substantially. Clearly, this argument relies on the idea that the average investor is significantly different than the average consumer violating the representative-agent assumption. Thus, it is important to characterize the joint distribution of consumption and wealth in the U.S. data to document that empirically this is indeed the case.

### 4.1. The empirical joint distribution of consumption and wealth

Table 2 reports the fraction of aggregate consumption and wealth accounted for by different percentiles of the wealth distribution calculated from the Panel Study of Income Dynamics (PSID) data set.<sup>5</sup> The first three rows display the size distributions of different measures of wealth, from the most comprehensive measure in the first row (net worth) to the most specific in the third (financial assets).<sup>6</sup> The intermediate measure, which we call

<sup>5</sup>Because we are interested in both wealth and consumption, we use the PSID instead of, for example, the Survey of Consumer Finances (SCF), which contains more detailed information about wealth but no consumption data, or the Consumer Expenditure Survey (CE), which has detailed consumption data but no comparable information on wealth holdings. One drawback of PSID is that consumption is limited to food expenditures and rent payments. We take the sum of these two components (which makes up about 40% of non-durables and services expenditures) as our consumption proxy in Table 2. The ratio of this proxy to the non-durables and services expenditures from the CE is roughly constant across income deciles, which suggests that it could provide a reasonable approximation to the distribution of this more general consumption measure. A data appendix, available from the author's website, provides further details about the construction of the consumption proxy and discusses how it compares to the CE data. Moreover, the size distributions of both wealth measures are quite similar to those reported by Wolff (2000) using data from the SCF.

<sup>6</sup>“Net worth” is defined as the current value of all marketable or fungible assets less the current value of debts. Specifically, it includes: (1) the net equity in owner-occupied housing; (2) other real estate; (3) cash and demand deposits; (4) time and saving deposits, CDs and money market accounts; (5) government bonds, corporate bonds and other financial securities; (6) cash surrender value of life insurance policies; (7) cash surrender values of

Table 1  
Baseline parameterization

Annual model		
Parameter		Value
$\beta$	Time discount rate	0.96
$\alpha^h$	Risk aversion of stockholders	3
$\alpha^n$	Risk aversion of non-stockholders	3
$\rho^h$	EIS of stockholders	1
$\rho^n$	EIS of non-stockholders	0.1
$\lambda$	Participation rate	0.3
$\pi_{11}$	Prob (good state   good state)	2/3
$\pi_{22}$	Prob (bad state   bad state)	2/3
$\sigma_Z$	Standard deviation of $Z$	0.031
$\theta$	Capital share	0.4
$\delta$	Depreciation rate	0.08
$\underline{B}^h$	Borrowing limit of stockholders	$4\bar{W}$
$\underline{B}^n$	Borrowing limit of non-stockholders	$0.3\bar{W}$

Note: The mean of the technology shock is a scaling parameter and is normalized to 1. The borrowing limits are indexed to the average wage rate,  $\bar{W}$ .

Table 2  
The concentration of wealth and consumption

Wealth percentile	U.S. data					Model	
	1–10	11–30	31–50	51–100	1–30	1–30	31–100
<i>Variable</i>							
Net worth	0.702	0.181	0.087	0.030	0.883	0.89	0.11
Productive wealth	0.742	0.201	0.054	0.002	0.942	0.89	0.11
Financial assets	0.832	0.210	0.026	–0.068	1.043	0.89	0.11
Consumption	0.169	0.251	0.205	0.373	0.429	0.37	0.63

Notes: The data are drawn from the 1989 PSID family file and wealth supplement. All variables are defined at household level. The consumption measure is a proxy constructed from food expenditures and rent as explained in footnote 5. Productive wealth is defined as net worth minus equity in owner-occupied housing, where net worth and financial assets are defined in footnote 6.

“productive wealth” for lack of a better term, is defined as net worth minus equity in owner-occupied housing. We use these three measures to show that the general conclusions drawn about the concentration of wealth are not sensitive to the specific definition of

(footnote continued)

pension plans, including IRAs, Keogh and 401(k) plans; (8) corporate stocks and mutual funds; (9) net equity in unincorporated businesses; and (10) equity in trust funds. From the sum of these assets we subtract consumer debt including auto loans and others. “Productive wealth” is the sum of (2)–(10). The narrowest definition is “financial wealth,” and is the sum of (3)–(8).

wealth used. Finally, the last row displays the share of aggregate consumption accounted for by households in different percentiles of the productive wealth distribution.

The main observation from this table is that all three measures of wealth display substantially higher concentration than consumption: households in the top 10% of the wealth distribution own three-quarters of aggregate productive wealth but account for only about 17% of total consumption in the U.S. data.<sup>7</sup> By also including the next decile (top 20%), this group of households can be thought of as the “investors” since they hold about 90% of capital and land, and virtually all financial assets. In contrast, observe the very gradual decline in consumption expenditures moving down the wealth distribution. The bottom 80% own only 12% of productive wealth, yet contribute almost 70% of total consumption. Thus this latter group can be thought of as the “consumers.” Expressing the same information in per capita terms makes this distinction even more striking: an average investor owns *29.3 times* the productive wealth of an average consumer, but consumes only *1.7 times* more. These two groups are also marked in Fig. 1 (which plots the joint distribution) to give a visual impression of their dramatically different contributions to aggregate consumption and wealth.<sup>8</sup> Moreover, the richest 30%—who own 99% of stocks before the 1990s, and hence whom we identify with the stockholders—held 88% of net worth and more than 100% of financial assets.

The fact that consumption is more evenly distributed than physical wealth is not totally surprising, since consumption is proportional to lifetime wealth inclusive of human capital, which constitutes a substantial part of lifetime wealth and is more evenly distributed than physical capital. As we shall see below, the limited participation model generates the same kind of wealth and consumption distributions and thus captures this key component in explaining the evidence on the EIS.

#### 4.2. Results from the baseline economy

Given how remarkably different the average consumer and the average investor are, it might seem almost obvious that each of these groups will (largely) determine different aggregates, and together with the heterogeneity in the EIS across these two groups, a reconciliation of the evidence on the EIS would seem to follow immediately. Why is it then not sufficient to merely document these two kinds of heterogeneity in the data to qualify as a full explanation?

The reason is the existence of trade in the bond market, which provides a channel through which the non-stockholders’ preferences can potentially influence the properties of aggregate quantities, including investment and output. For example, in a companion paper using essentially the same model (augmented with capital adjustment costs in production) we found that the non-stockholders’ preferences—and their low EIS in particular—play a key role in determining asset prices despite the fact that they hold very little wealth in that

---

<sup>7</sup>One drawback of existing micro-data sets containing consumption data is that they tend to underrepresent very rich households (i.e., those in the top 1% of the wealth distribution). As a result, stockholders’ consumption is likely to be somewhat understated to the extent that the consumption of these very rich households are not captured.

<sup>8</sup>Fig. 1 also makes clear that there is substantial heterogeneity among the very wealthy as well. Thus, while in the baseline parameterization we take the fraction of stockholders to be 30%, most of the wealth is in fact held by a subset of these households. In the next section we experiment with a lower participation rate of 15% to examine whether it makes a difference in our results.

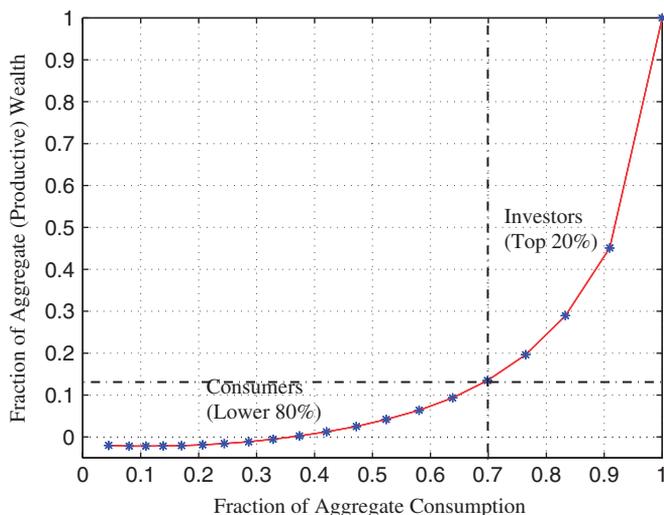


Fig. 1. The fraction of aggregate consumption accounted for by households who own a given fraction of aggregate wealth.

model as well (Guvenen, 2004). Thus, it is essential to investigate the behavior of aggregates allowing for equilibrium interactions through the bond market, which is done in this subsection.

We begin by comparing the cross-sectional distributions of consumption and wealth generated by the model to their empirical counterparts documented above. With only two types of agents, the model generates two-point distributions for each variable reported in the last two columns of Table 2. First, the wealth distribution is extremely skewed in the model, with 89% of aggregate wealth owned by the stockholders, similar to the share of wealth owned by the top 30% in the U.S. data. And second, consumption is much more evenly distributed, with 37% of aggregate consumption accounted for by the stockholders in the model, compared to 43% in the data. So, the model is able to generate the significantly different concentrations of these two variables, which is an essential element in our explanation of the evidence on the value of the EIS.

We now examine the implications of the limited participation model for business cycle statistics. The first row of Table 3 reports the standard deviations and the first-order autocorrelations of aggregate output (GDP), investment and consumption from the U.S. data. First, in order to see why a low EIS seems difficult to reconcile with aggregate fluctuations, consider the statistics reported on row 2 from a representative agent RBC model with an EIS of 0.1, which otherwise has the same parameterization as the baseline limited participation model. Notice that while the volatility of consumption matches its empirical counterpart, the volatilities of output and investment are overstated. The explanation is simple: because the agent desires a very smooth consumption path, investment has to absorb the shock to her income, making the former smooth at the expense of extra volatility in investment, and consequently, in output. Moreover, all three variables are too persistent.

The third row reports the results when the EIS of the representative agent is raised to 1.0. The implications of the model move closer to data: output and investment become less

Table 3  
Business cycle statistics

		Standard deviation (%)			Autocorrelation		
		<i>Y</i>	<i>I</i>	<i>C</i>	<i>Y</i>	<i>I</i>	<i>C</i>
(1)	U.S. data	2.2	8.5	1.5	0.52	0.36	0.57
Models		Panel A: Baseline parameterization					
(2)	RA ( <i>EIS</i> = 0.1)	3.6	10.8	1.5	0.65	0.49	0.98
(3)	RA ( <i>EIS</i> = 1.0)	3.0	8.7	1.7	0.52	0.33	0.94
(4)	Limited participation	3.0	8.8	1.5	0.53	0.36	0.96
		Panel B: Fraction holding stocks is 15%					
(5)	Limited participation	3.1	8.9	1.5	0.54	0.37	0.96
		Panel C: Business cycle duration is 8 years					
(6)	RA ( <i>EIS</i> = 0.1)	4.1	11.3	1.9	0.76	0.63	0.98
(7)	RA ( <i>EIS</i> = 1.0)	3.4	9.1	2.1	0.65	0.45	0.95
(8)	Limited participation	3.5	9.4	2.0	0.67	0.47	0.96

*Notes:* The empirical statistics of the U.S. economy are computed from the National Income and Product Accounts data at yearly frequencies covering 1959:1999. All variables are first logged and the trend is removed with Hodrick–Prescott filter with a smoothing parameter of 100. The consumption measure consists of expenditures on non-durables and services. In Panel B, the borrowing constraint of stockholders is increased to  $8 \times \bar{W}$  to keep the total wealth that can be accumulated by the non-stockholders unchanged from the baseline case.

volatile, without a major change in consumption volatility. The persistence of output and investment also come closer to their empirical counterparts. While the persistence of consumption falls slightly, it is still significantly higher than in the data.

The fourth row displays the moments from the baseline model. Comparing these statistics to those on the previous two rows, it is clear that the limited participation model behaves like the representative agent model with a high *EIS* (row 3) when output and investment are concerned. It is interesting that the existence of many households with a low *EIS* has almost no effect on the properties of investment and output. In fact, the next row reports the results when the participation rate is reduced to 15%, which shows no appreciable change in the behavior of output and investment. As for consumption, the unconditional moments examined here do not appear to be very sensitive to the *EIS* parameter.

To examine the robustness of this conclusion, we have experimented with a number of changes in the baseline parameterization such as increasing the persistence and variance of aggregate shocks, reducing the participation rate (discussed above), and relaxing the non-stockholders' borrowing constraints to 4 years' of expected income. As an example, rows 6–8 report the results when the business cycle duration is increased to 8 years from its baseline value of 6. The pattern seen here applies to the other exercises described: the properties of output and investment mainly reflect the high *EIS* of the stockholders, while the statistics for consumption are not very sensitive to the value of the *EIS*. However, the

co-movement of consumption growth with returns *is* informative about the value of the EIS parameter, which is further analyzed in Section 6.<sup>9</sup>

### 4.3. Heterogeneity in the EIS: further evidence

Our motivation for heterogeneity in the EIS came from a large empirical literature, which was discussed in Section 3. However, a potential caveat in that empirical evidence is pointed out by Laibson et al. (1998). Their argument can be summarized as follows. Suppose that there are two agents, a wealthy and a poor, who both have the same EIS. If there are borrowing constraints which frequently bind for the poor agent, her consumption will closely track her income and will not respond to the interest rate unlike the consumption of the wealthy (and unconstrained) agent. When the elasticities are recovered from consumption data on these two agents, the poor one will appear as if she has a low EIS. If true, then contrary to what we have assumed, there may not be any significant heterogeneity in the EIS in the first place.

To further examine the evidence on the heterogeneity in the EIS we first turn to cross-sectional data. An empirically well-documented difference between the stockholders and the non-stockholders is that the consumption growth of the former is more volatile than that of the latter: the ratio of the variances of consumption growth,  $\sigma^2(\Delta c^h)/\sigma^2(\Delta c^n)$ , ranges from 2 to 4 depending on the consumption measure and the threshold stockholding level used to identify the two groups.<sup>10</sup> This finding seems somewhat surprising though. Given that the stockholders are much wealthier than the rest and have access to financial markets, and low-income households typically have higher unemployment risk and labor earnings uncertainty, one might expect the former group's consumption to be smoother than that of the latter. In fact, it can easily be shown that a standard heterogenous-agent model with identical preferences and borrowing constraints (as in Aiyagari, 1994) would predict exactly that: because poor households cannot self-insure as effectively as the wealthy, they keep hitting their borrowing constraints. Thus their consumption will track their income and will be very volatile, contradicting this stylized fact.

In fact, a similar result is obtained in the limited participation model if preference heterogeneity is ignored: when we set  $\rho^h = \rho^n = 1$ , the ratio of the variances of consumption growth is 0.85, significantly lower than in the data. We have experimented with varying the shock size and its persistence, replacing the Markov structure with AR(1) shocks, and changing the calibration of constraints, without a noticeable increase in this ratio. Therefore, limited participation alone does not account for this fact. However, when we set  $\rho^n = 0.1$ , this picture changes. Now, the non-stockholders' stronger desire for a smooth consumption path causes them to trade vigorously in the bond market driving the interest rate down, and increasing the equity premium. Now, the non-stockholders'

<sup>9</sup>We do not examine a broader set of business cycle statistics, such as the behavior of labor hours, etc., because the model is not designed to address those issues. It is of course possible to extend the model to investigate those questions but that would require us to take a stand on further differences between the stockholders and the non-stockholders (such as whether or not they differ in their labor supply elasticities, etc.) which would distract from the main point of the paper.

<sup>10</sup>Mankiw and Zeldes (1991) and Poterba and Samwick (1995) document this finding from the PSID, Vissing-Jørgensen (1998) reports from the CEX, Attanasio et al. (2002) calculate from the Family Expenditure Survey data set from U.K.

consumption is smoother than that of the stockholders, who bear extra consumption volatility. As a result,  $\sigma^2(\Delta c^h)/\sigma^2(\Delta c^n)$  rises from 0.85 to 3.73.

Second, a number of financial statistics implied by the model are very sensitive to the EIS of the non-stockholders and hence make sharp predictions about its value. We study these asset pricing implications in Guvenen (2004). We find that, starting from an EIS value of 1 for both agents, the performance of the model improves dramatically as the non-stockholders' EIS is gradually reduced. For  $\rho^n \approx 0.05 - 0.2$ , the model is able to match a variety of asset pricing phenomena, including all the facts explained by the benchmark model of Campbell and Cochrane (1999). Examples in this list include a high equity premium, a low risk-free rate, the predictability of stock returns, the countercyclicality of the equity premium, among others. Given that many of these asset pricing phenomena have been considered puzzles, these findings provide further support to the thesis that the non-stockholders have a low EIS.

## 5. What is the effect of limited participation?

As noted earlier, the substantial wealth inequality is the driving force behind the results in this paper. This section examines the role played by limited participation in creating this wealth inequality. This is especially important because the severity of the restriction imposed by shutting some agents out of the stock market is not a priori obvious, since the only source of risk in the model is an aggregate shock with relatively low variance and persistence, *and* the non-stockholders still have access to the bond market. One could therefore suspect that some other feature of the model—such as the Epstein–Zin preferences, heterogeneity in the EIS, or borrowing constraints—is behind the wealth inequality.

To study the role of limited participation, in isolation, the following experiment is conducted. First, consider a simplified version of the baseline model where all the frictions and the preference heterogeneity is eliminated. More specifically, suppose that agents have CRRA utility with  $\rho^h = \rho^n = 1$ ; both agents have access to all markets and face no portfolio constraints. Clearly, this economy will reduce to a representative-agent model, and assuming that both agents start out at the same wealth level, there will be no trade in the bond market, no wealth inequality, and no heterogeneity at all.

Now consider introducing a single friction into this framework: suppose that the second group of agents are restricted from participating in the stock market (but otherwise impose no other constraint). As Table 4 shows, the effect of this simple change is striking. Suddenly the stockholders come to hold almost 80% of the aggregate wealth (compared to 30% without limited participation, since there was no wealth inequality in that case), or in per-capita figures, a stockholder now owns nearly eight times the average wealth of a non-stockholder. They also consume nearly 50% more per-capita than the non-stockholders. Furthermore, the two groups' wealth holdings are now virtually uncorrelated, down from perfect correlation with full participation. This example demonstrates the potential of limited participation for generating substantial cross-sectional heterogeneity.

### 5.1. Where does the wealth inequality come from?

It is useful to begin by discussing how the wealth inequality is *not* generated. First, it is not generated by the parameterization of preferences. Although preference heterogeneity

Table 4  
The effect of limited participation on cross-sectional heterogeneity

Statistics in a limited participation model with otherwise identical agents		
	Stockholders	Non-stockholders
Share of aggregate wealth	0.78	0.22
Per-capita wealth	8.27	1.00
Per-capital consumption	1.48	1.00
Std. of log consumption	0.018	0.019
Correlation of consumption with income	0.44	0.77
Cross-correlation between the two groups':		
Consumption	Income	Wealth
0.72	0.87	0.04

Notes: To make comparison across columns easier, the per-capita wealth and consumption of the non-stockholders are normalized to 1.

Table 5  
Wealth inequality for various parameter values

Parameters				Share of wealth held by	
$EIS^h$	$EIS^n$	$RA^h$	$RA^n$	Stockholders	Non-stockholders
1.0	0.1	3.0	3.0	0.88	0.12
1.0	1.0	3.0	3.0	0.85	0.15
1.25	0.25	2.0	7.0	0.87	0.13
0.33	0.33	2.0	2.0	0.85	0.15
0.33	0.14	2.0	4.0	0.86	0.14

Note:  $\rho$  and  $\alpha$  denote the EIS and relative risk aversion coefficients, respectively.

has some effect on the distribution of wealth, this effect is modest for plausible changes in the EIS parameter as well as in the risk aversion. This can be seen in Table 5 where the EIS varies from 0.1 up to 1.25, and the risk aversion varies from 2 to 7, without a substantial change in inequality.

Second, the wealth inequality is not generated by assumption, which is the case with a limited participation model in an exchange economy setting. In that framework, the stockholders are *endowed*, at time zero, with the entire future stream of dividends in addition to their labor income, and hence, are wealthier by assumption. In contrast, in the present model even when both agents start at the same wealth level, the stockholders *choose* to accumulate more wealth in equilibrium. Furthermore, there is nothing in principle that prevents the stockholders from having zero wealth in equilibrium despite the fact they must own all the capital stock: a stockholder could borrow in the bond market and invest all the proceeds in the firm, then pay the non-stockholder after production and consume the equity premium. So, the stockholders are not wealthier simply because they own the total capital stock. Then why are they?

The basic mechanism can be described as follows. For the sake of discussion, consider a simplified version of the limited participation model without preference heterogeneity, and

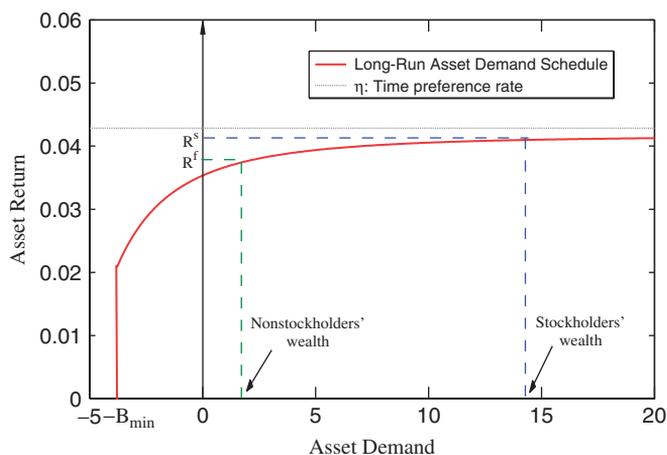


Fig. 2. Long-run wealth distribution in the presence of return differences (limited participation).

where both agents face the same portfolio constraints on their bond holdings.<sup>11</sup> These features are not primary determinants of wealth inequality as noted above, and this modification simplifies the following discussion.

In this environment, with infinite horizon and incomplete markets, both assets earn average returns below the time preference rate. A well-known result in this framework is that the wealth holdings of an agent will go to infinity if the (geometric) average of return is equal to the time preference rate.<sup>12</sup> This result also holds when asset returns are stochastic (Chamberlain and Wilson, 2000). Similarly, by a continuity argument, one can show that asset demand becomes unbounded as the average return approaches the time preference rate from below. Fig. 2 plots the typical long-run asset demand schedule for an agent in this environment and illustrates the extreme sensitivity of asset demand to small variations in returns near  $\eta$ .<sup>13</sup> The key feature of the baseline model is that both the stock and the bond return are close to the time preference rate, so agents are in this flat region of the asset demand schedule. Furthermore, the equilibrium stock return is much closer to the time preference rate than is the bond return:  $\eta - E(R^f) = 17.8 \times (\eta - E(R^s))$ . This last observation, combined with the fact that the equilibrium returns  $R^s$  and  $R^f$  are on the flat section of the asset demand schedule, explains why stockholders who have access to the slightly higher return,  $R^s$ , are willing to hold much more wealth than non-stockholders.

<sup>11</sup>It should be noted that borrowing constraints also contribute to the skewness of the wealth distribution. The non-stockholders can save only up to the maximum amount the stockholders are allowed to borrow. When the stockholders come closer to their borrowing constraints, this puts a downward pressure on bond return reducing the nonstockholders' demand for wealth. But as the example in Table 4 demonstrates, the concentration of wealth is still substantial without any constraints.

<sup>12</sup>See, for example, Aiyagari (1994) and references therein.

<sup>13</sup>Note that the stockholders and the non-stockholders will still have different asset demand schedules (despite having the same preferences and facing the same borrowing constraints) because of differences in their investment opportunity sets. For the parameterizations used in this paper, this difference in asset demand schedule turns out to be small, so we plot the same graph for both agents in Fig. 2.

Finally, the graph also makes clear that the wealth inequality cannot be simply explained by the argument that the stockholders have access to a positive equity premium and, hence, they are bound to become very wealthy in the long-run. This is because the equilibrium asset returns could very well lie in the steep region of the asset demand schedule, resulting in little wealth dispersion despite a positive equity premium.

A second, and probably more intuitive, way to explain this mechanism is as follows. With incomplete markets both agents want to accumulate precautionary wealth, and this incentive is stronger for the non-stockholders who only have access to one asset. However, the only way this group can accumulate wealth (bond) is if stockholders are willing to borrow. In contrast, the stockholders have access to capital accumulation, and they could smooth consumption even if the bond market was completely shut down (just as a representative agent would do in the standard RBC model). Furthermore, the asset demand of the non-stockholders is even more inelastic because, in addition, they have a low EIS and hence have a very strong desire for a smooth consumption profile. Therefore, trade in the bond market for consumption smoothing is more important for the non-stockholders than for the stockholders. As a result, the stockholders will only trade in the bond market if they can borrow at a low interest rate. This low interest rate in turn dampens the non-stockholders' demand for savings further, and they end up with little wealth in equilibrium (and the stockholders end up borrowing very little).

The described mechanism also corresponds to an argument commonly given for wealth inequality, but to our knowledge, which has not been formally investigated in a general equilibrium framework: the wealthy become wealthy because they face higher returns. This is exactly what happens in this model.

## 6. Another side of the puzzle: the empirical estimates of the EIS

We now look at the model economy through the lens of empirical macro-studies and show that the aggregate consumption data reveals the low EIS of the majority, i.e., the poor. Using simulated data we replicate existing studies which estimate the EIS from the log-linearized consumption Euler equation (Hansen and Singleton, 1983; Hall, 1988; Campbell and Mankiw, 1989, among others). The estimated equation is

$$\Delta c_{t+1} = k + \rho r_{t,t+1}^f + \varepsilon_{t+1}, \quad (6)$$

where small letters denote the natural logarithms of variables,  $\Delta$  is the difference operator,

$$k \equiv \rho \log(\beta) + \left[ \frac{1}{2\rho} \text{var}_t(\Delta c_{t+1}) \right],$$

and  $\varepsilon_{t+1}$  is the agent's forecast error with zero mean conditional on current information:  $E(\varepsilon_{t+1} | \Omega_t) = 0$ .

Eq. (6) is estimated via instrumental variables (IV) method using instruments commonly used in the literature. Specifically, our instrument set is an  $(8 \times 1)$  vector that consists of the lags of consumption growth and the risk-free rate:  $(1, \Delta c_{t-1}^i, \Delta c_{t-2}^i, \Delta c_{t-3}^i, \Delta c_{t-4}^i, r_{t-2}^f, r_{t-3}^f, r_{t-4}^f)$ , where  $i = h, n$  or  $A$  (referring to aggregate consumption) depending on whether Eq. (6) is estimated using the consumption of the stockholders, the non-stockholders, or the aggregate. Finally, we concentrate on large sample results, so a sample path of 63,000 observations is simulated and the first 3,000 periods are discarded.

Table 6  
Estimation of the log-linearized Euler equation with risk-free return

	Non-stockholder	Stockholders	Aggregate
<i>True value of EIS</i>	0.10	1.00	0.47
	$k^f = \text{constant}$		
<i>EIS</i> ( <i>t</i> -stats)	0.11 (22.15)	0.54 (51.1)	0.25 (36.82)
	$k^f = \text{constant} + (1/2\rho) \times \text{var}_t(\Delta c_{t+1})$		
<i>EIS</i> ( <i>t</i> -stats)	0.101 (17.02)	1.07 (27.96)	0.48 (21.84)

In the first step, the estimation is performed by assuming that the intercept term  $k$  is constant through time (or equivalently, that it is uncorrelated with the instruments). While this assumption is almost unanimously made in the empirical macro-literature, it will turn out to significantly affect the results. The first column of Table 6 reports the estimated EIS parameter of each group (estimated using that group's consumption data) as well as for the aggregate ( $\rho^A$ ). Although with incomplete markets there is no explicit mapping from individual elasticities into  $\rho^A$ , one can show that the latter will be close to a consumption-weighted average of each group's EIS:  $\rho^A \approx ((\lambda \bar{C}^h / \bar{C}^A) \rho^h + ((1 - \lambda) \bar{C}^n / \bar{C}^A) \rho^n)$ , where the bars indicate time averages of each variable.<sup>14</sup> This average in the baseline model is 0.47. However, the estimated value is 0.25, almost half of that figure.

One possible candidate for this bias is the omission of the conditional variance term  $k$  from the regression because in simulated data it is significantly (negatively) correlated with the instruments (lagged interest rates). Thus, we re-estimate (6) by properly including the time-varying conditional variance calculated from the model's optimal decision rules. The effect is dramatic: the aggregate EIS now jumps to 0.48, which is almost exactly the consumption-weighted elasticity. This exercise suggests that, to the extent that our model economy is able to capture the dynamics of these variables in the data, the previous estimates of the EIS are likely to be downward biased.<sup>15</sup>

<sup>14</sup>To see this simply note that  $\log(C_{t+1}^A / C_t^A) = \log(1 + (\lambda \Delta C_{t+1}^h + (1 - \lambda) \Delta C_{t+1}^n) / C_t^A) \approx \lambda \Delta C_{t+1}^h / C_t^A + (1 - \lambda) \Delta C_{t+1}^n / C_t^A \approx \Delta c_{t+1}^h (\lambda C_t^h / C_t^A) + \Delta c_{t+1}^n ((1 - \lambda) C_t^n / C_t^A)$ . This last expression is the weighted sum of the left hand side of Eq. (6) for the stockholders and the non-stockholders, where the weights are proportional to each group's contribution to aggregate consumption. If these weights do not change much over time, or are uncorrelated with the instruments, the estimated parameter on the right hand side will approximately be the same weighted average of the group-specific EISs, where the weights are replaced by their time averages.

<sup>15</sup>In principle, one can apply the same correction described here (by including the conditional variance of consumption growth) to the Euler equation estimations in the literature, such as Hall's (1988). In practice, however, computing the empirical conditional variance is not as straightforward as in the model (where we have access to individuals' explicit decision rules). To gain an idea about the potential empirical importance of this correction, we constructed a simple measure of  $\text{var}_t(\Delta c_{t+1})$  for each  $t$ , by calculating the rolling sample variance of consumption growth in the subsequent  $N$  periods. Including this term increases the estimate of the EIS from roughly zero to 0.3–0.6 range (depending on the choice of  $N$ , and the instrument set used) with an average slightly

## 7. Conclusion

In this paper we attempted to reconcile two seemingly contradictory views about the elasticity of intertemporal substitution. A number of observations on growth and aggregate fluctuations suggest a value of EIS close to 1. In contrast, the co-movement between aggregate consumption and interest rates—the focus of the empirical consumption literature—implies a very weak relationship between the two, suggesting an elasticity close to zero.

We showed that a simple and otherwise standard real business cycle model featuring limited participation in the stock market and heterogeneity in the elasticities is able to produce findings consistent with both capital and consumption fluctuations. In other words, an economy where the majority of households exhibit very low intertemporal substitution is consistent with aggregate capital and output fluctuations as long as most of the wealth is held by a small fraction of population with a high EIS. In this model limited participation in the stock market creates substantial wealth inequality similar to the distribution of wealth between the stockholders and the non-stockholders in the data. (This is in contrast to the majority of dynastic models which are known to generate little wealth dispersion, and hence suggests limited participation as an important factor in understanding distributional issues.) Consequently, the properties of output and investment are almost entirely determined by the high-elasticity stockholders, whereas consumption is strongly affected by the inelastic non-stockholders who contribute substantially. This dichotomy explains how the preferences of wealthy are hardly detectable in consumption data but still strongly influence the economy's aggregates.

The cross-sectional richness of the model allowed us to address related questions. For example, if the low EIS estimates of the poor were indeed due to severe borrowing restrictions instead of genuine preference heterogeneity (as argued by some papers), such a model would have implications inconsistent with cross-sectional and asset price data. On the other hand, the current model explains those facts when preference heterogeneity is properly acknowledged, providing further support to the thesis the poor have low EIS.

The idea that aggregates can be better explained by the interaction of heterogeneous agents rather than by a representative agent's intertemporal problem with a variety of frictions has important consequences. For example, in a representative-agent economy aggregate consumption and savings are determined by the very same preferences, whereas in the current model, they are significantly different from one another. We conclude that, as a result, economic analyses as well as policy discussions based on average elasticities may be seriously misguided.

## Appendix A. Policy implications

In the foregoing analysis our goal has been to demonstrate the interaction between wealth inequality and heterogeneity in the EIS from a positive perspective. We now conduct a policy experiment to demonstrate that one can reach misleading policy conclusions if this heterogeneity is ignored.

---

*(footnote continued)*

higher than 0.4. To save space, these further results are reported in an appendix available from the author's website.

It has long been recognized in the public finance literature that the welfare effects of capital income taxation critically depend on the degree of intertemporal substitution (Summers, 1981; King and Rebelo, 1990). Indeed, Hall (1988) concludes that his estimate of a small EIS also imply a weak response of savings to changes in interest rates. To the contrary, we argue that the effect of taxation on savings will be determined by the wealth-weighted average elasticity measure, which is—given the enormous wealth inequality—very close to that of the stockholders.

In order to demonstrate this point, we study a simple tax reform problem similar to the one studied by Lucas (1990). Imagine that initially the government imposes a flat-rate tax on capital income and returns the proceeds to households in a lump-sum fashion. At a certain date, capital income tax is completely eliminated and agents have not previously anticipated it. The initial tax rate is set equal to 36% which roughly corresponds to the average rate in the U.S. All aspects of the baseline model remain intact.

In order to provide a comparison to the existing literature, we first consider the welfare gain from this reform in a representative-agent framework. If the agent has  $\rho = 1.0$ , the welfare benefit of this policy is 0.93% of consumption per period—taking the transition path into account. Although it may not seem much, as Lucas argues, this is about 20 times the gain from eliminating the business cycle fluctuations, and two times the gain from eliminating 10% inflation rate. However, if we assume that the agent has  $\rho = 0.1$ , the welfare gain is reduced to 0.4% of consumption instead, mainly because now the transition takes about 110 years compared to about 23 years in the former case.

Now suppose that the limited participation economy is subjected to the same tax experiment. The welfare gain is 0.82% of total consumption.<sup>16</sup> In effect, this economy behaves as if it was populated only by agents with unit elasticity and non-stockholders' preferences virtually vanished from the problem.

There is an even more interesting side to this problem that transpires from explicitly modeling heterogeneity: based on policy experiments like the one above, some economists have argued in favor of eliminating capital income taxes. But, a representative-agent framework masks the question of “who gains and who loses from this reform?” In reality, all agents are not identical, and as we have shown so far, in some dimensions, they differ substantially. So, it is compelling to take this question seriously and break down the gains from this reform. It turns out that, in consumption terms, the stockholders gain by 5.4%, whereas the non-stockholders, who constitute 70% of the population, actually lose 2.1% of their consumption. Clearly, this is a different conclusion than what comes out from the representative-agent economy.

## References

- Aiyagari, R.S., 1994. Uninsured idiosyncratic shock and aggregate saving. *Quarterly Journal of Economics* 109, 659–684.
- Attanasio, O.P., Browning, M., 1995. Consumption over the life cycle and over the business cycle. *American Economic Review* 85, 1118–1137.
- Attanasio, O.P., Weber, G., 1989. Intertemporal substitution, risk aversion and the Euler equation for consumption. *The Economic Journal* 99, 59–73.

---

<sup>16</sup>This calculation is based on the assumption that there is a utilitarian government which tries to attain the same social welfare index as without taxes and makes transfers to agents in such a way to minimize the total amount of transfers.

- Attanasio, O.P., Banks, J., Tanner, S., 2002. Asset holding and consumption volatility. *Journal of Political Economy* 110, 771–793.
- Basak, S., Cuoco, D., 1998. An equilibrium model with restricted stock market participation. *Review of Financial Studies* 11, 309–341.
- Basu, S., Kimball, M., 2000. Long-run labor supply and the elasticity of the intertemporal substitution for consumption. Working Paper, University of Michigan.
- Blundell, R., Browning, M., Meghir, C., 1994. Consumer demand and the life-cycle allocation of household expenditures. *Review of Economic Studies* 61, 57–80.
- Browning, M., Crossley, T., 2000. Luxuries are easier to postpone: a proof. *Journal of Political Economy* 108, 1022–1026.
- Browning, M., Hansen, L.P., Heckman, J., 1999. Micro data and general equilibrium models. in: Taylor, J., Woodford, M. (Eds.), *Handbook of Macroeconomics*, vol. 1A. North-Holland, Amsterdam.
- Campbell, J.Y., Cochrane, J.H., 1999. By force of habit: a consumption based explanation of aggregate stock market behavior. *Journal of Political Economy* 107, 205–251.
- Campbell, J.Y., Mankiw, N.G., 1989. Consumption, income, and interest rates: reinterpreting the time series evidence. *NBER Macroeconomics Annual*. MIT Press, Cambridge, MA.
- Chamberlain, G., Wilson, C., 2000. Optimal intertemporal consumption under uncertainty. *Review of Economic Dynamics* 3, 365–395.
- Cooley, T., Prescott, E., 1995. Economic growth and business cycles. In: Cooley, T. (Ed.), *Frontiers of Business Cycle Research*. Princeton University Press, Princeton.
- Epstein, L.G., Zin, S.E., 1989. Substitution, risk aversion, and the temporal behavior of consumption and asset returns: a theoretical framework. *Econometrica* 57, 937–969.
- Guvenen, F., 2004. A parsimonious macroeconomic model for asset pricing: habit formation and cross-sectional heterogeneity? Working Paper, University of Rochester.
- Hall, R.E., 1988. Intertemporal substitution in consumption. *Journal of Political Economy* 96, 339–357.
- Hansen, L.P., Singleton, K.J., 1983. Stochastic consumption, risk aversion, and the temporal behavior of asset returns. *Journal of Political Economy* 91, 249–265.
- Investment Company Institute and the Securities Industry Association, 2002. *Equity Ownership in America*, downloadable at (<http://www.ici.org/pdf/>).
- Jones, L.E., Manuelli, R., Stacchetti, E., 1999. Technology (and policy) shocks in models of endogenous growth. Working Paper No. 7063, National Bureau of Economic Research.
- Jones, L.E., Manuelli, R., Siu, H., 2000. Growth and business cycles. Working Paper No. 7633, National Bureau of Economic Research.
- King, R.G., Rebelo, S.T., 1990. Public policy and economic growth: developing neoclassical implications. *Journal of Political Economy* 98, S127–S150.
- Krusell, P., Smith Jr., A.A., 1998. Income and wealth heterogeneity in the macroeconomy. *Journal of Political Economy* 106, 867–896.
- Laibson, D., Repetto, A., Tobacman, J., 1998. Self-control and saving for retirement. *Brooking Papers on Economic Activity*, pp. 91–172.
- Lucas Jr., R.E., 1990. Supply side economics: an analytical review. *Oxford Economic Papers* 42, 293–316.
- Mankiw, G.N., Zeldes, S., 1991. The consumption of stockholders and non-stockholders. *Journal of Financial Economics* 29, 97–112.
- Ogaki, M., Reinhart, C., 1998. Measuring intertemporal substitution: the role of durable goods. *Journal of Political Economy* 106, 1078–1098.
- Poterba, J., Samwick, A.A., 1995. Stock ownership patterns, stock market fluctuations, and consumption. *Brookings Papers on Economic Activity*, pp. 295–373.
- Summers, L.H., 1981. Tax policy, the rate of return, and savings. Manuscript No. 995, National Bureau of Economic Research.
- Vissing-Jørgensen, A., 1998. Limited asset market participation and the elasticity of intertemporal substitution. Working Paper, University of Chicago.
- Weil, P., 1989. The equity premium puzzle and the risk-free rate puzzle. *Journal of Monetary Economics* 24, 401–421.
- Wolff, E., 2000. Recent trends in wealth ownership, 1983–1998. *Levy Institute Working Paper*, No. 300.