

Rochester Center for

Economic Research

**Reconciling Conflicting Evidence on the Elasticity of
Intertemporal Substitution: A Macroeconomic Perspective**

M. Fatih Guvenen

Working Paper No. 491
May 2002

UNIVERSITY OF
ROCHESTER

Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective*

Fatih Guvenen[†]

January 20, 2003.

Abstract

This paper attempts to reconcile two opposing views about the elasticity of intertemporal substitution in consumption (EIS), a parameter that plays a key role in macroeconomic analysis. On the one hand, empirical studies using aggregate consumption data typically find that the EIS is close to zero (Hall, 1988). On the other hand, calibrated macroeconomic models designed to match growth and business cycle facts typically require that the EIS be close to one (Weil, 1989; Lucas, 1990). We show that this apparent contradiction arises from ignoring two kinds of heterogeneity across individuals. First, a large fraction of U.S. households do not participate in stock markets. Second, a variety of microeconomic studies using individual-level data conclude that an individual's EIS increases with his wealth. We study a dynamic macroeconomic model featuring these two realistic sources of heterogeneity which have been largely assumed away in macroeconomics to date. We find that limited participation creates substantial wealth inequality matching that in U.S. data. Consequently, the properties of aggregate variables directly linked to wealth, such as investment and output, are almost entirely determined by the (high-elasticity) stockholders. At the same time, since consumption is much more evenly distributed across households than is wealth, estimation using aggregate consumption uncovers the low EIS of the majority of households (i.e., the poor).

JEL Classification: E32, E44, E62

Keywords: Aggregate fluctuations, incomplete markets, wealth inequality, the elasticity of intertemporal substitution.

*For helpful conversations and comments, I thank Daron Acemoglu, Jeremy Greenwood, Lars Hansen, Per Krusell, Torsten Persson, David Laibson, Tony Smith, Nick Souleles, Stan Zin, and my discussant at the NBER Summer Institute, Robert Hall as well as the seminar participants at M.I.T, Chicago GSB, Pennsylvania, Rochester, UCSD, Carnegie-Mellon, Penn State, Queens, IIES at Stockholm University, Richmond FED, the 2001 Minnesota Macroeconomics Workshop, the NBER Economic Fluctuations Research Meeting in Boston, 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."

[†]Address: Department of Economics, 216 Harkness Hall, University of Rochester, Rochester, NY, 14627. E-mail: guvn@mail.rochester.edu. Phone: (585) 275-3968.

1 Introduction

A rational agent will seize intertemporal trade opportunities revealed in asset prices by adjusting his consumption growth, by an amount negatively related to his (counteracting) desire for a smooth consumption profile. The degree of this consumption growth response is called the elasticity of intertemporal substitution in consumption (EIS). The goal of this paper is to use dynamic macroeconomic analysis to reconcile seemingly contradictory evidence about the value of this parameter.

Research in many fields of macroeconomics has established the EIS as crucial for many questions ranging from government policy to the determinants of long-run growth. For example, the macroeconomic effects of capital income taxation critically depend on the magnitude of the EIS (see, for example, Summers, 1981; and King and Rebelo, 1990). In another line of research, Jones, Manuelli and Stachetti (1999) conclude that whether uncertainty boosts or slows down economic growth is determined, again, by the degree of intertemporal substitution. Finally, the effectiveness of the interest rate, commonly used by central banks as a monetary policy instrument also hinges on households' willingness to substitute consumption across time. These are just a few examples which demonstrate the significance of this parameter for economic analysis.

Given this substantial role, much effort has been devoted to accurately pin down the value of the EIS. On the one hand, macroeconomists generally use a large value, reflecting a common view that a high degree of intertemporal substitution is more consistent with aggregate data viewed through the lens of dynamic macroeconomic models. For example, in their seminal paper, Kydland and Prescott (1982) calibrated it to 0.66 and Lucas (1990) argued that even an elasticity of 0.5 appears too low when confronted with macro data. For a long time most of the real business cycle literature used a value around unity. In the next section we review some economic arguments which speak in favor of this practice.

An alternative approach is to directly estimate the EIS by focusing on consumption data, a method that received a lot of attention after the development of the rational expectations models of consumption behavior by Hall (1978), Grossman and Shiller (1981) and Hansen and Singleton (1982). 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, intertemporal elasticity is, in fact, very close to zero. The subsequent empirical macro literature, by and large, has confirmed his findings and provided further support (Campbell and Mankiw, 1989; and Patterson and Pesaran, 1992). Although this result has been challenged by some studies using micro data, the verdict for many macroeconomists seems to be that the average elasticity of substitution is close to zero.¹ Thus, there is an apparent contradiction between the macroeconomics literature and econometric studies which both use the same aggregate data.

¹A short list of these micro studies includes Attanasio and Weber (1993), Atkeson and Ogaki (1996), and Beaudry and Van Wincoop (1996) who use state-level data.

In this paper we offer a possible resolution to this puzzle. We argue that this apparent inconsistency is largely a consequence of the “representative-agent” perspective widely adopted in both literatures. More specifically, we show that the dynamics of different aggregates are disproportionately influenced by different groups in the economy, inconsistent with the representative-agent assumption.

We study a dynamic macroeconomic model which features two key sources of heterogeneity: limited participation in the stock market and heterogeneity in the intertemporal elasticity parameter. A large body of empirical evidence will be presented in the next section documenting these facts. Specifically, we consider an economy with neoclassical production and competitive markets. There are two types of agents. The majority of households (first type) do not participate in the stock market where claims to risky output are traded. However, a risk-free bond is available to all households, so non-stockholders can also accumulate wealth and smooth consumption intertemporally. Consistent with empirical evidence, we assume that stockholders have higher elasticity of substitution (around 1.0) than non-stockholders (around 0.1). This assumption is rather minor and this heterogeneity can also be generated endogenously within the model; see Section 9.

In this economy, the asymmetry in investment opportunities 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 investment, capital and output, are almost entirely determined by the (high-elasticity) stockholders who own virtually all the wealth. On the other hand, consumption turns out to be much more evenly distributed across households, as it is in the U.S. data, so aggregate consumption (and hence Euler equation estimations) mainly reveal the low elasticity of the majority, that is, the poor. In other words, there may not be a parameter that can properly be called “*the* elasticity of intertemporal substitution.” This asymmetry in the determination of macroeconomic aggregates explains how these two strands of literature viewed entirely different pictures of the same macroeconomy by concentrating on different aggregates implicitly assuming a representative-agent. This is the main result of the paper.

This explanation clearly relies on the premise that the average investor is very different than the average consumer. In Section 5.1, we present the joint distribution of consumption and wealth in the U.S. and document the remarkably different concentrations of these two variables in the population: basically, the top 20 percent of the population own 88 percent of all capital and land, but contribute less than 30 percent of aggregate consumption, whereas the lower 80 percent own only 12 percent of wealth but account for the remaining 70 percent of consumption (see Fig 1 in Section 5).

As noted above, our motivation for modeling heterogeneity in the EIS comes from a large empirical literature which we discuss in Section 2. However, some researchers suggested that frequently binding borrowing constraints may cause the poor to *appear* to have a low EIS even when their elasticity is as high as that of the rich (Laibson et. al 1998). In Section 5.2 we show that a general equilibrium model incorporating this feature (identical preferences and frequently binding constraints) has counterfactual implications for some key statistics as well as for asset

prices. On the other hand, our baseline model is able to explain those features of the data quite easily—as long as the poor are assumed to have a low EIS—suggesting that they indeed must have a low elasticity (Section 5.3).

A natural question is whether we should care about this “elasticity puzzle” for policy analysis. In the context of a well-known capital income taxation problem we show that the commonly used representative-agent model calibrated to the average EIS estimates yields substantially misleading conclusions (Section 7). Finally, in Section 9 we discuss how heterogeneity in the EIS can be generated as an endogenous outcome of the model by using a non-homothetic utility function. We also address issues related to endogenizing limited participation and the robustness of the results to such an extension.

This paper is related to two strands of literature. First, starting with an influential paper by Mankiw and Zeldes (1991), a literature in finance attempts to explain asset pricing puzzles by taking into account the well-documented limited stock market participation phenomenon. Basically, stockholding is extremely concentrated: until 1990s, less than 30 percent of households owned all the equity in the U.S. Even after the participation boom of the last decade half of all the stocks is held by 1 percent of the population.² Taking this observation as a starting point, a number of papers including Saito (1995), and Basak and Cuoco (1998) study these financial puzzles in a general equilibrium framework featuring stockholders and non-stockholders but abstract from production.³ 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 this new dimension of limited participation: we are interested in the determination of aggregate dynamics in this environment and apply the findings to resolve the elasticity puzzle. Two related macro papers are by Abel (2001) and Diamond and Geanakoplos (2001) who study policy issues in similar frameworks. They find that a government policy of investing social security funds in the stock market has significantly different consequences when there is limited participation.

A second related literature asks if market incompleteness matters quantitatively (Telmer, 1993; Rios-Rull, 1994; Krusell and Smith, 1998; and Levine and Zame, 2001). The main conclusion is that, under certain assumptions,⁴ it does not. So, in contrast to what one might initially think, appealing to heterogeneity does not always add extra richness to the model; the implications remain very close to that of the—complete markets—representative agent world. The interesting fact is that our model satisfies all those assumptions, so a priori, one would expect a similar

²For detailed information about stock market participation and recent trends, see the Investment Institute Company (2002) report.

³There are also a number of econometric studies, including Ait-Sahalia, Parker, and Yogo (2001), Attanasio, Banks and Tanner (2002), Brav, Constantinides and Geczy (2002), and Vissing-Jorgensen (2002) who estimate the Euler equations for the two groups separately and find that there is less evidence of asset pricing anomalies when data on stockholders is used.

⁴In calibrated examples, the common assumptions are: a long time horizon, stationary incomes and a risk-free bond available to all households. For further conditions and a discussion, see Section 6 and footnote 15.

conclusion to hold here. In this sense, the substantial wealth inequality implied by the model comes as a surprise. The novel feature here is that limited participation creates, what we call, “*asymmetric market incompleteness*” where one group of agents has extra insurance opportunities compared to the rest of the economy. This asymmetry has more striking consequences here compared to the exchange economy models mentioned above where aggregates are exogenously imposed. In this case market incompleteness does matter substantially. Moreover, many economic environments seem to be characterized by such an asymmetry suggesting that the model studied here can be applied to analyze a range of interesting economic questions. An example would be the interaction between wealthy countries and less developed countries with substantial differences in formal insurance opportunities.

2 Is there a puzzle?

In this section, we review the different views on the value of the EIS. We want to highlight two sides of the problem. First we ask, Do we really know the extent of intertemporal substitution in consumption? Second, and maybe more importantly, Do we really care about this parameter?

As we briefly mentioned above, a number of growth and fluctuations facts suggest a high degree of intertemporal substitution. For example, consider the familiar expression for the interest rate obtained by rearranging the consumption Euler equation:

$$R_t^f = \eta + \frac{1}{\rho} \log \left(\frac{C_{t+1}}{C_t} \right), \quad (1)$$

where η is the time preference rate, and ρ is the EIS parameter. For convenience we abstract from uncertainty. Given that the annual per-capita consumption growth in the U.S. is about 2 percent, an elasticity of 0.1 as estimated by Hall (1988) together with $\eta > 0$ imply a lower bound of 20 percent for the real interest rate! This observation is well-known as Weil’s (1989) risk-free rate puzzle. Alternatively we can invert the above equation to get a sense about a plausible value for ρ . Taking the average interest rate to be 3 percent (commercial paper rate), and a consumption growth of 2 percent requires the EIS to be at least 0.66 if η is to be positive.

In fact, a similar observation has led Lucas (1990) to rule out an elasticity below 0.5 as implausible (note that in his notation $\sigma = 1/\rho$):

If two countries have consumption growth rates differing by one percentage point, their interest rates must differ by σ percentage points (assuming similar time discount rates η). A value of σ 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.

where $\sigma = 2$ implies $\rho = 0.5$. This argument is robust to the introduction of uncertainty and

Table 1: BUSINESS CYCLE STATISTICS FOR DIFFERENT VALUES OF THE EIS

	EIS	$\sigma(Y)$	$\sigma(I_K)$	$\sigma(n)$	$\sigma(Y/n)$	$\sigma(C + I_K)$
US Data	—	2.60	6.05	2.00	1.24	2.00
Case						
1	1.11	3.91	6.39	3.54	1.47	3.24
2	1.00	3.03	4.58	2.19	1.24	2.60
3	0.80	2.29	3.01	1.07	1.33	2.07
4	0.66	2.03	2.56	0.66	1.42	1.89
5	0.57	1.90	2.28	0.45	1.47	1.80
6	0.50	1.82	2.11	0.33	1.51	1.74
7	0.40	1.73	1.91	0.18	1.55	1.69
8	0.33	1.67	1.79	0.10	1.58	1.65

The statistics are from the endogenous growth model of Jones, Manuelli and Siu (1999) reproduced from Table A3. The raw data is logged and detrended with a Hodrick-Prescott filter before the statistics are calculated. Each column reports the percentage standard deviation of the variable given in the top cell. Y denotes output, I_K is investment in physical capital, n is labor hours.

allowing for non-time-separability, because an equation very similar to (1) can still be derived under these conditions; see, for example, Attanasio and Weber (1989).⁵

As a second example, Jones, Manuelli, and Siu (2000) study business cycle fluctuations both in a neoclassical and in an endogenous growth framework. They focus on the effects of changing the EIS parameter on various measures of model performance such as the second-order moments of output, investment, consumption, hours, etc. For completeness we provide a summary of their results in Table 1. The moments from the U.S. data are reported in the top row and the model counterparts for different elasticity values are below. In each column, the values from the model which bracket the empirical counterpart are highlighted. As can be seen, for four out of five statistics an elasticity around unity seems to match the data the best. Consumption volatility seems to suggest a slightly lower EIS (and would imply an even lower elasticity if investment was left out) which we comment on later. Overall, though, comparing a large number of additional moments (including cross- and auto-correlations) they find that the endogenous growth model matches the data best when $\rho \approx 1$ and similarly the neoclassical model matches the data best when $0.8 < \rho < 1$.

For many macroeconomists economic reasoning like those above constitute strong, albeit indirect, evidence that the EIS is quite high, probably close to unity.

Note however that many of these inferences are based on unconditional information. For example, the risk-free rate puzzle is a statement about the *average* growth rate facts: that is, the

⁵In fact, with uncertainty the Euler equation is a parabola in the EIS and thus it is possible to obtain a low risk-free rate by assuming an extremely small elasticity of substitution. However, this fix gives rise to another problem: the interest rate becomes extremely volatile and reasonable values of the interest rate are achieved only as a knife-edge case.

average growth rate of consumption seems inconsistent with a low average interest rate if individuals are really very averse to substitution. However, one can potentially gain further insight by studying the conditional information embedded in the same Euler equation: the co-movement between consumption growth and the interest rate also identifies the elasticity parameter. As persuasive as the previous arguments may seem, this latter approach delivers a dramatically different result.

Consistent with the widely employed representative-agent approach, most of the empirical macro studies estimated a single elasticity parameter from aggregate consumption data, which is an average of individual EISs. These studies have, by and large, found the EIS to be very close to zero. For example, Campbell and Mankiw (1989) extended Hall's framework and obtained the same result, whereas Patterson and Pesaran (1992) reported similar estimates from the U.K. The following quote from Hall (1988) concisely states the conclusion that emerged from that literature:

All the estimates presented in this paper of the intertemporal elasticity of substitution are small. Most of them are also quite precise, supporting the strong conclusion that the elasticity is unlikely to be much above 0.1, and may well be zero.

In fact, many of the earlier econometric studies also obtained small—although not this small—values for the EIS. Davies (1981) reviews this earlier literature and concludes that a best guess is $\rho = 0.25$. Thus, one can say that econometric studies are, overall, in favor of a low degree of intertemporal substitution.⁶ Given that a plausible range for this parameter is probably $(0, 1)$ it is clear that the conclusions reached by studying the growth and fluctuations facts described above and by focusing on the co-movement between consumption and returns are at sharp contrast with each other.

This apparent contradiction has major consequences for normative questions. Among many policy issues that crucially depend on the magnitude of this parameter, we discuss the popular capital income taxation problem. For example, King and Rebelo (1990) find that, with a 10 percent increase in factor taxation, the growth rate of an economy falls from 2 percent all the way down to 0.48 percent accompanied by a substantial fall in consumption if $\rho = 1.0$ is assumed. In contrast, when $\rho = 0.2$, the same tax experiment reduces the growth rate from 2 percent only to 1.69 percent. This is a strikingly different picture than the one above. In a different model, Jones, Manuelli, and Rossi (1993) show that even changing the EIS slightly, from 0.66 to 0.4, reduces the welfare gain from switching to a Ramsey tax system from 37 percent down to 9 percent of consumption. These examples clearly show that a resolution to the elasticity puzzle is critical for satisfactory analyses of many policy questions.

⁶There are a number of exceptions, such as Hansen and Singleton (1982), Summers (1981), and Mankiw, Rotemberg and Summers (1985) who obtained estimates around 1.0. But Hall argued that their estimates were biased upward because of time aggregation. Mulligan (2001) recently argued that a high EIS is consistent with consumption fluctuations if the relevant measure of returns is the return on capital. Also, using micro data Attanasio and Weber (1993), and Blundell, Browning and Meghir (1994) have obtained larger values for the EIS.

2.1 What is the Explanation?

One feature common to all the arguments above is their reliance on the representative-agent assumption. In principle, aggregate consumption data yields information about an elasticity measure closer to the *average consumer's* preferences whereas investment and output data reveals (loosely speaking) the elasticity of the *average investor*. The representative-agent assumption implies that the average consumer and the average investor are the same and thus different macroeconomic time-series should yield comparable estimates of the EIS. To the extent that this assumption is violated it is entirely possible for capital and consumption fluctuations to imply very different values for the elasticity of substitution.

We emphasize two kinds of heterogeneity across households to reconcile the apparently inconsistent evidence on the elasticity of substitution. First, a well-known fact is that there is substantial wealth inequality in the U.S.: basically, 90 percent of all net worth⁷ is owned by 30 percent of the population (Table 3 presents the joint distribution of consumption and wealth). These wealthy households also hold 95 percent of productive wealth (net worth minus equity in owner-occupied housing) and 99 percent of all the equity. The concentration of capital is even more striking at the very top: the top 1 percent of the wealth distribution own 48 percent of all productive wealth. Naturally, then, the preferences of this small group will have a disproportionate impact on the properties of investment and capital (and, hence, of output). On the other hand this group's contribution to total consumption is much more modest: the top 1 percent wealthy account for less than 4 percent of aggregate consumption, and those in the top 10 percent contribute around 17 percent of the total. Thus the preferences of the wealthy are largely not revealed in consumption regressions.

Second, there are theoretical reasons which suggest that an individual's EIS is increasing with his wealth level. For example, if households consume a variety of goods with different income elasticities, the intertemporal elasticity of total consumption will be higher for wealthier households. This is because at lower wealth levels the share of necessity goods in a household's total consumption is large and the household is less willing to substitute these necessities across time. On the other hand, the consumption of a wealthy household will have a larger fraction of luxury goods, which are more easily substituted across time (see Browning and Crossley, 2000, for a proof of this statement.) Alternatively, non-homotheticity in preferences due to subsistence requirements, habit formation and so on also implies a higher EIS for the wealthy.

⁷Throughout the paper we adopt three definitions of wealth. The most comprehensive measure is "net worth" and is defined as the current value of all marketable or fungible assets less the current value of debts. Thus 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 pension plans, including IRAs, Keogh and 401(k) plans; (8) corporate stocks and mutual funds; (9) net equity in unincorporated businesses; (10) equity in trust funds. From the sum of these assets we subtract consumer debt including auto loans and others. The second definition of wealth is what we call "productive wealth" for the lack of a better term, and is intended to include all components of physical wealth that are productive. It is calculated as net worth minus equity in owner occupied housing. The narrowest definition is "financial wealth," and is the sum of (3) through (8).

Indeed, there is a fairly large and growing body of empirical work documenting this kind of heterogeneity. 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, Browning and Meghir (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. 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, Banks and Tanner (2002) use the Family Expenditure Survey dataset from the U.K. and consistently obtain elasticity values around 1 for stockholders, and between 0.1 to 0.2 for non-stockholders. Vissing-Jorgensen (1998) finds very similar estimates from the Consumption Expenditure Survey data on U.S. households.⁸

Finally, survey evidence also points in this direction. Barsky, et. al. (1997) recover preference parameters from survey data and also show that they are economically meaningful, in the sense that these parameters predict behavior in line with theory. They find that the population average for the elasticity is 0.2, with only less than 20 percent of households above 0.3. But a small group of individuals have elasticity even higher than 1.0.

The empirical evidence thus suggests that the majority of individuals display very low elasticity of substitution (around 0.1 – 0.2) but a wealthy stockholding minority are much more elastic than that (around 1.0). If we put these two pieces—wealth inequality and heterogeneity in the EIS— together, the following picture emerges. There is a small group of wealthy households who have significantly higher EIS than the rest. Consequently, macroeconomic aggregates directly linked to wealth, such as investment, saving, and output are almost entirely determined by the high elasticity of these stockholders who own almost all the capital in the economy. On the other hand, unlike wealth, consumption is quite evenly distributed with a Gini coefficient of 0.32 (compared to 0.82 for net worth and 0.90 for financial wealth).⁹ Hence, aggregate consumption data reveals mainly the low elasticity of the poor who contribute substantially. In the rest of the paper we will make

⁸The presented evidence raises some natural questions: Is the EIS also increasing over time since there is a secular increase in wealth? or, Do richer countries have higher elasticities than the poor ones? Note that, though there is some mild evidence in favor of both implications (see Atkeson and Ogaki 1996), neither one is necessarily implied by the presented results and would depend on the exact specification of preferences. We present one example of a utility function in Section 9 where agents’ utility depend on their own consumption as well as the average consumption in that country. With this specification neither of the implications above has to be true.

⁹c.f., Castañeda, Díaz-Giménez, and Ríos-Rull (2002).

this argument more rigorous and analyze its ramifications for the working of the macroeconomy.

3 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 necessary features discussed above. We consider an economy populated by two types of agents who live forever. The population is constant and is normalized to unity. Let λ ($0 < \lambda < 1$) denote the measure of the first type of agents (who will be called “stockholders” later) in total population.

PREFERENCES

Both agents value temporal consumption lotteries according to the following Epstein-Zin (1989) recursive utility function:

$$U_t^i = \left[(1 - \beta) (C_t)^{\varphi^i} + \beta \left(E_t (U_{t+1}^i)^{1-\alpha} \right)^{\frac{\varphi^i}{1-\alpha}} \right]^{\frac{1}{\varphi^i}} \quad 0 \neq (1 - \alpha), \varphi^i < 1 \quad (2)$$

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 $\beta = (\frac{1}{1+\eta})$, and α 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. We should note that this assumption is rather minor and it is possible to assume identical (non-homothetic) preferences for both agents and create heterogeneity in the elasticities endogenously in the model. This point will become clearer after we present the results. See Section 9 for further discussion of this point.

Remark: It is important to stress that our choice of recursive preferences is for clarity. It does not appreciably affect the outcome of the model, but by disentangling elasticity from risk aversion, it will make the exposition more transparent. Also, note that these preferences nest expected utility as a special case: when $\alpha = \varphi$, we obtain the familiar CRRA (constant relative risk aversion) expected utility function.

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 the familiar 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. The firm’s problem can simply be expressed as a static decision, and can be decomposed into a series of one-period profit maximization problems:

$$\underset{K_t, L_t}{Max} \quad \left[Z_t K_t^\theta L_t^{1-\theta} - (R_t^S + \delta) K_t - W_t L_t \right]$$

where R_t^S and W_t are the market return on capital and wage rate respectively, and δ is the

depreciation rate of capital.

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

$$\begin{aligned} R_t^S &= \theta Z_t (K_t/L_t)^{\theta-1} - \delta \\ W_t &= (1 - \theta) Z_t (K_t/L_t)^\theta. \end{aligned} \tag{3}$$

STOCKHOLDERS AND NON-STOCKHOLDERS

Both agents have one unit of time endowment in each period, which they supply inelastically to the firm. Besides the productive capital asset there is also a one-period risk-less household bond which is in zero net supply that is traded in this economy. The crucial difference between the two groups is in their investment opportunity sets: “non-stockholders” can freely trade in the bond, but as their name suggests, they are restricted from participating in the capital market. “Stockholders,” on the other hand, have access to both markets and hence are the sole capital owners in the economy. Finally, we impose portfolio constraints as a convenient way to prevent Ponzi schemes. For the main results of the paper, including the wealth inequality and the resolution of the EIS puzzle, these constraints need not even be binding; they can be as loose as possible.

The timing of events is as follows: each period starts with production; agents are paid their wages and asset returns are realized after production takes place. Then consumption and portfolio choice decisions are made and asset trading is carried out. Finally, consumption takes place and the period ends. Before we move to agents’ problem a final remark is in order.

Remark: Limited participation in the stock market is introduced exogenously here as in most of the existing literature. However, it is important to discuss the plausibility of this assumption and to understand how endogenizing it would affect the workings of the model. We address these issues in Section 9. We feel that for the purposes of this paper this is a reasonable assumption and the main conclusion is not likely to be overturned by this extension.

AGENTS’ DYNAMIC PROBLEM AND THE EQUILIBRIUM

To state the individual’s problem recursively, we need to specify the aggregate state-space for this economy. The Markov characteristic of the exogenous driving force naturally suggests concentrating on equilibria which 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 non-stockholders *after* production, B , and the technology level, Z . Let us denote the financial wealth of an agent in the current period by ω where we suppress superscripts for clarity of notation. Given the recursivity of the utility and the stationarity of the environment, maximization of (2) for the stockholders can be expressed as

the solution to the following dynamic programming problem:

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

where the expectation is conditional on the set Ω containing all the information at the time of decision, and b' and s' denote bond and stock (capital) choice of the agent respectively. The endogenous functions Γ_K and Γ_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 $V^i(\omega^i; K, B, Z)$, ($i = h, n$), bond holding decision rules for each agent $b^i(\omega^i; K, B, Z)$, stockholding decision for the stockholder, $s(\omega^h; K, B, Z)$, a bond pricing function, $q(K, B, Z)$, competitive factor prices, $R^S(K, Z)$, $W(K, Z)$, and laws of motion for aggregate capital and aggregate bond holdings of non-stockholders, $\Gamma_K(K, B, Z)$, $\Gamma_B(K, B, Z)$, 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 (equation (3) is satisfied)
- 3) Bond market clears: $\lambda b^h(\varpi^h; K, B, Z) + (1 - \lambda) b^n(\varpi^n; K, B, Z) = 0$, where ϖ^i denote the aggregate wealth of a given group; and labor market clears: $L = \lambda * 1 + (1 - \lambda) * 1 = 1$
- 4) Aggregates result from individual behavior:

$$K_{t+1} = \lambda s(\varpi^h, K_t, B_t, Z_t) \tag{4}$$

$$B_{t+1} = (1 - \lambda) b^n(\varpi^n, K_t, B_t, Z_t) \tag{5}$$

4 Numerical Solution and Calibration

Since an analytical solution is not possible, we solve for the equilibrium numerically. The equilibrium is solved globally—not as an approximation around a stochastic steady state—which allows us to perform transition experiments (as in Section 9) easily. To our knowledge, this is also the first attempt at numerically solving a dynamic programming problem with general recursive utility. Appendix A contains the algorithm as well as a discussion of the accuracy of the solution.

BASELINE PARAMETERIZATION

Following the tradition of the business cycle literature, we calibrate model parameters to replicate some 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 process¹⁰ with transition probabilities $\pi_{ij} = P(Z_{t+1} = j \mid Z_t = i)$ chosen such that business cycles are symmetric and last for 6 years. This condition implies $\pi_{11} = \pi_{22} = 2/3$. The percentage standard deviation, $\sigma(Z)$, is set equal to 3.1 percent which would be the (unconditional) variability in the yearly aggregate shock if the quarterly Solow residuals are assumed to follow an AR(1) process with persistence parameter of 0.95 and coefficient of variation of 1 percent. These numbers are quite standard and we also investigate the sensitivity of the results to alternative calibrations.

Participation rates: The stock market participation rate has gradually increased from around 5 percent in the 1950s to approximately 19 percent in 1982 (Survey of Consumer Finances, 1982). In the last decade, for many reasons ranging from the emergence of mutual funds and reduced costs of (on-line) trading to the retirement saving by baby-boomers, this trend has accelerated in the U.S. as in the rest of the world. As a result, in 2002 the stockholding rate has reached almost 50 percent.¹¹ Since in this paper, we are studying a stationary economy, we want to focus on the pre-1990s U.S. economy. Thus, the fraction of stockholders, λ , is set equal to the average participation rate in the 1980s, which is 30 percent. This choice is also motivated by data availability: most of the cross-sectional statistics that we use in this paper are calculated from the PSID (Panel Study of Income Dynamics) and the Survey of Consumer Finances for a period which centers around the 1980s.

Borrowing constraints are harder to measure and calibrate. We want to choose these bounds to reflect the fact that stockholders can potentially accumulate capital which can then be used as collateral for borrowing in the risk-free asset, whereas non-stockholders have to pay all their debt through future wages. For the baseline case, we allow the stockholders to borrow in bonds up to four years of expected labor income ($\underline{B}^h = 4 * E(W)$). As for non-stockholders, we calibrate their borrowing limit to 30 percent of one year's expected income, which is the average credit limit most short-term creditors, such as credit card companies, impose. Again, as will become clear in the next section, these constraints can be relaxed significantly without changing the main message of the paper.

Finally we specify the preference parameters. The subjective discount factor, β , 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).

¹⁰In a companion paper, Guvenen (2002), we calibrate the model to quarterly data and allow the technology shock to follow an AR(1) process with a quarterly persistence of 0.95 as in the RBC literature. The findings reported in that paper are very similar to what we obtain here, so we prefer the simpler Markov specification here.

¹¹In terms of wealth-weighted participation rates, participation boom is less pronounced because old stockholders still own most of the equity outstanding (see the "Equity Ownership in America" report by the Investment Company Institute, 2002).

Table 2: BASELINE PARAMETRIZATION

Yearly Model		
Parameter		Value
β	Time discount rate	0.96
α^h	Risk aversion of stockholders	3
α^n	Risk aversion of non-stockholders	3
ρ^h	EIS of stockholders	1
ρ^n	EIS of non-stockholders	0.1
λ	Participation rate	0.3
π_{11}	Probability of good state good state	2/3
π_{22}	Probability of bad state bad state	2/3
σ_ε	Standard deviation of shock	0.031
θ	Capital share	0.4
δ	Depreciation rate	0.08
\underline{B}^h	Borrowing limit of stockholders	$4\overline{W}$
\underline{B}^n	Borrowing limit non-stockholders	$0.3\overline{W}$

The mean of the technology shock is a scaling parameter and is normalized to one. The borrowing limits are indexed to the average wage rate, \overline{W} .

RECURSIVE UTILITY AND CALIBRATING THE EIS PARAMETER

The benefit of employing the Epstein-Zin preference specification becomes especially clear when it comes to calibrating the risk aversion and the EIS parameters. Unlike time-separable expected utility functions, recursive preferences disentangle these two conceptually different aspects of preferences. For transparency of results, we abstract from heterogeneity in risk aversion and set it equal to 3.0 for both agents, which is within the range viewed as plausible by most economists. This allows us to focus purely on the effects of the elasticity of substitution. We have also experimented with power utility and essentially got the same results when the elasticities were calibrated to the values chosen here.

In light of the evidence reviewed in Section 2, we set the EIS of non-stockholders to 0.1 and for stockholders, we set it equal to 1.0. To check the sensitivity of our results, we have also experimented with a wide range of elasticity values for both agents. Our impression is that the main message of this paper is robust to this choice, as long as stockholders are reasonably close to unit elasticity (0.8 – 1.25), and non-stockholders are much less elastic than that (0.05 – 0.25). Table 2 summarizes the baseline parameterization.

5 Macroeconomic Performance

In the rest of the paper the goal is to show that our explanation for the elasticity puzzle also holds quantitatively: basically, the majority of the population with low EIS has *quantitatively* no effect on aggregates directly linked to wealth, such as output and investment, which are determined

Table 3: THE CONCENTRATION OF WEALTH AND CONSUMPTION

	Percentage Share of Wealth or Consumption Held By					
	Wealth Percentiles					
	1–10	11–20	21–30	31–40	41–50	51–100
Productive Wealth	0.742	0.138	0.063	0.034	0.020	0.002
Financial Assets	0.832	0.155	0.055	0.018	0.008	−0.068
Consumption	0.169	0.134	0.117	0.108	0.097	0.373

by the preferences of the wealthy. On the other hand, aggregate consumption mainly reveals the preferences of the poor who contributes 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.

5.1 The Empirical Joint Distribution of Consumption and Wealth

Table 3 reports the fraction of aggregate consumption and wealth accounted for by different deciles of the wealth distribution calculated from the Panel Study of Income Dynamics (PSID).¹² The first two rows in the table show the size distribution of productive wealth (net worth minus equity in owner-occupied housing) and financial assets. Both measures of wealth display substantially higher concentration than consumption. Households in the top 10 percent of the wealth distribution own three-quarters of aggregate productive capital in the US economy but contribute only 17 percent of total consumption. By also including the next decile (top 20 percent), this group of households can be appropriately called “investors” since they hold about 90 percent of capital and land, and virtually all financial assets. In contrast, observe the very gradual decline in consumption expenditures as we move down the distribution. The bottom 80 percent own only 12 percent of productive wealth, yet contribute almost 70 percent of total consumption. Thus we call this latter group “consumers.” Expressing the same information in per-capita terms makes this distinction even more striking: an average investor owns 29.3 times the physical wealth of an average consumer, but consumes only 72 percent more. These two groups are also marked on Figure 1 (which plots the joint distribution) to give a visual impression of their dramatically different contributions to aggregate consumption and wealth.

This asymmetry is, of course, not totally surprising since consumption is proportional to life-time wealth inclusive of human capital which is the largest component and is distributed more

¹²It is hard to find a single household-level dataset with complete information on both consumption and wealth. PSID contains detailed information about wealth holdings every 5 years, but lacks a full consumption measure. Yet, it contains data on food expenditures and rent which we use in conjunction with the same variables from CEX to obtain a proxy for total consumption expenditures. Since we only want to characterize consumption for each wealth decile (and not at household level) the approximation turns out to be quite good. Moreover, the size distribution for both wealth measures are quite similar to those reported in Wolff (2000) using data from the Survey of Consumer Finances. Appendix B discusses the details of variable construction and calculations.




Figure 1: The Joint Distribution of Consumption and Wealth: Fraction of Consumption Accounted for by Households Who Own a Given Fraction of Aggregate Wealth

evenly 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 critical element in resolving the elasticity puzzle.

5.2 Results from the Baseline Economy

Given these remarkable differences between the average consumer and average investor, it might be tempting to conclude that the each group will largely determine different aggregates, and together with heterogeneity between the EIS of the two groups the resolution to the elasticity puzzle should follow immediately. In other words, why is it not sufficient to 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. For example, in a companion paper using essentially the same model we found that non-stockholders' preferences—and their EIS in particular—play a key role in determining asset prices even though they hold almost no wealth on

Table 4: THE DISTRIBUTION OF WEALTH AND CONSUMPTION ACROSS THE TWO GROUPS

	Productive Wealth		Consumption	
	Top 30%	Bottom 70%	Top 30%	Bottom 70%
US Data	0.94	0.06	0.41	0.59
Model	0.89	0.11	0.37	0.63

Table 5: BUSINESS CYCLE STATISTICS – BASELINE MODEL

	U.S. data	Representative Agent		Limited Participation
		$\rho = 0.1$	$\rho = 1.0$	
		Standard deviation (%)		
Output	2.2	3.6	3.0	3.0
Investment	8.5	10.8	8.7	8.8
Consumption	1.5	1.5	1.7	1.5
		Autocorrelation		
Output	0.52	0.65	0.52	0.53
Investment	0.36	0.49	0.33	0.36
Consumption	0.57	0.98	0.94	0.94

Notes: Empirical statistics for the U.S. economy are computed from the National Income and Products Account 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. Consumption measure includes non-durables and services.

average (Guenen 2002). Clearly, trade in the bond market provides a channel through which non-stockholders' preferences can *potentially* influence the properties of aggregate quantities, including investment and output, just as they affect asset prices. This necessitates an examination of the behavior of aggregates allowing for equilibrium interactions through the bond market, which we do in this subsection.

As a first step Table 4 displays the cross-sectional distribution generated by the model. Both the extreme skewness in the wealth distribution and the relative equity in the consumption distribution are captured very well.¹³ The substantial inequality of wealth has proved an especially tough challenge for many previous models using the infinite horizon framework, and given its central role for our results we discuss the underlying mechanism in Section 6.

The results of our baseline calibration is presented in Table 5 along with the corresponding statistics for the U.S. economy in the first column. In order to see why a low EIS seems difficult to be reconciled with macro data, we first compare two representative-agent (hereafter RA) real business cycle models with elasticities of 0.1 and 1.0 respectively.

When the EIS is equal to 0.1 (column 2), even though the volatility of consumption matches

¹³With a more detailed calibration, we would set stockholders' labor earnings to its empirical fraction of about 55 percent of aggregate earnings. But then we would also have to adjust for the fact that labor income tax is progressive and the burden of capital income taxes rest solely on stockholders. Moreover, in the current calibration we set capital's share of output, θ , to 0.4 to partly compensate for the low labor income assumed for stockholders, even though the empirical estimates range from 0.25 to 0.4. If instead, capital's share is set to 0.25 and we introduce the taxes described above, a back-of-the-envelope calculation (by setting consumption equal to average after-tax income for each group plus one-third of tax receipts as transfers) gives the share of consumption of stockholders to be 43 percent compared to 41 percent in the data. Although it is feasible to explicitly introduce these taxes, this would complicate the analysis without yielding any further insights, so we do not pursue it here. Also, what really matters is the fact that consumption is distributed evenly *in the data* which is already documented in the previous section.

Table 6: BUSINESS CYCLE STATISTICS WITH CYCLE DURATION CALIBRATED TO 8 YEARS

	Representative Agent		Limited Participation
	$\rho = 0.1$	$\rho = 1.0$	
	Standard deviation (%)		
Output	4.1	3.4	3.5
Investment	11.3	9.1	9.4
Consumption	1.9	2.1	2.0
	Autocorrelation		
Output	0.76	0.65	0.67
Investment	0.63	0.45	0.47
Consumption	0.98	0.95	0.96

the empirical value, the volatility 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 his income, making the former smoother at the expense of extra volatility in investment and consequently in output. Furthermore, all three variables are too persistent. In contrast, if we increase ρ to 1.0 (column 3), the outcome of the model economy moves in the right direction. The excess volatility of output decreases and the investment figure now matches data exactly. Also, the persistence of investment and output decreases to near the empirical values.

The last column displays the moments from our baseline economy. There are two points to notice. First, the limited participation model matches output and investment statistics much better than the low-elasticity RA model even though the majority have very low EIS. Second, and *more importantly*, the limited participation model behaves very much like the high-elasticity RA model when output and investment are concerned. This second point is what we want to emphasize. In other words, even though a standard calibration of the baseline model generates quite plausible aggregates, an in-depth analysis of business cycle statistics is not the focus of this paper, otherwise we would have incorporated standard features such as labor-leisure choice, a richer shock process, time-to-build, and so on. Our main goal here is to show that the existence of many low elasticity households have virtually no effect on the determination of investment and output, which transpires in this parsimonious framework.

In order to strengthen this result and to show that this conclusion does not rest on the particular parameterization, we change a number of key parameters including the variance and persistence of the aggregate shock, the fraction of stockholders, and display the corresponding statistics in Tables 6 to 8. As can be seen, output and investment continue to display very similar characteristics across the two models confirming our intuition. In this sense, the mechanism by which a small number of wealthy households control aggregate dynamics is quite robust.

On the other hand, unlike the unconditional moments of consumption we study in this section, the conditional information (revealed by Euler equation regressions) in Section 8 is very sensitive to the EIS and reflect the preferences of the poor. Overall, we believe that this analysis supports our intuition that aggregates do not necessarily reflect the preferences of the majority.

Table 7: BUSINESS CYCLE STATISTICS WITH FRACTION OF STOCKHOLDERS 15 PERCENT

	Representative Agent		Limited Participation
	$\rho = 0.1$	$\rho = 1.0$	
	Standard deviation (%)		
Output	3.8	3.2	3.3
Investment	11.0	9.3	9.6
Consumption	1.6	1.8	1.6
	Autocorrelation		
Output	0.64	0.52	0.54
Investment	0.49	0.33	0.35
Consumption	0.97	0.94	0.95

Table 8: BUSINESS CYCLE STATISTICS WITH SHOCK TWICE AS VOLATILE

	Representative Agent		Limited Participation
	$\rho = 0.1$	$\rho = 1.0$	
	Standard deviation (%)		
Output	6.9	6.0	6.1
Investment	22.4	18.7	19.3
Consumption	3.2	3.4	3.3
	Autocorrelation		
Output	0.71	0.52	0.54
Investment	0.50	0.34	0.36
Consumption	0.96	0.93	0.95

5.3 Heterogeneity in the EIS: Further Evidence

Heterogeneity in the EIS is one of the two building blocks of our model. Our motivation for this assumption came from a vast empirical literature some of which we discussed in Section 2. However, a potential caveat in those empirical studies is pointed out by Carroll (1997) and 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 elasticity of substitution. If there are borrowing constraints which frequently bind for the poor agent, his consumption will closely track his income and will not respond to the interest rate unlike the consumption of the wealthy agent who is totally unconstrained. When the elasticities are recovered from consumption data on these two agents, the poor agent will appear as if he has low EIS, even though binding constraints are in fact responsible for it. If true, then contrary to what we have assumed, there may not be any significant heterogeneity in the EIS in the first place.

The cross-sectional data help us evaluate if this possibility is likely to be important. We show that when the aggregate implications of this story is considered it is at odds with key features of cross-sectional data. Hence, the following macroeconomic inference should be viewed as complementary to and supporting the econometric studies cited in Section 2.

The first stylized fact is that stockholders' consumption is considerably more volatile than that of (the poor) non-stockholders; the ratio of the variances of log consumption ranges from 2 to 4 depending on the consumption measure and the threshold stockholding level used to identify the two groups.¹⁴ This finding seems quite surprising though. Given that stockholders are much wealthier than others and have access to financial markets, and low-income households typically have higher labor earnings uncertainty (Kydland, 1984), we would expect the former group's consumption to be smoother than that of the latter. In fact, a standard heterogenous-agent model with identical preferences and borrowing constraints (as in Aiyagari, 1994), which essentially incorporates the elements argued by Laibson et. al. (1998) into a general equilibrium context, would also 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.

Turning to the limited participation model, if we set $\rho^h = \rho^n = 1$, the results are also inconsistent with this observation: The ratio $\sigma^2(c^h)/\sigma^2(c^n)$ is just 0.78, much lower than even the lower bound of two. This finding could not be overturned by varying the shock size or its persistence, replacing the Markov structure with AR(1) shocks, or changing the calibration of constraints. Thus, limited participation alone does not account for this fact. However, when we set $\rho^n = 0.1$, this picture changes. Now, non-stockholders' stronger desire for a smooth consumption path causes

¹⁴A brief list of papers reporting this finding includes Mankiw and Zeldes (1991) and Poterba and Samwick (1995) from PSID, Atkeson and Ogaki (1997) from the Indian ICRISAT dataset, Vissing-Jorgensen (1998) from CEX, Attanasio, Banks and Tanner (2002) from the FES dataset from U.K. The last three papers measure consumption expenditures as non-durables and services and find the ratio of variances to be between 2 to 4. Ait-Sahalia, Parker and Yogo (2001) find that the sales of luxury (non-durable) goods are much more variable than the average variability of non-durables, which although indirect, provide further evidence.

them to trade vigorously in the bond market driving the interest rate down, and increasing the equity premium. Now, non-stockholders' consumption is smoother than that of stockholders, who bear the extra volatility and are paid a hefty (equity) premium. As a result $\sigma^2(c^h)/\sigma^2(c^n)$ rises from 0.78 to 1.96 (and the Sharpe ratio rises from 2 percent to about 20 percent). Furthermore, if the model is calibrated to quarterly data with AR(1) shocks with a persistence of 0.95, then the ratio of yearly variances is 3.4.

Second, a number of financial statistics implied by the model are quite sensitive to the elasticity of non-stockholders and make extremely sharp predictions about its value. We study these asset pricing implications in a companion paper (Guvenen 2002). We find that, starting from an elasticity of one for both agents, the performance of the model improves dramatically as non-stockholders' EIS is gradually reduced and is best when $\rho^n \approx 0.05 - 0.2$. With this calibration, the model is able to match a wide variety of asset pricing phenomena surprisingly well, including all the facts explained by the benchmark model of Campbell and Cochrane (1999) and some others on which that model is silent. Examples in this list include a high equity premium, a low risk-free rate as well as the volatilities of both returns, the procyclicality of stock prices, the countercyclicality of the equity premium, its volatility and the Sharpe ratio, among others. On the other hand, if borrowing constraints are tightened such that non-stockholders are constrained for about 5 percent of the time or more, many of the results, including the cyclical behavior of returns contradicts empirical observations. Given that many of these asset pricing phenomena have been considered puzzles, these findings seem to provide further evidence that the poor indeed have low EIS.

To sum up, a number of features of cross-sectional data and asset prices which are hard to explain in standard heterogenous-agent models can be easily accounted for by jointly considering preference heterogeneity and limited participation.

6 What is the Effect of Limited Participation?

As noted earlier, the substantial wealth inequality is an essential component and the driving force behind the results in this paper. Hence, we want to highlight the critical role played by limited participation in creating the wealth inequality.

This is especially important for two reasons. First, given other ingredients present in the model, it is instructive to see, in isolation, what limited participation contributes. Second, and more importantly, one might even be skeptical that limited participation, as a kind of market incompleteness, is an important mechanism for generating heterogeneity. The reason is that a number of recent papers have concluded that market incompleteness does not quantitatively matter if: (1) the time horizon is long, (2) agents have stationary income processes, and (3) there is a risk-free bond available to all in the economy. For example, Rios-Rull (1994) shows that the aggregate dynamics in an economy where agents cannot insure against all shocks is virtually indistinguishable from the complete markets version of the model.¹⁵ Noting that all the

¹⁵In other examples, Telmer (1993), Heaton and Lucas (1996), Krusell and Smith (1998) study economies which

Table 9: BASIC STATISTICS IN A LIMITED PARTICIPATION MODEL WITH IDENTICAL AGENTS AND NO BORROWING CONSTRAINTS

	Stockholder	Nonstockholder
Fraction of population	0.30	0.70
Share of aggregate wealth	0.77	0.23
Per-capita wealth	8.10	1.00
Per-capital consumption	1.48	1.00
Std. of log consumption	.018	.019
Correlation of consumption and income	0.44	0.75
	Correlation across groups	
Consumption	0.71	
Income	0.87	
Wealth	0.03	

Notes: Per-capita wealth and consumption are normalized. Std refers to the standard deviation of log consumption

above assumptions are also satisfied in our model, one might want to question the severity of the restriction imposed by shutting some agents out of the stock market since they still have access to the bond market.

Hence, in order to show that limited stock market participation *does* matter, we conduct the following experiment. First consider a simplified version of our baseline economy where we eliminate all the frictions and the preference heterogeneity. More specifically, suppose that $\rho^h = \rho^n = 1$, 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 suppose that we introduce a single friction: we restrict the second group of agents from participating in the stock market, but otherwise there are still no borrowing constraints are imposed. As Table 9 displays, the effect of this simple change is striking. Suddenly stockholders come to hold almost 80 percent of the aggregate wealth (compared to 30 percent just a moment ago), or in per-capita figures, a stockholder now owns nearly eight times the average wealth of a non-stockholder. They also consume nearly 50 percent more per-capita than non-stockholders. Furthermore, the two groups' wealth holdings are now virtually uncorrelated, down from perfect correlation! And all of this is happening without any idiosyncratic risk, borrowing constraint or

satisfy these common assumptions but differ in many other respects. They all find that market incompleteness has very little effect quantitatively. Finally, Levine and Zame (2001) prove a similar result in the limit as $\beta \rightarrow 1$ and when there are no aggregate shocks. So, one can say that there is substantial evidence that market incompleteness does not quantitatively matter, at least in this general framework.

Table 10: WEALTH INEQUALITY FOR VARIOUS PARAMETER VALUES

Parameters				Share of Wealth Held By:	
ρ^h	ρ^n	α^h	α^n	Stockholders	Nonstockholders
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: ρ and α denote the EIS and relative risk aversion coefficients respectively.

preference heterogeneity. This example demonstrates the potential of limited participation for generating substantial heterogeneity with important economic consequences.

A natural question is: How does limited participation generate significant heterogeneity? How is this model any different than the previous work we discussed above? The crucial difference is that limited participation gives rise to *asymmetric* market incompleteness, in the sense that some agents (i.e., non-stockholders) face more severe market incompleteness than others since they have access to fewer assets for insurance. Thus, this model provides an example where market incompleteness does matter quantitatively because it affects individuals asymmetrically.

6.1 Where Does the Wealth Inequality Come From?

First, notice that extreme skewness in wealth distribution is robust to different parametrizations of preferences as displayed in the table above, which shows that preference heterogeneity is not the main source of inequality.¹⁶

Let us first discuss how the inequality is *not* generated. It is not generated by design, which is the case with a limited participation model in an exchange economy setting. In that case, stockholders are *endowed*, at time zero, with the entire future stream of dividends on top their labor income, and hence, are wealthier by assumption. In contrast, here, even for identical initial wealth levels, the stockholder *chooses* to accumulate more wealth in equilibrium. For example, it is entirely possible that a stockholder has zero wealth in equilibrium: she could borrow in the bond market and invest all the proceeds in the firm, payoff the non-stockholder after production and consume the equity premium. So, stockholders are not wealthier simply because they own the total capital stock. Then why are they?

In fact, the basic mechanism is quite familiar. First, with infinite horizon both assets earn

¹⁶It should be noted that borrowing constraints also contribute to the skewness of the wealth distribution. Non-stockholders can save only up to the maximum amount stockholders are allowed to borrow. When stockholders come closer to their borrowing constraints, this puts a downward pressure on bond return reducing nonstockholders demand for wealth. But as the example in Table 9 demonstrates, the concentration of wealth is still substantial without any constraints.

average returns below the time preference rate. A well-known result 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.¹⁷ This result also holds when asset returns are stochastic (Chamberlain and Wilson 1984). Similarly, by a continuity argument, one can show that asset demand will become unbounded as the average return approaches the time preference rate from below. Figure 2 plots the typical long-run asset demand schedule for an individual in this environment and illustrates the extreme sensitivity of demand to return variations close to η . In the current model even though the equity premium is quite small in an absolute sense, it is still very large compared to the difference between R^S and the time preference rate: $\eta - E(R^f) = 17.8 * (\eta - E(R^S))$. This 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.¹⁸

Finally, the graph also makes clear that the significant wealth of stockholders cannot be trivially explained by the fact that they have access to a positive equity premium and 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 which would result in very little wealth dispersion.¹⁹

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-stockholder who has only access to one asset. However, the only way this group can accumulate bond is if stockholders are willing to borrow. In contrast, stockholders have access to capital accumulation, and they could smooth consumption even if the bond market was completely shut down (just as they do in the standard RBC model). This asymmetry in insurance opportunities between the two groups puts non-stockholders in an unfavorable position in the bond market. This disadvantage is further exacerbated and non-stockholders have an even more inelastic demand if, in addition, they have low elasticity of substitution and hence are desperate for a smooth consumption profile. In this case, stockholders will only trade if they are paid well in the bond market, which means a high equity premium. (The Sharpe ratio is around 20 percent in the baseline economy which is about 2/3 of the empirical value). The low return to saving for non-stockholders in turn dampens their demand and they end up with little wealth in equilibrium.

¹⁷See Clarida (1991), Aiyagari (1994), and references therein. Tobin (1967) also contains a diagram similar to Figure 2 in this paper.

¹⁸To be precise, this argument holds exactly if both agents have identical asset demand schedules, which is *not* the case here because even without borrowing constraints and differences in preferences, the non-stockholder has fewer assets to use for insurance and hence has more precautionary demand. But these factors shift the asset demand schedule by very little exactly because of the reasons found in the previous literature: a second asset does not help a lot for self-insurance at a given interest rate. So, the actual asset demand schedules for each agent are very close to what is displayed in the figure.

¹⁹It is worth noting that the discussion so far, including Figure 2, is about the stationary wealth distribution. In particular, it seems possible a priori that with a small equity premium it will take very long to reach the stationary distribution which would then put pressure on the infinite-life assumption. Fortunately that is not the case. The big effect on wealth dispersion comes from a large wedge between the saving rates of the two groups caused by the modest equity premium. Thus, 90 percent of the adjustment is completed in less than 40 years in the baseline case.

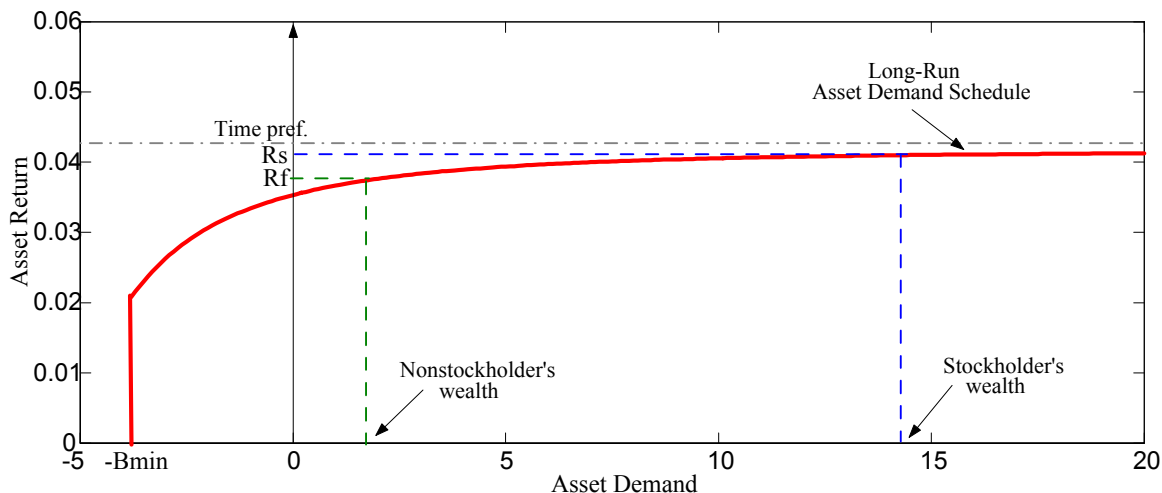


Figure 2: Wealth Inequality as a Result of Return Differential (Limited Participation)

The described effect also corresponds to an explanation commonly given, but to our knowledge which has never been investigated in a general equilibrium framework: the wealthy become wealthy because they face higher returns (in the stock market). This is exactly what happens in this model.

7 Should We Care About the Average EIS? A Policy Experiment

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.

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). We imagine that initially the government imposes a flat-rate tax on capital income and returns the proceeds to households in a lump-sum fashion. Suppose that, at a certain date, capital income tax is completely eliminated and agents have not previously anticipated it. We set the initial tax rate $\tau^k = 36\%$ 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 percent 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 percent inflation rate. However, if we assume that the agent has $\rho = 0.1$, the welfare gain is reduced to 0.4 percent of consumption instead, mainly because now the transition takes approximately 100 years compared to 20 years in the former case.

Now we subject the limited participation economy to the same tax experiment. The welfare gain is 0.82 percent of total consumption.²⁰ 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 modelling heterogeneity: based on policy experiments like the one above, many 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, stockholders gain by 5.4 percent, whereas non-stockholders, who constitute 70 percent of the population, actually lose 2.1 percent of their consumption!* Clearly, this is a different conclusion than what comes out from the representative-agent economy. Moreover, if there is an exogenous government expenditure stream that needs to be financed, eliminating the capital income tax can be viewed as imposing a consumption tax instead. Hence, this result suggests that consumption taxes must be progressive from an equity point of view.

8 Another side of the Puzzle: The Econometrics 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 estimating the EIS from the log-linearized consumption Euler equation which is the most extensively used method in the literature (Hansen and Singleton, 1983, Hall, 1988, Campbell and Mankiw, 1989, among others). The estimated equation is:

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

where small letters denote the natural logarithms of variables, Δ is the difference operator, $k \equiv \rho \log(\beta) + \left[\frac{1}{2\rho} \text{var}_t(\Delta c_{t+1}) \right]$, and ϵ_{t+1} is the expectation error with zero mean conditional on current information: $E(\epsilon_{t+1} | \Omega_t) = 0$.

We perform instrumental variables (IV) estimation of equation (6) and choose instruments commonly used in the literature which include the lags of consumption growth and the risk-free

²⁰This 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.

Table 11: ESTIMATION OF THE LOG-LINEARIZED EULER EQUATION WITH RISK-FREE RETURN

<i>True Value</i>	Nonstockholder	Stockholder	Aggregate
	0.10	1.00	0.47
	$k^f = \text{constant}$		
<i>EIS</i> (t-stats)	0.11 (22.15)	0.54 (51.1)	0.25 (36.82)
	$k^f = \left(\frac{1}{2\rho}\right) \text{var}_t(\Delta c_{t+1})$		
<i>EIS</i> (t-stats)	0.101 (17.02)	1.07 (27.96)	0.48 (21.84)

rate: $(1, \Delta c_{t-1}, \Delta c_{t-2}, \Delta c_{t-3}, \Delta c_{t-4}, r_{t-2}^f, r_{t-3}^f, r_{t-4}^f)$ for each agent as well as for the estimation from aggregates.

First, we run the estimation by assuming that the intercept term k is constant through time or is uncorrelated with the commonly used instruments. This seemingly innocuous assumption is unanimously made in the empirical macro literature but will turn out to significantly affect the results. We concentrate on large sample results, so a sample of 63,000 periods is simulated where the first 3,000 periods are discarded.

Table 11 reports the elasticity parameters recovered for each group (from that group's consumption data) as well as for the aggregate (ρ^A). We expected ρ^A to be close to a weighted average of individual elasticities where the weights would be the consumption shares of each group. This average is 0.47 in this model. 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). We re-estimate (6) by properly including the time-varying conditional variance calculated from the model's optimal decision rules. The improvement 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 seriously downward biased.²¹

Replicating Hall's regression using U.S. time-series data indicates that this issue is likely to be significant in practice. Including a simple measure of consumption volatility (the squared sum of realized consumption growth rates in the next N periods) into the regression increases the estimate of the EIS from around zero up to 0.4–0.5 range. Although we do not report the entire estimation here to save space, we summarize this result in Appendix C for completeness. This exercise is also a nice example of how quantitative theory can help us understand empirical studies and possibly check for their reliability in a sound way.

²¹A number of recent papers have looked at similar types of biases ranging from omitting durables (Ogaki and Reinhart 1998) to ignoring nonseparability of leisure and consumption in the utility function (Basu and Kimball 2000).

9 Discussion and Extensions

HETEROGENEITY IN THE EIS: Although the empirical evidence reviewed in Section 2 reveals that the wealthy have higher EIS, there is no presumption that this heterogeneity is exogenous. In this paper, we chose to state the baseline model by making preference heterogeneity exogenously given. This modeling choice was made for expositional purposes because it allowed us to use Epstein-Zin preferences and disentangle risk aversion from the elasticity in a clear way. However, it is easy to accommodate endogenous differences in the elasticities of each group by slightly changing the model set up. In this alternative case, agents are assumed to have *identical* non-homothetic preferences exhibiting “keeping up with Joneses”:

$$U^h(C, X) = U^n(C, X) = \frac{(C - \varphi X)^{1-\alpha}}{1-\alpha}$$

where C is individual consumption, X denotes aggregate consumption and φ is a parameter which determines how much agents want to “keep-up.”

The main results of the paper survive under this alternative scenario. The reason is as follows: from Table 9, it is clear that limited participation alone creates substantial wealth inequality even without any other frictions or heterogeneity in the model. We also argued that this inequality is quite robust to changes in preference parameters. Moreover, because the EIS is increasing in the wealth level with “habit” utility functions, in equilibrium the wealthy stockholders (endogenously) exhibit higher intertemporal elasticity compared to non-stockholders.²² This interpretation has the nice implication that all heterogeneity, including that in the EIS, is coming from a single friction (limited participation) which is empirically very well established.

LIMITED PARTICIPATION: First, in terms of comparability of our results, limited participation is also assumed exogenously in all the studies discussed in this paper.²³ This is mainly because a satisfactory and tractable model of the participation decision is not currently available. Given that we are interested in the consequences of this phenomenon which are largely unexplored in a macroeconomic context, we believe that this is a reasonable assumption for our purposes.

In a way, the exogeneity assumption even has certain advantages because it allows us to separate the effect of limited participation from self-selection of certain types into the stock market. Take the wealth inequality as an example. We observe that in the U.S. data, on average, a stockholder (or a household in the top 30 percent) owns twenty times the wealth of a non-stockholder. A potential explanation will have two components which may also interact with each other. First, agents with certain characteristics (high risk tolerance, high initial wealth, lower discount rate, etc.) will be those who self-select into the stock market. Hence, these characteristics may be the

²²Note that in this case the elasticity is time-varying as a function of wealth. However, the elasticity of each agent remains confined within a narrow interval; they are separated from each other and centered around the high and low ends of the (0,1) interval which is what we need.

²³For some recent attempts to explain the participation decision (in partial equilibrium contexts) see Davis, Kubler and Willen (2002) and the references therein.

main source of their high wealth and they would be almost as wealthy even if they were non-stockholders. This is the “selection effect.” In this case, it is not entirely clear how the difference between investing in the stock market versus investing in a risk-free asset changes the saving incentives and contributes to generating heterogeneity (as an “amplification effect”) since it is entangled with these other factors. On the other hand, here we investigate the other extreme. We take two identical agents and place one in the stock market and leave one out. Then, also allowing for general equilibrium interactions, we investigate purely the contribution of limited participation to generating heterogeneity. For example, given that we are able to account for most of the wealth inequality, it seems like the stock market is an important amplification mechanism for generating wealth inequality.

Finally, if the participation decision is endogenized, more risk-tolerant agents would be more willing to enter the stock market (Yaron and Zhang, 2000). Overall, stockholders would have lower risk aversion than non-stockholders. As we commented before, the results are robust to using power utility with risk aversions of 1 and 10 for stockholders and non-holders respectively.

Thus, while acknowledging the importance of understanding the reasons behind incomplete participation, we still believe that a reduced form specification for non-participation, in order to analyze its implications for the macroeconomy, is acceptable for the purpose of this paper. We believe that the results of this paper which underscores the importance of limited participation provides further motivation to study how this phenomenon arises in the first place.

10 Conclusion

In this paper we attempted to reconcile two seemingly contradictory views about the elasticity of intertemporal substitution, a crucial parameter in macroeconomics. A number of observations on growth and aggregate fluctuations suggest a value of EIS close to one. In contrast, the comovement between aggregate consumption and interest rates—the focus of the recent empirical consumption literature—implies a very weak relationship between the two, suggesting an elasticity close to zero. A satisfactory answer to this question is also critical for determining the effectiveness as well as the consequences of widely used monetary and fiscal policies.

We showed that a simple and otherwise standard macroeconomic 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 and resolve this apparent puzzle. 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 and the resulting “asymmetric market incompleteness” creates substantial wealth inequality matching the actual distribution between stockholders and 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 under-

standing 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 clarify other related questions: For example, we found that 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 contradict important stylized facts. On the other hand the current model is successfully able to explain those facts when preference heterogeneity is properly acknowledged which provides further support that the poor indeed have low EIS.

The idea that aggregates can be better explained by the interaction of heterogeneous agents rather than by a single 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 show that, as a result, economic analysis as well as policy discussions based on average elasticities may be seriously misguided.

To sum up, we conclude that aggregate consumption is more sensitive to interest rates than suggested by Hall (1988) and others because of a downward bias in their estimates. More importantly, aggregate savings will be even more responsive than revealed by the consumption regression, because wealth is controlled by wealthy households with a high elasticity of substitution.

Overall, we believe that a view of the macroeconomy based on heterogeneity across agents in investment opportunity sets and preferences provides a rich description of the data as well as enabling a better understanding of the determination of aggregate dynamics.

A Appendix: Numerical Solution and Accuracy

This Appendix describes the numerical solution of the model introduced in Section 3 and related accuracy issues. Let X denote the aggregate state (K, B, Z) throughout this Appendix. Solving the model amounts to finding the following functions which are part of the stationary recursive equilibrium:

1. Value functions $(V^i(\omega; X))$ and decision rules: $b^i(\omega; X)$ for each agent $i = h, n$, and $s'(\omega; X)$.
2. Equilibrium bond pricing function, $q(X)$, which clears the bond market.
3. Equilibrium laws of motion $\Gamma_K(X)$, $\Gamma_B(X)$ consistent with individual decision rules.

Note that there is an interdependence between the functions in 1 to 3 above. There are a number of nonlinear functional equations to solve in order to obtain these functions, so instead of attempting to solve them simultaneously, we use an iterative algorithm.

Specifically, we first solve each agent's dynamic programming problem with initial guesses for $q(X)$ and laws of motion. Note that $R^S(K, Z)$ and $W(K, Z)$ are easily determined by the FOCs of the firm's problem. Moreover, to our knowledge, this is the first attempt to numerically solve a dynamic program with Epstein-Zin preferences. Then we use the decision rules to find a bond pricing function which clears the market and update the old value of $q(X)$. Similarly, we update the laws of motion as will be described below. We go back to the agent's problem and solve it with the updated values of the equilibrium aggregate functions and continue the procedure until convergence. The details of the algorithm are as follows:

Step 0: *Initialization:*

- (a) Choose a grid for individual wealth levels: ω^h, ω^n . We used 80 grid points for each wealth variable for the baseline case discussed in the text; using as few as 60 points and as many as 100 points did not noticeably affect the results. There is not much curvature in the equilibrium functions in K direction, so eight equally-spaced grid points gave sufficiently accurate results. On the other hand, the bond price exhibits substantial variation in B direction close to the borrowing constraints, so we took 30 grid points. We chose the locations of the grid points corresponding to Chebyshev roots which oversamples near boundaries.
- (b) Take an initial guess for $q(X)$, $\Gamma_K(X)$ and $\Gamma_B(X)$: we set $q^0(X)$ such that neither asset dominates the other in return state-by-state. For $\Gamma_K^0(X)$, we set it equal to the equilibrium law of motion for capital obtained from a representative-agent economy with the same calibration of relevant variables, and preferences are calibrated to that of the stockholder. $\Gamma_B^0(X)$ is set such that initially, $B' = B$.
- (c) Set $V^i \equiv c > 0$, $i = h, n$, for some constant c .

Step 1: *Solve each agent's dynamic problem:*

- (a) Let l and j index grid points and iteration number respectively. For each grid point $(\omega_l; X_l)$ we find the optimal consumption and portfolio choice for the individual. The portfolio choice of the stockholder is a two-dimensional maximization problem with a very flat objective function given the small equity premium. The standard optimization packages which rely on Jacobian or Hessian matrix information, such as the "npsol" subroutine of the NAG library, or similar routines from the IMSL library fail very frequently. Hence, instead we used a constrained optimization algorithm similar to the one described in Krusell and Smith (1997) which is not very fast but is very robust. To evaluate the value function off the grid points we use interpolation methods. One advantage of Epstein-Zin preferences is that without borrowing constraints the value function is linear in individual wealth. Our experience is that, in our model, it is also almost linear except in the close neighborhood of the constraint. So we were able use linear interpolation in ω direction, and we used cubic spline interpolation in K and B directions.
- (b) After decision rules are obtained, we apply Howard's policy iteration algorithm to speed up convergence. This amounts to updating the value function by assuming that the agent uses the same decision rule for t periods, where we used $t = 20$.
- (c) We iterate on a-b until the maximum percentage deviation in each decision rule is less than 10^{-5} for the stockholder and 10^{-6} for the non-stockholder.

Step 2: *Update the bond pricing function:* In iteration j , at each grid point for current state X_l , we want to find the new bond price $q^j(X_l)$ which clears the markets today, when agents take $q^{j-1}(X)$ to apply to all future dates. More specifically, we first solve the following maximization problem for the stockholder and with $s' \equiv 0$ for the non-stockholder:

$$\begin{aligned}
 J(\omega; K, B, Z, \hat{q}) &= \max_{b', s'} \left((1 - \beta)(C)^\varphi + \beta (E_t [V(\omega'; K', B', Z')^\alpha])^{\frac{\varphi}{\alpha}} \right)^{\frac{1}{\varphi}} \\
 &\quad s.t. \\
 C + \hat{q} * b' + s' &\leq \omega + W(K, Z) \\
 &\text{and equations 2 to 5 in the text.}
 \end{aligned}$$

Note that this is not a functional equation. This problem will give rise to bond holding rules $f_B^h(\omega; K, B, Z, \hat{q})$ and $f_B^n(\omega; K, B, Z, \hat{q})$ as a function of the current bond price \hat{q} . Then, at each grid point X_t , we search over the bond price \hat{q} to find q_t^* such that the bond market clears: $\text{ABS}(\lambda * f_B^h(\omega; K, B, Z, q_t^*) + (1 - \lambda) * f_B^n(\omega; K, B, Z, q_t^*)) < 10^{-8}$. We set $q^j(X_t) = q^*(X_t)$.

Step 3: Now we update the laws of motion using the updated decision rules: $K_{t+1} = \Gamma_K^j(X) = \lambda s^j(\omega^h, K_t, B_t, Z_t)$ where $\omega^h = (K(1 + R^e(K, Z)) - B) / \lambda$, and $B_{t+1} = \Gamma_B^j(X) = (1 - \lambda)b^{j,n}(\omega^n, K_t, B_t, Z_t)$ where $\omega^n = B / (1 - \lambda)$.

Step 4: Iterate on steps 1 to 3 until convergence. We require the maximum deviation in consecutive updates to be less than 10^{-6} for bond pricing function, and 10^{-5} for the aggregate laws of motion.

We then simulate the model through time and assume that the economy reaches a stationary distribution after 3,000 periods. All the statistics in the paper are averages over 60,000 periods of simulation.

B Data Appendix

This appendix describes the data and the construction of variables for the joint distribution of wealth and consumption reported in Figure 1 and Table 3. The data is from the 1984 wave of the Panel Study of Income Dynamics and its wealth supplement. We choose the 1984 wave because it is one of the four waves (along with 1989, 1994 and 1999) which include the wealth supplement used to construct net worth and financial asset variables. Food data (used in the construction of consumption) is missing in the 1989 wave rendering it unusable. The latter two waves come only in early release form with very little documentation. Our attempts to construct consistent definitions of variables proved problematic. Also, due to increasing reliance on secondary respondents when the household head could not be contacted (after 1993) data quality seems to have been negatively affected (See Haider 1997 for a discussion).

The basic economic unit in all the calculations is a household. The definition of net worth and financial wealth is the same as given in footnote 7. These definitions also correspond to Wolff's (2000) variables from the Survey of Consumer Finances making comparison easier. For consumption we take expenditures of non-durables and services (denoted C_{ns}) which is the measure used by Euler equation studies discussed in the text. We construct the CEX measure of C_{ns} that we take as benchmark following Attanasio and Weber (1995) which is also very close to the National Income and Products Accounts definition. We exclude durables (vehicle purchases, household furnishings and equipment), apparel and education expenses since they are likely to have a significant durable component.

One problem with using PSID is that it does not include a comprehensive measure of consumption but rather has food expenditures (the sum of food at home, food away from home and the value of food stamps), and the rent value. On the other hand the Consumer Expenditure Survey has high quality consumption data but no detailed information on wealth. Using these two variables we construct a proxy for non-durables and services. Note that we are interested in the distribution of consumption rather than its absolute level, so we can construct a reasonable proxy as long as our variables (rent and food) constitute a reasonably fixed fraction of total consumption for most households. Figure 3 plots the ratio of rent to C_{ns} for different income brackets from the CEX. Rent accounts for about 23 to 25 percent of C_{ns} and is pretty stable across different household groups. Second, together with food expenditures (the dashed line) they account for 41 to 43 percent of total and this ratio is also pretty much fixed across income groups. We use this second measure to calculate the corresponding proxy from the PSID.

One last issue is about the calculation of a rental equivalent for homeowners in PSID. To obtain the rental equivalent from the reported house values for homeowners we compute the user cost of housing, $C = [(1 - \tau)(i + \tau_p) - \pi + \delta]V \equiv \Theta V$, where τ is the personal income tax rate, τ_p is the property tax, π is capital gains on house value, δ is the maintenance cost, and V is the house value (c.f., Hendershott and Slemrod, 1983). The Bureau of Economic Analysis (BEA) releases annual data that measure the value of the stock of owner-occupied nonfarm housing as well as the imputed rent for owner-occupied non-farm housing. (The imputed rent is calculated using actual rents on comparable dwellings both from the Census of Housing Survey and the CPI housing survey.) Dividing the two numbers yields a measure of Θ equal to 8.8 percent which we use to calculate the rental equivalents from house values. To see if possible variation in Θ across income levels might bias the results, we also constructed the same consumption proxy using only renters' information. In this case, the share of consumption of the top 20 percent was 34.2 percent. We also tried using total consumption expenditures (instead of non-durables and services) as the measure of consumption. In this case the share of consumption of the top 20 percent increased from 30.3 percent to 33.6 percent.

C Appendix: Replicating Hall (1988) with Time-varying Conditional Variances

To gain some insight into how serious this problem is, we conducted the following simple experiment with actual data. We re-ran Hall's regression, closely matching the definition of variables, time period and construction of variables to existing studies. As a first step, we calculate the correlation of the variance of consumption growth with

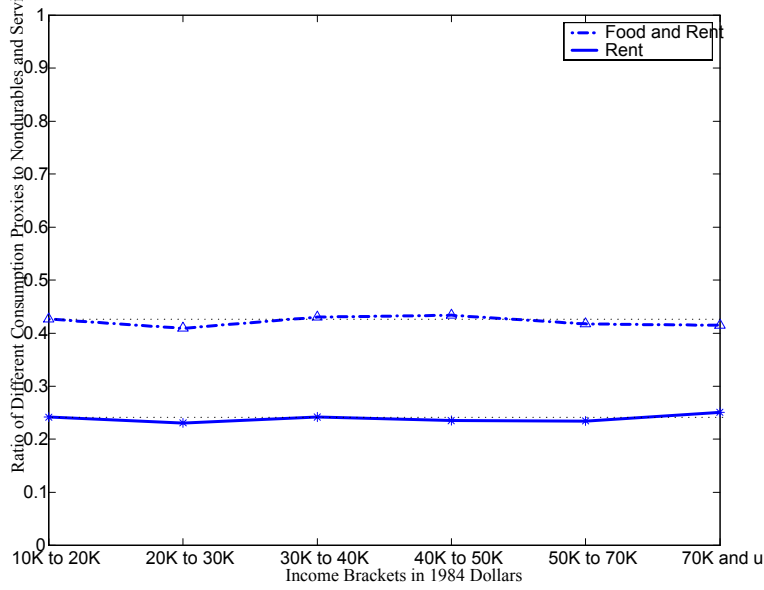


Figure 3: The Ratio of Consumption Proxies to Non-Durables and Services from the CEX, 2000

Table 12: ESTIMATION OF THE LOG-LINEARIZED EULER EQUATION (6) WITH POST-WAR U.S. DATA WITH AND WITHOUT CORRECTING FOR TIME VARIATION

$Var_t(\Delta c_{t+1})$ included?	No	Yes			
		4	8	12	16
N (quarters)	—	4	8	12	16
EIS estim. with \mathbf{Z}_1	0.11	0.29	0.63	0.48	0.41
(t -stat)	(0.94)	(1.21)	(1.67)	(1.94)	(1.49)
EIS estim. with \mathbf{Z}_2	-0.09	0.39	0.57	0.40	0.51
(t -stat)	(-0.97)	(1.82)	(1.92)	(2.28)	(2.12)

once lagged interest rate, and surprisingly it is -0.2 , closely matching that in the model. This also provides some further support to the model in terms of capturing the rich dynamics of data.

The consumption measure is the real quarterly consumption of non-durables from the National Income and Products Account covering 1951.1 to 1984.4. The interest rate is the 3-month T-bill rate from the FRED database, deflated by the corresponding NIPA non-durable consumption prices deflator. Finally, $var_t(\Delta c_{t+1})$ is calculated as follows: for each quarter t , we computed the sample variance using the realized consumption changes in the next N quarters ($\Delta c_{t+1}, \dots, \Delta c_{t+N}$). Results are reported for a range of values of N . The first instrument set is $\mathbf{Z}_1 = (1, r_{t-2}^f, r_{t-3}^f, r_{t-4}^f, \Delta c_{t-2}, \Delta c_{t-3}, \Delta c_{t-4})$ and the second one is $\mathbf{Z}_2 = (\mathbf{Z}_1, i_{t-2}, i_{t-3}, i_{t-4})$ where i_t denotes the inflation rate between t and $t + 1$.

Table 12 presents the results. When the intercept k , is assumed to be constant (column 1) the EIS is estimated to be 0.11 and insignificant, in line with most of the previous literature. However, with $var_t(\Delta c_{t+1})$ appropriately included in the rest of the table, all of a sudden the estimates of EIS jump to around 0.4! Repeating the same experiment with a larger instrument set only strengthens our finding, if anything we get more precise estimates which cluster around 0.4 to 0.5.

References

- [1] Abel, Andrew (2001): “The Effects of Investing Social Security Funds in the Stock Market When Fixed Costs Prevent Some Households from Holding Stocks,” *American Economic Review*, Vol. 91, No.1, pp. 128-148.
- [2] Aït-Sahalia, Yacine, Jonathan Parker and Motohiro Yogo (2001): “Luxury Goods and the Equity Premium,” NBER working paper No. 8417
- [3] Aiyagari, Rao S. (1994): “Uninsured Idiosyncratic Shock and Aggregate Saving,” *Quarterly Journal of Economics*, Vol. 109, No.3, pp. 659-684.
- [4] Atkeson, Andrew, and Masao Ogaki (1996): “Wealth-Varying Intertemporal Elasticities of Substitution: Evidence from Panel Data and Aggregate Data,” *Journal of Monetary Economics*, Vol.38, pp. 507-534.
- [5] Attanasio, Orazio P., and Martin Browning (1995): “Consumption over the Life Cycle and over the Business Cycle,” *American Economic Review*, Vol 85, No. 5, pp. 1118-1137.
- [6] Attanasio, Orazio P., and Guglielmo Weber (1993): “Consumption Growth, the Interest Rate and Aggregation,” *Review of Economic Studies*, Vol 60, pp. 631-649.
- [7] Attanasio, Orazio P., and Guglielmo Weber (1989): “Intertemporal Substitution, Risk Aversion and the Euler Equation for Consumption,” *The Economic Journal*, Vol. 99, pp. 59-73.
- [8] Attanasio, Orazio P., James Banks and Sarah Tanner (2002): “Asset Holding and Consumption Volatility,” *Journal of Political Economy*, forthcoming.
- [9] Basak, Suleyman and Domenico Cuoco (1998): “An Equilibrium Model with Restricted Stock Market Participation,” *Review of Financial Studies*, Vol. 11, No.2, pp. 309-341.
- [10] Basu, Susanto, and Miles Kimball (2000): “Long-Run Labor Supply and the Elasticity of the Intertemporal Substitution for Consumption,” Working paper, University of Michigan.
- [11] Barsky, Robert, Thomas Juster, Miles Kimball, and Matthew Shapiro (1997): “Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey,” Working Paper, National Bureau of Economic Research.
- [12] Beaudry, Paul, and Eric Van Wincoop (1996): “The Intertemporal Elasticity of Substitution: An Exploration Using a U.S. Panel of State Data,” *Economica*, Vol. 63, No. 251
- [13] Boldrin, Michele, Lawrence Christiano, and Jonas Fisher (1999): “Habit Persistence, Asset Returns and the Business Cycle,” Working Paper, Northwestern University.
- [14] Brav, Alon, George Constantinides, and Chris Geczy (2001): “Asset Pricing with Heterogeneous Consumers and Limited Participation: Empirical Evidence,” *Journal of Political Economy*, forthcoming.
- [15] Browning, Martin and Thomas Crossley (2000): “Luxuries are easier to postpone: A proof,” *Journal of Political Economy*, Vol 108, pp. 1022-1026.
- [16] Browning, Martin, Lars P. Hansen and James Heckman (1999): “Micro Data and General Equilibrium Models” in Taylor, J., M. Woodford (eds.): *Handbook of Macroeconomics*, Vol. 1A. Amsterdam, North Holland, 1999

- [17] Campbell, John Y., Andrew W. Lo, and Craig MacKinlay (1997): “The Econometrics of Financial Markets,” Princeton University Press, Princeton, New Jersey.
- [18] Campbell, John Y. and N. Gregory Mankiw (1989): “Consumption, Income, and Interest Rates: Reinterpreting the Time Series Evidence,” *NBER Macroeconomics Annual*, MIT Press.
- [19] Castañeda, Ana, Javier Díaz-Giménez and José-Víctor Ríos-Rull (2002): “Accounting for the U.S. Earnings and Wealth Inequality,” *Journal of Political Economy*, forthcoming.
- [20] Chamberlain, Gary and Charles Wilson (2000): “Optimal Intertemporal Consumption Under Uncertainty,” *Review of Economics Dynamics*, forthcoming.
- [21] Clarida, Richard (1991): “International Lending and Borrowing in a Stochastic Stationary Equilibrium,” *International Economic Review*, Vol. 31, pp. 543-558.
- [22] Cooley, Thomas, and Edward Prescott (1995): “Economic Growth and Business Cycles,” in *Frontiers of Business Cycle Research*, Thomas Cooley editor, Princeton University Press.
- [23] Diamond, Peter, and John Geanakoplos (2001): “Social Security Investment in Equities I: Linear Case,” NBER Working Paper No. W7103
- [24] Davies, James B. (1981): “Uncertain Lifetime, Consumption, and Saving in Retirement,” *Journal of Political Economy*, Vol. 89, No. 3, pp. 561-577.
- [25] Davis, Steve J., Felix Kubler and Paul Willen (2002): “Borrowing Costs and the Demand for Equity Over the Life-Cycle,” NBER Working Paper No 9331.
- [26] Epstein, Larry G. , Stanley E. Zin (1989): “Substitution, Risk Aversion, and the Temporal Behavior of Consumption and Asset Returns: A Theoretical Framework,” *Econometrica*, Vol. 57, No. 4., pp. 937-969.
- [27] Grossman, Sanford J., Robert J. Shiller (1981): “The Determinants of the Variability of Stock Market Prices,” *American Economic Review*, Vol. 71, No. 2, pp. 222-227.
- [28] Guvenen, M. Fatih (2002): “A Parsimonious Macroeconomic Model for Asset Pricing: Habit Formation and Cross-Sectional Heterogeneity?” Working Paper, University of Rochester.
- [29] Hall, Robert E (1978): “Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence,” *Journal of Political Economy*, Vol. 86, No. 6., pp. 971-987.
- [30] Hall, Robert E. (1988): “Intertemporal Substitution in Consumption,” *Journal of Political Economy*, Vol. 96, No.2, pp. 339-357.
- [31] Hansen, Lars P., and Kenneth J. Singleton (1982): “Generalized Instrumental Variables Estimation of Nonlinear Rational Expectations Models,” *Econometrica*, Vol. 50, No. 5, pp. 1269-1286.
- [32] Hansen, Lars P., and Kenneth J. Singleton (1983): “Stochastic Consumption, Risk Aversion, and the Temporal Behavior of Asset Returns,” *Journal of Political Economy*, Vol 91, No.2, pp. 249-265.
- [33] Heaton, John and Deborah J. Lucas (1996): “Evaluating the Effects of Incomplete Markets on Risk Sharing and Asset Pricing,” *Journal of Political Economy*, Vol. 104, No.3, pp. 443-487.

- [34] Hendershott, Patric, and Joel Slemrod (1983): "Taxes and the User Cost of Capital for Owner-Occupied Housing," *AREUEA Journal*, pp. 376-394.
- [35] Investment Company Institute and the Securities Industry Association (2002): "Equity Ownership in America," downloadable at <http://www.ici.org/pdf/>
- [36] Jones, Larry E., Rodolfo Manuelli, and Ennio Stachetti (1999): "Technology (and Policy) Shocks in Models of Endogenous Growth," Working Paper No. 7063, National Bureau of Economic Research.
- [37] Jones, Larry E., Rodolfo Manuelli, and Henry Siu (2000): "Growth and Business Cycles," Working Paper No.7633, National Bureau of Economic Research.
- [38] King, Robert G., and Sergio T. Rebelo (1990): "Public Policy and Economic Growth: Developing Neoclassical Implications" *Journal of Political Economy*, Vol. 98, No. 5, pp. S127-S150.
- [39] Krusell, Per, and Anthony A. Smith Jr. (1998): "Income and Wealth Heterogeneity in the Macroeconomy," *Journal of Political Economy*, Vol. 106, pp. 867-896.
- [40] Krusell, Per, and Anthony A. Smith Jr. (1997): "Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns," *Macroeconomic Dynamics*, Vol 1, pp. 387-422.
- [41] Kydland, Finn E., and Edward C. Prescott (1982): "Time to Build and Aggregate Fluctuations," *Econometrica*, Vol. 50, No. 6, pp. 1345-1370.
- [42] Kydland, Finn E. (1984): "Labor Force Heterogeneity and the Business Cycle," *Carnegie-Rochester Conference Series on Public Policy*, 21, 173-209.
- [43] Laibson, David, Andrea Repetto and Jeremy Tobacman (1998): Self-Control and Saving for Retirement," *Brooking Papers on Economic Activity*, pp 91-172.
- [44] Levine, David K., and William R. Zame (2001): "Does Market Incompleteness Matter?" *Econometrica*, forthcoming.
- [45] Lucas, Robert E. Jr. (1990): "Supply Side Economics: An Analytical Review," *Oxford Economic Papers*, 42, pp. 293-316.
- [46] Mankiw, Gregory N, and Stephen Zeldes (1991): "The Consumption of Stockholders and Non-stockholders," *Journal of Financial Economics*, Vol 29, pp. 97-112.
- [47] Mankiw, Gregory N., Julio J. Rotemberg and Lawrence H. Summers (1985): "Intertemporal Substitution in Macroeconomics," *Quarterly Journal of Economics*, Vol. 100, No. 1. (Feb., 1985), pp. 225-251.
- [48] Mulligan, Casey (2002): "Capital, Interest, and Aggregate Intertemporal Substitution," NBER Working Paper, No. W9373.
- [49] Ogaki, Masao and Carmen Reinhart (1998): "Measuring Intertemporal Substitution: The Role of Durable Goods," *Journal of Political Economy*, Vol. 106, No. 5, pp. 1078-1098.
- [50] Patterson, Kerry D. and Bahram Pesaran (1992): "The Intertemporal Elasticity of Substitution in Consumption in the U.S. and in the U.K.," *Review of Economics and Statistics*, Vol. 74, No.4, pp 573-584.

- [51] Poterba, James, and Andrew A. Samwick (1995): “Stock Ownership Patterns, Stock Market Fluctuations, and Consumption,” *Brookings Papers on Economic Activity*, pp. 295-373.
- [52] Ríos-Rull, José-Víctor (1994): “On the Quantitative Importance of Market Completeness,” *Journal of Monetary Economics*, Vol. 34, pp. 462-496.
- [53] Saito, Makoto (1995): “Limited Participation and Asset Pricing,” Working Paper, University of British Columbia.
- [54] Summers, Lawrence H. (1981): “Tax Policy, the Rate of Return, and Savings,” Manuscript, National Bureau of Economic Research, No. 995.
- [55] Telmer, Chris I. (1993): “Asset-Pricing Puzzles and Incomplete Markets,” *Journal of Finance*, Vol. 48, No. 5, pp. 1803-1832.
- [56] Tobin James (1967): “Life Cycle Saving and Balanced Growth,” in *Ten Economic Studies in the Tradition of Irving Fisher*, W. Fellner (eds.), New York, Wiley.
- [57] Vissing-Jørgensen, Annette (2001): “Limited asset market participation and the elasticity of intertemporal substitution,” *Journal of Political Economy*, forthcoming.
- [58] Weil, Philippe (1989): “The Equity Premium Puzzle and the Risk-Free Rate Puzzle,” *Journal of Monetary Economics*, Vol 24, pp. 401-421.
- [59] Wolff, Edward (2000): “Recent Trends in Wealth Ownership, 1983-1998,” Levy Institute Working Paper, No 300.
- [60] Yaron, Amir and Harold Zhang (2000): “Fixed Costs and Asset Market Participation,” GSIA Working Paper, Carnegie Mellon University.