

Do Stockholders Share Risk More Effectively Than Non-stockholders?*

Fatih Guvenen[†]

July 11, 2005

Abstract

This paper analyzes the extent of risk-sharing among stockholders and among non-stockholders. Wealthy households play a crucial role in many economic problems due to the substantial concentration of asset holdings in the U.S. data. Hence, to evaluate the empirical importance of market incompleteness, it is essential to determine if idiosyncratic shocks are important for the wealthy, who have access to better insurance opportunities, but also face different risks, than the average household. We study a model where each period households decide whether to participate in the stock market by paying a fixed cost. Due to this endogenous entry decision, the testable implications of perfect risk-sharing take the form of a sample selection model, which we estimate and test using a semi-parametric GMM estimator proposed by Kyriazidou (2001). Using data from PSID we strongly reject perfect risk-sharing among stockholders, but perhaps surprisingly, do not find evidence against it among non-stockholders. These results appear to be robust to several extensions we considered. These findings indicate that market incompleteness may be more important for the wealthy, and suggest further focus on risk factors that primarily affect this group, such as entrepreneurial income risk.

Keywords: Perfect risk-sharing, incomplete markets, semiparametric estimation, Generalized Method of Moments, limited stock market participation.

JEL Classification Codes: C33, G11, D52.

*For helpful comments and suggestions, I would like to thank Chris Flinn, John Heaton, Boyan Jovanovic, Shakeeb Khan, Katerina Kyriazidou, Mico Mrkaic, Werner Ploberger, Tony Smith, Harold Zhang, and the seminar participants at NYU, University of Rochester, the 2002 SED meetings, and the North American summer meetings of the Econometric Society. Financial support from William Larimer Mellon foundation is gratefully acknowledged. The previous version of this paper was circulated under the title "Does Stockholding Provide Perfect Risk-Sharing?" The usual disclaimer applies.

[†]Mailing 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

In the past several years, models with incomplete markets and uninsurable idiosyncratic shocks have achieved a central place in many fields of economics. These models are now used to study a wide range of economic questions, such as business cycle dynamics, fiscal policy, wealth inequality, asset prices, among many others (c.f., Aiyagari (1994); Constantinides and Duffie (1996); Krusell and Smith (1998); Heaton and Lucas (2000); Storesletten, et. al (2001)).

A major motivation for these studies has been the decisive empirical rejection of perfect risk-sharing—the hypothesis that individuals are able to insure against all idiosyncratic shocks, and consequently, that their marginal utilities move in lockstep with each other. A number of empirical studies have found individuals’ marginal utility growth (sometimes proxied by consumption growth) to be correlated with certain idiosyncratic shocks, violating the premise of perfect insurance (Cochrane (1991); Nelson (1994); Townsend (1994); and Attanasio and Davis (1996)).

An important point to note is that these studies test if perfect risk sharing (PRS) holds among *all* households in the population. However, given that asset holdings and wealth are extremely concentrated—basically 90 percent of non-housing wealth and 98 percent of stocks is held by the top 20 percent of the U.S. population—wealthy households play a crucial role in many economic interactions.¹ Thus, for a satisfactory analysis of the issues mentioned above, it is especially important to determine the extent of risk-sharing among these *wealthy households*.

On the one hand, there are good reasons to suspect that wealthy households stand a better chance of achieving perfect risk-sharing than the rest of the population. These households almost exclusively trade in stock markets, and therefore have access to the arguably most sophisticated market-based risk sharing mechanism. Thus, it seems possible that the empirical rejection of PRS among all households could be driven by the lack of insurance among the poor, and may not provide a justification that idiosyncratic risk is important for the wealthy. But, on the other hand, wealthy households are exposed to certain risks to a much larger extent than the rest of the population. For example, private capital—which is roughly as large as the capital in publicly traded companies, and is potentially difficult to insure due to asymmetric information problems—is concentrated among the wealthy, exposing them to entrepreneurial income risk not faced by other households.² Furthermore, investors in financial markets face several trading frictions (such as transactions costs, margin requirements, costs of information acquisition, etc.) which could prevent optimal diversification and expose stockholders to the idiosyncratic risk of the stocks in their portfolio. The first goal of this paper is then to formally investigate if the (wealthy) stockholders—who face these various risks

¹See Guvenen (2005a) for empirical evidence on the inequality of wealth and asset holdings in the U.S. Because households in the top 20 percent own almost all stocks outstanding as well as 90 percent of the wealth, we also refer to stockholders as “the wealthy,” and to non-stockholders as “the poor.”

²Gentry and Hubbard (1998), and Heaton and Lucas (2000) document the extreme concentration of private capital and the business risks faced by entrepreneurs.

and have access to various risk sharing mechanisms—are able to share risk effectively.

As emphasized in the literature, however, trading in financial assets is not the only channel for risk sharing. There are a number of informal (or non-market) insurance mechanisms available to most households, such as redistributive government and social programs, gifts and loans between family members, financial assistance provided by charities, long-term contracts and labor hoarding by companies, and so on. Given also that the less wealthy households (non-stockholders) may not be exposed to certain types of risks as noted above, it is conceivable that they could also be sharing risk effectively among themselves. To investigate this possibility then, we also test for perfect risk sharing among the less wealthy (non-stockholders). Applying the test of PRS to this latter group also has the advantage of providing a benchmark to compare the results for stockholders.

We consider an economy with a complete set of financial assets traded in a stock market. Thus all shocks are *potentially* insurable. Each period households decide whether to participate in the stock market by paying a per-period participation cost, or to stay out and trade a single risk-free asset. Households also make optimal portfolio and labor supply decisions. In addition, since perfect risk sharing imposes restrictions on marginal utilities, a sufficiently flexible parameterization of the utility function is crucial for our purposes. To this end, we allow for non-separabilities between consumption and the leisure times of head and spouse, and incorporate household-specific preference shifters. This specification is more general than most of those adopted in previous studies (with the exception of Altug and Miller (1990) and Hayashi, Altonji and Kotlikoff (1996) which are similar to ours).

Due to the endogenous nature of stock market participation in this model, the testable implications of the risk-sharing hypothesis for stockholders take the form of a sample selection model, where the participation decision serves as the selection equation. To eliminate the selection bias, we implement a semiparametric GMM estimator recently proposed by Kyriazidou (1997, 2001) for panel data models, which does not require strong distributional assumptions about the error terms. To our knowledge, this is the first implementation of this estimator.

Our findings can be summarized as follows. Using data from the Panel Study of Income Dynamics (PSID) we strongly reject perfect risk-sharing among stockholders, but perhaps surprisingly, do not find similar evidence against it among non-stockholders. This result is robust to several extensions we considered, such as including future wages into the instrument set (as advocated by Hayashi et. al (1996)), and using different moment conditions implied by PRS, among others. To interpret this finding, it would seem hard to argue that non-stockholders have access to better risk sharing opportunities than stockholders, suggesting that the difference (in the tests of PRS) is likely to be due to the *additional uninsurable risks* faced by stockholders. Finally, we also strongly reject PRS for the whole population consistent with the existing literature, suggesting that those earlier rejections could also be driven by the lack of risk sharing among wealthy households.

These results have important implications for modeling choices made in heterogenous-agent models regarding the types of uninsurable risks. For example, the most common sources of uncertainty in these models are labor income risk, unemployment risk, health shocks, and mortality risk. However, it is also common to assume that different groups in the population—such as the wealthy and the poor—are exposed to these uninsurable risks to similar extents, which is hard to reconcile with our findings.³ In contrast, some recent studies have begun incorporating entrepreneurial income risk—a risk that primarily affects the wealthy—into incomplete markets models (c.f., Angeletos (2005); Cagetti and De Nardi (2004); Chari, Golosov and Tsyvinski (2002)). Our results support this emphasis and suggest further focus on risks that fall disproportionately on the wealthy.

In terms of approach this paper is most closely related to the literature which test for PRS among smaller groups in the population. The discouraging rejection of PRS in the whole population by the studies mentioned above led researchers to focus on smaller economic units who have strong ties with the hope of uncovering full insurance within these groups. Examples include households living in the same geographical regions (Hess and Shin, 2000), inhabitants of small villages in various underdeveloped countries (Townsend (1994); and Ogaki and Zhang (2001)), and finally, family members (Hayashi, Altonji and Kotlikoff (1996)). However, these studies treat participation in these groups as exogenous, and hence do not address selection bias.

The paper is organized as follows. In the next section we set up the model and derive the testable implications of PRS. In Section 3 we develop the econometric techniques to analyze this problem in the presence of selection bias. Then in Section 4 we describe the data, and discuss the estimation of the selection equation. Section 5 presents the results and discusses robustness. We then discuss some potential explanations for the findings, and conclude in Section 6.

2 The Model

There are a finite number of households, each with a life span of T periods. Households derive utility from consumption as well as from husband’s and spouse’s leisure times. Specifically, the lifetime utility of household i is given by

$$E_0 \left[\sum_{t=1}^T \beta^t u(C_{it}, L_{1it}, L_{2it}, \mathbf{X}_{it}) \right], \quad (1)$$

³See, for example, Hubbard, Skinner and Zeldes (1994) and Carroll and Samwick (1997), who estimate very similar labor income processes for different education groups in the population. These estimates are extensively used to calibrate heterogeneous-agent models. Unemployment risk is sometimes assumed to vary across demographic groups consistent with empirical evidence (high skill workers have lower unemployment rate). But this assumption implies more risks for the poor in these models, in contrast to our results.

where β is the subjective time discount rate; C_{it} is consumption in period t ; L_{1it} and L_{2it} are the leisure times of head and spouse respectively in period t , and \mathbf{X}_{it} is a vector of household-specific preference shifters that captures all the heterogeneity across households relevant for preferences.⁴ The period utility function, u , is continuously differentiable and concave in the choice variables for each value of \mathbf{X}_{it} .

FINANCIAL MARKETS AND PARTICIPATION

Let s_t denote a date-event pair (state), which constitutes a complete description of uncertainty for all the economy that is realized in t , and let $\mathbf{s}^t = (s_1, s_2, \dots, s_t)$ be the history of states realized up through period t . For example, s_t will contain the realization of wages of *all* households, the return on all assets in the economy in period t , etc. Each node s_t branches out into S possible states (successor nodes) in the next period. There is a complete set (S) of Arrow securities (one-period contingent claims) available in every state, each paying one unit of consumption good in exactly one state of the world tomorrow. From this description, it is clear that all shocks are potentially insurable.

In addition, there is also a risk-free bond available in the economy. While this bond can be traded freely by all households (that is, without incurring any fixed or proportional transaction costs), the same does not hold for the Arrow securities. Households must pay a fixed cost of Ψ^P in every period they participate in financial markets (which we will also refer to as the “stock market”) where these securities are traded.⁵

To better understand the choices facing a typical household, it is useful to express the decision problem recursively. Each period a household decides whether to participate in the stock market in the current period by paying Ψ^P , or to stay outside and trade the risk-free bond only. Define $\mathbf{q}(s)$ to be the $(1 \times S)$ price vector of the Arrow securities when the current state is s , and q_0 to be the bond price. Similarly, let the $(S \times 1)$ vector $\mathbf{k}_i \equiv (k_{1i}, k_{2i}, \dots, k_{Si})$ denote a current stockholder’s portfolio choice vector of Arrow securities, and k_{0i} be the bond holdings of a non-stockholder. We drop the time subscript, and denote next period’s variables by primes. Then a household’s problem is:

$$v(\omega_i, \mathbf{X}_i; \mathbf{s}) = \max \left[v^h(\omega_i, \mathbf{X}_i; \mathbf{s}), v^n(\omega_i, \mathbf{X}_i; \mathbf{s}) \right]$$

⁴Throughout the paper bold letters are used to denote vectors in order to distinguish them from scalars.

⁵This per-period fee is intended to capture several explicit and implicit costs of trading in financial markets, such as the time and effort it takes to monitor and rebalance one’s portfolio, the additional time it takes to file tax returns, etc. Vissing-Jorgensen (2000) estimates these costs from micro data and concludes that even modest fixed costs (about \$150 per year) are sufficient to keep a large fraction of households out of the stock market. However, it seems plausible that certain types of participation costs, (especially those associated with learning how financial markets work, etc.) are incurred only once. This would suggest that there could be a one-time fixed cost as well as a per-period fixed cost. In the working paper version (Guisen 2003) we allowed for both of these costs. This extension made the participation decision dynamic (and the model more complicated), but had no substantive effect on our results. Thus, we abstract from the one-time cost in the current version.

where

$$\begin{aligned}
v^h(\omega_i, \mathbf{X}_i; \mathbf{s}) &= \max_{C_i, L_{1i}, L_{2i}, \mathbf{k}_i} \left\{ u(C_i, L_{1i}, L_{2i}, \mathbf{X}_i) + \beta E(v(\omega'_i, \mathbf{X}'_i; \mathbf{s}') | \mathbf{s}) \right\} \\
&\quad s.t \\
C_i + \mathbf{q}(\mathbf{s}) \mathbf{k}_i &\leq \omega_i + \sum_{j=1,2} (1 - L_{ji}) W_{ji}(\mathbf{s}) - \Psi^P \\
\omega'_i &= k_{s'i},
\end{aligned}$$

and

$$\begin{aligned}
v^n(\omega_i, \mathbf{X}_i; \mathbf{s}) &= \max_{C_i, L_{1i}, L_{2i}, k_{0i}} \left\{ u(C_i, L_{1i}, L_{2i}, \mathbf{X}_i) + \beta E(v(\omega'_i, \mathbf{X}'_i; \mathbf{s}') | \mathbf{s}) \right\} \\
&\quad s.t \\
C_i + q_0(\mathbf{s}) k_{0i} &\leq \omega_i + \sum_{j=1,2} (1 - L_{ji}) W_{ji}(\mathbf{s}) \\
\omega'_i &= k_{0i},
\end{aligned}$$

where v^h and v^n are the value functions of current stockholders and non-stockholders respectively; $W_{1i}(\mathbf{s})$, and $W_{2i}(\mathbf{s})$ denote the wages of the head and the spouse respectively after history \mathbf{s} , and ω_i denotes financial wealth. Finally, although the arguments of the choice variables are suppressed for clarity of notation, they are all functions of the state vector $(\omega_i, \mathbf{X}_i; \mathbf{s})$.

Note the difference between the budget constraints of current stockholders and non-stockholders. In particular, stockholders choose an unrestricted $(S \times 1)$ portfolio vector implying that they can transfer any (budget-feasible) amount of wealth to a given state in the next period. Thus markets are dynamically complete within the stock market community. In contrast, non-stockholders are restricted to choosing a constant wealth level, k_0 , for the next period.

2.1 Perfect risk-sharing

In the model presented above households optimally enter and exit the stock market in different periods, and thus face complete markets and incomplete markets at different points in time. This is in contrast to the canonical model used to test for PRS in the previous literature where (under the null hypothesis) households face complete markets every period. Despite this difference, the main testable implication of PRS derived in the canonical model continues to hold in our model, but only for households who participate in the stock market in two consecutive periods.⁶ For clarity of exposition then, we present the derivation of the PRS condition in the canonical model, which can be obtained from our general framework by simply setting $\Psi^P \equiv 0$.

⁶As shown below, precisely because the PRS condition holds for households who self-select into the stock market, PRS must be tested jointly with the selection equation.

In this case, without the participation cost, a household faces complete markets in every period, and hence maximizes the preferences given in (1) subject to a single lifetime budget constraint:

$$E_0 \left[\sum_{t=0}^T \beta^t \lambda(s_t) (C_{it} - \sum_{j=1,2} (1 - L_{jit}) W_{jit}) \right] \leq 0, \quad (2)$$

where $\beta^t \lambda(s_t)$ is the time-zero price of one unit of consumption in state s_t . This budget constraint equates the discounted lifetime value of expenditures to the discounted lifetime value of wage income. Let ϕ_i be the multiplier associated with constraint (2). The first order condition for consumption choice is given by

$$\frac{1}{\phi_i} u_1(C_{it}, L_{1it}, L_{2it}, \mathbf{X}_{it}) = \lambda_t. \quad (3)$$

where the subscript of u denotes the partial derivative with respect to the indicated argument. Similarly, the labor supply decision of the head (assuming an interior solution) is determined by:⁷

$$\frac{1}{\phi_i} u_2(C_{it}, L_{1it}, L_{2it}, \mathbf{X}_{it}) = \lambda_t W_{1it}. \quad (4)$$

Taking the ratio of (3) for periods $t - 1$ and t to eliminate the unobservable component, ϕ_i , we get

$$\frac{u_1(C_{it}, L_{1it}, L_{2it}, \mathbf{X}_{it})}{u_1(C_{it-1}, L_{1it-1}, L_{2it-1}, \mathbf{X}_{it-1})} = \frac{\lambda_t}{\lambda_{t-1}}. \quad (5)$$

This last equation clearly illustrates the main implication of PRS: the marginal utility growth of a given household (on the left hand side) should only be a function of aggregate variables (the right hand side), and hence, should not be correlated with any idiosyncratic variable. In contrast, with *incomplete markets*, there is no reason to expect marginal utility growth to be equated across households. The main test of PRS is then to estimate the relationship in (5) and then see if any idiosyncratic variable is correlated with the resulting error term. Because of the stock market participation decision in our framework, the PRS condition holds for households who are in the stock market in periods $t - 1$ and t .⁸ Consequently, this self-selection must be taken into account when testing condition (5).

2.2 The participation decision

In general a closed form solution for the participation decision is not available, although it is easy to see that a numerical solution can be obtained by solving backwards starting from the last period

⁷A condition analogous to (4) holds for the female labor supply decision, although we do not use it in our analysis.

⁸See the working paper version (Güvenen 2003, Appendix A), which shows that equation (5) indeed holds for the present model with $\Psi^P > 0$.

of a household’s life. The focus of this paper is on risk-sharing, and our interest in the participation decision is mainly for having a good specification of the characteristics of households who self-select into the stock market. Thus, rather than explicitly solve for the participation decision, we seek variables that determine this choice, which can be obtained from the optimization problem above.

A household enters the market in period t if $v^h(\omega_{it}, \mathbf{X}_{it}; \mathbf{s}^t) > v^n(\omega_{it}, \mathbf{X}_{it}; \mathbf{s}^t)$ and stays outside if the reverse holds, implying that the decision rule will be a function of the the state vector: $(\omega_{it}, \mathbf{X}_{it}; \mathbf{s}^t)$. Although in general \mathbf{s}^t is a high-dimensional vector containing the entire history of asset prices and wages in the economy, it can be simplified substantially by making two observations. First, the efficient markets hypothesis holds that asset returns are not predictable. Although there is some evidence of predictability of stock returns over long horizons, at shorter horizons (one to two years) stock returns are difficult to predict (see Guvenen (2005b, table 4)). Consistent with this evidence, we assume that asset returns are *i.i.d* over time, implying that that the history of asset prices can be excluded from \mathbf{s} . Second, there is substantial empirical evidence indicating that the persistence in individual wage dynamics can be adequately represented by an AR(1) process (which includes a random walk as a special case. See Browning, Hansen and Heckman (1999) for a detailed review of this evidence). This implies that the current wage (W_{1it}, W_{2it}) is a sufficient state variable for predicting the future evolution of a household’s wage process. Putting these two pieces together, we obtain $\mathbf{s} = (W_{1it}, W_{2it})$.

To sum up, the participation decision rule for a typical agent can be written as

$$d_{it} = 1 \{ \pi(\omega_{it}, W_{1it}, W_{2it}, \mathbf{X}_{it}) \geq 0 \},$$

where $1 \{ \cdot \}$ is an indicator function, and $\pi(\cdot)$ is determined by the solution to the problem.⁹ Clearly, the set of variables included in $\pi(\cdot)$ represent all the *potential* determinants of stockholding, and it is likely that empirically only a subset of them are significant factors in participation choice. For example, variables which affect v^h and v^n symmetrically will leave $v^h - v^n$ unchanged and will have no impact on participation. Thus, identifying the significant determinants of stockholding is ultimately an empirical question, which we address in Section 4.2.

2.3 Empirical Specification

Since perfect risk-sharing imposes restrictions on marginal utilities, the specification of the utility function is especially important for the purposes of this paper. Following Altug and Miller (1990)

⁹Notice that other parameters of the model (including Ψ^P , ρ_0 , etc.) will also play a role in the participation choice. But these parameters are assumed to be identical across the population—except those already summarized in \mathbf{X}_{it} —so they are all soaked up into the functional form of π .

we assume the following period utility function

$$u(C, L_1, L_2, \mathbf{X}) \equiv \delta_0(\mathbf{X})C^{\rho_0}L_2^{\rho_2} + \delta_1(\mathbf{X})L_1^{\rho_1}L_2^{\rho_3}. \quad (6)$$

This specification is quite flexible, and indeed, it is more general than most of those considered in the previous literature. One limitation of the PSID consumption data that we use in our empirical work is that it consists of only food expenditures. We thus interpret the specification in (6) as the first two sub-utilities of a more general utility function in which non-food consumption enters in a separable manner. (This has been the maintained assumption in all previous studies using PSID. See Section 4.1 for further discussion). Then the first sub-utility can be interpreted as a Cobb-Douglas home-production function where food expenditures serves as capital and female leisure hours as labor input.

The second sub-utility captures the possible non-separability between the leisure times of head and spouse. Non-separable specifications in both sub-utilities have empirical support (Browning and Meghir (1991), Altug and Miller (1990)). Another possibility is to have male leisure also enter the first sub-utility. But first, Hayashi et al. (1996) test for this possibility and do not find support for it. Second, if in fact male leisure and consumption are non-separable, tests based on equation (5) are invalid due to observational equivalence (See Attanasio and Davis (1996, page 1235) for a discussion of this point).

Some components of the vector \mathbf{X}_{it} may not be observable to the econometrician. Hence, it is convenient to write $\mathbf{X}_{it} = (\mathbf{x}_{it}, \varepsilon_{0i}, \varepsilon_{1i}, \varepsilon_{0it}, \varepsilon_{1it})$, where \mathbf{x}_{it} is a vector representing the observable component, and the remaining elements denote the unobservables. Each sub-utility is weighted by indices which are log-linear functions of \mathbf{X}_{it} :

$$\delta_m(\mathbf{X}_{it}) = \rho_m^{-1} \exp(\mathbf{b}_m \mathbf{x}_{it} + \varepsilon_{mi} + \varepsilon_{mit}), \quad m = 0, 1. \quad (7)$$

Here b_m is a fixed vector of coefficients; ε_{mi} represents the fixed household-effect, and ε_{mit} is a zero-mean disturbance term which varies both over time and across households. Further assumptions on the error terms will be stated in the next section. Note that each subutility is scaled by ρ_0 and ρ_1 .

For tractability, we specialize the selection function $\pi(\cdot)$ to a linear form, which allows us to write the decision rule as a standard binary choice equation. Substituting the observable and unobservable parts of \mathbf{X}_{it} , we obtain

$$d_{it} = 1\{\boldsymbol{\theta} \mathbf{y}_{it} + \eta_i - \eta_{it} \geq 0\}, \quad (8)$$

where $\boldsymbol{\theta}$ is a fixed vector of coefficients; $\mathbf{y}_{it} \equiv (\omega_{it}, W_{1it}, W_{2it}, \mathbf{x}_{it})$; $\eta_i \equiv \varepsilon_{0i} + \varepsilon_{1i}$, and $\eta_{it} \equiv -(\varepsilon_{0it} + \varepsilon_{1it})$.

2.4 Moment Conditions

Using the parameterization for preferences, the risk-sharing condition for stockholders (5) yields our first moment condition. After taking logarithms, first differencing, and rearranging, we obtain

$$(\rho_0 - 1)\Delta c_{it} = \Delta \ln(\lambda_t) - \mathbf{b}_0\Delta \mathbf{x}_{it} - \rho_2\Delta l_{2it} + \Delta \varepsilon_{0it}, \quad (9)$$

where Δ_t denotes the difference operator between t and $t - 1$, and lower-case letters denote the natural logarithms of their upper-case counterparts (except for \mathbf{x}_{it}).

Although this equation by itself is sufficient to test for risk-sharing, it cannot identify all the structural parameters of the model. For that purpose, we add another moment condition which is valid for both the stockholders and the non-stockholders. Using equations (3) and (4), we get

$$\frac{u_2(C_{it}, L_{1it}, L_{2it}, \mathbf{X}_{it})}{u_1(C_{it}, L_{1it}, L_{2it}, \mathbf{X}_{it})} = W_{1it}. \quad (10)$$

This is the familiar intra-temporal efficiency condition equating the marginal rate of substitution between consumption and leisure to the wage rate.¹⁰ Notice that this is a static condition, which does not depend on market completeness, so it holds for non-stockholders as well. It is convenient to take logarithms, then difference (10) to obtain:

$$(\rho_0 - 1)\Delta c_{it} = -\Delta w_{1it} + (\mathbf{b}_1 - \mathbf{b}_0)\Delta \mathbf{x}_{it} - (\rho_2 - \rho_3)\Delta l_{2it} + (\rho_1 - 1)\Delta l_{1it} + (\Delta \varepsilon_{1it} - \Delta \varepsilon_{0it}). \quad (11)$$

3 Econometric Method

Since the disturbance terms, $\Delta \varepsilon_{0it}$ and $(\Delta \varepsilon_{1it} - \Delta \varepsilon_{0it})$, in the equations above have zero mean by construction, it might seem reasonable to look at (9) and (11) as defining orthogonality conditions, which could then be estimated using GMM. However, this strategy is not directly applicable in this framework due to sample selection bias. To clearly see this, first consider the PRS condition (9). Under the null hypothesis, only stockholders are able to share risk perfectly, so the appropriate moment condition is:

$$E[\Delta \varepsilon_{0it} \mid d_{it}d_{i,t-1} = 1] = 0, \quad (12)$$

¹⁰Notice that the MRS equation holds as an equality only when the head is working in a given period, which can potentially cause another selection problem. Altug and Miller (1990) estimate a Tobit specification for selection into the labor market and find that the error term in the selection equation has a small and insignificant correlation with the error in the MRS equation. Moreover, we eliminate far fewer households compared to these authors, since we only require the head to work for two consecutive years to be included in the estimation (whereas they require this for 14 consecutive years), so this problem is probably even less critical in this case.

which requires

$$\begin{aligned} E[\varepsilon_{0it} \mid d_{it} = 1] &= E[\varepsilon_{0it-1} \mid d_{i,t-1} = 1] \Rightarrow \\ E[\varepsilon_{0it} \mid \eta_{it} \leq \boldsymbol{\theta}\mathbf{y}_{it} + \eta_i] &= E[\varepsilon_{0it-1} \mid \eta_{it-1} \leq \boldsymbol{\theta}\mathbf{y}_{it-1} + \eta_i] \end{aligned} \quad (13)$$

to hold. In general both sides of this equation will be non-zero, because ε_{0it} is correlated with η_{it} , which we earlier defined as $-(\varepsilon_{0it} + \varepsilon_{1it})$. In other words, the unobservable preference shock that affects the risk-sharing condition, ε_{0it} , also influences the stock market participation decision, creating a selection bias. Furthermore, while it is possible that these conditional expectations are non-zero, but still equal to each other, this is not likely to be the case either given that these expectations are functions of $\boldsymbol{\theta}\mathbf{y}_{it}$, and will vary over time as this selection index changes.

This last observation, however, suggests a way of eliminating the selection bias, and forms the basis of the estimator proposed by Kyriazidou (1997, 2001). To explain the basic idea of this estimator, we first make two assumptions.

Let \mathbf{Z}_{it} be a vector of instruments.

Assumption A1. $\{(\varepsilon_{0it}, \varepsilon_{1it})\}_{t=1}^T$ is *i.i.d* over time for all i , conditional on $\zeta_i \equiv \{\eta_i, \mathbf{Z}_{i0}, (\mathbf{y}_{i1}, \dots, \mathbf{y}_{iT})\}$.

Assumption A2. $(\varepsilon_{0it}, \varepsilon_{1it})$ is independent of $\mathbf{Z}_{i\tau}$ for all $\tau < t$, and for all i conditional on ζ_i .

The first assumption is the same as condition (A1') in Kyriazidou (2001). An important implication of this assumption is that all regressors in the selection equation must be exogenous. The second assumption is a slight weakening of her assumption (A2') that allows us to have endogenous variables in the main equation and to instrument for them using lagged dependent variables.¹¹ Because η_{it} is the sum of ε_{0it} and ε_{1it} , it also satisfies both assumptions above.

The idea behind Kyriazidou's estimator can be explained as follows. From the discussion above, it is clear that the term $E[\varepsilon_{0it} \mid \eta_{it} \leq \boldsymbol{\theta}\mathbf{y}_{it} + \eta_i]$ will remain unchanged if the selection index, $\boldsymbol{\theta}\mathbf{y}_{it}$, (and consequently the conditioning set) is constant in two consecutive periods. Thus, under assumptions A1 and A2, the following modified version of condition (12) holds:

$$E[\Delta\varepsilon_{0it} \mid d_{it}d_{i,t-1} = 1, \Delta\boldsymbol{\theta}\mathbf{y}_{it} = 0] = 0. \quad (14)$$

One immediately observes, however, that if \mathbf{y}_{it} contains any continuous variables, the set of households that satisfy $\{\boldsymbol{\theta}\mathbf{y}_{it} = \boldsymbol{\theta}\mathbf{y}_{it-1}\}$ may be very small, or even empty. One strategy is then to assign a weight to each observation which is inversely proportional to the change in the index

¹¹This is the only difference between Kyriazidou (2001)'s original estimator and the one used here. Kyriazidou considers the case where all regressors in the main equation are either strictly exogenous or lagged endogenous variables. Instead in our case, consumption and leisure are likely to be correlated with contemporaneous preference shocks, and so we instrument for these variables. Assumption A2 ensures that such instruments exist. Furthermore, there are a number of additional regularity assumptions that are required for the estimator discussed in Appendix B.

of that household, $\Delta\theta\mathbf{y}_{it}$, such that asymptotically only observations with constant indices are included in the estimation.¹²

The estimator can then be implemented as follows. In the first step, the selection equation (8) is estimated to obtain an estimate of θ (denoted by $\hat{\theta}$). Then, using this estimate we construct weights which we take to be “kernel density” functions of the following form:

$$\psi_{it}^N = \frac{1}{h_N} K\left(\frac{\Delta\hat{\theta}\mathbf{y}_{it}}{h_N}\right), \quad (15)$$

where $K(\cdot)$ is a scalar density function which satisfies certain regularity conditions (described in Appendix B), and h_N is a sequence of “bandwidths” which tends to zero as the sample size $N \rightarrow \infty$. For a fixed value of $\Delta\hat{\theta}\mathbf{y}_{it}$ the weight ψ_{it}^N shrinks as N increases, while for fixed N , a larger deviation in the selection index corresponds to a smaller weight.

The kernel-weighted GMM estimator is constructed as follows. Let $\mathbf{f}(\boldsymbol{\alpha})$ denote a column vector of orthogonality conditions, $\mathbf{f}(\boldsymbol{\alpha}, i)$ be its sample counterpart for the i th observation, and $\boldsymbol{\alpha}$ be the vector of identifiable parameters in that system. For example, $\mathbf{f}(\boldsymbol{\alpha})$ could be obtained by interacting $\Delta\varepsilon_{0it}$ in equation (14) with some appropriate instruments from the set \mathbf{Z}_{it} . The key difference from a standard GMM estimator is in the construction of the sample counterparts of the moment conditions, which are multiplied by the kernel weights in our case: $\mathbf{G}_N(\boldsymbol{\alpha}) \equiv \frac{1}{N} \sum_{i=1}^N \psi_{it}^N \mathbf{f}(\boldsymbol{\alpha}, i)$. Once $\mathbf{G}_N(\boldsymbol{\alpha})$ is obtained, we proceed as in the case of a standard GMM estimation:

$$\hat{\boldsymbol{\alpha}}_N = \arg \min \left[\mathbf{G}_N(\boldsymbol{\alpha})^T \Phi_N^T \Phi_N \mathbf{G}_N(\boldsymbol{\alpha}) \right],$$

where Φ_N be a stochastic matrix that converges in probability to a finite non-stochastic limit Φ_0 , and the superscript T denotes the transpose of a matrix. This estimator is consistent and asymptotically normal with $\sqrt{Nh_N}$ convergence rate (Kyriazidou, 2001). Further details of the estimation method are contained in Appendix B.

¹²It is clear that the same approach can be used to correct for the selection bias among non-stockholders, which we do when testing for PRS among this group. In addition, a similar selection problem also exists in the estimation of the MRS condition (11) because our sample selection procedure described in next section eliminates households who change their stockholding status during the sample period: $d_{it} \neq d_{i,t-1}$. Even though this moment condition holds for the whole population unlike the PRS condition, the error term ($\Delta\varepsilon_{1it} - \Delta\varepsilon_{0it}$) has zero mean over the entire population, whereas we need this expectation to be zero over the sample that we observe: $d_{it} = d_{i,t-1}$

4 Estimation

4.1 The PSID data

This section briefly describes the data and the variables used in the empirical analysis. Further details are provided in Appendix A. The data are drawn from the Panel Study of Income Dynamics (PSID) covering the period from 1982 to 1993. Although PSID data is available from 1968 onwards, data on stock ownership was not collected until 1984, making the earlier period unsuitable for our purposes. (Data from 1982-83 is used for constructing instruments.) The availability of income and consumption data¹³, along with detailed demographic information has made PSID attractive for studying perfect risk-sharing, and not surprisingly it has been used extensively for this purpose in the previous literature (among others, Altug and Miller (1990); Cochrane (1991); Hayashi, et al. (1996); and Hess and Shin (2000)). To the extent possible, we follow these studies in our sample selection criteria.

We use households from the core sample of PSID, which is a representative sample of U.S. households. In addition, we select a family into the sample in year t if the head: (i) was in the study for four consecutive years including 1984 or 1989; (ii) was married to the same spouse in t and $t-1$; and (iii) had positive labor hours in t and $t-1$. We further eliminate households who had missing or inconsistent data on some key variables as described in Appendix A. Filtering out these observations leaves a total of 8941 household-years (observations) that can be used in estimation.

For each household $i \in \{1, \dots, N\}$ we have data on (a) annual leisure hours of head and spouse; (b) real average hourly earnings of head and spouse; (c) age of head and spouse, denoted age_{1it} and age_{2it} respectively; (d) real household food consumption expenditures (which is the sum of “food at home,” “food away from home,” and “the cash value of food stamps”); (e) number of household members, hsz_{it} (f) completed education of head, E_{it} (g) a dummy indicating whether the household is a stockholder, d_{it} . These variables are available in every year from 1982 to 1993, except for food data which is missing in 1988 and 1989, and stockholding variable which is only available in 1984 and 1989.

Table 1 provides the summary statistics of the data for both groups. First, note that there are roughly twice as many non-stockholders as there are stockholders in our sample (65 percent versus 35 percent). This point is important to keep in mind when interpreting the relative power of the tests of PRS in the two subsamples. Second, the average annual work hours, both for males and females, are similar across the two groups. Third, the average hourly earnings of stockholders is higher—by 70 percent for males, and by 50 percent for females—than that of non-stockholders. And fourth, the average food consumption is 18 percent higher for stockholders. However, despite

¹³Although consumption data is restricted to food and a few other expenditure items. We discuss this limitation below.

Table 1: A LIST OF KEY VARIABLES AND THEIR SIMPLE STATISTICS

	Stockholders	Non-stockholders	All
(i) Number of Observations	3178	5763	8941
(ii) Percentage of sample	34.8%	65.2%	100%
Hours and Earnings			
(i) Average annual hours of husband	2213 (646.1)	2177 (686.5)	2189 (672.1)
(ii) Average annual hours of spouse	1451 (741.5)	1501 (706.8)	1483 (718.2)
(iii) Average hourly earnings of husband	\$17.83 (13.79)	\$10.41 (7.45)	\$12.99 (9.76)
(iv) Average hourly earnings of spouse	\$10.10 (9.08)	\$6.82 (5.71)	\$7.96 (6.88)
(v) Average annual Food consumption	\$5249 (2806)	\$4419 (2253)	\$4708 (2445)
Demographic Variables			
(i) Average age of husband	43.8 (11.3)	39.9 (11.4)	41.2 (11.3)
(ii) Average Education of Head	6.07 (1.54)	4.9 (1.62)	5.31 (1.59)
(iii) Average household size	3.3 (1.13)	3.6 (1.21)	3.5 (1.18)

Note: Standard deviations are in parenthesis. The sample selection criteria are detailed in Appendix A.

these differences in levels, there is no significant difference between the dispersion of any of these variables across the two groups: the coefficient of variations of hours, earnings and consumption are similar for stockholders and non-stockholders.

Finally, there are a few points concerning the use of food expenditures as the measure of consumption that should be addressed. First, separability between food and non-food has been the maintained assumption in all studies on risk-sharing using PSID data, which makes our results comparable (Altug and Miller (1990); Cochrane (1991); Hayashi, et. al. (1996); Hess and Shin (2000), etc.). Second, Atkeson and Ogaki (1996) provide evidence that food and non-food consumption are separable.¹⁴ Moreover, Ogaki and Zhang (2001) find virtually the same results regarding PRS when they replicate their tests using non-durable consumption instead of food expenditures. Third, a possible concern could be that food consumption may not be sufficiently variable causing risk-sharing tests to have low power. But, if anything, the volatility of food consumption (from PSID) is higher than the volatility of non-durables consumption calculated from the Consumption Ex-

¹⁴See also, however, Attanasio and Browning (1995) who argue against separability between food and non-food consumption.

penditure Survey. This is true for both stockholders and non-stockholders. Finally, in section 5.2 we derive and test another implication of the perfect risk sharing hypothesis that does not rely on consumption data. As we further discuss there, the results of that test confirm our findings using food expenditures.

4.2 First step: The selection equation

The first step in the procedure is to estimate the parameter vector θ in the participation equation. This question has received a lot of attention in the recent literature, and as a result, there now exists a wealth of information on the empirical determinants of stock market participation (c.f., Haliassos and Bertaut (1995), Hurst et al. (1998), Vissing-Jorgensen (2000), Guiso et al. (2001), Curcuru et al. (2005)). Many of these papers use data sets containing detailed wealth and portfolio information, such as the Survey of Consumer Finances (SCF), that are more suitable for this estimation than the PSID. Given this existing body of work, we do not re-estimate the selection equation in this paper. Rather, we construct the kernel weights based on the findings of this literature, which we discuss here.

In an early paper Haliassos and Bertaut (1995) estimate a logit model for stockholding choice using data from the SCF with a large number of explanatory variables that include demographics, preferences towards risk, income, wealth and occupation. In particular, their list of variables contains all the variables included in our specification of the selection equation (8), except that they use labor income instead of wages. They find (1) race, (2) education, (3) risk aversion measures, (4) labor income, (5) financial net worth, and (6) “whether the individual has a managerial occupation,” to be significant determinants of stockholding. Subsequent studies mentioned above have confirmed these findings using various alternative specifications and different data sets.

An important point to observe about these findings is that except for labor income and financial wealth, these explanatory variables represent individual (or household) characteristics that show little or no change over time. Since kernel weights are constructed based on the *time change* in the selection index, all but two of these regressors become redundant. Consequently, the coefficient estimates on these fixed characteristics do not affect the second step estimation, as they are always differenced out.

The conclusion we draw from this analysis is that, in order to correct for the selectivity bias, we need to mainly consider movements over time in labor income and financial wealth. Nevertheless, labor income and financial wealth are both endogenous variables and are thus likely to be correlated with the preference shifters included in the main equation (e.g., the MRS and PRS conditions), which is not allowed by Assumption A1 above. In our work then, we proxy for both variables with the sum of the head’s wage and spouse’s wage, $W_{1it} + W_{2it}$, and construct a univariate kernel with this variable only (see Guvenen (2003) for further justifications for using this proxy.)

4.3 Second step: The main equation

There are a number of different ways the risk-sharing hypothesis can be tested. The first and most obvious one is to estimate (9) alone for stockholders and use Hansen’s J -test as a model specification test. If stockholders are not able to share risk perfectly, then their marginal utility growth cannot be explained by aggregate shocks alone, and the residuals will be correlated with idiosyncratic variables. By including household-level variables in the instrument set this correlation will be caught by the J -test as a model specification error. This idea forms the basis of the previous tests implemented in the literature.

A second method, whose advantage will become clear in a moment, is the following: First estimate the MRS condition (11), which holds for the entire population. Then append (9) multiplied by d_{it} as an additional moment condition and estimate the two jointly, and test for PRS as an overidentifying restriction of the model. Specifically, if the additional orthogonality condition imposes p_1 extra restrictions and identifies p_2 additional parameters (and $p_1 > p_2$) then Nh_N (effective sample size) times the increment in the GMM criterion function has a χ^2 distribution with $(p_1 - p_2)$ degrees of freedom. The second approach has the advantage of exploiting more information thereby increasing the efficiency of the estimator and the power of our hypothesis test. We will report the results using both approaches.

The choice of optimal bandwidth.—A standard Gaussian density is used for the kernel function $K(\cdot)$. As noted by Kyriazidou (2001), the choice of bandwidth is potentially more important for the performance of the estimator than the choice of kernel. The optimal bandwidth is determined with a cross-validation procedure described in appendix xx. This procedure yields an optimal bandwidth of $h_N^* = 0.24$. However, the cross-validation criterion function is quite flat between 0.2 to 0.5, although it increases steeply outside this region. Due to the exponential nature of weights, small differences in the value of h_N in this range results in large variations in kernel weights. For example, a household whose selection index changes by 50% between two periods is weighted by 0.61, 0.13 and 0.004 for $h_N = 0.5$, 0.24 and 0.15 respectively. To make sure that our conclusions are robust to values of h_N in this range, in the next section we will also report the results for $h_N = 0.5$ as well.

5 Results

In this section we report our empirical findings. First, we present the results obtained from tests of risk-sharing for stockholders, non-stockholders and the whole population. Then we consider several extensions and alternative tests to examine for the robustness of these results. We conclude by discussing parameter estimates.

In PSID consumption (food) data is not available in 1988 and 1989, which leaves us with six

time differences that can be used in estimation: 84-85, 85-86, 86-87, 90-91, 91-92, 92-93.

Instruments. Our main set includes the following nine variables for the estimation using the time difference between $t - 1$ and t : a constant; age of head at time t , age_{1it} ; age of spouse at time t , age_{2it} ; the contemporaneous change in household size, Δhsz_{it} ; twice lagged change in log consumption, $\Delta c_{i,t-2}$; twice lagged change in log wage of husband, $\Delta w_{1i,t-2}$; and the twice lagged change in log spouse's wage, $\Delta w_{2i,t-2}$.

This instrument set is used for the PRS condition. Notice that we have not included the first lags of variables which are susceptible to measurement error, such as consumption and wages, because the resulting correlation with variables in the PRS equation would make them invalid. For the MRS condition, we use the same instruments as in the PRS condition, with two additions (to increase the precision of the estimates of preference parameters): we include the dummy for stockholding status, d_{it} ; and we use the levels of household size in $t - 1$ and t , hsz_{it-1} and hsz_{it} , instead of the change.

For the empirical specification of \mathbf{x}_{it} , we choose a square and a cubic polynomial of husband's age, age_{1it}^2 , age_{1it}^3 , and the household size, hsz_{it} .

5.1 Tests of risk-sharing

Table 2 presents the main result of the paper. We first estimate the PRS condition for stockholders only, using the optimal bandwidth ($h_N^* = 0.24$). If perfect risk sharing holds, this equation should adequately describe the marginal utility growth for stockholders and the model specification J -test should not reject the estimated moment condition. However, this is not the case: in column 1, the PRS condition has a p -value of 0.004, strongly rejecting perfect risk sharing among stockholders.

Next we turn to non-stockholders. As discussed in the introduction, non-stockholders also have access to informal or non-market insurance mechanisms, and may face different (or even fewer types of) risks than stockholders. Therefore, it is conceivable that they might be able to share risk effectively among themselves. To investigate this possibility then, we now test for PRS only among non-stockholders, which is reported in the second column. The p -value of the test is 70 percent, showing no evidence against perfect risk-sharing among non-stockholders.

To examine if this result is sensitive to the assumed bandwidth for the kernel function, we repeat the same test with a wider bandwidth: $h_N = 0.5$ (columns 4 and 5). The p -value for stockholders slightly increases to 0.009, still indicating a rejection at 1% level. The p -value for non-stockholders falls to 0.421, but is still far away from rejection. Similarly, using a tighter bandwidth of $h_N = 0.15$ (columns 6 and 7) only has a small effect on these results: the p -value for stockholders is 0.025, and for non-stockholders it rises to 0.66. As noted earlier because of the exponential nature of the kernel weights, the effective sample size shrinks quickly with a smaller bandwidth, which might be partly responsible for the slightly higher p -values for the stockholders in this latter case.

Table 2: TESTS OF RISK SHARING USING THE PRS EQUATION ONLY

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	$h_N^* = 0.24$			$h_N = 0.50$		$h_N = 0.15$	
Group	H	N	All	H	N	H	N
Moment Conditions	PRS	PRS	MRS	PRS	PRS	PRS	PRS
Test Statistics							
χ^2	52.1	24.3	54.1	49.8	29.9	45.7	25.3
df	29	29	45	29	29	29	29
p -value	0.004	0.709	0.166	0.009	0.421	0.025	0.66

Notes: H and N denote stockholders and Non-stockholders respectively. P-value (PRS) refers to the significance level associated with the PRS moment condition df is the degrees of freedom for the moment conditions in a given column. The instrument set for the PRS equation includes a constant, age of head, age of spouse, change in household size, consumption growth lagged twice, and husband's and spouse's wage growth lagged twice. The instrument set for MRS adds a dummy indicating stock ownership to the previous list, and uses the levels of household size instead of its change.

5.2 PRS tests: Robustness

Before looking for interpretations for these findings, it is important to address several issues regarding the robustness of these results. The first question is if the model specification test could be rejecting the PRS moment condition because of invalid instruments. One way to check this possibility is to estimate the MRS equation with instruments that were also included in the instrument set for the PRS equation. This will be informative because under the null hypothesis of perfect risk sharing, the error term in the PRS equation is $\Delta\varepsilon_{0it}$, which also appears in the MRS equation (compare equations 9 and 11). So, if some instruments are invalid, they are also likely to result in the rejection of the MRS equation. In Column 3 of Table 2 we estimate the MRS equation using data on all households in the sample. The J – test has a p-value of 0.166 showing no evidence against the MRS equation, and consequently, against the validity of the instruments. Moreover, if instruments were indeed invalid, it is not clear why this would affect the moment conditions of the stockholders significantly while not being revealed in the non-stockholders' moment conditions.

Second, could these results be explained by the poor finite sample properties of the kernel-weighted GMM estimator? For example, if the model specification test tends to over-reject in small samples, it will be more likely to reject the null for stockholders than for non-stockholders (since the sample size is smaller for the former group). However, note that the effective sample size for the estimator is Nh_N . So, for example, the stockholders' effective sample size when $h_N = 0.5$ ($3178 \times 0.5 = 1589$) is close to that of non-stockholders when $h_N = 0.24$ ($5763 \times 0.24 = 1383$). Similarly, the effective sample size of non-stockholders' when $h_N = 0.15$ is close to that of stockholders

Table 3: TEST OF RISK-SHARING USING LEAD INSTRUMENTS

Group	$h_N^* = 0.24$		$h_N = 0.50$	
	H	N	H	N
Moment Condition	PRS	PRS	PRS	PRS
Test Statistics				
χ^2	48.7	33.8	52.2	36.4
df	29	29	29	32
p-value (PRS)	0.012	0.247	0.005	0.162

Notes: The instrument set includes a constant, age of head, age of spouse, change in household size, change in head's log wage from $t - 3$ to $t + 1$, and from $t - 2$ to $t + 1$.

when $h_N = 0.24$. Since the results of PRS tests are the same for all these bandwidth sizes, this particular concern does not appear to be critical. Furthermore, consistent with this finding, the Monte Carlo evidence in Kyriazidou (1997, 2001) suggests that the small sample properties are quite well-behaved for sample sizes around those considered in this paper.¹⁵

Third, is it possible that we fail to reject PRS in non-stockholders' sample because there is more unobserved variability (and hence less information) in that sample, whereas the reverse is true in stockholders' sample leading to a rejection? This does not seem likely to be the case, because non-stockholders' sample is twice as large as that of stockholders (table 1), so everything else equal specification tests should have more power to reject in the former sample. As another reflection of this fact, in the next section we find that the estimates obtained from non-stockholders' sample are more precise than those from stockholders' sample. Thus, it is unlikely that less information (or larger variances) in non-stockholders' sample accounts for differences in the PRS test results.

Fourth, as noted by Hayashi et. al. (1996), tests of perfect risk sharing may not have high power against the alternative of self-insurance if the instrument set only includes lagged values of variables such as wages and consumption.¹⁶ To investigate this possibility, we replace the lagged wage changes ($\Delta w_{1i,t-2}$ and $\Delta w_{2i,t-2}$) with the bracketing wage changes of the head from $t - 2$ to $t + 1$ and from $t - 3$ to $t + 1$ in the main instrument set used before, as suggested by Hayashi et al. As Table 3 displays, PRS is rejected for stockholders as before, with p-values between 0.005 and 0.012. For the non-stockholders, although the p-value is lower than before (0.162 and 0.247), the

¹⁵Another well-known problem is that small sample properties could deteriorate when the number of instruments is large. To address this possibility, the working paper version (Güvenen (2003, table 4)) reports additional tests of PRS, where we eliminate several instruments to reduce the degrees of freedom. The results reported there shows that this has no appreciable effect on the test results.

¹⁶This is because even with incomplete markets, the permanent income hypothesis implies that lagged endogenous variables will be uncorrelated with current forecast errors. If, further, the forecast error can be written as the sum of an aggregate and an idiosyncratic component, then lagged variables will have zero correlation with the idiosyncratic component even when markets are incomplete. See Hayashi et al (1996) for further discussion.

Table 4: TESTS OF RISK SHARING USING THE MARGINAL UTILITY OF LEISURE EQUATION

	$h_N^* = 0.24$		$h_N = 0.50$	
Group	H	N	H	N
Moment Conditions	PRS	PRS	PRS	PRS
Test Statistics				
χ^2	49.7	28.9	51.7	30.1
df	29	29	29	29
p -value (<i>PRS</i>)	0.009	0.471	0.006	0.409

Notes: The instrument set is the same one as the one used in Table 2.

PRS cannot be rejected at conventional significance levels.

A final concern discussed earlier is about the use of food expenditures as a measure of consumption. It is possible to test for PRS without relying on consumption data. To develop such a test note that PRS also imposes structure on the cross-section of marginal utility of leisure growth through equation (4). After taking logs, differencing and rearranging this equation (for household heads that are working in years t and $t - 1$), we get

$$\rho_2 \Delta l_{2nt} = \Delta w_{1nt} + \Delta \ln(\lambda_t) - \mathbf{b}_0 \Delta x_{nt} - (\rho_1 - 1) \Delta l_{1nt} - \Delta \varepsilon_{0nt} \quad (16)$$

The results reported in Table 4 are similar to those found above: using the optimal bandwidth, risk-sharing is rejected for stockholders with a p -value of 0.009, but not for non-stockholders (p -value = 0.471). Increasing h_N to 0.50 has little effect as in previous cases.

Before closing this section we compare our findings to existing work reviewed in the Introduction, which strongly rejected perfect insurance in the whole population. We repeat our main test of PRS for the whole population, and append the MRS equation to increase the power of the test. In column 1, the overidentifying restriction for PRS has a p -value of 0.013 rejecting the null of perfect risk sharing for the whole population.

To further increase the power of the test, we add a simple wage equation (20)—which expresses the wage of the head as a function of individual-characteristics—as an additional moment condition (see appendix C for the description of this equation). This equation acts as a “seemingly unrelated regression” and is also used to obtain more precise estimates in the next section. In the next column, we report the results with this wage equation added: the p -value of PRS now goes further down to 4×10^{-4} . Considering a higher bandwidth value of $h_N = 0.5$ makes the rejection even stronger with p -values never higher than 10^{-4} . Thus, our results are consistent with earlier results indicating strong rejection of perfect risk sharing for the whole population. In light of these findings, it seems that the rejection of PRS in the whole population obtained in the literature is likely to be due to

Table 5: TESTS OF RISK SHARING IN THE WHOLE POPULATION

	(1)	(2)	(3)	(4)
	$h_N^* = 0.24$		$h_N = 0.5$	
Group	$H+N$	$H+N$	$H+N$	$H+N$
Moment Conditions	MRS & PRS	MRS & PRS & WAGE	MRS & PRS	MRS & PRS & WAGE
Test Statistics				
χ^2	102.6	137.67	149.2	196.0
df	74	114	74	114
p -value (<i>Model</i>)	0.016	0.003	0.000	0.000
p -value (<i>PRS</i>)	0.013	0.000	0.000	0.000

Notes: The instrument set is the same one used as in Table 2. For the wage equation, (20), we exclude wife's twice lagged wage change, and twice lagged consumption growth, but add the education of head, $E1nt$ to the instrument set above.

the rejection of PRS not among the non-stockholders, but instead among the stockholders.

5.3 Parameter estimates

All the structural parameters of the model can be identified by jointly estimating: (1) the MRS equation, and (2) the PRS condition for either stockholders or non-stockholders. The wage equation (20) is added as a third moment condition to obtain more precise estimates. Because the PRS equation is strongly rejected for stockholders, but not for non-stockholders, in Table 6 we report the estimates of structural parameters obtained by estimating the PRS equation for non-stockholders only, in addition to the MRS and wage equation estimated for the whole population.

The first column presents the parameters obtained when the optimal bandwidth value is used. First, the curvature of consumption, ρ_0 , is estimated to be 0.27 and is not significantly different from zero, suggesting logarithmic consumption preferences. In the next column, which reports the results when the bandwidth is $h_N = 0.50$, the estimate of ρ_0 is only slightly higher (0.49) and still not significantly different from zero. This result is consistent with earlier studies using PSID data (c.f., Altug and Miller (1990)).

Second, the curvature of male leisure is estimated to be -6.16 and is statistically significant at 1 percent level. The implied elasticity of male labor supply with respect to wages (holding the marginal utility of wealth constant) is $L_{1it} ((1 - L_{1it}) (1 - \rho_1))^{-1}$, which can easily be derived from the first order condition for labor choice, equation (4). Given that the average time spent at work is approximately 2200 hours per year in our sample (see Table 1), assuming 16 hours of discretionary time per day, we get $L_{1it} = 0.37$. Then the implied elasticity is 0.08. Similarly, if we take the estimate of $\rho_1 = -4.83$ from the column 2, the implied elasticity would be 0.10. These values are

Table 6: ESTIMATES OF THE STRUCTURAL PARAMETERS

<i>Bandwidth</i>	Utility function:		
	$\delta_0(\mathbf{X})C^{\rho_0}L_2^{\rho_2} + \delta_1(\mathbf{X})L_1^{\rho_1}L_2^{\rho_3}$		
	$h_N^* = 0.24$	$h_N = 0.50$	$h_N = \infty$
<i>Model Rejected at 5%?</i>	No	No	Yes
<i>Curvature Parameters</i>			
ρ_0	0.27 (0.43)	0.49 (0.33)	0.99 (3.63)
ρ_1	-6.16 (1.96)	-4.83 (1.10)	-6.26 (41.54)
ρ_2	-24.1 (31.2)	-16.32 (28.1)	-0.87 (265.8)
ρ_3	-42.0 (27.4)	-29.90 (23.7)	-34.2 (336.1)
<i>Demographic Effects</i>			
b_0^1 (age-squared)	-0.88 (2.55)	-0.42 (1.47)	-0.01 (2.47)
b_0^2 (age-cubed)	-2.50 (1.24)	-1.18 (0.77)	-0.03 (3.64)
b_0^3 (family size)	0.67 (1.05)	0.32 (0.68)	-0.01 (3.88)
b_1^1 (age-squared)	-1.16 (2.65)	-0.54 (1.48)	-0.43 (4.12)
b_1^2 (age-cubed)	-2.30 (2.63)	-1.12 (1.39)	0.17 (3.65)
b_1^3 (family size)	1.24 (2.54)	0.92 (1.44)	1.54 (13.64)
<i>Test Statistics</i>			
χ^2 (<i>Model</i>)	124.3	120.5	149.2
<i>df</i> (<i>Model</i>)	114	114	114
<i>p</i> -value (<i>Model</i>)	0.239	0.320	0.014
<i>p</i> -value (<i>PRS</i>)	0.566	0.762	0.022

Notes: The moment conditions used are the MRS and Wage equations (11 and 20) for the whole population, and the PRS equation (9) for non-stockholders. The structural parameters are exactly identified and the standard error for parameter estimates are in parenthesis.

within the range found in the previous literature, reported for example in Browning, Hansen and Heckman (1999, table 3.3). As for the curvature parameters of female leisure, ρ_2 and ρ_3 , while the point estimates are negative, neither one is statistically different from zero. The negative signs are consistent with the estimates of Altug and Miller (1990) and Hayashi et al. (1996).

The estimated coefficients of household characteristics seem reasonable. The coefficients on the age polynomial are all negative indicating that, everything else held constant, the utility derived from both sub-utilities decreases with age. The structural coefficient of hsz_{it} is positive which means that an increase in household size (which is mostly due to a new child, since our sample contains only married couples) increases both subutilities. The implied Frisch elasticity of male labor supply with respect to family size is $(b_1^3 hsz_{it}) L_{1it} ((1 - L_{1it})(\rho_1 - 1))^{-1}$. Evaluating this formula using the estimate of $b_1^3 = 1.24$ from the first column, and the average family size of 3.5 from Table 1, yields an elasticity of -0.36 . The negative value suggests that a new child reduces the work hours of the head, most probably in response to the increased home production demand associated with childrearing.

Finally, the last column reports the parameter estimates when no kernel is used: $h_N = \infty$. First, notice in the last row that the model is rejected at any significance level higher than 1.4%. Second, while the parameter estimates are not dramatically different, the standard errors are substantially higher in most cases. Similarly, we found that when stockholders' PRS condition is used (instead of non-stockholders') the parameter estimates are erratic, and often have the wrong sign with large standard errors (available upon request). As noted above, in this case the joint moment conditions (MRS, PRS and wage equations) are always rejected regardless of kernel bandwidth.

6 Discussion and Conclusion

In this paper we found strong evidence against perfect risk-sharing among stockholders, but were unable to reject it among non-stockholders. Furthermore, risk-sharing in the whole population is also strongly rejected. Overall, these findings suggest that the failure of PRS in the whole population, also found by many previous studies, is likely to be due to the failure of the wealthy to insure the additional risks they face.

One potential source of risk faced primarily by the wealthy (stockholders) is entrepreneurial income risk.¹⁷ The literature offers several reasons—based on asymmetric information (agency) problems—for why entrepreneurial income is difficult to insure (Bitler, Moskowitz and Vissing-Jorgensen (2005)). Thus, while stockholders may have access to additional insurance opportunities, their income is harder to insure as well.¹⁸ In contrast, the main sources of income for

¹⁷ According to the 1995 Survey of Consumer Finances the fraction of business owners (with a business value above \$10,000) among non-stockholders is only 3 percent.

¹⁸ Although it is beyond the scope of this paper, in principle it is possible to identify the role of entrepreneurial

non-stockholders are wages and salaries, which already include several implicit or explicit sources of insurance (unemployment and disability insurance, long-term contracts, labor hoarding behaviour by firms, etc.) In addition, a number of informal risk sharing mechanisms (e.g., inter-vivos transfers, charitable donations, and borrowing and lending) further eliminate the risks faced by most households. Note that many of these insurance opportunities may not be effective in insuring the potentially large losses experienced by business owners.

A second potential source of additional risk for stockholders could be the stock market itself. If stockholders face certain trading frictions, such as information acquisition costs (for each stock they trade), transactions costs, and short-selling constraints, this could easily prevent them from forming optimal portfolios. As a result, if each stockholder only holds a few stocks instead of a well-diversified optimal portfolio, there is no reason to expect that they will be equating their marginal utility growth since each investor is exposed to substantial idiosyncratic uncertainty of the stocks in their portfolio. Indeed, empirical studies have shown that, especially before 1990's, more than 1/3 of stockholders were seriously under-diversified, holding a portfolio containing 5 stocks or less (see Curcuru et al. (2005) for a detailed review of these empirical facts). However, during the 1990's, investors increasingly switched from holding individual stocks to mutual funds, improving diversification (although still far from perfect). Although it would be interesting to explore if this trend has improved risk-sharing among stockholders, unfortunately PSID does not contain detailed portfolio information to conduct such an analysis.

Accounting for differences in idiosyncratic risks between the wealthy and the poor could potentially help explain several differences in behavior between these groups, such as why the wealthy have a higher savings rate, and why they demand such a high return for holding risky assets. Overall these results underscore the need to focus on risks faced by wealthy households as important sources of market incompleteness. Further work is needed to pinpoint the exact types of risks that are uninsurable for wealthy households.

A Appendix: The Data

Starting with the PSID family files covering 1982-93 waves, we use the following sample selection criteria to select our main sample. Specifically, we include household-years in t and $(t - 1)$ in estimation if the head of the family:

- (i) is in the study for at least four consecutive years $(t - 3, \dots, t)$ including 1984 or 1989,
- (ii) is married to the same spouse at least in the last two years $(t - 1, t)$ of the same period,
- (iii) has a positive labor income at least in the last two years $(t - 1, t)$ of the same period.

These criteria produced a sample of 2350 households who were in the study between 1984 and 1993, not necessarily for all years. We further eliminate a household-year if:

risk for stockholders. In particular, looking at stockholders only, one can test the PRS separately for those who own private businesses and those who do not. If risk-sharing is rejected for the former group but not for the latter, then this would suggest entrepreneurial risk as an important source of market incompleteness. However, entrepreneurship introduces another selection equation which must be dealt with as well. So, we leave this question for future research.

- (iv) annual family food consumption expenditure is less than \$150,
- (v) the head’s education variable is missing for the last two years ($t - 1, t$) of this period,
- (vi) if the head’s or spouse’s reported annual labor hours exceeded 4860 hours,
- (vii) if the head or spouse had positive annual labor hours but zero annual labor income, or vice versa.

Conditions (iv)-(vii) are similar to those used to eliminate irregular observations in the literature (c.f., Altug and Miller (1990), Hayashi et al (1996))

- (viii) Finally, if a household changed its stockholding status from 1984 to 1989, we eliminate that observation from estimation between these two dates.

An important concern is coding errors. To eliminate potential outliers, we first isolated observations on total consumption, head’s wage and spouse’s wage, if the following bound was violated: $E(x_t) - 2std(x_t) \leq x_t \leq 2E(x_t) + 2std(x_t)$, where x_t denotes the variable. There were a total of 46 observations violating this bound for at least one of the three variables. Upon closer inspection of the time-series of these variables, we eliminated 41 observations which had small standard deviations and the outlier was significantly away from the sample average. These criteria produced a total of 8941 of observations. The breakdown for each moment condition is as follows: 1292 observations for the 83 – 84 moments, 1289 for 84 – 85, 1302 for 85 – 86, 1761 for 89 – 90, 1709 for 90 – 91, and 1588 for 91 – 92.

Wages: The average hourly labor earnings (wages) of head and spouse reported in PSID and adopted in this paper are calculated from the sum of the following types of income and total annual hours: V19127, Labor Part of Farm Income; V19128, Labor Part of Business Income; V19129, Salary Income; V19131, Bonuses, Overtime, Commissions; V19132, Income from Professional Practice or Trade; V19133, Labor Part of Market Gardening Income; V19134, Labor Part of Roomers and Boarders Income.

Stockholding: The definition of stockholding adopted in this paper includes ownership of shares of stock in publicly held corporations, mutual funds, investment trusts, including stocks in IRA’s. This definition corresponds to PSID variables V10912 for 1984 and V17325 for 1989. All households who indicate they do not own any of these assets are considered non-stockholders that year.

This definition of stockholding does not include indirect ownership of stocks through pension funds. Notice that, first, indirect holding was more modest during the 1980’s and has become much more popular in the 1990’s (see Investment Company Institute (2002) for direct and indirect stock ownership rates). And second, pension funds impose several restrictions on the use of funds, which is not consistent with our null hypothesis that stockholders can optimally use their assets to insure against shocks to their budget sets.

PSID collects stock ownership data every five years (from 1984 on) whereas for our empirical work we need this information for every year. We identify a household as a stockholder (alternatively, non-stockholder) in every year between 1984 to 1989, if he is present in the sample in both years as a stockholder (non-stockholder). Second, if a household switches between these two groups from 1984 to 1989, we eliminate those observations from the sample between these two dates since we are not able to determine when the switch exactly happens. Clearly, this step creates another selection bias, which the econometric method is able to handle as explained in the text. Finally, for years after 1989 we take the status of a household as it is given in 1989. This identification scheme is clearly not perfect, but notice that the estimation method asymptotically assigns zero weight to an observation if the probability of being a stockholder changes ($\Delta\theta_{y_{it}} \neq 0$). Thus, a household who moves in, or who moves out of the stock market between 1984 and 1989 will receive a small (and asymptotically zero) weight in estimation since this move is likely to be accompanied by a change in the selection index.

B Appendix: Estimation

We describe the procedure for the estimation of the PRS condition (9) for the stockholders; the case for the MRS equation and the wage equation (eq (20) described in the next appendix) are analogous. First, in order to construct $\mathbf{f}(\boldsymbol{\alpha}, i)$, for each year t we pick $(r \times 1)$ dimensional vectors of instruments \mathbf{Z}_{it} satisfying

$$E[\mathbf{Z}_{it}\Delta\varepsilon_{0it}|\Delta\theta_{y_{it}} = 0, d_{it}d_{i,t-1} = 1] = 0, \tag{17}$$

Denote these moment conditions for year t by $\mathbf{f}_t(\boldsymbol{\alpha}, i)$. From the moment conditions (9) and (11), it is clear that for testing PRS and identifying the structural parameters of the model, one only needs data in two consecutive periods. Thus we reduce the panel data estimation into cross-section by forming the following T^*r dimensional vector where T^* is our panel length:

$$\mathbf{f}(\boldsymbol{\alpha}, i) = (\mathbf{f}_1(\boldsymbol{\alpha}, i), \dots, \mathbf{f}_{T^*}(\boldsymbol{\alpha}, i))' \tag{18}$$

$T^* = 6$, and $r = 39$ for the PRS equation ($r = 51$ for the MRS equation, and 47 for the wage equation). We then construct the sample counterparts by weighting each observation by its corresponding kernel weight: $\mathbf{G}_N(\boldsymbol{\alpha}) \equiv \frac{1}{N} \sum_{i=1}^N \psi_{it}^N \mathbf{f}(\boldsymbol{\alpha}, i)$, which is then used to construct the GMM objective function:

$$\widehat{\boldsymbol{\alpha}}_N = \arg \min \left[\mathbf{G}_N(\boldsymbol{\alpha})^T \Phi_N^T \Phi_N \mathbf{G}_N(\boldsymbol{\alpha}) \right]$$

Under certain regularity conditions outlined in Kyriazidou (2001), this estimator is consistent and asymptotically normal with $\sqrt{Nh_N}$ convergence rate. In the absence of a formula for the optimal weighting matrix, we will choose Φ_0^* such that: $\Phi_0^{*T} \Phi_0^* = E \left((\widehat{\psi}_{nt}^N)^2 \mathbf{f}(\boldsymbol{\alpha}, n)^T \mathbf{f}(\boldsymbol{\alpha}, n) \right)^{-1}$ which is optimal in the standard GMM case.

Unlike in the standard GMM case, the J -test for the kernel-weighted estimator is non-central chi-squared (because of the asymptotic bias of the estimator) with the non-centrality parameter (NCP) equal to the squared mean of $(1/(\sqrt{Nh_n})) \sum_{n=1}^N \widehat{\psi}_N f_j(\alpha_j, n)$. Even though this quantity can be estimated in principle, this is very difficult in practice (see Bierens (1987) for a detailed discussion). The Monte Carlo experiments in Kyriazidou (1997, 2001) suggest that this bias is small in general, suggesting that the NCP is also likely to be small. We use the central chi-squared distribution to perform the hypothesis tests. In the worst case, this will bias the results towards rejection if the NCP is large (see Davidson and MacKinnon 1993, pp. 412-414).

Kernel function.—We assume that the K satisfies the standard regularity conditions. In particular, $\int |K(v)| dv < \infty$, $\int K(v) dv = 1$, and we consider symmetric kernels: $\int vK(v) dv = 0$. Moreover the smoothness of the kernel affects the asymptotic convergence rate which imposes restrictions on the empirical choice of the function $K(\cdot)$. We work with a Gaussian kernel which satisfies these conditions. Because asymptotically optimal kernel functions perform only slightly better even in the limit, normal density is a reasonable choice in practice.

The choice of optimal bandwidth.—The first step is to choose h_N . As is usually the case with semiparametric methods, asymptotically optimal methods for selecting the bandwidth provide little guidance for practical implementation with a fixed sample size. However, observing that the estimated weighting function, ψ_{it}^N , has a structure similar to a kernel density estimator, a sensible approach is to select h_N as the cross-validated value for the estimation of the density of the selection index, $\boldsymbol{\theta}_{y_{it}}$. Hence, the bandwidth is chosen by minimizing the mean integrated squared error of the kernel density estimator.

C Appendix: The Wage Equation

Provided that the labor market is competitive, the wage rate of a given household head can be written as follows:

$$W_{1it} = \delta_w(\mathbf{X}_{wit}) \overline{W}_t, \quad (19)$$

where $\delta_w(\mathbf{X}_{wit})$ is an efficiency index function; \mathbf{X}_{wit} is a vector of household characteristics possibly containing some elements not included in \mathbf{X}_{it} , and \overline{W}_t is the market wage rate. This wage equation is used in some specifications (described in the text) to increase the asymptotic efficiency of the estimator through the correlation of the error terms. We assume that $\mathbf{X}_{wit} = (\mathbf{x}_{wit}, \varepsilon_{wit})$, where \mathbf{x}_{wit} and ε_{wit} denote the observable and unobservable components. Moreover similarly to the specification in (7) we assume that $\delta_w(\mathbf{X}_{wit}) = \exp(b_w \mathbf{x}_{wit} + \varepsilon_{wit})$. After taking logs, first differencing and rearranging we have:

$$\Delta w_{1it} = \Delta \ln(\overline{W}_t) + \mathbf{b}_w \Delta \mathbf{x}_{wit} + \Delta \varepsilon_{wit}. \quad (20)$$

The variable \mathbf{x}_{wit} includes a constructed experience variable: $(A_{1nt} + E_{nt}^2)$ where E_{nt} is the completed education of the individual.

References

- [1] Aiyagari, S. Rao (1994): “Uninsured Idiosyncratic Shock and Aggregate Saving,” *Quarterly Journal of Economics*, Vol. 109, No.3, pp. 659-684.
- [2] Altug, S., and R. A. Miller (1990): “Household Choices in Equilibrium,” *Econometrica*, 58, 543-570.
- [3] Angeletos, G. M. (2005): “Uninsured Idiosyncratic Investment Risk and Aggregate Saving,” Working paper, Massachusetts Institute of Technology.

- [4] Atkeson A., and M. Ogaki (1996): "Wealth-Varying EIS: Evidence from Panel and Aggregate Data," *Journal of Monetary Economics*, Vol. 38, No. 3, pp. 507-534.
- [5] Attanasio, O. P., and M. Browning (1995): "Consumption over the Life Cycle and over the Business Cycle," *American Economic Review*, Vol 85, No. 5, pp. 1118-1137.
- [6] Attanasio, O. and S. Davis (1996): "Relative Wage Movements and the Distribution of Consumption," *Journal of Political Economy*, Vol. 104, No. 6, pp. 1227-1262.
- [7] Bierens, H. J.(1987): "Kernel Estimators of Regression Functions," in *Advances in Econometrics: Fifth World Congress*, Vol. 1, Ed. by T. F. Bewley, Cambridge: Cambridge University Press.
- [8] Bitler M., T. Moskowitz and A. Vissing-Jorgensen (2005): "Testing Agency Theory With Entrepreneur Effort and Wealth," *Journal of Finance*, forthcoming.
- [9] Browning, M. and C. Meghir (1991): "The Effects of Male and Female Labor Supply on Commodity Demands," *Econometrica*, Vol. 59, No. 4, pp. 925-951.
- [10] Browning, M., L. P. Hansen and J. J. Heckman (1999): "Micro Data and General Equilibrium Models," *Handbook of Macroeconomics*, Vol 1, J.B. Taylor and M. Woodford, eds., Elsevier Science.
- [11] Cagetti, M., and M. DeNardi (2004): "Entrepreneurship, Frictions, and the Distribution of Wealth," Federal Reserve Bank of Minneapolis Staff Report No 322.
- [12] Carroll, Christopher, and Andrew Samwick (1997): "The Nature of Precautionary Wealth," *Journal of Monetary Economics*, 40, No. 1.
- [13] Chari, V. V., M. Golosov, and A. Tsyvinski (2002): "Business Start-ups, the Lock-in Effect, and Capital Gains Taxation," Working Paper, University of Minnesota.
- [14] Cochrane, J.(1991): "A Simple Test of Consumption Insurance," *Journal of Political Economy*, 99 , 957-976.
- [15] Constantinides, G. and D. Duffie (1996): "Asset Pricing with Heterogeneous Consumers," *Journal of Political Economy*, Vol. 104, No. 2., pp. 210-240
- [16] Curcuro, S., J. Heaton, D. Lucas, and D. Moore (2005): "Heterogeneity and Portfolio Choice: Theory and Evidence," Working Paper, University of Chicago.
- [17] Davidson, R. and J. G. MacKinnon (1993): "Estimation and Inference in Econometrics," Oxford University Press, New York.
- [18] Gentry, W., and G. Hubbard (1998): "Why Do the Wealthy Save So Much? Saving and Investment Decision of Entrepreneurs," Working Paper, Columbia University.
- [19] Gourinchas P. O., and J. Parker (2002): "Consumption over the life cycle," *Econometrica*, Vol 70, No. 1, pp. 47-90
- [20] Guiso, L., M. Haliassos, T. Japelli (2002): "Household Stockholding in Europe: Where Do We Stand and Where Do We Go?" *Economic Policy*, forthcoming.

- [21] Guvenen, Fatih (2005a): “Reconciling Conflicting Evidence on the Elasticity of Intertemporal Substitution: A Macroeconomic Perspective,” *Journal of Monetary Economics*, forthcoming.
- [22] Guvenen, Fatih (2005b): “A Parsimonious Macroeconomic Model for Asset Pricing: Habit Formation of Cross-sectional Heterogeneity?” Working Paper, University of Rochester.
- [23] Haliassos, M., and C. Bertaut (1995): “Why Do So Few Hold Stocks?,” *The Economic Journal*, Vol, 105, 1110-1129.
- [24] Hayashi, F., J. Altonji, and L. Kotlikoff (1996): “Risk-Sharing Between and Within Families,” *Econometrica*, Vol. 64, pp. 261-294.
- [25] Heaton, J., and D. Lucas (2000): “Asset Pricing and Portfolio Choice: The Importance of Entrepreneurial Risk,” *Journal of Finance*, No 3, pp 1163-97.
- [26] Hess G., and K. Shin (2000): “Risk-sharing by Households within and across Regions and Industries,” *Journal of Monetary Economics*, Vol. 45, pp. 533-560.
- [27] Hubbard, R Glenn, Jonathan Skinner and Steve P. Zeldes (1994): “Precautionary Saving and Social Insurance,” *Journal of Political Economy*, Vol. 103, pp. 360-399.
- [28] Hurst, E., M. C. Luoh, and F. Stafford (1998): “The Wealth Dynamics of American Families, 1984-94,” *Brooking Papers on Economic Activity*, No.1, pp. 267-337.
- [29] Investment Company Institute (2002): “Equity Ownership in America”, available online at <http://www.ici.org/statements/res>.
- [30] Krusell, P., and A. A. Smith (1998): “Income and Wealth Heterogeneity in the Macroeconomy,” *Journal of Political Economy*, Vol 106, 867–896.
- [31] Kyriazidou, E. (1997): “Estimation of a Panel Data Sample Selection Model,” *Econometrica*, 65, 1335-1364.
- [32] Kyriazidou, E. (2001): “Estimation of Dynamic Panel Data Sample Selection Models,” *Review of Economic Studies*, Vol. 68, No. 3, pp. 543-572.
- [33] Nelson, J. (1994): “On Testing for Full Insurance using Consumer Expenditure Survey Data,” *Journal of Political Economy*, Vol. 102, No. 2, pp. 384-394.
- [34] Ogaki, M. and Q. Zhang (2001): “Decreasing Relative Risk Aversion and Tests of risk-sharing,” *Econometrica* Vol. 69, No. 2, pp. 515-526.
- [35] Storesletten, K., C. Telmer, and A Yaron (2001): “Asset Pricing with Idiosyncratic Risk and Overlapping Generations,” Working Paper, Carnegie-Mellon University
- [36] Townsend, R.(1994): “Risk and Insurance in Village India,” *Econometrica*, Vol. 62, pp. 539-590.
- [37] Vissing-Jorgensen, A. (2000): “Towards an Explanation of Household Portfolio Choice Heterogeneity: Nonfinancial Income and Participation Cost Structures,” Working Paper, University of Chicago.