

Household Finance

JOHN Y. CAMPBELL*

ABSTRACT

The study of household finance is challenging because household behavior is difficult to measure, and households face constraints not captured by textbook models. Evidence on participation, diversification, and mortgage refinancing suggests that many households invest effectively, but a minority make significant mistakes. This minority appears to be poorer and less well educated than the majority of more successful investors. There is some evidence that households understand their own limitations and avoid financial strategies for which they feel unqualified. Some financial products involve a cross-subsidy from naive to sophisticated households, and this can inhibit welfare-improving financial innovation.

A PRESIDENTIAL ADDRESS IS A PRIVILEGED OPPORTUNITY to ask questions without answering them, and to suggest answers without proving them. I use this opportunity to explore a field, household finance, that has attracted much recent interest but still lacks definition and status within our profession. Teaching and research are presently organized primarily around the traditional fields of asset pricing and corporate finance. Economists in the former field ask how asset prices are determined in capital markets and how average asset returns reflect risk. Economists in the latter field ask how business enterprises use financial instruments to further the interests of their owners, and in particular to resolve agency problems.

By analogy with corporate finance, household finance asks how households use financial instruments to attain their objectives. Household financial problems have many special features that give the field its character. Households must plan over long but finite horizons; they have important nontraded assets, notably their human capital; they hold illiquid assets, notably housing; they face constraints on their ability to borrow; and they are subject to complex taxation.

*Department of Economics, Harvard University and NBER. This paper was delivered as the Presidential Address to the American Finance Association on January 7, 2006. It reflects the intellectual contributions of colleagues, coauthors, and students too numerous to thank individually. I would like to acknowledge, however, the special influence of my dissertation advisers at Yale, Robert Shiller and the late James Tobin, and comments received from Robert Barro, Dan Bergstresser, Steve Cecchetti, Karine de Medeiros, Xavier Gabaix, Michael Haliassos, David Laibson, Anna Lusardi, Greg Mankiw, James Poterba, Tarun Ramadorai, Robert Shiller, Andrei Shleifer, Nick Souleles, Jeremy Stein, Sam Thompson, Luis Viceira, Tuomo Vuolteenaho, Moto Yogo, Luigi Zingales, and seminar participants at Harvard. I am grateful to Laurent Calvet and Paolo Sodini for allowing me to draw on our joint research in Section III. Allie Schwartz provided able research assistance for Sections II and IV. This material is based upon work supported by the National Science Foundation under Grant No. 0214061.



John Y. Campbell

Of course, household asset demands are important in asset pricing too. In that context, however, wealthy and risk-tolerant households have a disproportionate impact on equilibrium asset returns while in the household finance context the behavior of typical households is of greater interest.

Research in finance, as in other parts of economics, can be positive or normative. Positive research describes what economic agents actually do, while normative research prescribes what they should do. Economists have often assumed either that actual and ideal behavior coincide, or that they can be made to coincide by the selection of an appropriately rich model of agents' beliefs and preferences. Revealed preference theory (Samuelson (1938)), for example, shows how one can work backwards from a household's choices over multiple consumption goods to the implied preferences of the household. The revealed preference perspective leaves no room for normative economics as distinct from positive economics.¹

Household finance poses a particular challenge to this framework. Many households seek advice from financial planners and other experts, yet some households make decisions that are hard to reconcile with this advice or with any standard model. One response to this observation is to maintain the assumption that actual and ideal behavior coincide, but to consider nonstandard behavioral models of preferences that incorporate phenomena such as loss aversion and mental accounting. An alternative response is to abandon the framework of revealed preference and to consider the possibility that households may not express their preferences optimally. This response leads to the view that behavioral finance theory describes the choices households currently make, whereas standard finance theory describes the choices that maximize household welfare, and that households can be educated to make.

In this address I adopt the second response. I compare what we know about what households actually do—*positive household finance*—with our body of knowledge about what households should do—*normative household finance*. The comparison is not trivial to make. First, positive household finance requires high-quality data that are hard to obtain. Second, normative household finance requires significant extensions of textbook financial theory. I show that while for many households, the discrepancies between observed and ideal behavior have relatively minor consequences and can easily be rationalized by small frictions that are ignored in standard finance theory, for a minority of households, particularly poorer and less educated households, there are larger discrepancies with potentially serious consequences. I argue that these discrepancies, or *investment mistakes*, are central to the field of household finance.

It should not be surprising that some households make investment mistakes, given the complexity of their financial planning problem and the often confusing financial products that are offered to them. An important question is what determines the set of financial products available to households. This part of the field might be called *equilibrium household finance*. I suggest that retail

¹ My colleague Robert Barro summarized this view with characteristic sharpness when he told me that normative economics is "what you do when your model fails to fit the data."

financial innovation is slowed by the cost of advertising and educating households, together with the weakness of patent protection for financial products. In addition, using my presidential privilege I speculate that the existence of naive households permits an equilibrium of the sort described by Gabaix and Laibson (2006), in which confusing financial products generate a cross-subsidy from naive to sophisticated households, and in which no market participant has an incentive to eliminate this cross-subsidy.

If households make investment mistakes, it may be possible for financial economists to offer remedies that reduce the incidence and welfare costs of these mistakes. As a financial educator, I am tempted to call for an expansion of financial education. However, academic finance may have more to offer by influencing consumer regulation, disclosure rules, and the provision of investment default options. Work on these topics, which might be called *household financial engineering*, offers a powerful practical rationale for the study of household finance.

The organization of the paper is as follows. Section I summarizes the empirical and theoretical challenges faced by researchers studying household finance. Section II discusses household participation and asset allocation decisions, and Section III studies diversification of risky asset holdings. Section IV uses the choice of a mortgage, one of the most important financial decisions a household must make, to illustrate the themes of household finance. This treatment of mortgages contrasts with the body of research conducted by real estate specialists or asset pricing economists interested in the valuation and hedging of mortgage-backed securities. Section V considers barriers to innovation in retail financial markets, and Section VI concludes.

I. Two Challenges of Household Finance

A. Measurement

Positive household finance asks how households actually invest. While this is a conceptually straightforward question, it is hard to answer because the necessary data are hard to obtain. One reason is that households tend to guard their financial privacy jealously: Indeed, it may be more unusual today for people to reveal intimate details of their financial affairs than to reveal details of their intimate affairs. In addition, many households have complicated finances, with multiple accounts at different financial institutions that have different tax status and include both mutual funds and individual stocks and bonds. Even households that wish to provide data may have some difficulty answering detailed questions accurately.

The ideal data set for positive household finance would have at least five characteristics. First, it would cover a representative sample of the entire population. It is particularly important to have good coverage by both age and wealth, since many aspects of financial behavior vary with these characteristics. Second, for each household the data set would measure both total wealth and an exhaustive breakdown of wealth into relevant categories. Third, these categories

would be sufficiently disaggregated to distinguish among asset classes, and ideally would capture specific individual assets so that one could measure household diversification within asset classes. Fourth, the data would be reported with a high level of accuracy. Finally, the data set would follow households over time; that is, it would be a panel data set rather than a series of cross-sections.

Most work on household portfolio choice relies on surveys. The U.S. survey with the best data on financial wealth is generally thought to be the Survey of Consumer Finances (SCF).² The SCF scores highly on the first two criteria listed above. With respect to the first criterion, the SCF has good coverage and it oversamples the wealthy, who have a disproportionate influence on asset demands. With respect to the second criterion, it covers all aspects of wealth including both liquid and illiquid assets. However the SCF is less satisfactory in the other three respects. It is only disaggregated enough to address questions of asset allocation, and cannot shed light on diversification, because it does not report holdings of individual assets. Moreover, like any survey, the SCF relies on the willingness of households to participate and the accuracy of the answers they give when they do participate. Kennickell (1998) reports that in 1995, about one-third of households chosen for the standard sample refused to participate in the survey. The refusal rate was higher in the high-wealth sample: 56% for households in the lowest high-wealth stratum, and 87% for extremely wealthy households in the highest stratum. Even households that do participate in the survey may fail to report certain items. In 1995, for example, 64% of stockholding households reported a numerical value for their stockholdings, 21% stated that their stockholdings fell within a range, and 15% did not report any value. Finally, since 1989 the SCF has not followed households over time; rather, it interviews a fresh sample of households every 3 years.³

There are several ways to improve and verify the quality of survey data. For instance, the problem of nonresponse can be mitigated by offering households the opportunity to state a range for asset holdings rather than a precise answer, and possibly using follow-up questions to narrow the initially given range. This approach has been used with considerable success in the Health and Retirement Survey, which offers an attractive alternative to the SCF for measuring the behavior of older households (Juster and Smith (1997), Juster, Smith, and Stafford (1999)). To verify data quality, it is sometimes possible to cross-check survey evidence against objective external data. For example, the 1975 Wharton Survey of individual stockholders cross-checks self-reported share ownership data against corporate records compiled by the New York Stock Exchange (Blume

² Recent studies that use the SCF include Bergstresser and Poterba (2004), Bertaut and Starr-McCluer (2002), Carroll (2002), Heaton and Lucas (2000), Poterba and Samwick (2003), and Tracy, Schneider and Chan (1999).

³ Other surveys have similar problems. The Panel Study of Income Dynamics (PSID) asks questions on wealth every 5 years, but financial assets are divided into only three broad categories that correspond roughly to cash, bonds, and stocks (Mankiw and Zeldes (1991), Banks, Blundell, and Smith (2005)). Taking a different approach, the UBS/Gallup survey (Vissing-Jorgensen (2003), Graham, Harvey, and Huang (2005)) relies on telephone interviews; however, this limits the complexity of the questions that can be asked.

and Friend (1978)), and the decennial Residential Finance Survey interviews residents of housing units and then contacts their mortgage lenders to verify the data they provide on their mortgages.

Given the deficiencies of even the best survey data, recent work explores alternative data sources. One approach uses the records of companies that act as custodians for households or lend to households. Schlarbaum, Lewellen, and Lease (1978), who pioneer this approach, study approximately 3,000 retail brokerage accounts held during the 1960s. Similarly, Odean (1998, 1999) studies 10,000 discount brokerage accounts held during the 1990s and Barber and Odean (2000) expand the sample to 78,000 accounts at the same discount brokerage.⁴ While these brokerage records are highly accurate reports of holdings and trades in individual stocks, they sample customers of the brokerage house rather than the entire population. Further, these records do not necessarily represent the total wealth of even these customers, as they may have other accounts elsewhere. Similar difficulties afflict recent studies of asset allocation in 401(k) accounts and other tax-favored retirement accounts.⁵

Another approach exploits the fact that some countries maintain centralized registers of share ownership. For example, Grinblatt and Keloharju (2000, 2001) use data from the Finnish Central Securities Depository to measure daily transactions in and holdings of Finnish equities by Finnish individuals and institutions and foreign investors. However, while these data provide a comprehensive picture of Finnish household trading behavior in individual stocks, they do not reveal households' indirect holdings through mutual funds, their holdings of foreign stocks, or their allocations to other asset classes.

Finally, government tax records provide a tried-and-true source of financial data, which are highly accurate in countries that effectively limit tax evasion. Blume and Friend (1975), for example, use dividend payments reported on U.S. income tax returns, together with dividend-price ratios, to infer taxpayers' holdings of individual stocks. Unfortunately, this method of inferring wealth from reported income gives only a partial picture of household assets. The U.S. tax system requires reporting of wealth itself only in connection with the estate tax, which is levied only on the holdings of the very rich at the date of death. Blume and Friend (1978) and Kopczuk and Saez (2004) use U.S. estate tax records to study household asset allocation.

In Sweden, by contrast, households are liable to pay a wealth tax throughout their lives. As a consequence the Swedish government has detailed records of households' financial assets. Calvet, Campbell, and Sodini (2006) use these records to construct a panel of wealth and income data covering the entire population of Sweden (almost 5 million households). The data set provides highly disaggregated information on the income, wealth, demographic composition,

⁴ Other recent papers that use the large discount brokerage data set include Barber and Odean (2001), Zhu (2002), Ivkovic and Weisbenner (2003), Ivkovic, Sialm, and Weisbenner (2004), and Goetzmann and Kumar (2004). Feng and Seasholes (2004) analyze a similar data set from China.

⁵ Recent studies of such accounts include Agnew, Balduzzi, and Sunden (2003), Ameriks and Zeldes (2004), Benartzi and Thaler (2001), Choi et al. (2002, 2004a,b), and Madrian and Shea (2001).

education, and location of all households. Financial asset, mutual fund, and real estate portfolios are measured at the single property or security level, and each individual can be followed over time. The income data begin in 1983 and the wealth data in 1999. This data set affords the unique opportunity to analyze the financial behavior of the entire population of an industrialized country.⁶

B. Modeling

Normative household finance asks how households should invest. This is a challenging question to address because household decision problems involve many complications that are neglected by standard textbooks. Perhaps most obviously, households must plan their financial strategies over a lifetime rather than over a single short period. The seminal work of Merton (1971, 1973) introduces a conceptual framework for long-term financial planning given time-varying investment opportunities. Merton emphasizes that long-term investors must consider not only risks to their wealth, but also risks to the productivity of their wealth, that is, the rate of return at which wealth can be reinvested. This implies that long-term investors should hedge not only shocks to wealth itself, but also shocks to any state variable that predicts expected returns on wealth.

Because the Merton framework is much harder to work with than the traditional mean–variance analysis of portfolio choice, it was not until the 1990s that empirically usable versions of the Merton model emerged. One branch of the literature concentrates on shocks to the real interest rate, assuming that all movements in investment opportunities are captured by this variable—that is, assuming that risk premia are constant through time (Campbell and Viceira (2001), Wachter (2003)). A second branch of the literature concentrates on the equity premium, assuming that it follows an exogenous time-series process such as an AR(1) (Kim and Omberg (1996), Campbell and Viceira (1999)). More recent work allows for general multivariate processes that drive both interest rates and risk premia on multiple assets (Brennan, Schwartz, and Lagnado (1997), Lynch (2001), Campbell and Viceira (2003)).

An appealing feature of these models is that they can explain some obvious discrepancies between the predictions of mean-variance analysis and the financial planning advice that is usually offered to households. For example, Canner, Mankiw, and Weil (1997) point out that, contrary to the mutual fund theorem of Tobin (1958), financial planners typically advise conservative investors to hold more bonds relative to stocks in their risky portfolios. Campbell and Viceira (2001) show that this advice makes sense if bonds are hedges against time-variation in interest rates.

⁶ Massa and Simonov (2006) also study the portfolios of Swedish households by merging survey data on income, wealth, and asset allocation with data on individual stock ownership of Swedish companies from 1995 to 2000. Stock ownership data are available for this period since Swedish companies were legally required to report the identity of most of their shareholders.

An important theme in this work is the distinction between real and nominal magnitudes. The risk properties of long-term nominal bonds, for example, depend critically on the properties that one assumes for inflation. If inflation is well controlled, then nominal bonds are safe assets for long-term investors; otherwise, these assets are highly risky and poor proxies for inflation-indexed bonds. Because inflation shocks are persistent, the distinction between real and nominal bonds is much more important than the distinction between real and nominal bills in short-horizon models.

Models in the Merton tradition assume that all wealth is held in a liquid, easily tradable form. However, the largest component of wealth for most households is human capital, which is nontradable. Put differently, households receive labor income but cannot sell claims to that income. If labor income is perfectly correlated with traded assets, and if households can short those assets, then households can hedge their labor income risk and undo the effects of labor income on their total portfolio (Bodie Merton, and Samuelson (1992)). In practice, however, much of the risk in labor income is idiosyncratic and therefore unhedgeable. This background risk increases effective risk aversion and leads households to invest more cautiously (Heaton and Lucas (2000), Viceira (2001)). On the other hand, to the extent that some households can increase their labor supply in response to poor investment returns, either by increasing hours worked or by delaying their retirement, this added flexibility increases households' willingness to take financial risks (Bodie et al. (1992), Farhi and Panageas (2005)).⁷

There is a debate in the literature about the risk properties of labor income. Some authors find that labor income is similar to an implicit holding of safe assets, stimulating investment in risky financial assets (Cocco, Gomes, and Maenhout (2005)); others argue that labor income and capital income covary in the long run (Benzoni, Collin-Dufresne, and Goldstein (2005)), or that the volatility of idiosyncratic labor income risk covaries negatively with stock returns (Lynch and Tan (2004), Storesletten, Telmer, and Yaron (2004)), leading labor income to crowd out stock market investments.

Housing is an asset class of dominant importance for middle-class homeowners. Houses are long-term assets that deliver a stream of housing services to their owners; in this sense they are like long-term bonds and can be used to hedge changes in the relative price of housing and nonhousing consumption (Pelizzon and Weber (2005), Sinai and Souleles (2005)). But houses are also illiquid assets, so homeowners find it costly to adjust their consumption of housing services in response to economic shocks. This illiquidity may discourage homeownership and financial risk taking by homeowners.⁸

⁷ A small recent literature builds on the human capital literature in labor economics and treats education as a risky investment that is chosen jointly with risky financial assets (Palacios-Huerta (2003), Saks and Shore (2005)).

⁸ See Cocco (2005), Flavin and Nakagawa (2004), Fratantoni (2001), Shore and Sinai (2005), and Yao and Zhang (2005). Davidoff (2006) finds that the covariance between labor income and house prices affects the size of a household's position in owner-occupied housing.

Housing, unlike labor income, provides collateral that can be used to facilitate borrowing. Another important aspect of household finance is the existence of borrowing constraints.⁹ Households must consider the fact that their future consumption may be determined not just by their wealth and investment opportunities, but also by their net future income if they are borrowing constrained. Financial investments that do poorly when income is temporarily low may be unattractive for this reason.

Borrowing constraints are likely to be more important for young households than for older households that have built up some retirement savings. Life cycle aspects of household finance also complicate the normative theory, because one cannot use stationary infinite horizon models but instead must use more complicated finite horizon models that capture the evolution of financial strategy as households age and accumulate financial assets.¹⁰

Finally, household decisions must take into account the complexities and non-neutralities of the tax code. Relevant complications include the taxation of nominal rather than real interest, the availability of tax-favored retirement accounts, the tax deductibility of mortgage interest, the taxation of capital gains only when these gains are realized through asset sales, and the adjustment of the capital gains tax basis at death.¹¹

A particularly important example that illustrates how these considerations can interact with one another is the choice between an adjustable-rate mortgage (ARM) and a fixed-rate mortgage (FRM). An ARM is effectively a floating-rate note issued by a household, while a FRM is a long-term nominal bond, typically with a call option that allows the household to repay its loan at face value and refinance the mortgage if interest rate movements make it desirable to do so. In textbook financial theory, a floating-rate note is a safer instrument than a long-term nominal bond; it has a stable value that is almost unaffected by movements in interest rates, while the value of a long-term bond is highly sensitive to interest rates. Yet financial planners typically describe ARMs as risky for households.¹²

This apparent paradox can be resolved by taking into account two special characteristics of the household financial problem. First, the household is planning over a long horizon. If real interest rates vary, then an ARM exposes the

⁹ A variety of studies find that consumption responds to predictable changes in income in a manner that suggests the relevance of borrowing constraints for many households. Particularly convincing are recent microeconomic studies of social security tax withholdings (Parker (1999)), income tax refunds (Souleles (1999), Johnson, Parker, and Souleles (2004)), and paycheck receipts (Stephens (2006)).

¹⁰ Gourinchas and Parker (2002) provide a workhorse model of saving over the life cycle in the presence of risky labor income. Cocco et al. (2005) extend the model to allow for portfolio choice. Davis, Kubler, and Willen (2006) argue that the lowest and most realistic equity demands result from borrowing that is expensive (costing the average rate of return on equity) rather than impossible. The possibility of expensive borrowing reduces precautionary saving and thus equity demand later in life.

¹¹ Poterba (2002) surveys the vast literature on taxation and optimal portfolio choice. Recent contributions that emphasize life cycle effects include Dammon, Spatt, and Zhang (2001, 2004) and Gomes, Michaelides, and Polkovnichenko (2004).

¹² Fisher and Shelly (2002), for example, write "An ARM can pay off, but it's a gamble. Sometimes there's a lot to be said for something that's safe and dependable, like a FRM," (p. 319).

household to the risk that real borrowing costs will increase. The household may wish to hedge this risk, as in the Merton framework, by using a long-term FRM. The ideal instrument for this purpose would be an inflation-indexed mortgage, but if inflation risk is modest a nominal FRM may be an adequate proxy.

Second, the household faces the risk that borrowing constraints may bind in future periods. If future income declines temporarily, the household may wish to borrow; it may be unable to do so, however, if future housing prices have fallen such that collateral is unavailable. If future borrowing constraints are a concern, then ARMs are relatively risky even when real interest rates are constant. To see this, consider what happens when expected future inflation increases. The nominal interest rate, and hence the required monthly payments on an ARM, increase even though the price level has not yet increased. This accelerated repayment of the loan compensates the lender for future inflation. It has no effect on a household that can borrow to make the accelerated payments required by the ARM, but it reduces the consumption of a borrowing-constrained household.

Campbell and Cocco (2003) solve a numerical model of household mortgage choice and show that ARMs should be attractive to unconstrained households when inflation risk is large relative to real interest rate risk and to potentially borrowing-constrained households with low risk aversion; they should be unattractive to risk-averse borrowing-constrained households, particularly those that have high mortgage debt relative to their income. In this paper, Appendix A presents a simple analytical model in which the same points can be understood.

A fundamental issue that confronts the normative literature is how to specify the household utility function. It is common to assume that households have time-separable power utility or Epstein–Zin (1989) utility, so that their relative risk aversion does not vary with their wealth. Asset pricing models with this feature capture the stability of interest rates and asset valuation ratios in the face of long-run economic growth. However, some aspects of short-run asset price behavior and cross-sectional variation in risk taking suggest that relative risk aversion declines with wealth. Carroll (2002), for instance, proposes a model in which bequest utility has lower curvature than consumption utility, so that risk aversion falls as households accumulate wealth over the life cycle. Models of habit formation (Campbell and Cochrane (1999)) or consumption commitments (Chetty and Szeidl (2005)) imply further that risk aversion fluctuates with short-term movements in wealth. Ultimately it should be possible to assess these alternative models by their consistency with the behavior of households that appear to be more sophisticated, or with the advice of financial planners (Canner et al. (1997)). Until some consensus is reached, normative household finance should emphasize results that are robust to alternative specifications of household utility.

II. Participation and Asset Allocation

How do households allocate their assets across broad categories such as money market instruments, bonds, equities, and real estate? More specifically:




Figure 1. The U.S. wealth distribution. The cross-sectional distribution of total assets, financial assets, and net worth in the 2001 Survey of Consumer Finances.

How many households participate in these markets at all? Given that they have decided to participate, what fraction of their assets do they allocate to each category? How does household behavior vary with age, wealth, and other household characteristics? Because these questions can be answered without having detailed information on individual asset holdings, the data problems described in Section I.A are not as serious in this context.

Following Bertaut and Starr-McCluer (2002), Haliassos and Bertaut (1995), and Tracy, Schneider, and Chan (1999), I now summarize the information in the 2001 SCF that relates to these questions. Figure 1 presents the cross-sectional wealth distribution. The horizontal axis in this figure shows the percentiles of the distribution of total assets, defined broadly to include both financial assets and nonfinancial assets (durable goods, real estate, and private business equity, but not defined benefit pension plans or human capital). The vertical axis reports dollars on a log scale. The three lines in the figure show the average levels of total assets, financial assets, and net worth (total assets less debts, including mortgages, home equity loans, credit card debt, and other debt) at each percentile of the total assets distribution. It is clear from the figure that many households have negligible financial assets. Even the median household has financial assets of only \$35,000, net worth of \$86,000, and total assets of \$135,000.

The figure also shows the extreme skewness of the wealth distribution. Wealthy households at the right of the figure have an overwhelming influence on aggregate statistics. To the extent that these households behave differently from households in the middle of the wealth distribution, the aggregates tell




Figure 2. Participation rates by asset class. The cross-sectional distribution of asset class participation rates in the 2001 Survey of Consumer Finances.

us very little about the financial decision making of a typical household: Again, while the behavior of wealthy households is disproportionately important for asset pricing models, household finance is more concerned with the behavior and welfare of typical households.

A. Wealth Effects

Figure 2 illustrates the participation decisions of households with different levels of wealth. The horizontal axis is the same as in Figure 1, but the vertical axis now shows the fraction of households that participate in particular asset classes. In this figure I aggregate the SCF asset data into several broad categories, namely, safe assets, vehicles, real estate, public equity, private business assets, and bonds.¹³

Given the negligible financial assets held by households at the left of the figure, it should not be surprising that these households often fail to participate in risky financial markets. Standard financial theory predicts that households should take at least some amount of any gamble with a positive expected return, but this result ignores fixed costs of participation, which can easily overwhelm the gain from participation at low levels of wealth. Figure 2 shows that most

¹³ Safe assets include checking, saving, money market, and call accounts, CDs, and U.S. savings bonds. Public equity includes stocks and mutual funds held in taxable or retirement accounts or trusts. Bonds include government bonds other than U.S. savings bonds, municipal, corporate, foreign, and mortgage-backed bonds, cash-value life insurance, and amounts in mutual funds, retirement accounts, trusts, and other managed assets that are not invested in stock.

households in the bottom quartile of the wealth distribution hold only liquid assets and vehicles, with a minority participating in real estate through homeownership.

As we move to the right in the figure we see that an increasing fraction of households participate in equity markets. Participation is far from universal, however, even among quite wealthy households. This finding has also been emphasized by Mankiw and Zeldes (1991), Haliassos and Bertaut (1995), and Heaton and Lucas (2000). Limited participation among the wealthy poses a significant challenge to financial theory and is one of the main stylized facts of household finance. At the 80th percentile of the wealth distribution, for example, a typical household has about \$200,000 in financial assets, but almost 20% of these households own no public equity.

Many wealthy households have significant private business assets. Gentry and Hubbard (2004) report that private business owners hold as much as 40% of total net worth even though they comprise less than 10% of the population, implying that these households are particularly important for aggregate asset demands and hence for asset pricing. Figure 2 shows that the fraction of business owners increases from 22% at the 80th percentile of the wealth distribution to 70% at the right tail of the distribution. Heaton and Lucas (2000) emphasize that private business assets substitute for public equity in the portfolios of some wealthy households. The fraction of households at the 80th percentile of the wealth distribution that hold neither private business assets nor public equity is just under 10%. Thus, private business assets can explain much of the nonparticipation in public equity markets by wealthy households, but there remains a significant number of these households who have no exposure to equity risk of any kind.

Figure 3 illustrates the asset allocation decisions of households with different levels of wealth. The horizontal axis is the same as in the two previous figures, but the vertical axis now depicts the weight of an asset class in the aggregate portfolio of households at each level of wealth (equivalently, the wealth-weighted mean share of the asset class, which is almost identical to the unweighted mean share for households within a given wealth range). The figure illustrates the dominant role of liquid assets and vehicles for the poor, and real estate—primarily owner-occupied housing—for middle-class households. Equity has some importance for the middle class, but represents the largest portfolio share only for wealthier households at the right of the figure.¹⁴

Figure 3 shows that wealthy households are willing to take greater risk in their portfolios. This is partly the result of greater participation in risky asset classes by wealthy households, but also partly the result of higher portfolio shares conditional on participation. Carroll (2002) emphasizes this

¹⁴ Tracy et al. (1999) show a similar figure for the median portfolio share rather than the wealth-weighted mean portfolio share. Joe Tracy calls their figure a “whale chart” because the real estate line defines the body of a whale, while the equity line traces out its tail. Kopeczuk and Saez (2004), using estate tax returns to look at the extreme right tail of the wealth distribution, find that the real estate share continues to decline, while the equity share continues to increase, within the top 0.5% of the population.




Figure 3. Asset class shares in household portfolios. The share of each asset class in the aggregate portfolio of households at each point in the wealth distribution, in the 2001 Survey of Consumer Finances.

phenomenon and shows that similar patterns are obtained in several European countries.¹⁵

B. Demographic Effects

Wealth is not the only household characteristic that may predict its willingness to take financial risk. Income, age, race, education, and self-reported attitudes to risk may also be important.

Before one can understand the relative importance of these effects, one must confront a fundamental identification problem (Heckman and Robb (1985), Ameriks and Zeldes (2004)). At any time t a person born in year b is a_t years old, where $a_t = t - b$. Thus, it is inherently impossible to separately identify age effects, time effects, and cohort (birth year) effects on portfolio choice. Even if one has complete panel data on portfolios of households over time, any pattern in the data can be fit equally well by age and time effects, age and cohort effects, or time and cohort effects.

Theory suggests that there should be time effects on portfolio choice if households perceive changes over time in the risks or expected excess returns of risky assets. Theory also suggests that there should be age effects on portfolio choice if older investors have shorter horizons than younger investors and investment

¹⁵ King and Leape (1998) capture the same phenomenon by estimating wealth elasticities of demand for different asset classes. They find that risky assets tend to be luxury goods with high wealth elasticities.

opportunities are time-varying, or if older investors have less human wealth relative to financial wealth than younger investors (Bodie et al. (1992), Campbell and Viceira (2002)). Thus, it seems hard to rule out either time or age effects in studying portfolio choice. Cohort effects are more problematic. In principle cohort effects could be caused by different labor market experiences that affect the ratio of human to financial wealth held by a cohort at each age, but this effect is unlikely to be strong in modern U.S. conditions. Cohort effects could also arise from differences in preferences, perhaps driven by different asset market experiences. Such effects cannot be identified by the data without modeling them (or age or time effects) in some way. Accordingly, I follow Heaton and Lucas (2000) and most other studies by setting cohort effects to zero. Under this assumption age effects can be estimated in any cross-section.

Table I summarizes demographic effects on asset allocation in the 2001 SCF. The left panel of the table reports logit regressions of asset class participation on household income, wealth, and demographic characteristics.¹⁶ The right panel reports regressions of portfolio shares on the same variables, conditional on participation. Within each panel, the first regression looks at public equity (including equity held in retirement accounts), and the second regression looks at private business assets. Standard errors are reported below each coefficient, and coefficients significant at the 10% level or better are indicated with stars. To illustrate the quantitative importance of each effect in the logit regressions, the table also reports the participation probability for a reference household, and the change in this probability caused by a change in each dummy variable from zero to one, or a one-standard-deviation change in each continuous variable.

The table shows that in the United States in 2001, there was a weak negative age effect on participation in public equity markets.¹⁷ This result is presumably due to increased participation by younger households during the 1990s and the fact that the regression controls for wealth and income, which tend to be higher for middle-aged households. Unsurprisingly, households that report they have no tolerance for investment risk are less likely to hold public equity. There are strong positive effects of education, income, and wealth on public equity participation.

The results are somewhat different for private business ownership. In this case the age effect is hump-shaped, reflecting the tendency for younger households to acquire and older households to sell off private businesses. Income has a U-shaped effect on the incidence of private business ownership, with a minimum at an income of \$250,000, and wealth has an extremely powerful, but imprecisely estimated, quadratically increasing effect. Both these variables capture the strong tendency for the richest and highest-income households to

¹⁶ The influence of outliers is limited by truncating income and wealth at the 1st and 99th percentiles of the cross-sectional distribution. The regressions use logs of income and wealth, but results are similar using the level rather than the log of income.

¹⁷ Stronger age effects on equity market participation are reported by Bertaut and Starr-McCluer (2002) for earlier SCF data, by Banks and Tanner (2002) for the U.K., by Guiso and Jappelli (2002) for Italy, and by Iwaisako (2003) for Japan.

Table I
Equity Participation and Portfolio Share

The table reports demographic determinants of participation in public equity and private business ownership (logit regressions, left panel) and of the portfolio share among participants (OLS regressions, right panel) for households in the 2001 Survey of Consumer Finances. Standard errors are reported underneath the coefficients in parentheses. Coefficients significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***. All regressions control for gender and marital status. In the reference household, the household head is a married white male with no high school diploma. The column headed "Probability Estimates" reports the probability of participation for the reference household, and the change in this probability caused by a unit change in a binary variable and a one-standard-deviation change in a continuous variable.

Dependent Variable	Whether a Household Owns a Given Asset				Portfolio Share for Participants	
	Public Equity		Private Business		Public Equity Coefficients	Private Business Coefficients
	Coefficients	Probability Estimates (%)	Coefficients	Probability Estimates (%)		
<i>Reference household</i>						
No tolerance for investment risk	-1.173*** (0.121)	56.5 -40.9	0.189 (0.167)	0.7 0.5	-0.032** (0.014)	0.061* (0.032)
Age	-0.034* (0.019)	-20.2	0.094*** (0.031)	-0.6	0.002 (0.002)	-0.003 (0.008)
Age squared	3.9×10^{-5} (1.8×10^{-4})		-1.3×10^{-3} *** (3.0×10^{-4})		-1.7×10^{-5} (2.4×10^{-5})	-5.6×10^{-6} (7.6×10^{-5})
White	0.133 (0.128)	5.3	0.335* (0.198)	0.4	0.032** (0.014)	0.000 (0.038)
High school diploma	0.401*** (0.157)	14.9	-0.086 (0.236)	-0.2	0.009 (0.020)	0.028 (0.048)
College diploma	0.848*** (0.186)	27.9	-0.211 (0.275)	-0.3	0.037* (0.021)	-0.041 (0.052)
Graduate school	1.051*** (0.239)	32.2	-0.210 (0.280)	-0.3	0.040* (0.023)	-0.038 (0.051)
Number of children	-0.010 (0.049)	-0.5	-0.068 (0.058)	-0.1	-0.011** (0.005)	0.003 (0.011)
Ln(income)	2.650* (1.452)	17.3	-2.507** (1.036)	-0.4	0.038 (0.153)	0.106 (0.283)
Ln(income) squared	-0.101 (0.070)		0.101** (0.047)		-0.002 (0.007)	-0.007 (0.012)
Ln(wealth)	0.094 (0.372)	37.3	0.315 (0.583)	30.6	-0.312*** (0.047)	-0.409*** (0.120)
Ln(wealth) squared	0.026 (0.017)		0.029 (0.023)		0.014*** (0.002)	0.018*** (0.005)
Sample size	4,304		4,304		2,822	1,245

own private businesses. White households are more likely to own private businesses, but there are no significant effects of education.

Turning to portfolio shares for participants, the main influence is wealth. For both public equity and private business assets there is a quadratic pattern with a minimum share at \$70,000 for public equity and \$85,000 for private equity. This pattern reflects the fact that low-wealth households are likely to hold large portfolio shares if they participate at all, but in the upper part of the wealth distribution portfolio shares are strongly increasing in wealth. White and educated households have higher portfolio shares in public equity than other households.

The regressions in Table I omit some variables that have been found to be important in other studies. Bertaut and Starr-McCluer (2002) show that defined benefit pension rights increase the allocation to risky assets, while self-employment decreases it. Using the Health and Retirement Study, Rosen and Wu (2004) show that poor self-reported health decreases the allocation to risky assets. These effects work strongly through the participation decision, and also to some extent through asset allocation conditional on participation. Poterba and Samwick (2003) show that households with higher marginal tax rates are more likely to hold tax-advantaged assets such as stock and tax-exempt bonds, and are more likely to hold assets in tax-deferred accounts.

These results describe household asset allocation at a point in time, but I do not attempt to follow households through time to see how their asset allocations evolve. A small recent literature finds strong evidence for inertia in asset allocation. Participants in retirement savings plans rarely alter the allocations of their contributions or rebalance their portfolios, and default options have long-lasting effects on these portfolios (Agnew et al. (2003), Ameriks and Zeldes (2004), Choi et al. (2002, 2004b), Madrian and Shea (2001)). Capital gains and losses also generate little portfolio rebalancing in U.S. survey data studied by Brunnermeier and Nagel (2005).

C. Interpretation

How can we make sense of the above empirical results? Textbook financial theory implies that all households, no matter how risk averse, should hold some equities as long as the equity premium is positive. It follows that limited participation in the equity market must be due to a failure of one of the standard assumptions.

One possibility is that some households are unaware of the existence of stocks as an asset class: Over 35% of Italian households reported that they were unaware of stocks in the late 1990s (Guiso and Jappelli (2005)); however, this proportion is likely to be much smaller in the United States. Alternatively, households may have nonstandard preferences or may face fixed costs, for example one-time entry costs or ongoing participation costs. Fixed costs can explain why participation increases with wealth, since a larger portfolio is more likely to justify the payment of a fixed cost to increase return. One-time entry costs imply strong positive age effects, since a household will continue to participate

once a fixed entry cost has been paid; ongoing participation costs produce limited participation with weaker age effects. Haliassos and Bertaut (1995) and Vissing-Jorgensen (2003) show that moderate ongoing participation costs can explain the nonparticipation of many U.S. households, although not the richest households.

Fixed participation costs can be interpreted in different ways. One approach is to think of fixed costs as capturing time and money that must be spent in order to invest in the stock market. Vissing-Jorgensen (2003), for example, points out that equity ownership often complicates the preparation of tax returns. Alternatively, fixed costs may be an economist's description of psychological factors that make equity ownership uncomfortable for some households. Hong, Kubik, and Stein (2004), for example, find that households that interact less with other households in their community are less likely to own stocks, suggesting that households prefer to follow financial practices that they know they share with others. Similarly, Guiso, Sapienza, and Zingales (2005) find that households that express reluctance to trust others are less likely to own stocks. According to this interpretation, nonparticipation can be regarded as an investment mistake, one that households with high fixed costs are more likely to make.¹⁸

Both interpretations must confront the important effect of education on equity ownership. Table I shows that education directly predicts equity ownership, even controlling for age, income, and wealth. Guiso et al. (2005) report that the effect of trust on equity ownership is weaker for educated households. It is tempting to conclude that educated households have learned that nonparticipation is an investment mistake, or as Haliassos and Bertaut (1995) put it, that "education and the free acquisition of information are important in overcoming the barrier to stockholding erected by ignorance and misperceptions." While this conclusion is probably correct, it is also plausible that education reduces the objective costs of stock market participation. Related to this idea, in the next section I report that educated households in Sweden diversify their portfolios more efficiently, and therefore can expect to earn higher returns per unit of risk if they do participate.

An interesting question is whether stock market participants are more risk tolerant than nonparticipants. If nonparticipants are relatively risk averse, then small fixed costs suffice to deter them from participation. Carroll (2002) proposes a model in which all agents have a common utility function with declining relative risk aversion, and argues that this model explains the high participation rate and more aggressive asset allocation of wealthy households. However, Haliassos and Michaelides (2003) and Gomes and Michaelides (2005) argue that risk-averse households have a strong precautionary saving motive, which leads them to accumulate more wealth. If there is exogenous cross-sectional variation in risk aversion and the precautionary saving effect is sufficiently strong, those households that are wealthy enough to pay the fixed costs

¹⁸ A similar interpretation issue arises if one accounts for inertia in asset allocation by invoking fixed costs of portfolio rebalancing. Such fixed costs may be objective, or they may simply be another way to describe an investment mistake.

of stock market participation may actually be more risk averse than nonparticipating households.

Other features of the data can be explained by the effect of background risk on portfolio choice. Self-employed households and households with significant private business assets are exposed to private business risk that increases their effective risk aversion even if it is uncorrelated with returns on publicly traded equities. Private business risk has an even stronger discouraging effect on equity ownership if, as seems plausible, it is positively correlated with public equity risk. When I include a dummy for private business ownership in the public equity participation regression of Table I, it enters negatively although not significantly. The effect of poor health on asset allocation can also be understood as a background risk effect; in this case the risk is to spending needs rather than to income.

III. Diversification

A second major topic in household finance is how households construct their portfolios within each asset class. Accurate measurement is significantly more challenging in this context because we would ideally like to measure the holdings of each individual asset, and survey data do not generally give this much detail. However, a large and ingenious empirical literature explores the composition of household stock portfolios extracting information from surveys (for example, on the decision to hold any individual stocks, and the number of individual stocks held), tax returns (which list dividends by payer and thus reveal individual stockholdings), and brokerage accounts. The main conclusions of this literature are as follows.

First, many households own relatively few individual stocks. Analyses of the SCF find that among households that hold individual stocks directly, the median number of stocks held was two until 2001, when it rose to three (Blume and Friend (1975), Kelly (1995), Polkovnichenko (2005)). Of course, many households own equity indirectly, through mutual funds or retirement accounts, and these indirect holdings tend to be much better diversified.¹⁹ Thus, it is not clear that concentrated individual stockholdings have a large effect on household portfolio risk.

Second, the local bias or preference for local securities that has been found in aggregate data (French and Poterba (1991), Cooper and Kaplanis (1994)) shows up in household-level data both with respect to domestic versus foreign investments, and with respect to regional versus nonregional companies. Using company records, Huberman (2001) finds that individual investors prefer to own the stocks of their local telecommunications company. Using brokerage

¹⁹ Curcuru et al. (2004) define households to be undiversified if they hold more than 50% of their equity in a brokerage account with fewer than 10 stocks. They find that the fraction of undiversified households has been declining, from almost a third of stockholding households, owning 21% of equity, in the 1989 SCF to 14% of households, owning 12% of equity, in the 2001 SCF. Polkovnichenko (2005) also emphasizes the co-existence of mutual funds and individual stocks in household portfolios.

account data, Zhu (2002) finds that regional bias is stronger among investors who do not own international stocks, suggesting a connection between the two forms of local bias. Feng and Seasholes (2004) find that Chinese investors overweight not only local companies, but also companies that are traded on a local exchange, suggesting that familiarity drives local bias.

Third, many U.S. households have large holdings in the stock of their employer, particularly within their 401(k) retirement savings accounts (Mitchell and Utkus (2003)). While some of these holdings result from employer policies, Benartzi (2001) shows that a substantial fraction of unrestricted employee contributions go to employer stock rather than diversified alternatives. This is especially true when the employer stock has performed well over the previous decade, suggesting that households extrapolate the past performance of the employer.²⁰

Fourth, discount brokerage customers trade intensively (Odean (1999), Barber and Odean (2000)). This finding contrasts with the inertia in asset allocation found in studies of retirement savings plans, probably because discount brokerage customers tend to be households with a particular interest in equity trading. Brokerage customers also display a disposition effect, that is, a tendency to sell winners and hold losers. Odean (1998) shows that the propensity among brokerage customers to realize gains is substantially greater than the propensity to realize losses, except in December when tax-loss selling reverses the relationship. Selling winners can be a way to restore diversification to a portfolio that has become excessively concentrated, but the tendency to hold losers is hard to rationalize, both because it is tax-inefficient and because it lowers pre-tax returns to the extent that stocks display momentum.

Finally, there is heterogeneity in the strength of these effects across households. Puri and Robinson (2005), for example, show that households with optimistic beliefs about their life expectancy place a higher portfolio weight on individual stocks even though they do not place a higher overall weight on equities. Further, Graham, Harvey, and Huang (2005) find that investors who claim to be comfortable with investment products also tend to trade more frequently and to be more diversified internationally.

There is an active debate about the performance of concentrated portfolios. Odean (1999) finds that the stocks purchased by discount brokerage customers tend to underperform the stocks sold by these households. Zhu (2002) argues that households with a relatively weak local bias—which tend to have higher income and professional status—outperform households with a stronger local bias. On the other hand Ivković and Weisbenner (2003) find that households' local investments outperform their non-local ones, and Ivković et al. (2004) find that among wealthier households, concentrated portfolios have higher average returns than diversified portfolios although they also have higher risk and lower Sharpe ratios. All these studies use account-level data from a discount brokerage.

²⁰ However Choi et al. (2004a) find that 401(k) plan participants reallocate their portfolios by selling company stock after the stock rises, consistent with the disposition effect.

A. Household Risk Exposures in Sweden

A weakness of this literature is that it cannot directly measure households' risk exposures. Surveys do not identify individual stocks or mutual funds, and brokerage accounts do not reveal total portfolios. In joint research with Laurent Calvet and Paolo Sodini (Calvet et al. (2006)), we use Swedish data to look more directly at the idiosyncratic risk in Swedish household portfolios. We adopt the perspective that systematic risk is compensated and idiosyncratic risk is not, so that taking idiosyncratic risk is an investment mistake. Since the time dimension of our data set is short, we do not attempt to measure the performance of Swedish household portfolios directly.

The Swedish data appear to be broadly consistent with U.S. data as regards asset allocation: At the aggregate level, real estate accounts for over 70% of household assets, bank deposits and money market funds for 11%, directly held stocks and mutual funds for 6% each, and bonds, derivatives, and capital insurance products for the remainder. At the end of 2002, 62% of households participated in financial markets by holding financial assets other than bank deposits and money market funds.

We construct a sample of 100,000 households and measure the composition of their portfolios at the end of 2002, down to the level of individual stocks and mutual funds. We calculate the risk properties of these portfolios by estimating a variance–covariance matrix Σ for the returns of all stocks and mutual funds held by Swedish households. Then, if a household h has portfolio weight vector ω_h , the variance of its portfolio return is estimated as $\omega_h' \Sigma \omega_h$. This procedure captures the risk in household portfolios at a point in time; it does not track the trading decisions of households within the year.

The median household in our sample has a risky portfolio standard deviation of about 20%. Part of this standard deviation comes from exposure to systematic risk in the world equity market, and part comes from unsystematic risk. As a measure of systematic risk, we calculate the standard deviation of the fitted value in a regression of each household's portfolio return on the dollar excess return of the MSCI All World Index. For the median household, this systematic standard deviation is slightly less than the standard deviation of the residual, a measure of unsystematic risk, implying that more than half of the median Swedish household's portfolio variance is idiosyncratic.

Although Swedish households can obtain the dollar excess return on international stocks by hedging their currency exposure when they invest internationally, this may be an unrealistic benchmark given that international equity funds widely marketed in Sweden are not currency hedged. If we repeat the above exercise with the Swedish Krona excess return on the MSCI index, we find that slightly less than half of the median household's portfolio variance is idiosyncratic.

While the median standard deviation of the risky portfolio is about 20%, there is wide variation in this number across households. Some households take low risk and hold primarily bond funds; others take high risk. The 95th percentile of the risky portfolio standard deviation is about 50%, and the 99th percentile

is 70%. Portfolios with this level of risk tend to have betas above one, but they also have extremely high shares of idiosyncratic as opposed to systematic risk.

To analyze idiosyncratic risk in a portfolio, it is helpful to consider a stylized symmetrical model in which all the assets in the portfolio of household h have the same idiosyncratic variance σ_{ah}^2 and the same correlation ρ_{ah} . It is straightforward to show that the idiosyncratic variance of the household portfolio, σ_{ih}^2 , satisfies

$$\sigma_{ih}^2 = C_{ah}\sigma_{ah}^2 + (1 - C_{ah})\rho_{ah}\sigma_{ah}^2, \quad (1)$$

where $C_{ah} = \omega'_h \omega_h$ is a measure of the concentration of the overall portfolio. Let $\bar{c}_a$ denote the average value of $\log C_{ah}$ in the population, and $\bar{C}_a = \exp(\bar{c}_a)$. A log linearization of (1) around $\rho_{ah} = 0$ and $c_{ah} = \bar{c}_a$ implies

$$\log(\sigma_{ih}) \approx \log \sigma_{ah} + \frac{1}{2} \log C_{ah} + \frac{1}{2} \left(\frac{1}{\bar{C}_a} - 1 \right) \rho_{ah}. \quad (2)$$

This decomposition relates log idiosyncratic portfolio standard deviation to the log of the average idiosyncratic standard deviation of assets in the portfolio, the log concentration of the portfolio, and the average correlation across assets in the portfolio.

The above analysis treats all assets in the portfolio equally, whether they are stocks or mutual funds. An alternative approach is to assume that mutual funds are fully diversified, with zero idiosyncratic risk. Let D_h denote the share of directly held stocks in the risky portfolio, and let $C_{sh} = C_{ah}/D_h^2$ denote the concentration of the stock portion of the portfolio. Then

$$\log(\sigma_{ih}) \approx \log D_h + \log \sigma_{sh} + \frac{1}{2} \log C_{sh} + \frac{1}{2} \left(\frac{1}{\bar{C}_s} - 1 \right) \rho_{sh}, \quad (3)$$

where s subscripts denote the characteristics of the directly held stocks in the portfolio. This alternative decomposition attributes idiosyncratic risk to a high share of stocks rather than mutual funds in the portfolio, volatile stocks, a concentrated stock portfolio, and correlated stocks.

In the Swedish data, we find that portfolios with high idiosyncratic risk tend to have high shares of directly owned stocks, and the directly owned portfolios tend to be concentrated in one or two volatile stocks. Concentration, however, can be a misleading statistic; many portfolios with low idiosyncratic risk also contain one or two directly owned stocks, but these portfolios are dominated by mutual funds and contain only a small share of directly owned stocks. This pattern illustrates the danger of looking only at the number of directly held stocks in a portfolio without considering the broader context within which those stocks are held. Correlation across stocks in the portfolio contributes very little to the cross-sectional risk pattern in Swedish portfolios.

In order to evaluate the consequences of underdiversification for household welfare, we assume that mean returns on stocks and mutual funds obey an

international asset pricing model (either the CAPM or the three-factor Fama–French (1993) model, estimated in dollars). This assumption avoids the difficult task of estimating average returns on individual stocks and mutual funds from short historical time series, while enabling us to plot Swedish household portfolios on a mean-standard deviation diagram. By assumption, all portfolios must fall below the efficient frontier, which in the case of the international CAPM is a straight line connecting the riskless rate to the currency-hedged return on the MSCI world index. We find that many household portfolios come close to the Sharpe ratio of the unhedged world index (which we estimate at 35%), but almost none attain the efficient Sharpe ratio of the currency-hedged world index, which we estimate at 45%. The Swedish domestic equity index, with an estimated Sharpe ratio of 27%, lies within the middle of the distribution of household portfolios.

There are several ways to measure portfolio inefficiency within this framework. One is to calculate the percentage difference between a household portfolio's Sharpe ratio S_h and the Sharpe ratio of a benchmark index S_B , $1 - S_h/S_B$. A second approach is to calculate the return lost, at the portfolio's given standard deviation, by the lower Sharpe ratio of the household portfolio. This is $w_h(S_B\sigma_h - \mu_h)$, where w_h is the portfolio's weight in risky assets, σ_h is the standard deviation of the household's risky portfolio return, and μ_h is the mean of that return. A third approach is to calculate the utility lost by a household that correctly perceives its own Sharpe ratio, and chooses its risk optimally given its risk aversion, but fails to understand that a higher Sharpe ratio is available by investing efficiently. This utility loss is equivalent to a decrease in the riskless interest rate of $w_h\sigma_h(S_B^2 - S_h^2)/2S_h$.

According to the first measure of portfolio inefficiency, the median Swedish household gives up slightly more than a third of the maximum available Sharpe ratio if the international CAPM holds, and slightly less than a third if the international Fama–French model holds. The difference is caused by the fact that Swedish household portfolios are tilted towards small stocks and value stocks, which earn higher returns in the Fama–French model than in the CAPM. The Sharpe ratio loss is reduced by more than half if we take as our benchmark the world index in Swedish kronas rather than the currency-hedged world index. The median Swedish household portfolio has a higher Sharpe ratio than the Swedish equity index, reflecting the fact that many Swedish households hold global equity mutual funds.

Reductions in Sharpe ratios have little effect on portfolio returns if households invest conservatively. Return loss, the second measure of portfolio inefficiency, places greater weight on low Sharpe ratios that are accompanied by aggressive investment strategies. If converted to dollars by multiplying by portfolio value, it also places greater weight on large portfolios. The median Swedish household loses not much more than 1% or \$100 per year relative to the currency-hedged world index under the CAPM. Relative to the unhedged world index, the median household loses only one-quarter as much. Clearly portfolio underdiversification has only modest effects on the welfare of the median Swedish household.

Once again, however, there is wide variation in these numbers across households. At the right tail of the distribution of return losses, these losses are substantial. The 95th percentile of the return loss, relative to the hedged world index, is almost 5 times greater than the median in percentage units, and over 15 times greater in dollar units. In dollar units, the 95th percentile of the loss is over \$2,200 per year relative to the hedged world index, and almost \$850 per year relative to the unhedged index.

These numbers suggest that underdiversification is a problem for a minority of households. A natural next question is which households lose the most by inefficient investing. If the CAPM holds, the overall return loss can be written as the product of three household-specific and one market-wide component

$$w_h(S_B\sigma_h - \mu_h) = w_h\beta_h \left(\frac{S_B}{S_h} - 1 \right) Er_m^e. \quad (4)$$

Here, w_h is the share of the household's portfolio invested in risky assets, β_h is the beta of those risky assets with the benchmark portfolio, $(S_B/S_h - 1)$ is a transformed measure of the relative Sharpe ratio, and Er_m^e is the expected excess return on the world market portfolio.

In Calvet et al. (2006) we take logs of both sides of this equation and then regress the log return loss and its three household-specific components onto demographic characteristics of households. We find offsetting effects on return losses. On the one hand, financially sophisticated households with high disposable income, wealth, education, private pension savings, and financial liabilities tend to invest more aggressively. They invest a higher fraction of their wealth in risky assets, and those assets have higher betas. On the other hand, these households also tend to invest more efficiently, consistent with the findings of Goetzmann and Kumar (2004) for U.S. brokerage account data. In the Swedish data we find that the first effect dominates, so financially sophisticated households actually have higher overall return losses.²¹

These results have two important limitations. First, we assume that mutual fund returns to investors obey the CAPM. If mutual funds hold stocks that obey the CAPM, and if they charge fees to investors, mutual funds will deliver returns with negative alphas reflecting the fee drag. This effect is likely to be significant, as Hortacsu and Syverson (2004) report average fees for equity funds ranging from almost 100 basis points for S&P 500 index funds to over 225 basis points for global and international funds, with wide dispersion across individual funds. A priority for future research is thus to measure the fees charged by each mutual fund available to Swedish investors and the effect of these fees on household portfolio performance.

Second, we treat the financial portfolio in isolation, abstracting from the possibility that financial assets are used to hedge households' labor income

²¹ We also find an age effect on investment performance. Consistent with the findings of Korniotis and Kumar (2006) in U.S. brokerage account data, older households invest more cautiously but less efficiently. These two effects work against each other and almost cancel one another, so that overall return losses are almost invariant to age.

risk. Massa and Simonov (2006) explore this issue and find that while investors in general hold stocks that are positively correlated with their labor income, possibly because these stocks are familiar to them, wealthy investors have a greater tendency to pick negatively correlated stocks that can hedge their labor income risk. Massa and Simonov's results are consistent with the theme of this paper that sophisticated households come closer to the investment strategies recommended by standard financial theory.

An important remaining question is to what extent the results for Sweden describe household behavior in other countries. There are several reasons to think that Swedish households may diversify more effectively than households elsewhere. First, Sweden is a country with a well-educated population and an unusually high stock market participation rate. Second, it is a small country, so Swedish investors are used to the idea that they must diversify internationally. Third, Swedish households were exposed to a national financial education campaign in the late 1990s as part of a reform of the pension system.

An additional priority for future research is to try to understand the implications of underdiversification for the wealth distribution. Household underdiversification has the potential to explain the puzzling dispersion of wealth at retirement reported in U.S. data by Venti and Wise (2001). Venti and Wise argue that differences in lifetime earnings or asset allocation do not explain dispersion, and conclude that it must be caused by differences in savings propensities. However, poorly diversified stock investments could also explain a great deal of dispersion.

B. Diversification and Participation

The demographic predictors of portfolio inefficiency in Sweden are strikingly similar to the demographic predictors of nonparticipation and cautious investing in both Sweden and the United States (see Table I). This suggests that some households may fail to invest in stocks, or may invest only cautiously in stocks, in part because they are aware that they lack the skills to invest efficiently. They may correctly calculate lower welfare benefits of participation given their investment skills, or they may simply feel uncomfortable participating in an activity for which they are poorly prepared (this can be interpreted as a higher psychological fixed cost of participation). To demonstrate the relevance of the first channel, Calvet et al. calculate the extra return that stock market participation can provide a household with the typical demographic characteristics of nonparticipants, assuming that this household invests with the efficiency predicted by a demographic regression. The implied increase in portfolio return is slightly lower than the increase for a household that invests with the average efficiency of Swedish households, and only about half the increase for a household that invests fully efficiently. Put another way, the fixed costs that are needed to deter participation are smaller when households correctly anticipate their own limitations as investors.

There is other, more direct, evidence for a relation between skills, knowledge, and investment behavior. Lusardi and Mitchell (2006) find that people who incorrectly answer simple questions about investing are less likely to plan for

retirement, while Graham, Harvey, and Huang (2005) find that investors who claim to be comfortable about their “ability to understand investment products, alternatives, and opportunities” trade more frequently and are more internationally diversified. Of course, there could be reverse causality from investment activity to understanding of investments, but this cannot explain the finding of Benjamin, Brown, and Shapiro (2006) that people with low cognitive ability, as measured by standardized tests administered in youth, are less likely to participate in financial markets or accumulate assets during their subsequent adult life. In the next section of this address, I use demographic evidence to argue that skills and knowledge are also important in another financial context, the financing of housing by mortgage debt.

IV. Household Mortgage Decisions

The data on asset allocation reveal the importance of housing and the associated mortgage debt for typical households. Yet there has been surprisingly little work on mortgage decisions from the perspective of the household. Instead, most research on mortgages has been conducted by real estate or fixed-income securities specialists who are interested in pricing mortgage-backed derivatives.²²

Mortgage contracts take a bewildering variety of forms, but the two main types are FRMs and ARMs. In the United States, FRMs predominate, accounting for 72% of newly issued mortgages on average over the period from 1985 to 2005 according to the Federal Housing Finance Board (FHFB). Since FRMs have a longer average life before repayment, they account for an even larger fraction of the stock of mortgages at any point in time. Most FRMs have 30-year maturities at origination, must be refinanced when a borrower moves (that is, they cannot be assumed by a new borrower), and can be refinanced at the borrower’s discretion without penalty at any time.

There is interesting variation in the use of ARMs over time. Figure 4, which is an update of a similar figure in Campbell and Cocco (2003), plots the FRM share of new mortgages originated by major lenders over the period from 1985 to 2005. The figure also plots the typical ARM and 30-year FRM rates, and the spread between them. The figure shows that homeowners are more likely to use ARMs when the FRM rate has recently increased, and are more likely to use FRMs when the FRM rate has recently decreased. There is also some evidence that homeowners respond to the spread between the ARM and FRM rates, but this does not seem to explain all the movements in the FRM share. Over 1990 to 1993 and 2001 to 2002, for example, the spread widened yet homeowners did not shift toward ARMs. The correlation between the FRM share and the spread is -0.42 , while the correlation between the FRM share and the lagged 1-year change in the FRM rate is -0.51 . If one regresses the FRM share on the spread

²² The following classic articles on prepayment of FRMs all concentrate on the implications of prepayment for the valuation of mortgage-backed securities, and all include the word “valuation” in their titles: Dunn and McConnell (1981), Schwartz and Torous (1989), McConnell and Singh (1994), Stanton (1995), LeRoy (1996), and Longstaff (2004).




Figure 4. Fixed rate mortgage share and mortgage rates. The share of newly originated mortgages with fixed rates, reported in the monthly survey of the Federal Housing Finance Board, together with 30-year fixed and 1-year adjustable mortgage rates reported by the Federal Home Loan Mortgage Corporation (FHLMC) and the spread between them, from 1985 to 2005.

and the lagged change in the FRM rate, both variables are highly significant. Further, the lagged change in the FRM rate remains significant even if one includes one lag of the spread to capture inertia in mortgage decisions.

Choosing an optimal mortgage contract is a complex problem, as illustrated by the stylized model of Appendix A and the numerical model solved by Campbell and Cocco (2003). Households must take into consideration real interest rate risk, inflation risk, borrowing constraints today and the possibility of borrowing constraints in the future, their risk aversion, their probability of moving, and their ability to refinance a FRM optimally. However it is hard to rationalize the time-series variation shown in Figure 4 using a standard model of household optimization. A wide spread between short-term and long-term interest rates should make ARMs attractive to households that are currently borrowing-constrained or likely to move in the near future, but the recent movement of the long-term interest rate should not be a relevant state variable except insofar as it predicts future movements in long-term rates. The data do not suggest that changes in long-term interest rates are highly autocorrelated, and it would be surprising if they were as this would imply large predictable variation in bond returns. In summary, some households appear to choose between FRMs and ARMs as if they irrationally believe that long-term interest rates are mean-reverting.²³

²³ Some personal finance books encourage the belief that long rates are predictable. Steinmetz (2002), for example, advises "If you think rates are going up, get a FRM" (p.48), and Irwin (1996) implicitly assumes mean-reversion when he writes "When interest rates are low, get a FRM and lock in the low rate" (p. 143).

A. Refinancing

The option to refinance a FRM means that households do not have to pay a fixed rate that greatly exceeds the current level of mortgage rates. When interest rates fall, households have an incentive to refinance their mortgages, either reducing their monthly payments for a fixed level of mortgage debt, or increasing their debt while maintaining the same monthly payments. The latter practice is known as home equity extraction, and has attracted the attention of the Federal Reserve Board for its possible impact on consumer spending (Greenspan and Kennedy (2005)).

Refinancing incurs a substantial one-time cost and thus is not optimal unless the spread between a household's existing mortgage rate and the currently available rate is large enough to cover this cost.²⁴ Since interest rates are volatile, the option to delay refinancing is valuable and the spread must also cover the loss in option value caused by refinancing. Agarwal et al. (2006) estimate the spread that justifies refinancing at 1.1–1.4% for mortgages between \$100,000 and \$200,000 in size.

Declining interest rates in recent years have created large incentives to refinance, and it is generally thought that households have become more responsive to such incentives (Bennett, Peach, and Peristiani (2001)). Despite this, a large minority of households pay interest rates on old FRMs that greatly exceed the currently available rate. Figure 5 summarizes the distribution of rates paid on 30-year FRMs in 1997, 1999, 2001, and 2003. The underlying data are taken from the American Housing Survey (AHS), a data source described in Appendix B and analyzed by Schwartz (2006a). For each spread over the current mortgage rate, the figure shows the fraction of households that pay more than this rate. The data show a striking difference between the 2003 survey and the surveys conducted in 1997, 1999, and 2001. In the three earlier years, rates had been stable or gently declining in the 2 years prior to the survey. Around a quarter of households were paying mortgage rates more than 1% above the currently prevailing rate, 15–20% were paying rates more than 1.5% above, 12–14% were paying more than 2% above, and 6–8% were paying rates more than 3% above the current rate. High mortgage rates tended to be paid on slightly smaller mortgages, so the shares of mortgage value that paid high rates were somewhat lower; for example, 9–12% of mortgage dollar value was paying more than 2% above the currently available rate.

In 2003, fixed mortgage rates had declined precipitously in the 2 years prior to the survey, as illustrated in Figure 4. Although refinancing increased dramatically in 2002 and 2003, by the time of the 2003 survey this had not caught up with the lower available mortgage rate, leading the distribution of mortgage spreads to shift to the right. In 2003 more than half the households surveyed were paying a spread of more than 1%, more than a third were paying more than 1.5%, a quarter were paying more than 2%, and an eighth were paying more

²⁴ The refinancing cost includes the lender's application and attorney review fees, appraisal and home inspection fees, title search and title insurance fees, hazard insurance, the loan origination fee, and mortgage insurance. Agarwal, Driscoll, and Laibson (2006) estimate the refinancing cost during the 1990s at \$2,000 plus 1% of mortgage value.




Figure 5. Distribution of mortgage spreads. The distribution of self-reported mortgage rates, corrected for implausibly low rates, in the 1997, 1999, 2001, and 2003 American Housing Surveys, relative to the currently prevailing FHLMC rate in the year of each survey.

than 3%. Again these numbers are slightly lower when calculated as shares of dollar value, but almost 20% of mortgage dollars were paying spreads of more than 2% in 2003.

The sluggishness of household refinancing has been a major theme of the literature on valuation of mortgage-backed securities. Early work by Dunn and McConnell (1981) assumed costless refinancing, and failed to fit the behavior of mortgage-backed securities prices. The subsequent literature either works with reduced-form econometric models of prepayment rates (Schwartz and Torous (1989)) or specifies a cross-sectional distribution of refinancing costs across households to account for the willingness of some households to pay high mortgage spreads. Further realism can be achieved by adding an exogenous delay to refinancing (McConnell and Singh (1994), Stanton (1995)). The parameters of these models themselves evolve over time, generating random variations in prepayment speed that create an unhedgeable risk in the cash flows of mortgage-backed securities. Gabaix, Krishnamurthy, and Vigneron (2006) argue that this accounts for the option-adjusted spread commonly used to value these securities.

This work generally does not explore the underlying economic causes of slow household refinancing. Some households may be prevented from refinancing by declines in their house value that have eroded their home equity, or by declines in their income that have decreased their creditworthiness. The lock-in effect of declining house prices is emphasized by Caplin, Freeman, and Tracy (1997), and Archer, Ling, and McGill (1996) find that households with insufficient collateral value or income are less likely to respond to refinancing incentives. But it is

also plausible that sluggish refinancing is an investment mistake, on a par with nonparticipation or underdiversification.²⁵

In support of this view, Table II reports characteristics of households that did not move but refinanced their FRMs between 2001 and 2003, a period when rates fell dramatically and created large incentives to refinance. The first panel of the table reports a probit regression of a refinancing dummy on household characteristics, using data from the AHS. The regression includes dummy variables for the mortgage origination year to control for variation in the incentive to refinance created by time-series variation in prevailing mortgage rates. As in Table I, the refinancing probability for a reference household is reported along with the change in this probability caused by a change in each dummy variable, or a one-standard-deviation change in each continuous variable.

The first two explanatory variables in Table II capture constraints that may prevent some borrowers from refinancing. The variable *loan problem* is a dummy variable that equals one if the current principal value of the mortgage exceeds 90% of the self-reported value of the house. The variable *income problem* is a dummy variable that equals one if the mortgage payments that would be required by a refinanced 30-year FRM exceed 28% of current self-reported income. These dummies are used by Archer et al. (1996), who emphasize their important effects on mortgage refinancing during the 1980s. In Table II they enter with the theoretically predicted negative signs, but only the loan problem is economically or statistically significant. The weakness of these effects compared with those reported by Archer et al. may be due to relaxed standards for mortgage lending in recent years, or to the rise in house prices that has reduced the fraction of households with insufficient home equity.

The remaining variables in Table II capture the demographic characteristics of each household along with its income, the value of its house, and the size of its mortgage. Quadratic terms are included in the last three variables to capture possible nonlinearities, but are not of major importance. The results show that younger, smaller, better educated, better off, white households with more expensive houses were more likely to refinance their mortgages between 2001 and 2003. These patterns suggest that prompt refinancing requires financial sophistication.

A household may rationally fail to refinance its mortgage if it expects to move. Households that expect to move and actually do move are dropped from the refinancing regression, while households that expect to move and do not move are recorded as nonrefinancers. Thus, one explanation for the results in the first panel of Table II could be that more sophisticated households are less mobile. The second panel of Table II estimates the determinants of mobility within a larger sample of households surveyed in 2001. Those that moved between 2001 and 2003 (as identified by a different household answering survey questions

²⁵ Consistent with this interpretation, sluggish refinancers are known in the mortgage industry as “woodheads.” Agarwal et al. (2006) point out that some households make the opposite mistake, refinancing when the mortgage spread is just enough to cover the fixed cost of refinancing but ignoring the option to delay and thus refinancing too quickly.

Table II
Mortgage Refinancing and Mortgage Rates

The table reports demographic determinants of mortgage behavior for households with 30-year fixed mortgages in the 2001 and 2003 American Housing Surveys. The first three columns report probit regressions of refinancing between 2001 and 2003, confirmed changes of address between 2001 and 2003, and 2001 self-reporting of implausibly low mortgage rates more than 2% below standard FHLMC rates prevailing at origination, onto demographic variables. The last two columns report OLS regressions of self-reported 30-year fixed mortgage rates in 2001 and 2003, adjusted to correct for implausibly low self-reported rates, onto demographic variables. Standard errors are reported underneath the coefficients in parentheses. Coefficients significant at the 10% level are denoted by *, at the 5% level by **, and at the 1% level by ***. All regressions include year and region fixed effects, and control for gender, marital status, and urban house location. In the reference household, the household head is a married white male with no high school diploma. The column headed "Probability Estimates" reports the probability for the reference household, and the change in this probability caused by a unit change in a binary variable and a one-standard-deviation change in a continuous variable.

Dependent Variable	Whether a Household Refinanced a 30-year Fixed Mortgage in 2001–2003		Whether a Household Moved between 2001 and 2003		Whether a Household Reported a Rate 2 or More % Points Below FHLMC in 2001		Mortgage Rate for 30-year Fixed Mortgages in 2001		Mortgage Rate for 30-year Fixed Mortgages in 2003	
	Coefficients	Probability Estimates (%)	Coefficients	Probability Estimates (%)	Coefficients	Probability Estimates (%)	Coefficients	Probability Estimates (%)	Coefficients	Probability Estimates (%)
<i>Reference household</i>										
Loan problem	−0.179*** (0.071)	27.9 −4.6	−0.050 (0.088)	4.5 −0.5	0.233*** (0.082)	1.3 2.2	0.097** (0.048)	0.060** (0.031)		
Income problem	−0.076 (0.099)	−0.2	0.132 (0.130)	1.4	−0.064 (0.107)	−0.6	0.139** (0.063)	0.364*** (0.065)		
First house	0.104** (0.043)	3.6	0.001 (0.054)	0.0	−0.037 (0.053)	0.2	−0.039 (0.035)	−0.057** (0.029)		
Age	−0.026** (0.011)	−4.3	−0.078*** (0.012)	−1.6	0.003 (0.013)	−0.1	0.007 (0.008)	0.004 (0.007)		
Age squared	1.5×10^{-4} (1.1×10^{-4})		6.8×10^{-4} *** (1.3×10^{-4})		-2.3×10^{-5} (1.2×10^{-4})		-4.6×10^{-5} (9.2×10^{-5})	-4.1×10^{-6} (8.3×10^{-5})		
White	0.250*** (0.056)	7.7	0.251*** (0.081)	1.9	0.071 (0.071)	−0.1	0.002 (0.041)	−0.202*** (0.038)		

High school diploma	0.163** (0.077)	5.4	0.293*** (0.115)	3.5	-0.272*** (0.076)	-0.6	-0.156** (0.062)	-0.170*** (0.052)
College diploma	0.268*** (0.083)	9.2	0.364*** (0.122)	4.6	-0.203** (0.091)	-0.5	-0.095 (0.067)	-0.195*** (0.055)
Graduate school	0.261*** (0.091)	8.9	0.316** (0.132)	3.9	-0.296*** (0.104)	-0.7	-0.070 (0.074)	-0.154** (0.063)
Number of children	-0.006 (0.017)	-0.3	-0.060** (0.024)	-0.6	2.5×10^{-4} (0.023)	0.1	0.006 (0.013)	0.004 (0.011)
Number of adults	-0.094*** (0.028)	-2.6	-0.124*** (0.041)	-0.8	-0.005 (0.031)	-0.1	-0.015 (0.022)	0.053*** (0.020)
Ln(income)	0.274** (0.140)	1.3	-0.528** (0.224)	0.9	0.049 (0.107)	-0.3	-0.037 (0.110)	0.018 (0.095)
Ln(income) squared	-0.011* (0.007)		0.028*** (0.011)		-0.007 (0.006)		0.000 (0.005)	0.000 (0.005)
Ln(house value)	1.259* (0.747)	6.6	1.589* (0.960)	-0.2	1.422* (0.822)	0.1	1.101** (0.464)	-0.233 (0.440)
Ln(house value) squared	-0.042 (0.031)		-0.068* (0.040)		-0.056* (0.034)		-0.045** (0.019)	0.010 (0.018)
Ln(mortgage)	-0.024 (0.136)	-0.1	-0.024 (0.045)	0.0	-0.057* (0.031)	0.0	0.335*** (0.079)	0.447*** (0.054)
Ln(mortgage) squared	0.003 (0.007)		0.002 (0.002)		-0.003* (0.002)		-0.027*** (0.004)	-0.038*** (0.003)
Sample size	5,190		5,735		7,610		7,610	9,749

at the same address in 2003, about 9% of the sample) are recorded as movers, while those that responded to both the 2001 and 2003 surveys at the same address are recorded as nonmovers. The regressions show that almost all the determinants of refinancing probability, notably age, race, education, family size, and home value, affect mobility in the same direction. Thus, mobility does not seem to be a plausible explanation for the cross-sectional variation in the propensity to refinance.²⁶

As further support for the interpretation that sluggish refinancing is an investment mistake, the third panel of Table II reports demographic determinants of implausible self-reported mortgage rates. As discussed in Appendix B, some households in the AHS report that their current mortgage rates are implausibly low, over two percentage points below the average mortgage rate prevailing during the origination period. If these households do not understand that they are paying high rates, it is understandable that they might fail to take advantage of a refinancing opportunity. About 7% of households report implausibly low mortgage rates, and most of these households took out their mortgages more than 10 years ago. The third panel of Table II shows that this reporting error is much more common among less-educated households.²⁷

The last two columns of Table II push this investigation further by examining the determinants of the mortgage rate that households pay at any point in time. The regressions are run separately for the 2001 and 2003 surveys. In each case the dependent variable is the self-reported mortgage rate, adjusted to correct for implausible rates as described in Appendix B. The regression includes dummy variables for the year of house purchase.

The variables that predict prompt refinancing behavior also generally predict low mortgage rates. The loan problem and income problem dummies from Archer et al. (1996) have a positive and strongly significant effect on mortgage rates paid in both years. Education has a negative effect in 2001, which strengthens in 2003, reflecting the importance of education in driving prompt refinancing in 2002 and 2003. The effect of race also becomes significant in 2003. Income and home value have relatively weak effects on mortgage rates paid, but the value of the mortgage has a strong effect; a one-standard-deviation increase in the mortgage size above the sample mean lowers the mortgage rate by 44 basis points in 2001, and 49 basis points in 2003. Presumably the mortgage variable captures both mortgage size as a proxy for household wealth, and a reverse causality effect that households with the credit quality and sophistication to obtain cheaper mortgages tend to take out larger mortgages.

²⁶ Some households may have moved as an endogenous response to declining mortgage rates, refinancing in order to buy a larger house. To the extent that this behavior is more common among sophisticated households, it will lead the regression in the first panel of Table II, which includes only nonmoving households, to understate the effect of sophistication on refinancing.

²⁷ The reporting error is relatively rare for the reference household in Table II because this household has a mortgage of average age rather than an old mortgage. When the reporting error dummy is included directly in the regression predicting refinancing, it enters negatively but is not statistically significant. Bucks and Pence (2006) present complementary evidence on ARM borrowers. They find that many households, especially lower-income households, do not understand the potential increase in their ARM rate that can be caused by rising interest rates.

One difficulty with these regressions is that they confound the effects of refinancing decisions with the mortgage rate that a household can obtain at a point in time. It may well be the case that better-educated households have better credit quality and can obtain mortgages on more favorable terms.²⁸ But if one eliminates cross-sectional variation in rates available at a point in time by replacing self-reported mortgage rates with average Federal Home Loan Mortgage Corporation (FHLMC) rates prevailing at the mortgage origination date, the results are similar to those reported in Table II. In particular, the effects of race and education remain significant.

An interesting question is whether less sophisticated households can anticipate their inability to refinance FRMs optimally. If so, a rational response might be to use an ARM instead. Schwartz (2006b) looks at the determinants of ARM usage in the AHS and finds no evidence that less sophisticated households use ARMs. During the earlier survey years, ARMs were favored by younger households financing their first houses and with relatively small mortgages. This is consistent with the model of Campbell and Cocco (2003) and Appendix A, in which mobile and currently borrowing-constrained households should be attracted by the low initial cost of an ARM unless a large mortgage makes its interest rate risk unacceptable. During the later survey years, and particularly in 2003, ARMs were favored by better-educated households. It is striking that 2003 was a year of an unusually wide spread between FRM and ARM rates, even though FRM rates were low; thus, sophisticated households should have been attracted to ARMs while unsophisticated households may have anticipated rapid mean-reversion in long rates and may have avoided ARMs for that reason. This interpretation is speculative, but can be tested by a more systematic investigation of the response of households with different education levels to movements in mortgage rates.

Overall, it does not seem that households that lack the knowledge to refinance FRMs substitute away from these mortgage contracts in a way that would be analogous to nonparticipation as a response to lack of knowledge about the stock market. Presumably the lack of a simple alternative is a barrier to this response in the mortgage market.

V. Equilibrium in Retail Financial Markets

Given the complexity of the household financial optimization problem, it may not be surprising that some households make mistakes. For example, to refinance a FRM optimally one must solve an irreversible investment problem and this is a difficult task. What is perhaps surprising is that so many of the contracts available to households reward sophisticated decision making and

²⁸ Moore (2003) surveys mortgage borrowers in Washington State, including “victims” who borrowed from a predatory lender. She finds that the victims were more likely to lack basic financial knowledge, suggesting that they failed to understand the cost of their mortgage loans. However, she also finds that the victims were more likely to be in financial distress, implying that they might have had difficulty obtaining more favorable loans.

continuous monitoring of financial markets. One might expect that simpler contracts would be offered that would leave less room for expensive mistakes. Economists often recommend such instruments. Specifically, economists have often recommended mortgages that adjust interest and principal payments for inflation, thereby combining the best features of nominal FRMs and ARMs (Statman (1982), Alm and Follain (1984), McCulloch (1986)). More recently, Flesaker and Ronn (1993) and Nalebuff and Ayres (2003) have proposed an automatically refinancing nominal FRM that would eliminate sluggish refinancing and also save consumers' considerable costs of current refinancing procedures.²⁹

Despite these recommendations, financial innovation in retail markets often appears to proceed slowly. There is considerable inertia in the general form of mortgage contracts, despite robust competition by mortgage lenders and changes in credit standards over time. A related puzzle is that different types of mortgages are standard in different countries. In the United Kingdom, for example, ARMs are standard and FRMs of the U.S. variety are almost unknown (Miles (2003), Green and Wachter (2005)).

In this section I argue that the existence of unsophisticated households helps to explain these phenomena. A first and perhaps obvious point is that unsophisticated households tend to use whatever financial contracts are standard in a particular country, possibly because they follow the lead of relatives and neighbors. It is expensive for would-be financial innovators to reach such households, particularly if they need to explain a complex new financial product. Second, the absence of effective patent protection in the financial industry makes it hard for financial innovators to recoup the costs of advertising and financial education needed to establish a new product.

Even if these two points are valid, why can financial innovators not establish a foothold by attracting sophisticated households who understand the benefits of a new financial product? Once sophisticated households adopt the product, other households might follow. The explanation here may be that existing products often involve a cross-subsidy from naive to sophisticated households. A refinanceable FRM, for example, offers a low rate in part because many households do not optimally refinance. Sophisticated households gain by pooling with naive households, and will not be attracted to a new mortgage if it is only taken up by other sophisticated households.

To understand this possibility, consider a market in which a fraction α of the population is naive (denoted by N) and the remainder is sophisticated (S). An existing mortgage contract can be provided at overall cost C_0 , but is structured in such a way that sophisticated households pay a lower cost than naive households. For example, there may be a refinancing option that only sophisticated households exercise. In competitive equilibrium with full participation by both groups, their costs must be related by

$$\alpha C_{0N} + (1 - \alpha)C_{0S} = C_0. \quad (5)$$

²⁹ Automatic refinancing has a close parallel in the automatically refunding "ratchet" bond issued by the Tennessee Valley Authority in 1998 (Kalotay and Abreo (1999)).

Write x for the cross-subsidy from the naive to the sophisticated: $x = C_{0N} - C_{0S}$. Then

$$C_{0S} = C_0 - \alpha x. \quad (6)$$

The mortgage costs of sophisticated households fall with the fraction of naive households and the size of the cross-subsidy.

Now suppose that a new mortgage contract is invented, which provides the same benefits at a lower cost $C_1 = C_0 - g$. The new contract might be, for example, an automatically refinancing mortgage that reduces the costs of refinancing, or an inflation-indexed mortgage that reduces the need to refinance. Assume that the new contract is made available initially to a negligible fraction of the population, so that its introduction does not perturb the pricing of the existing contract. In the simplest setting only sophisticated households can understand the new product, but all households can be reached costlessly. Sophisticated households that opt for the new product gain directly from its lower cost, but lose indirectly by giving up the cross-subsidy. They therefore switch to the new product only if it offers a social gain—a cost saving—larger than the per capita cross-subsidy from naive households using the existing product

$$g > \alpha x. \quad (7)$$

New products with smaller social gains will not be adopted.

More realistically, suppose that a new product can be advertised at per capita cost k , but only sophisticated households understand the advertisement. Such households switch to the new product if it offers them a cost $C_{0S} - s = C_1 + g - \alpha x - s$, where s is a switching cost. If a monopolist offering the new product charges this cost, it will make a profit of $g - \alpha x - s$ on each customer. It will be worth advertising the new product only if advertising attracts enough customers, that is, if

$$(1 - \alpha)(g - \alpha x - s) > k. \quad (8)$$

Even a product that appeals to sophisticated consumers may not gain a foothold if such consumers are difficult to reach. In this case financial education—a reduction in the fraction of naive households—has a positive effect on financial innovation both by increasing the effectiveness of advertising the new product, and by reducing the cross-subsidy that each sophisticated household receives in the existing product.

Is it worthwhile for financial innovators to offer financial education privately, converting a small number of naive households into sophisticated households? It may not be, for the same reason that sophisticated households may not be attracted to new products at a price that covers the cost of recruiting them. Newly educated households understand the cross-subsidy they receive from naive households, and may refuse to switch to a new product offered along with financial education. Instead, innovators may have an incentive to mislead naive households by offering confusing products with high fees. The possibility

of such perverse financial innovation depends on the details of naive investors' behavior, and deserves further theoretical and empirical research.

These effects are an example of "shrouded equilibrium" (Ellison (2005), Gabaix and Laibson (2006)), in which some consumers are unaware of hidden costs associated with certain products. It may not pay competitors to reveal these hidden costs if sophisticated consumers have the ability to avoid them while still purchasing the products, which are cheaper because of the revenue provided by naive consumers. Gabaix and Laibson concentrate on examples in which the hidden costs are associated with expensive add-ons (such as cartridges for ink-jet printers or telephone calls from hotel rooms), whereas in the example of a refinanceable mortgage the hidden costs arise from consumer failure to understand an advantageous option.³⁰

Miles (2003) emphasizes the effect of hidden costs on the equilibrium of a retail financial market. David Miles was commissioned by the British Chancellor of the Exchequer to review the state of the U.K. mortgage market, and in particular to understand why FRMs are so much less popular in the United Kingdom than in the United States and many European countries. His 2003 interim report argues that in the U.K. mortgage market, ARMs are often sold with discounted initial rates (sometimes called teaser rates) that automatically adjust to a much higher "standard variable rate" after 2 years. Borrowers have the right to refinance their mortgages without penalty after 2 years, yet the teaser rates are extremely attractive relative to prevailing money market rates.³¹ Miles concludes that mortgage lenders can only offer these rates because many households (close to a third of borrowers in 2003) fail to refinance their mortgages and end up paying standard variable rate. He argues that the resulting cross-subsidy from naive to sophisticated households inhibits the development of a FRM market in the United Kingdom. By contrast with the United States, where FRMs are standard and are widely used by naive households, in the United Kingdom they are considered only by sophisticated borrowers who are reluctant to give up the attractive rates available in the adjustable-rate market.

In the United States, a similar mechanism may keep down the cost of standard refinanceable FRMs. The excess mortgage interest shown in Figure 5 is large enough to have a noticeable effect on the mortgage rates offered by competitive lenders. In 2001, for example, the total payments made by households in the AHS were higher by 0.66% of mortgage value than they would have

³⁰ Hidden costs may also be important in the mutual fund industry. Barber, Odean, and Zheng (2003) show that operating expenses have a lower effect on mutual fund flows than more visible front-end loads. In this case sophisticated investors do not have the ability to avoid such fees so they do not receive a cross-subsidy from naive investors; however, high-fee mutual funds may still survive if naive investors are costly to reach through financial advertising (Sirri and Tufano (1998), Hortacsu and Syverson (2004)).

³¹ In October 2003, for example, one-month sterling LIBOR was 3.63% and the discounted adjustable rate was only 7 basis points higher at 3.70%. Two-year LIBOR was 4.51%, and the discounted initial rate on a mortgage with a two-year fixed interest period was 2 basis points lower at 4.49%. The adjustable rate for mortgages with no initial discount was 4.51% and the standard variable rate was 5.42%, 88 basis points and 179 basis points, respectively, above one-month LIBOR.

been if interest were capped at 1% above the current mortgage rate. The excess interest was somewhat lower in 1997 and 1999 at 53 and 43 basis points, respectively, but much higher in 2003 at 107 basis points. The numbers are similar if one eliminates cross-sectional variation in credit quality by using the FHLMC rate prevailing at each mortgage origination date rather than the self-reported mortgage rate. Thus, excess interest reflects the failure of naive borrowers to refinance their mortgages, rather than cross-sectional variation in credit quality, and can be thought of as a hidden cost of the sort discussed by Gabaix and Laibson. Of course, these numbers are only suggestive as they reflect outcomes along a single, declining path for mortgage rates rather than a probability-weighted average across all possible rate scenarios, and they are measured relative to current rates rather than minimum rates.

Hidden costs may also be important in another aspect of U.S. mortgage markets. Borrowers have the option of paying mortgage origination costs, including perhaps a fee to a mortgage broker, in the form of cash up front ("points") or by paying a higher interest rate on the mortgage. The conventional analysis of this arrangement (e.g., Stanton and Wallace (1998)) emphasizes that it is a way for mortgage lenders to separate households by moving probability. Households that expect to move soon are unwilling to pay points and can therefore be distinguished from relatively immobile households; they receive a favorable rate because the refinancing option is less valuable for these households, and thus less costly for mortgage lenders to provide to them. Woodward (2003), however, argues that points also increase the opportunity for mortgage brokers to confuse borrowers. In a sample of 2,700 mortgage loans with average mortgage broker fees of almost \$2,500, she finds that households tend to pay higher broker fees on mortgages with points, and that college education is associated with a remarkable \$1,500 reduction in average broker fees. In a competitive market for mortgage brokerage services, broker fees may be lower for sophisticated households because of the high fees paid by naive households.

An important question is whether public policy can improve outcomes in financial markets with naive households. Financial education is certainly helpful, and an audience of financial educators is likely to agree on its importance. However, one should not overestimate the power of education; in particular, the education variables in demographic regressions are endogenous, proxying for cognitive ability and other omitted factors, and so they may overstate the effects of exogenous increases in education on investment behavior.

Regulation to prohibit "predatory lending" and excessive cross-subsidization may also be helpful, but as Gabaix and Laibson emphasize, it is difficult to design regulations that protect naive households without also inhibiting helpful innovation. For example, regulations that prevent negative amortization in mortgage contracts, which are intended to protect households from incurring unmanageable debts, also prevent inflation-adjustment of principal if amortization is defined in nominal terms.

Public policy can subsidize suitable financial instruments through tax incentives or credit guarantees. Swensen (2005) argues that low-cost passive mutual funds should be subsidized in this way. In mortgage markets, the U.S. government plays a major role through the Government National Mortgage

Association (GNMA) and its sponsorship of the Federal National Mortgage Association (FNMA) and FHLMC. The ability of these agencies to issue low-cost debt has likely reduced the cost of the mortgages they hold (Green and Wachter (2005)). The agencies have traditionally favored nominal FRMs, and it is plausible that they could encourage the use of mortgages with inflation adjustment or automatic refinancing features.

Disclosure requirements can reduce the incidence of investment mistakes, but here too they must be designed appropriately. Miles (2004) points out that in the United Kingdom the annual percentage rate (APR) on a mortgage—the main cost measure considered by prospective borrowers—is calculated under the assumption that interest rates will remain unchanged during the life of the mortgage. This assumption tends to understate the cost of an ARM when the yield curve is upward sloping, for then forward rates exceed spot rates, suggesting that the bond market expects interest rates to rise. Miles recommends that the APR should be calculated under the assumption that spot rates at all future dates will equal the corresponding forward rates prevailing at the disclosure date. Disclosure requirements will be increasingly important if households come to rely more heavily on mortgage calculators and other financial websites to compare financial products.

Finally, recent research in behavioral finance finds that default options—standard choices that households believe to be recommended by authoritative bodies—have a powerful effect on household behavior (Choi et al. 2002, 2004b). The mortgage policies of GNMA, FNMA, and FHLMC may influence household mortgage decisions through this channel as well as by driving down the cost of standard mortgage contracts.

VI. Conclusion

In this paper I outline the field of household finance. I argue that although many households find adequate solutions to the complex investment problems they face, some households make serious investment mistakes. These mistakes come in a variety of forms. I emphasize nonparticipation in risky asset markets, underdiversification of risky portfolios, and failure to exercise options to refinance mortgages.

Investment mistakes have a number of interesting characteristics that make them central to the study of household finance. First, it appears that poorer and less educated households are more likely to make mistakes than wealthier and better educated households. This pattern reinforces the interpretation of nonstandard behavior as reflecting mistakes rather than nonstandard preferences.

Second, some mistakes may result from efforts to avoid others. The same types of households that tend to invest poorly are more likely not to participate in risky asset markets at all. Nonparticipating households may be aware of their limited investment skill and may react by withdrawing from risky markets altogether. Other households, wishing to delegate financial decisions to professionals, may pay high fees to financial planners, mutual funds, or banks.

Third, the presence of households that make investment mistakes may inhibit financial innovation. Many investment products allow some degree of cross-subsidization from naive households to sophisticated households that optimally exploit embedded options. It may be difficult for new investment products to gain acceptance if sophisticated households, who are the natural early adopters, must give up the benefit of a cross-subsidy when they move from an existing product to a new product.

Inevitably I neglect many important issues. For instance, I discuss portfolio choice but not household savings decisions, and mortgage debt but not credit card debt (Agarwal et al. (2005), Bertaut and Haliassos (2006), Zinman (2005)). I do not review the large literature on the fees charged by mutual funds, banks, and other financial intermediaries, nor do I discuss evidence that some households fail to exploit tax incentives for retirement saving (Amromin, Huang, and Sialm (2005), Choi, Laibson, and Madrian (2005)). I emphasize financial markets that are relevant for middle-class households, but I say nothing about payday lending and other forms of credit that are used by poor households (Bolton and Rosenthal (2005)). I likewise ignore issues that are important for extremely wealthy households, such as estate tax management and the role of hedge funds. I treat households as unitary entities and do not consider the possibility that bargaining between family members influences household decisions.

I also sidestep the issue of whether investment mistakes influence the pattern of returns available in asset markets. This issue is important because it may affect the welfare cost of mistakes. For example, nonparticipation in equity markets may increase the equity premium and worsen the welfare loss caused by this mistake. There is an active debate about the magnitude of such effects. On the one hand, asset prices are ultimately determined by supply and demand, and household investment mistakes are surely relevant for household asset demands. On the other hand, asset prices are disproportionately influenced by the demands of wealthy, risk-tolerant investors and professional arbitrageurs, so investment mistakes may be less important in asset pricing than they are in household finance.

Even if asset prices are set efficiently, investment mistakes can have large welfare costs for households. Since investment mistakes are particularly likely when new financial markets are created or when households are asked to take on new financial planning responsibilities, they may greatly reduce the welfare gains that can be realized from the current period of financial innovation and from proposed new financial instruments (Shiller (2003)). If household finance can achieve a good understanding of the sources of investment mistakes, it may be possible for the field to contribute ideas to limit the costs of these mistakes. For example, we can try to define the core elements of financial literacy that make it possible for households to undertake financial planning (Bernheim (1998), Lusardi and Mitchell (2006)). We can also propose more informative disclosures, structure the customized advice that is offered by financial planning websites, suggest appropriate default investment options, or encourage public provision or tax subsidy of simple financial products such as well designed U.S. savings bonds (Tufano and Schneider (2005)). Work of this sort extends the innovative spirit of financial engineering to the retail marketplace.

The possibility that household finance may be able to improve welfare is an inspiring one. Keynes (1932) wrote that he looked forward to a distant future when economists would be “thought of as humble, competent people, on a level with dentists.” Today, dentists spend much of their time delivering advice and easy-to-use products that promote oral hygiene; economists for their part can deliver, or at least design, advice and innovations that promote financial hygiene.

Appendix A: A Normative Model of Mortgage Choice

To clarify the issues that arise in mortgage finance, I consider a simplified example. Campbell and Cocco (2003) present a richer model with a similar spirit and solve it numerically.

I assume that at an initial date 0, a household buys a house and finances it with a mortgage that has face value M . For simplicity, I assume that the house is also worth M , so the loan-to-value ratio is 100%. At date 1, the household pays interest on the mortgage. At date 2, the household sells the house and repays the mortgage with interest. The household also receives income and chooses nonhousing consumption at each date.

Mortgage contracts are specified in nominal terms, so the debt M is nominal. The house, however, is a real asset whose nominal value grows with inflation. I assume that inflation can take one of two values, low or high, and uncertainty about inflation is resolved between dates 0 and 1. Thus, inflation between date 0 and date 1 is either $(1 + \pi_L)$ or $(1 + \pi_H)$, and cumulative inflation between dates 0 and 2 is either $(1 + \pi_L)^2$ or $(1 + \pi_H)^2$. This structure captures the historical tendency for inflation movements to be highly persistent.

The household can choose between two standard mortgage contracts. A nominal FRM requires a nominal payment of $R_F M$ at date 1, and a final repayment of $(1 + R_F)M$ at date 2. The fixed mortgage rate R_F is set at date 0, before inflation is observed. An ARM requires nominal payments of $R_{AH}M$ at date 1 and $(1 + R_{AH})M$ at date 2 if inflation is high, and $R_{AL}M$ at date 1 and $(1 + R_{AL})M$ at date 2 if inflation is low.

These mortgages do not correspond exactly to real-world mortgages, because they repay principal in one lump sum at maturity (like Treasury bonds), rather than amortizing the debt to produce level payments over the life of the mortgage. Level-payment mortgages, however, repay principal slowly at first and more rapidly later, and the assumptions of the model capture this pattern within a simplified two-period structure.

I assume that the riskless real interest rate is a constant r , and I ignore default risk in mortgage lending. Thus, the interest rate on an adjustable mortgage is

$$R_{AH} = (1 + r)(1 + \pi_H) - 1 \approx r + \pi_H$$

if inflation is high, and

$$R_{AL} = (1 + r)(1 + \pi_L) - 1 \approx r + \pi_L$$

if inflation is low. The required real mortgage payment in period 1 is

$$\frac{R_{Ai}}{1 + \pi_i} M = \left(r + \frac{\pi_i}{1 + \pi_i} \right) M$$

for $i = H, L$. The required real payment in period 2 is

$$\frac{(1 + R_{Ai})}{(1 + \pi_i)^2} M = (1 + r) \left(1 - \frac{\pi_i}{1 + \pi_i} \right) M$$

for $i = H, L$. The sum of all real payments, discounted back to date 0 at the real interest rate r , does not depend on inflation and is always equal to M . In this sense an ARM is a riskless liability, just as a floating-rate note is a riskless investment.

Note, however, that the *timing* of required payments does depend on inflation. An increase in inflation raises the nominal interest rate and accelerates the repayment of the mortgage debt. Each percentage point increase in inflation requires an increased mortgage payment of about 1% of the face value of the mortgage. This is compensated by a reduced real payment when the mortgage matures. In other words, an ARM is a shorter-duration liability when inflation is high.

A FRM has a very different sensitivity to inflation. The interest rate on a FRM is R_F , regardless of whether inflation is high or low. The required real mortgage payment in period 1 is

$$\frac{R_F}{1 + \pi_i} M$$

for $i = H, L$, and the required real payment in period 2 is

$$\frac{(1 + R_F)}{(1 + \pi_i)^2} M.$$

The sum of all real payments, discounted back to date 0 at the real interest rate r , is declining in inflation; thus, a FRM benefits the borrower when inflation is high.

What is the optimal mortgage contract in this model? I will assume that the household has real income Y each period. If the household can borrow or lend freely at the riskless real interest rate, then the household does not care about the timing of mortgage payments and is concerned only with the expected level and variability of lifetime resources. The ARM gives the household lifetime resources, discounted to time 0, of

$$\left[\frac{1}{1 + r} + \left(\frac{1}{1 + r} \right)^2 \right] (Y - rM),$$

allowing the household to consume a riskless $(Y - rM)$ each period.

If mortgage rates are set by lenders who are risk neutral with respect to inflation risk, then the fixed mortgage rate will be set to equate the expected

present values of fixed and adjustable mortgage payments.³² In other words, the expected value of household lifetime resources will be the same under fixed and adjustable mortgages. However, lifetime resources are random under a FRM, so a risk-averse household will always prefer an ARM. This conclusion is only strengthened if mortgage lenders are averse to inflation risk and charge a premium rate for bearing it.

How, then, can we understand the observed predominance of long-term nominal FRMs in the U.S. mortgage market? First, it is important to consider the fact that many households are borrowing constrained, particularly in the early years of homeownership. Borrowing-constrained households care not just about lifetime resources, but about the resources available in each period. A borrowing-constrained household with an ARM will have real consumption in period 1 of

$$C_{A1i} = Y - \left(r + \frac{\pi_i}{1 + \pi_i} \right) M$$

and real consumption in period 2 of

$$C_{A2i} = Y - \left(r - \frac{\pi_i}{1 + \pi_i} (1 + r) \right) M.$$

A borrowing-constrained household with a FRM will have real consumption in period 1 of

$$C_{F1i} = Y - \frac{R_F}{1 + \pi_i} M$$

and real consumption in period 2 of

$$C_{A2i} = Y + M \left(1 - \frac{1 + R_F}{(1 + \pi_i)^2} \right).$$

This consumption profile has less variability in period 1, because inflation shocks affect period-1 consumption only in proportion to the product of the nominal interest rate and the face value of the mortgage, and not in proportion to the whole face value of the mortgage. A Taylor expansion around the average inflation rate, π , shows that for small inflation volatility, the standard

³² Note that this does not mean the fixed mortgage interest rate will equal the gross expected inflation rate times the gross real interest rate. The present value formula is convex in inflation, so fixed-rate lenders gain more from low inflation than they lose from high inflation. This convexity effect lowers the equilibrium fixed mortgage rate. The fixed mortgage rate satisfies the equation

$$\frac{1}{1+r} E \left[\frac{1}{1+\pi} \right] R_F + \left(\frac{1}{1+r} \right)^2 E \left[\left(\frac{1}{1+\pi} \right)^2 \right] (1 + R_F) = 1.$$

The solution to this equation is lower than the rate R_F^* that would equate expected first-period mortgage payments with those of an ARM, which in turn is lower than the gross real interest rate times the expected gross inflation rate.

deviation of period-1 consumption is $M\sigma_\pi/(1+\pi)^2$ for the ARM, and only R_F times as large for the FRM. On the other hand, the fixed-rate consumption profile has greater variability in period 2 because inflation shocks affect the whole value of the mortgage and are not compensated by interest rate variations. A similar Taylor expansion shows that for small inflation volatility, the standard deviation of period-2 consumption is $(1+r)M\sigma_\pi/(1+\pi)^2$ for the ARM, and $2(1+R_F)/(1+r)(1+\pi)$ times as large for the FRM.

What about the average levels of real consumption in period 1 and period 2? If inflation is deterministic, then FRMs and ARMs deliver the same time path of payments and thus the same average levels of consumption in the two periods. If inflation is random, the convexity of FRMs lowers the fixed mortgage rate below the level that would equate expected real payments in period 1 with those required by an ARM. FRMs are thus back-loaded relative to ARMs.

A borrowing-constrained household will find the back-loading of a FRM attractive as a way to increase expected period-1 consumption relative to period-2 consumption. If the household is risk averse, it will also find the reduced volatility of period-1 consumption attractive. For both reasons, a borrowing-constrained household with a sufficiently high time discount rate will prefer a FRM.

Refinancing

Thus far I ignore an important feature of FRMs, namely, that they give the household an option to refinance. This option affects both the interest rate and the risk characteristics of a FRM. To model refinancing, suppose that the FRM can be repaid without penalty in period 1, after the period-1 fixed interest payment has been made. The household will find it worthwhile to refinance, and take out a new mortgage in period 1, if inflation turns out to be low so that nominal rates are lower in period 1 than they were in period 0. Thus, a refinaneable FRM requires fixed payments if inflation is high, but if the mortgage is refinanced because inflation is low, payments become

$$\frac{R_F}{1+\pi_L}M$$

in period 1, and

$$\frac{(1+r)}{(1+\pi_L)}M$$

in period 2. The household now avoids making high period-1 interest payments when inflation turns out to be high, and has partial protection against a high real debt burden when inflation turns out to be low. (The refinancing option protects against one period of low inflation but not two, since a period elapses between initial financing and refinancing.)

Of course, the option to refinance does not come for free. The rate on a refinaneable FRM must be higher than on a FRM that prohibits refinancing. A

refinanceable FRM must also require higher period-1 payments than an ARM. If lenders are risk neutral, then the expected present value of real lifetime payments is the same for all mortgages. A refinanceable FRM has lower expected real payments in period 2 than does an ARM, thus the refinanceable FRM must have higher expected payments in period 1. This means that a borrowing-constrained household that is risk neutral will prefer an ARM, because such a household wishes to increase average period-1 consumption relative to period-2 consumption, and does not care about the randomness of mortgage payments.³³

A risk-averse-constrained household, however, may prefer a refinanceable FRM because the reduction in period-1 consumption risk may outweigh both the increase in period-2 consumption risk and the reduction in the average level of period-1 consumption. This is a striking reversal of the analysis for unconstrained households, which find ARMs to be safer than FRMs. The lesson is that perceptions of risk can be profoundly affected by the presence of borrowing constraints.

The nature of risk aversion in this model deserves some discussion. If the household has a utility function defined over consumption in each period, then under standard assumptions about the utility function, the derived risk aversion with respect to mortgage payment risk is increasing in the level of consumption risk aversion. It is also increasing with respect to randomness in real income that is uncorrelated with inflation. Such randomness creates a "background risk" that makes the consumer more averse to mortgage risk. Finally, the randomness of mortgage payments is relatively more important the larger the mortgage relative to the household's income. Thus, the model implies that a household should be more likely to prefer a refinanceable FRM if it is naturally conservative, has a volatile income, or has a large mortgage relative to its income. Campbell and Cocco (2003) obtain these results in their multiperiod numerical model.

It is possible to enrich this model further to allow for refinancing that is driven by decisions to move rather than by movements in interest rates. Households that move must refinance their mortgages even if interest rates are high rather than low. Refinancing of this sort reduces the period-2 benefits of refinanceable FRMs to borrowers, and thereby lowers the equilibrium interest rate on these mortgages. For a given moving probability in the population as a whole, and thus a given refinanceable FRM rate, a household that is more likely to move gets a lower benefit from a refinanceable FRM and is more likely to prefer an ARM.

Finally, one can allow for the fact that many households do not refinance their mortgages even when it is optimal to do so. A household that fails to refinance a FRM when the interest rate falls pays additional mortgage interest, so the presence of such households in the population reduces the equilibrium mortgage

³³ This result is consistent with statements in some personal finance guides. Orman (1999), for example, writes that "ARMs are best utilized . . . when your cash flow is currently tight but you expect it to increase as time goes on" (p. 254), while Irwin (1996) writes "Sometimes ARMs have lower initial loan costs. If cash is a big consideration for you, look into them" (p. 144).

rate. Given the aggregate prepayment rate, and thus the refinanceable FRM rate, a household that believes itself to be more likely than average to refinance optimally will realize a higher benefit from a refinanceable FRM and is more likely to prefer this type of mortgage.

Appendix B: Household-Level Mortgage Data

The household-level mortgage data studied in this paper and in Schwartz (2006) come from two data sources: the American Housing Survey (AHS), conducted by the U.S. Department of Housing and Urban Development, and the Residential Finance Survey (RFS), conducted by the Census Bureau. Both surveys provide basic information on housing units, their residents, and their mortgages. However, they have different strengths and weaknesses.

The AHS is conducted in every odd year. The survey follows housing units rather than households, with interviewers returning to the same housing units since 1985 and adding new units to reflect new housing development. Thus, the AHS is a panel data set with a long time dimension. Interviewers ask detailed questions about the residents of each housing unit, including demographic and educational information, whether residents own their housing unit, whether they have a mortgage, what is the form of the mortgage, and what rate is being paid on the mortgage.

The main weaknesses of the AHS have to do with data quality. Participation in the survey is voluntary, and many residents either refuse to participate or give only partial answers to survey questions. Even a simple question such as whether a housing unit's residents lived in the unit at the time of the previous survey is answered unreliably. To establish whether a household is the same as the one that responded to the previous survey, one can combine the answer to this question with demographic data.

Mortgage data are even more problematic. A detailed review of the AHS by Lam and Kaul (2003) shows that as mortgages age, households' reports of the original principal value, monthly payments, and mortgage rates can change from one survey to the next. Schwartz (2006a) handles this problem by following households through time, assuming at each point in time that their mortgage was originated at the most recent date they have ever previously reported as a mortgage origination date, and assuming that the terms of the mortgage are those they reported in the most recent survey following the origination date. The motivation for this procedure is that households are more likely to report the terms of their mortgage accurately if they recently took out the mortgage. Mortgage rates derived in this manner track the historical rates reported by the FHLMC much more accurately than the raw reported rates. However, there are still some implausibly low rates, and I replace any rate that is more than two percentage points below the average FHLMC rate prevailing during the mortgage origination period with the average FHLMC rate less 2%. The results are robust to reasonable variations in this truncation procedure, for example, replacing self-reported mortgage rates with FHLMC rates.

The RFS has been conducted every 10 years since 1951 as part of the decennial census. Each survey draws a different subsample from the census population, so the RFS is a series of cross-sections and not a panel. The RFS asks homeowners to identify their mortgage lenders and then cross-checks mortgage data with the lenders, resulting in much greater accuracy of mortgage terms. The main weakness of the RFS is that it contains relatively little information on homeowners; in particular, it lacks the educational information contained in the AHS. I use the RFS as a cross-check on the mortgage rates reported in the AHS. I find that the distribution of mortgage spreads in the 2001 RFS closely matches the distribution extracted from the 2001 AHS and reported in Figure 5.

REFERENCES

- Agarwal, Sumit, Souphala Chomsisengphet, Chunlin Liu, and Nicholas S. Souleles, 2005, Do consumers choose the right credit contracts? Working paper, University of Pennsylvania.
- Agarwal, Sumit, John C. Driscoll, and David I. Laibson, 2006, When should borrowers refinance their mortgages? Working paper, Brown University and Harvard University.
- Agnew, Julie, Pierluigi Balduzzi, and Annika Sunden, 2003, Portfolio choice and trading in a large 401(k) plan, *American Economic Review* 93, 193–215.
- Alm, James, and James R. Follain, Jr., 1984, Alternative mortgage instruments, the tilt problem, and consumer welfare, *Journal of Financial and Quantitative Analysis* 19, 113–126.
- Ameriks, John, and Stephen P. Zeldes, 2004, How do household portfolio shares vary with age? Working paper, Columbia University.
- Amromin, Gene, Jennifer Huang, and Clemens Sialm, 2005, Responsible fools? The tradeoff between mortgage prepayments and tax-deferred retirement savings, Working paper, Federal Reserve Bank of Chicago, University of Texas, and University of Michigan.
- Archer, Wayne R., David C. Ling, and Gary A. McGill, 1996, The effect of income and collateral constraints on residential mortgage terminations, *Regional Science and Urban Economics* 26, 235–261.
- Banks, James, Richard Blundell, and James P. Smith, 2005, Wealth portfolios in the United Kingdom and the United States, in David A. Wise, ed.: *Perspectives on the Economics of Aging* (University of Chicago Press, Chicago, IL).
- Banks, James, and Sarah Tanner, 2002, Household portfolios in the United Kingdom, in Luigi Guiso, Michael Haliassos, and Tullio Jappelli, eds.: *Household Portfolios* (MIT Press, Cambridge, MA).
- Barber, Brad M., and Terrance Odean, 2000, Trading is hazardous to your wealth: The common stock investment performance of individual investors, *Journal of Finance* 55, 773–806.
- Barber, Brad M., and Terrance Odean, 2001, Boys will be boys: Gender, overconfidence, and common stock investment, *Quarterly Journal of Economics* 116, 261–292.
- Barber, Brad M., Terrance Odean, and Lu Zheng, 2003, Out of sight, out of mind: The effects of expenses on mutual fund flows, Working paper, University of California and University of Michigan.
- Benartzi, Shlomo, 2001, Excessive extrapolation and the allocation of 401(k) accounts to company stock, *Journal of Finance* 56, 1747–1764.
- Benartzi, Shlomo, and Richard Thaler, 2001, Naive diversification strategies in defined contribution saving plans, *American Economic Review* 91, 79–98.
- Benjamin, Daniel J., Sebastian A. Brown, and Jesse M. Shapiro, 2006, Who is “behavioral”? Cognitive ability and anomalous preferences, Working paper, Harvard University and University of Chicago.
- Bennett, Paul, Richard Peach, and Stavros Peristiani, 2001, Structural change in the mortgage market and the propensity to refinance, *Journal of Money, Credit, and Banking* 33, 955–975.

- Benzoni, Luca, Pierre Collin-Dufresne, and Robert S. Goldstein, 2005, Portfolio choice over the life-cycle in the presence of 'trickle down' labor income, NBER working paper no. 11247.
- Bergstresser, Daniel, and James Poterba, 2004, Asset allocation and asset location: Household evidence from the Survey of Consumer Finances, *Journal of Public Economics* 88, 1893–1915.
- Bernheim, B. Douglas, 1998, Financial illiteracy, education, and retirement saving, in Olivia Mitchell and Sylvester Schieber, eds.: *Living with Defined Contribution Pensions* (University of Pennsylvania Press, Philadelphia, PA).
- Bertaut, Carol C., and Michael Haliassos, 2006, Credit cards: Facts and theories (forthcoming), in G. Bertola, R. Disney, and C. Grant, eds.: *The Economics of Consumer Credit Demand and Supply* (MIT Press, Cambridge, MA).
- Bertaut, Carol C., and Martha Starr-McCluer, 2002, Household portfolios in the United States, in Luigi Guiso, Michael Haliassos, and Tullio Jappelli, eds.: *Household Portfolios* (MIT Press, Cambridge, MA).
- Blume, Marshall, and Irwin Friend, 1975, The asset structure of individual portfolios and some implications for utility functions, *Journal of Finance* 30, 585–603.
- Blume, Marshall, and Irwin Friend, 1978, *The Changing Role of the Individual Investor* (John Wiley, New York, NY).
- Bodie, Zvi, Robert C. Merton, and William Samuelson, 1992, Labor supply flexibility and portfolio choice in a life cycle model, *Journal of Economic Dynamics and Control* 16, 427–449.
- Bolton, Patrick, and Howard Rosenthal, eds.: 2005, *Credit Markets for the Poor* (Russell Sage Foundation, New York, NY).
- Brennan, Michael J., Eduardo S. Schwartz, and Ronald Lagnado, 1997, Strategic asset allocation, *Journal of Economic Dynamics and Control* 21, 1377–1403.
- Brunnermeier, Markus, and Stefan Nagel, 2005, Do wealth fluctuations generate time-varying risk aversion? Micro-evidence on individuals' asset allocation, Working paper, Princeton University and Stanford University.
- Bucks, Brian, and Karen Pence, 2006, Do homeowners know their house values and mortgage terms? Working paper, Federal Reserve Board of Governors.
- Calvet, Laurent, John Y. Campbell, and Paolo Sodini, 2006, Down or out: Assessing the welfare costs of household investment mistakes, NBER working paper no. 12030.
- Campbell, John Y., and João F. Cocco, 2003, Household risk management and optimal mortgage choice, *Quarterly Journal of Economics* 118, 1449–1494.
- Campbell, John Y., and John H. Cochrane, 1999, By force of habit: A consumption-based explanation of aggregate stock market behavior, *Journal of Political Economy* 107, 205–251.
- Campbell, John Y., and Luis M. Viceira, 1999, Consumption and portfolio decisions when expected returns are time varying, *Quarterly Journal of Economics* 114, 433–495.
- Campbell, John Y., and Luis M. Viceira, 2001, Who should buy long-term bonds? *American Economic Review* 91, 99–127.
- Campbell, John Y., and Luis M. Viceira, 2002, *Strategic Asset Allocation: Portfolio Choice for Long-Term Investors* (Oxford University Press, New York, NY).
- Campbell, John Y., and Luis M. Viceira, 2003, A multivariate model of strategic asset allocation, *Journal of Financial Economics* 67, 41–80.
- Canner, Niko, N. Gregory Mankiw, and David Weil, 1997, An asset allocation puzzle, *American Economic Review* 87, 181–191.
- Caplin, Andrew, Charles Freeman, and Joseph Tracy, 1997, Collateral damage: Refinancing constraints and regional recessions, *Journal of Money, Credit, and Banking* 29, 496–516.
- Carroll, Christopher D., 2002, Portfolios of the rich, in Luigi Guiso, Michael Haliassos, and Tullio Jappelli, eds.: *Household Portfolios* (MIT Press, Cambridge, MA).
- Chetty, Raj, and Adam Szeidl, 2005, Consumption commitments: Neoclassical foundations for habit formation, Working paper, University of California Berkeley.
- Choi, James J., David Laibson, and Brigitte Madrian, 2005, \$100 bills on the sidewalk: Suboptimal saving in 401(k) plans, Working paper, Yale University, Harvard University, and University of Pennsylvania.
- Choi, James J., David Laibson, Brigitte Madrian, and Andrew Metrick, 2002, Defined contribution pensions: Plan rules, participant decisions, and the path of least resistance, in James Poterba, ed.: *Tax Policy and the Economy* 16, 67–113.

- Choi, James J., David Laibson, Brigitte Madrian, and Andrew Metrick, 2004a, Employees' investment decisions about company stock, in Olivia S. Mitchell and Stephen P. Utkus, eds.: *Pension Design and Structure: New Lessons from Behavioral Finance* (Oxford University Press, New York, NY).
- Choi, James J., David Laibson, Brigitte Madrian, and Andrew Metrick, 2004b, For better or for worse: Default effects and 401(k) savings behavior, in David Wise, ed.: *Perspectives on the Economics of Aging* (University of Chicago Press, Chicago, IL).
- Cocco, João F., 2005, Portfolio choice in the presence of housing, *Review of Financial Studies* 18, 535–567.
- Cocco, João F., Francisco J. Gomes, and Pascal J. Maenhout, 2005, Consumption and portfolio choice over the life cycle, *Review of Financial Studies* 18, 491–533.
- Cooper, Ian, and Evi Kaplanis, 1994, Home bias in equity portfolios, inflation hedging and international capital market equilibrium, *Review of Financial Studies* 7, 45–60.
- Curcuro, Stephanie, John Heaton, Deborah Lucas, and Damien Moore, 2006, Heterogeneity and portfolio choice: Theory and evidence (forthcoming), in Yacine Ait-Sahalia and Lars P. Hansen, eds.: *Handbook of Financial Econometrics* (Elsevier Science, Amsterdam).
- Dammon, Robert M., Chester S. Spatt, and Harold H. Zhang, 2001, Optimal consumption and investment with capital gains taxes, *Review of Financial Studies* 14, 583–616.
- Dammon, Robert M., Chester S. Spatt, and Harold H. Zhang, 2004, Optimal asset location and allocation with taxable and tax-deferred investing, *Journal of Finance* 59, 999–1037.
- Davidoff, Thomas, 2006, Labor income, housing prices, and homeownership, *Journal of Urban Economics* 59, 209–235.
- Davis, Steven J., Felix Kubler, and Paul Willen, 2006, Borrowing costs and the demand for equity over the life cycle, *Review of Economics and Statistics* 88, 348–362.
- Dunn, Kenneth B., and John J. McConnell, 1981, Valuation of mortgage-backed securities, *Journal of Finance* 36, 599–617.
- Ellison, Glenn, 2005, A model of add-on pricing, *Quarterly Journal of Economics* 120, 585–638.
- Epstein, Lawrence, and Stanley Zin, 1989, Substitution, risk aversion, and the temporal behavior of consumption and asset returns: A theoretical framework, *Econometrica* 57, 937–969.
- Fama, Eugene, and Kenneth French, 1993, Common risk factors in the returns on stocks and bonds, *Journal of Financial Economics* 33, 3–56.
- Farhi, Emmanuel, and Stavaros Panageas, 2005, Saving and investing for early retirement: A theoretical analysis, Working paper, MIT and University of Pennsylvania.
- Feng, Lei, and Mark S. Seasholes, 2004, Portfolio choice and location of trade, Working paper, University of California Berkeley.
- Fisher, Sarah, and Susan Shelly, 2002, *The Complete Idiot's Guide to Personal Finance in Your 20's and 30's* (Alpha Books, Indianapolis, IN), second edition.
- Flavin, Marjorie, and Shinobu Nakagawa, 2004, A model of housing in the presence of adjustment costs: A structural interpretation of habit persistence, NBER working paper no. 10458.
- Flesaker, Bjorn, and Ehud I. Ronn, 1993, The pricing of FIREARMs ("falling interest rate adjustable-rate mortgages"), *Journal of Real Estate Finance and Economics* 6, 251–275.
- Fratantoni, Michael C., 2001, Homeownership, committed expenditure risk, and the stockholding puzzle, *Oxford Economic Papers* 53, 241–259.
- French, Kenneth, and James M. Poterba, 1991, Investor diversification and international equity markets, *American Economic Review* 81, 222–226.
- Gabaix, Xavier, Arvind Krishnamurthy, and Olivier Vigneron, 2006, Limits of arbitrage: Theory and evidence from the mortgage-backed securities market (forthcoming), *Journal of Finance*.
- Gabaix, Xavier, and David Laibson, 2006, Shrouded attributes, consumer myopia, and information suppression in competitive markets, *Quarterly Journal of Economics* 121, 505–540.
- Gentry, William, and Glenn Hubbard, 2004, Entrepreneurship and household saving, *Advances in Economic Analysis and Policy* 4, Article 8.
- Goetzmann, William N., and Alok Kumar, 2004, Why do individual investors hold under-diversified portfolios? Working paper, Yale University and University of Notre Dame.
- Gomes, Francisco, and Alexander Michaelides, 2005, Optimal life-cycle asset allocation: Understanding the empirical evidence, *Journal of Finance* 60, 869–904.

- Gomes, Francisco, Alexander Michaelides, and Valery Polkovnichenko, 2004, Portfolio choice and wealth accumulation with taxable and tax-deferred accounts, Working paper, London Business School, London School of Economics, and University of Minnesota.
- Gourinchas, Pierre-Olivier, and Jonathan Parker, 2002, Consumption over the life cycle, *Econometrica* 70, 47–89.
- Graham, John R., Campbell R. Harvey, and Hai Huang, 2005, Investor competence, trading frequency, and home bias, NBER working paper no. 11426.
- Green, Richard K., and Susan M. Wachter, 2005, The American mortgage in historical and international context, *Journal of Economic Perspectives* 19, 93–114.
- Greenspan, Alan, and James Kennedy, 2005, Estimates of home mortgage originations, repayments, and debt on one-to-four family residences, Finance and Economics Discussion Series 2005-41, Federal Reserve Board, Washington, DC.
- Grinblatt, Mark, and Matti Keloharju, 2000, The investment behavior and performance of various investor-types: A study of Finland's unique data set, *Journal of Financial Economics* 55, 43–67.
- Grinblatt, Mark, and Matti Keloharju, 2001, How distance, language, and culture influence stock holdings and trades, *Journal of Finance* 56, 1053–1073.
- Guiso, Luigi, Michael Haliassos, and Tullio Jappelli, eds.: 2002, *Household Portfolios* (MIT Press, Cambridge, MA).
- Guiso, Luigi, and Tullio Jappelli, 2002, Household portfolios in Italy, in Luigi Guiso, Michael Haliassos, and Tullio Jappelli, eds.: *Household Portfolios* (MIT Press, Cambridge, MA).
- Guiso, Luigi, and Tullio Jappelli, 2005, Awareness and stock market participation, *Review of Finance* 9, 537–567.
- Guiso, Luigi, Paola Sapienza, and Luigi Zingales, 2005, Trusting the stock market, Working paper, University of Chicago and Northwestern University.
- Haliassos, Michael, and Carol Bertaut, 1995, Why do so few hold stocks? *Economic Journal* 105, 1110–1129.
- Haliassos, Michael, and Alexander Michaelides, 2003, Portfolio choice and liquidity constraints, *International Economic Review* 44, 143–177.
- Heaton, John, and Deborah Lucas, 2000, Portfolio choice and asset prices: The importance of entrepreneurial risk, *Journal of Finance* 55, 1163–1198.
- Heckman, James, and R. Robb, 1985, Using longitudinal data to estimate age, period, and cohort effects in earnings equations, in W. Mason and S. Fienberg, eds.: *Cohort Analysis in Social Research: Beyond the Identification Problem*, 137–150 (Springer-Verlag, New York, NY).
- Hong, Harrison, Jeffrey D. Kubik, and Jeremy Stein, 2004, Social interaction and stock-market participation, *Journal of Finance* 59, 137–163.
- Hortacsu, Ali, and Chad Syverson, 2004, Product differentiation, search costs, and competition in the mutual fund industry: A case study of S&P 500 index funds, *Quarterly Journal of Economics* 119, 403–456.
- Huberman, Gur, 2001, Familiarity breeds investment, *Review of Financial Studies* 14, 659–680.
- Irwin, Robert, 1996, *Tips and Traps When Buying a Home* (McGraw-Hill, New York, NY).
- Ivković, Zoran, Clemens Sialm, and Scott Weisbenner, 2004, Portfolio concentration and the performance of individual investors, NBER working paper no. 10675.
- Ivković, Zoran, and Scott Weisbenner, 2003, Local does as local is: Information content of the geography of individual investors' common stock investments, NBER working paper no. 9685.
- Iwaisako, Tokuo, 2003, Household portfolios in Japan, NBER working paper no. 9647.
- Johnson, David, Jonathan Parker, and Nicholas Souleles, 2004, Household expenditure and the income tax rebates of 2001, NBER working paper no. 10784.
- Juster, Thomas, and James Smith, 1997, Improving the quality of economic data: Lessons from the HRS and AHEAD, *Journal of the American Statistical Association* 92, 1268–1278.
- Juster, Thomas, James Smith, and Frank Stafford, 1999, The measurement and structure of household wealth, *Labor Economics* 6, 253–275.
- Kalotay, Andrew, and Leslie Abreo, 1999, Ratchet bonds: Maximum refunding efficiency at minimum transaction cost, *Journal of Applied Corporate Finance* 12, 40–47.

- Kelly, Morgan, 1995, All their eggs in one basket: Portfolio diversification of U.S. households, *Journal of Economic Behavior and Organization* 27, 87–96.
- Kennickell, Arthur B., 1998, Multiple imputation in the Survey of Consumer Finances, Working paper, Federal Reserve Board, available at: <http://www.federalreserve.gov/pubs/oss/oss2/method.html>.
- Keynes, John Maynard, 1932, Economic possibilities for our grandchildren, in *Essays in Persuasion* (Harcourt Brace, New York, NY).
- Kim, Tong Suk, and Edward Omberg, 1996, Dynamic nonmyopic portfolio behavior, *Review of Financial Studies* 9, 141–161.
- King, Mervyn A., and Jonathan I. Leape, 1998, Wealth and portfolio composition: Theory and evidence, *Journal of Public Economics* 69, 155–193.
- Kopczuk, Wojciech, and Emmanuel Saez, 2004, Top wealth shares in the United States, 1916–2000: Evidence from estate tax returns, NBER working paper no. 10399.
- Korniotis, George, and Alok Kumar, 2006, Does investment skill decline due to cognitive aging or improve with experience? Working paper, University of Notre Dame.
- Lam, Ken, and Bulbul Kaul, 2003, *Analysis of Housing Finance Issues Using the American Housing Survey (AHS)*, U.S. Department of Housing and Urban Development, Office of Policy Development and Research, Washington, DC.
- LeRoy, Stephen F., 1996, Mortgage valuation under optimal prepayment, *Review of Financial Studies* 9, 817–844.
- Longstaff, Francis A., 2004, Optimal recursive refinancing and the valuation of mortgage-backed securities, NBER working paper no. 10422.
- Lusardi, Annamaria, and Olivia S. Mitchell, 2006, Financial literacy and planning: Implications for retirement well-being, Working paper, Dartmouth College and University of Pennsylvania.
- Lynch, Anthony W., 2001, Portfolio choice and equity characteristics: Characterizing the hedging demands induced by return predictability, *Journal of Financial Economics* 62, 67–130.
- Lynch, Anthony W., and Sinan Tan, 2004, Labor income dynamics at business cycle frequencies: Implications for portfolio choice NBER working paper no. 11010.
- Madrian, Brigitte, and Dennis Shea, 2001, The power of suggestion: Inertia in 401(k) participation and savings behavior, *Quarterly Journal of Economics* 66, 1149–1188.
- Mankiw, N. Gregory, and Stephen P. Zeldes, 1991, The consumption of stockholders and nonstockholders, *Journal of Financial Economics* 29, 97–112.
- Massa, Massimo, and Andrei Simonov, 2006, Hedging, familiarity and portfolio choice, *Review of Financial Studies* 19, 633–685.
- McConnell, John J., and Manoj Singh, 1994, Rational prepayments and the valuation of collateralized mortgage obligations, *Journal of Finance* 49, 891–921.
- McCulloch, J. Huston, 1986, Risk characteristics and underwriting standards for price level adjusted mortgages versus other mortgage instruments, *Housing Finance Review* 5, 65–97.
- Merton, Robert C., 1971, Optimum consumption and portfolio rules in a continuous-time model, *Journal of Economic Theory* 3, 373–413.
- Merton, Robert C., 1973, An intertemporal capital asset pricing model, *Econometrica* 41, 867–887.
- Miles, David, 2003, *The U.K. Mortgage Market: Taking a Longer-Term View, Interim Report: Information, Incentives, and Pricing* (HM Treasury, London).
- Miles, David, 2004, *The U.K. Mortgage Market: Taking a Longer-Term View, Final Report and Recommendations* (HM Treasury, London).
- Mitchell, Olivia S., and Stephen P. Utkus, 2003, The role of company stock in defined contribution plans, in Olivia S. Mitchell and Kent Smetters, eds.: *The Pension Challenge: Risk Transfers and Retirement Income Security* (Oxford University Press, Oxford).
- Moore, Danna, 2003, Survey of financial literacy in Washington State: Knowledge, behavior, attitudes, and experiences, Technical Report 03-39, Social and Economic Sciences Research Center, Washington State University.
- Nalebuff, Barry, and Ian Ayres, 2003, *Why Not? How to Use Ordinary Ingenuity to Solve Problems Big and Small* (Harvard Business School Press, Boston, MA).

- Odean, Terrance, 1998, Are investors reluctant to realize their losses? *Journal of Finance* 53, 1775–1798.
- Odean, Terrance, 1999, Do investors trade too much? *American Economic Review* 89, 1279–1298.
- Orman, Suze, 1999, *The Courage to Be Rich* (Riverhead Books, New York, NY).
- Palacios-Huerta, Ignacio, 2003, An empirical analysis of the risk properties of human capital returns, *American Economic Review* 93, 948–964.
- Parker, Jonathan A., 1999, The reaction of household consumption to predictable changes in social security taxes, *American Economic Review* 89, 959–973.
- Pelizzon, Lorian, and Guglielmo Weber, 2005, Efficient portfolios when housing is a hedge against rent risk, Working paper, University of Venice and University of Padua.
- Polkovnichenko, Valery, 2005, Household portfolio diversification: A case for rank dependent preferences, *Review of Financial Studies* 18, 1467–1502.
- Poterba, James M., 2002, Taxation, risk-taking, and household portfolio behavior, in Alan Auerbach and Martin Feldstein, eds.: *Handbook of Public Economics Vol. 3* (North-Holland, Amsterdam).
- Poterba, James M., and Andrew Samwick, 2003, Taxation and household portfolio composition: Evidence from tax reforms in the 1980s and 1990s, *Journal of Public Economics* 87, 5–39.
- Puri, Manju, and David Robinson, 2005, Optimism and economic choice, NBER working paper no. 11361.
- Rosen, Harvey S., and Stephen Wu, 2004, Portfolio choice and health status, *Journal of Financial Economics* 72, 457–484.
- Saks, Raven E., and Stephen H. Shore, 2005, Risk and career choice, *Advances in Economic Analysis and Policy* 5, Article 7.
- Samuelson, Paul A., 1938, A note on the pure theory of consumer's behavior, *Economica* 5, 61–71.
- Schlarbaum, Gary G., Wilbur G. Lewellen, and Ronald C. Lease, 1978, The common-stock-portfolio performance record of individual investors: 1964–1970, *Journal of Finance* 33, 429–441.
- Schwartz, Allie, 2006a, Household refinancing behavior in fixed rate mortgages, Working paper, Harvard University.
- Schwartz, Allie, 2006b, Who takes out adjustable rate mortgages? Working paper, Harvard University.
- Schwartz, Eduardo S., and Walter N. Torous, 1989, Prepayment and the valuation of mortgage-backed securities, *Journal of Finance* 44, 375–392.
- Shiller, Robert J., 2003, *The New Financial Order: Risk in the 21st Century* (Princeton University Press, Princeton, NJ).
- Shore, Stephen H., and Todd Sinai, 2005, Commitment, risk, and consumption: Do birds of a feather have bigger nests? Working paper, University of Pennsylvania.
- Sinai, Todd, and Nicholas S. Souleles, 2005, Owner-occupied housing as a hedge against rent risk, *Quarterly Journal of Economics* 120, 763–789.
- Sirri, Erik R., and Peter Tufano, 1998, Costly search and mutual fund flows, *Journal of Finance* 53, 1589–1622.
- Souleles, Nicholas S., 1999, The response of household consumption to income tax refunds, *American Economic Review* 89, 947–958.
- Stanton, Richard, 1995, Rational prepayment and the valuation of mortgage-backed securities, *Review of Financial Studies* 8, 677–708.
- Stanton, Richard, and Nancy Wallace, 1998, Mortgage choice: What's the point?, *Real Estate Economics* 26, 173–105.
- Statman, Meir, 1982, Fixed rate or index-linked mortgages from the borrower's point of view: A note, *Journal of Financial and Quantitative Analysis* 17, 451–457.
- Steinmetz, Thomas C., 2002, *The Mortgage Kit* (Dearborn Trade Publishing, Chicago, IL).
- Stephens Jr., Melvin, 2006, Paycheck receipt and the timing of consumption (forthcoming), *Economic Journal*.
- Storesletten, Kjetil, Chris Telmer, and Amir Yaron, 2004, Cyclical dynamics in idiosyncratic labor-market risk, *Journal of Political Economy* 112, 695–717.
- Swensen, David F., 2005, *Unconventional Success: A Fundamental Approach to Personal Investment* (Free Press, New York, NY).

- Tobin, James, 1958, Liquidity preference as behavior towards risk, *Review of Economic Studies* 25, 68–85.
- Tracy, Joseph, Henry Schneider, and Sewin Chan, 1999, Are stocks overtaking real estate in household portfolios? *Federal Reserve Bank of New York Current Issues in Economics and Finance* 5(5), 1–6.
- Tufano, Peter, and Daniel Schneider, 2005, Reinventing savings bonds, Working paper, Harvard Business School.
- Venti, Steven F., and David A. Wise, 2001, Choice, chance, and wealth dispersion at retirement, in S. Ogura, T. Tachibanaki, and D. Wise, eds.: *Aging Issues in the United States and Japan* (University of Chicago Press, Chicago, IL).
- Viceira, Luis M., 2001, Optimal portfolio choice for long-horizon investors with nontradable labor income, *Journal of Finance* 56, 433–470.
- Vissing-Jorgensen, Annette, 2003, Perspectives on behavioral finance: Does “irrationality” disappear with wealth? Evidence from expectations and actions, in Mark Gertler and Kenneth Rogoff, eds.: *NBER Macroeconomics Annual 2003* (MIT Press, Cambridge, MA).
- Wachter, Jessica, 2003, Risk aversion and allocation to long-term bonds, *Journal of Economic Theory* 112, 325–333.
- Woodward, Susan, 2003, Consumer confusion in the mortgage market, Working paper, Sand Hill Econometrics, available at: <http://www.sandhillecon.com/pdf/consumer.confusion.pdf>.
- Yao, Rui, and Harold H. Zhang, 2005, Optimal consumption and portfolio choices with risky housing and borrowing constraints, *Review of Financial Studies* 18, 197–239.
- Zhu, Ning, 2002, The local bias of individual investors, Working paper, Yale University.
- Zinman, Jonathan, 2005, Piecing together a portfolio puzzle: Accounting for why households borrow “high” and lend “low,” Working paper, Dartmouth College.