



Research Statement

Stijn Van Nieuwerburgh

Stern School of Business, New York University and NBER

June 2008

For an average household, at least two-thirds of total wealth is human wealth, the present value of future labor income. Of the remainder, about one-third is housing. Yet, many theories of households' financial choices overlook the role of these assets as well as the role of imperfect information in households' financial decision making. The field of *household finance* studies households' financial choices and their effects on asset prices and the economy at large. My research contributes three new considerations to this literature:

1. How does housing affect asset markets and how do asset markets affect household mortgage choices?
2. How do the risk-return properties of human wealth differ from those of financial wealth?
3. What is the role of information frictions in financial decision making?

Standard finance theory has historically had a difficult time accounting for most of the time-series fluctuations and cross-sectional differences in asset prices. One reaction is to call each failure a mispricing and argue that it represents a profit opportunity. My work considers these failures of existing theory and uses them to decipher what new model features might be productively incorporated. Specifically, incorporating housing wealth, human wealth, or information frictions helps to diminish the gap between theory and data.

I. Housing Wealth and Asset Markets

A first part of my research agenda focuses on real estate, the largest financial asset for most households. One important feature of housing is that it serves as collateral to borrow against. In the US, aggregate household residential real estate wealth is currently \$20 trillion and mortgage debt stands at \$10 trillion. I argue that the extent to which households can insure against income shocks is limited by the value of housing collateral they can borrow against. By controlling the ability to share risk, the value of housing affects households' willingness to hold risky assets, such as stocks and bonds. Thus fluctuations in housing collateral wealth can help explain puzzling features of stock returns and dispersion in households' consumption. The following four papers derive these predictions from this heterogeneous-agent general equilibrium model and evaluate them empirically.

"Housing Collateral, Consumption Insurance, and Risk Premia: an Empirical Investigation" with H. Lustig (Journal of Finance, 2005), predicts that households are less keen to take on financial risks, and therefore demand a higher return for bearing these risks, when housing collateral is scarce. In the data, we show that a decrease in housing collateral is followed by higher future stock returns, in excess of the risk-free rate and that this relationship is statistically significant. The cross-sectional prediction of the model is that assets whose returns covary more with the value of housing must offer their investors higher returns relative to other assets, particularly when housing

collateral is scarce. In contrast, assets whose value increases when house prices decline are a valuable hedge against the risk of being borrowing-constrained. This additional benefit induces the holders of these assets to accept lower returns. In the data, this mechanism explains more than 80% of the cross-sectional difference between average returns on value and growth stocks. Its pricing errors compare favorably to those of competing asset pricing models.

One of the biggest challenges facing the asset pricing literature is explaining why equities pay returns that are so much higher than risk-free returns. While many mechanisms generate some of this equity premium, few are quantitatively powerful enough to account for its observed magnitude, without an excessively high degree of risk aversion. Even fewer speak to its variation over time. “Can Housing Collateral Explain Long-Run Swings in Asset Returns?” with H. Lustig (2007, mimeo) builds on the housing collateral model sketched in the previous paper, solves for the equilibrium prices, calibrates the model and shows that it can generate an equity premium of 6% with a moderate level of risk aversion. The reason this model can generate the observed equity premium is that, in equilibrium, stocks offer low returns exactly when housing collateral values are low and households are unable to hedge a drop in income by borrowing. Because holding stocks carries this risk, stocks must offer high returns to entice investors. This risk of being poor and unable to borrow also induces households to save more and pushes the return on risk-free assets down to a realistic level. To see if the model can explain the equity premium’s time-variation, we feed the model annual data on housing collateral between 1930 and 2005. The resulting asset return predictions match three long-run trends: a decline in the equity risk premium, a decline in the volatility of stock and bond returns, and a positive correlation between housing collateral and the risk-free rate.

While the first paper showed that value stocks should offer higher returns than growth stocks because those returns covary more with housing collateral, it did not explain the origin of this covariance. “Exploring the Link between Housing and the Value Premium” with H. Lustig (2007, mimeo) considers assets whose cash-flows are relatively front-loaded, such as value stocks, and assets whose cash-flows are relatively back-loaded, such as growth stocks. While the former are exposed to both business cycle shocks to consumption and long-run shocks to the value of housing collateral, the latter are exposed only to the housing collateral shocks. Using the same model calibration as the previous paper, we show that the collateral mechanism is strong enough to account for the observed value premium of 6%.

The last part of the housing collateral research project looks for direct evidence that the scarcity of housing collateral inhibits risk-sharing by investigating consumption differences across US metropolitan areas. “How Much Does Housing Collateral Constrain Regional Risk Sharing?” with H. Lustig (2007, R&R Review of Economic Dynamics) measures the degree of risk sharing as the cross-sectional variance of

consumption relative to the cross-sectional variance of income. The model aggregates heterogeneous, borrowing-constrained households into regions characterized by a common housing market and solves for the equilibrium consumption dynamics. It generates a higher degree of risk sharing when housing collateral is scarce to an extent similar to what we find in the data. This result provides direct support for the risk-sharing mechanism that drives the results in the other three papers.

My interest in regional variations across housing markets led to a new project that explores why house prices differ across regions and over time. The spatial location model in "Why Has House Price Dispersion Gone Up?" with P.-O. Weill (2007, R&R at Review of Economic Studies) is the first dynamic version of the seminal Rosen-Roback model in urban economics. Regions differ in their productivity levels and therefore the wages paid to their resident workers. Since workers are free to move, house prices must adjust to make them indifferent between living in any region. Regions which experience fast wage growth attract new households who bid up house prices. Housing supply regulation constrains the number of new units that can be built per period in each area. Muting the response of quantities amplifies price changes. By feeding realized regional wages into a calibrated version of the model, we can explain the magnitude of the increase in average house prices and the increase in the dispersion of house prices across regions. An increase in housing supply regulation by itself (an alternative candidate explanation) does not generate much of an increase in the price level or its dispersion.

Rather than explore how house prices affect financial asset prices, "Mortgage Timing" with R. Koijen and O. Van Hemert (2008, R&R Journal of Financial Economics) looks at how financial asset prices affect households' mortgage choices. Since one of the most important financial decisions a household makes is what kind of mortgage to choose, understanding mortgage choice is an important component of household finance. The paper starts by documenting large fluctuations in the fraction of mortgages that are adjustable-rate (ARM), as opposed to fixed-rate. We build a simple model of utility-maximizing households, a competitive mortgage market, and a term structure of interest rates with time-varying bond risk premia. The analysis reveals that the relative attractiveness of an ARM is determined primarily by the long-term bond risk premium. This is the expected return on a strategy that buys and holds a long-term bond to maturity and finances it by rolling over a short term bond. When the bond risk premium is high, fixed-rate mortgages are relatively expensive and households prefer the ARM contract. To test the model, we form several measures of the bond risk premium and show that they predict the ARM share in aggregate, state-level, and loan-level data. We conclude that the time variation in households' mortgage choices is consistent with a rational model.

One of the underlying ideas in the work on housing collateral and asset prices is that long-run structural changes in the economy, in this case the changes in the value of housing collateral, can alter the relationship between the return on equity and its price-dividend ratio. “Reconciling the Return Predictability Evidence” with M. Lettau (2008, Review of Financial Studies) extends this idea of structural changes in the price-dividend ratio to consider any source of such structural change. One of the big questions in the asset pricing literature is the extent to which returns are predictable. Campbell and Shiller (1989) noted that in any model that links the price of a stock to the present value of its cash-flows, the price-dividend ratio must forecast either future returns or future dividend growth or both.¹ Yet, the existing literature has found that the price-dividend ratio is a statistically unreliable predictor of stock returns and an even worse predictor of dividend growth. We show that allowing for structural breaks reconciles the present-value model with the weak predictability results found by an econometrician who is unaware of the break. The data point to a structural break in the mean of the price-dividend ratio in the early 1990s. After implementing our simple structural break correction, the break-adjusted price-dividend ratio is a strong predictor of stock returns. Dividend growth is not predictable.

II. Human Wealth and Asset Markets

A typical US household’s total wealth is made up largely of human wealth. One reason that most research has not incorporated this important asset is that its value is difficult to measure. To address this problem, my work takes two approaches: use asset pricing data to infer what properties human wealth must have or use contract theory to predict labor income, and in turn human wealth.

“The Returns on Human Capital: Good News on Wall Street is Bad News on Main Street” with H. Lustig (2008, Review of Financial Studies) addresses the question, what do asset prices tell us about expected returns on human wealth? Our measurement approach uses the first-order condition and budget constraint of a representative investor with Epstein-Zin preferences to back out the human wealth returns that reconcile the data on aggregate consumption and aggregate labor income with that on financial wealth returns. Financial wealth returns are either measured as stock market returns or a more comprehensive return on all claims issued by the entire corporate sector (using Flow of Funds data). The main finding is that the low volatility of aggregate consumption growth and its modest covariance with financial wealth returns implies that human wealth return innovations must covary negatively with financial wealth return innovations.

¹ In “Financial Economics, Market Efficiency and Return Predictability” with R. Koijen (2008), we survey the stock return predictability literature and emphasize the importance of imposing the cross-equation restrictions implied by the present-value model.

Further investigation uncovered evidence of negative covariance between news about future labor income and financial income growth rates at longer horizons. From the perspective of standard production-based asset pricing with Cobb-Douglas preferences, such a negative correlation is a puzzle. This became the motivation for “IT, Corporate Payouts, and the Growing Inequality in Managerial Compensation” with H. Lustig and C. Syverson (2008, mimeo). The paper documents two additional facts: a large increase in total payouts to U.S. corporations’ bond- and stockholders: from 2% of value-added in the mid-1970s to 9% in the mid-2000s; at the same time, the dispersion of managerial pay across industries soared.² An optimal contract model, where the manager has the option to leave his firm and transfer some of his “organizational capital” to a new firm, can explain all three facts. This limited commitment friction is similar to the one that affects borrowing contracts in my housing collateral work. Our explanation rests on the idea that the arrival of information technology in the 1970s slowed the rate at which firms became obsolete. This allowed successful firms to grow larger in size, increasing dispersion in firm size. As firm size becomes more dispersed, managers of larger firms earn higher wages because the organizational capital they can walk away with is more valuable. Because successful managers are tempted to leave, owners must offer all managers less insurance. Less insurance results in more wage dispersion and a lower share of profits being paid out to the average manager. Hence, the model generates a negative correlation between the average pay of owners and managers.

“The Wealth-Consumption Ratio” with H. Lustig and A. Verdelhan (2008, mimeo) makes further progress on measuring (the returns on) human and total wealth. The idea is to use stock and bond prices to learn about the prices of aggregate risk in the economy. The no-arbitrage framework from the term-structure literature allows us to estimate these risk prices in a flexible way. We show that they can be used to price human wealth, which is a claim on a sequence of aggregate payments to labor. Likewise, the price of a claim to consumption is total wealth; its price-dividend ratio the wealth-consumption ratio. We find a high average wealth-consumption ratio (87) and correspondingly a low average risk premium on the consumption claim of 2.2%. Total wealth has the risk return characteristics of a real bond, rather than a stock. This finding is important because it provides a new measurement of the return to total wealth, the main asset pricing factor in many standard asset pricing models, such as the CAPM. Our work also suggests that modeling richer cash-flow processes for stocks may be needed to generate the pronounced difference in the risk-return properties between stocks and total wealth that we find.

² We measure managerial pay using executive compensation data from Execucomp, occupational data for managers from the Current Population Survey, and establishment-level data from the Bureau of Labor Statistics.

III. Information Frictions and Seemingly-Irrational Household Behavior

All previous papers assumed that households made decisions based on full knowledge of the true distribution of future financial payoffs, and stayed perfectly abreast of all newly revealed information. The third area of my research agenda relaxes this perfect information assumption and allows me to tackle some important questions in household finance. The first four papers of this section take a theoretical approach to this question, asking what information households would likely have access to or choose to acquire, while the last two papers take an empirical approach, using survey data to disentangle preference-based from uncertainty-based effects.

To determine if information-based theories are a viable alternative to behavioral, preference-based theories, we need to know what information sets households are basing their decisions on. We cannot observe these information sets directly, but we can ask what information an investor with limited ability to acquire or process information would choose before investing. While information acquisition with one risky asset has been studied extensively, extending this framework to consider how to choose among multiple risky assets to learn about is the more realistic case and delivers new insights. “Information Acquisition and Under-Diversification” with L. Veldkamp (2008, R&R Review of Economic Studies) shows that a household wants to learn about assets they expect to hold and prefers to hold more of assets whose payoffs are less uncertain to them. The interaction of these two effects creates increasing returns to learning about and investing in a set of correlated assets. The resulting asset portfolio appears under-diversified from the perspective of standard theory, but is optimal, given the information constraints.

One of the greatest challenges to household finance, is rationalizing households’ tendency to overweight home assets in the equity portfolio, otherwise known as “home bias.” If households in different countries all acquire the same information, even if they specialize in that information, symmetric information sets will lead them to hold symmetric portfolios in equilibrium. “Information Immobility and the Home Bias Puzzle” with L. Veldkamp (2009, Journal of Finance) extends the one-investor model in the previous paper to an equilibrium setting and shows that households want to make their information sets as different as possible from those of other investors. The most efficient way to do this is to learn about assets that they initially know more about, such as home-country assets. Adding information choice to the standard equilibrium portfolio model allows it to explain a wide variety of puzzling facts. In addition to the home bias, it explains local bias, geographic patterns of analyst forecast accuracy, patterns of portfolio out-performance, the under-diversification within the foreign asset component of investors’ portfolios, and asset price changes following changes in the information environment, such as disclosure laws.

“Inside Information and the Own Company Stock Puzzle” with L. Veldkamp (2006, Journal of the European Economic Association) explores how accounting for human wealth interacts with the information choice problem studied in the previous two papers. A household wants to construct a financial portfolio that hedges its labor income risk, which makes own-company stock undesirable. At the same time, labor income is a source of private information about shocks to the firm. This information advantage makes own-company stock more desirable to learn more about and less risky to hold. Thus, the information effect counter-acts the standard diversification effect. It contributes to an explanation of why households might hold their own-company’s stock.

Another seemingly-irrational pattern is that households react more quickly to bad news about the business cycle than to good news. “Learning Asymmetries in Real Business Cycles” with L. Veldkamp (2006, Journal of Monetary Economics) argues that if a representative household does not know the current state of the business cycle, but only how much it invests, how hard it works and at the end of each quarter, observes how much output was produced, such behavior can be rationalized. When households believe that productivity is high, they work and invest more. Because high production amplifies productivity changes but leaves the variance of noise in output unchanged, it makes output a clearer signal about productivity. Since productivity is mean-reverting, most bad news is observed in times when production is high and news is easiest to observe. Thus households react more to bad news because they are more certain of its informational content. The resulting asymmetries in macroeconomic time series are consistent with U.S. data.

While the previous papers explored the idea that households are imperfectly informed about the return on financial or real investments, the last two papers investigate what households believe about their financial needs in retirement. One of the puzzling aspects of retirement behavior is that most households do not run down their assets late in life. One reason a household might not spend all their wealth is to bequeath some of it to their children. Alternatively, it could be a precaution, in case of a large medical expense such as long-term care. “The Joy of Giving or Assisted Living? Using Strategic Surveys to Separate Bequest and Precautionary Motives” with J. Ameriks, A. Caplin, and S. Laufer (2008, R&R Journal of Finance) proposes a novel tool, strategic survey questions, which places the respondent in a hypothetical scenario which is revealing of the relative strength of bequest and precautionary motives. Vanguard financed our survey which collected answers to strategic survey questions from 1000 retirees, as well as data on their wealth, income, and consumption. Combining the survey data and a standard life-cycle consumption-savings model allowed us to disentangle the bequest and precautionary savings motives. We uncovered a large degree of preference heterogeneity, but also a much stronger average bequest motive than previously believed.

One reason that it is important to understand (and separate) savings motives is to properly design and price new retirement products or government programs for retirees. "Annuity Valuation Given Long-term Care Concerns and Bequest Motives", with J. Ameriks, A. Caplin, and S. Lauffer, (2007, Pension Research Council Conference Volume), uses the results from the previous paper to calculate our survey respondents' willingness to pay for standard retirement products, such as annuities. We found low willingness to pay, in line with the low observed take-up rates. We then used the data and the model to ask which retirement products would be most valuable. Of the various products we investigated, long-term care insurance and annuities with increased payments when households are in need of long-term care, would improve welfare of retirees the most.

IV Future Research

My work so far makes the case that three often-overlooked considerations are important in understanding household financial decision-making: their housing wealth, non-financial risk from labor income or health, and the information households use in their decision-making. My work in progress further extends each of these themes.

One important question for academics, financial market participants, and monetary policy makers is how much households increase spending when house prices rise and what makes that response change. Relatively little is known because empirical estimates based on aggregate data are plagued by endogeneity issues. With S. Ludvigson and J. Favilukis, we are studying households' consumption, housing purchase, and housing financing decisions in a heterogeneous-agent general equilibrium model. Starting from exogenous aggregate and idiosyncratic productivity shocks, the model allows us to quantify the aggregate propensity to consume out of housing wealth and to understand its composition across age and income groups.

Still in the realm of housing, another question I plan to work on is how households' mortgage and refinancing choices affect the term structure of interest rates in a general equilibrium model. The existing literature treats mortgages essentially as redundant assets, whose prices are implied by the prices of other financial assets. For example, the mortgage prepayment literature has focused on explaining observed mortgage refinancing decision in empirical, reduced-form models that take the interest rate as given. The sheer size of the mortgage-backed securities market, which exceeds that of the Treasury market, suggests that it may be important to understand feedback relationships from both initial mortgage origination decisions and refinancing choices on interest rates. Put differently, reduced-form prepayment regressions have endogenous variables on the right-hand side.

In the area of retirement finance, I am working with R. Koijen and M. Yogo on financial product design for retirees. Rather than exogenously fixing the menu of assets households have access to, and investigating whether an additional retirement product is valuable, we aim to proceed more systematically. We consider a world where households have access to a full set of state-contingent claims, but which is subject to adverse selection frictions that arise in insurance markets for longevity risk, health risk, and house price risk. In this world, what retirement product maximizes household welfare? The solution will identify which risks are most valuable to insure. A comparison with the allocation that can be achieved with currently available products will allow us to gauge the welfare improvements from financial innovation.

Finally, I am continuing my work on information frictions with L. Veldkamp to better understand the asset management industry. Households delegate their investment decisions to portfolio managers who collect and process information on their behalf. While previous work has considered partial equilibrium models of delegated portfolio management, our general equilibrium model allows us to understand how a fund's investment choices affect asset prices. It also allows us to ask whether a highly skilled manager, one with a large information capacity, will run a large fund which charges low fees or a small fund with high fees. This should deepen our understanding of how the trade-off between fund size and performance depends on the manager's skill and the assets he invests in. It also allows us to think about whether actively-managed funds are socially beneficial. Actively-managed funds discover private information which reduces uncertainty for their investors. The tradeoff is that this activity increases asset prices and therefore the cost of investing for everyone.