

UC Berkeley

UC Berkeley Electronic Theses and Dissertations

Title

Essays on the Consumption and Investment Decisions of Households in the Presence of Housing and Human Capital

Permalink

<https://escholarship.org/uc/item/38g4q4vm>

Author

Betermier, Sebastien

Publication Date

2010

Peer reviewed|Thesis/dissertation

**Essays on the Consumption and Investment Decisions of Households in
the Presence of Housing and Human Capital**

by

Sebastien Jean-Michel Betermier

A dissertation submitted in partial satisfaction of the
requirements for the degree of
Doctor of Philosophy

in

Business Administration

in the

GRADUATE DIVISION
of the
UNIVERSITY OF CALIFORNIA, BERKELEY

Committee in charge:
Professor Richard Stanton, Chair
Professor Martin Lettau
Professor Adam Szeidl
Professor Johan Walden
Professor Nancy Wallace

Spring 2010

**Essays on the Consumption and Investment Decisions of Households in
the Presence of Housing and Human Capital**

Copyright 2010

by

Sebastien Jean-Michel Betermier

Abstract

Essays on the Consumption and Investment Decisions of Households in the Presence
of Housing and Human Capital

by

Sebastien Jean-Michel Betermier

Doctor of Philosophy in Business Administration

University of California, Berkeley

Professor Richard Stanton, Chair

This dissertation consists of three essays in which I study the consumption and investment decisions of households in the presence of two major asset classes: housing and human capital.

In the first essay, I analyze how the dual consumption-investment nature of housing affects the consumption dynamics of households. A key feature of the housing market is that for most households, the consumption and investment benefits of housing are non-separable. I propose a tractable theoretical framework to understand the impact of this constraint on the consumption allocation of homeowners who would ideally like to own just a fraction of their home. For these homeowners, the relative cost of living in their home is not just the imputed rental cost. It also includes an opportunity cost of having an unbalanced financial portfolio. This cost varies substantially over time, and it is especially high in “good” times, when available investment opportunities yield high returns and homeowners allocate a high fraction of their wealth to current consumption. As a result, this cost dampens variations in the level of their housing consumption, and it amplifies variations in both their level of non-housing consumption and the composition of their consumption baskets.

I then test empirically this theory in the second essay. Using household-level data from the Panel Study of Income Dynamics (PSID), I test the hypothesis that homeowners who face a high opportunity cost choose *ceteris paribus* a low housing con-

sumption volatility. I also develop a method to identify these constrained homeowners by comparing their characteristics to those of a subset of unconstrained homeowners: the landlords. The results are consistent with the predictions of the model. First, the characteristics of homeowners that determine how constrained they are in the model are strong predictors of those homeowners who choose to be landlords in the data. For example, homeowners with a low level of risk aversion, little value for housing consumption, and a long horizon are relatively more likely to be landlords. Second, I find evidence that the more constrained homeowners adjust their level of housing consumption much less over time.

In the third essay, which was developed in collaboration with Thomas Jansson, Christine Parlour, and Johan Walden, we investigate the relationship between workers' labor income and their investment decisions. Using a detailed Swedish data set on employment and portfolio holdings we estimate wage volatility and labor productivity for Swedish industries and, motivated by theory, we show that highly labor productive industries are more likely to pay workers variable wages. We also find that both levels and changes in wage volatility are significant in explaining changes in household investment portfolios. A household going from an industry with low wage volatility to one with high volatility will *ceteris paribus* decrease its portfolio share of risky assets by 25%, i.e., 7,750 USD. Similarly, a household that switches from a low labor productivity industry to one with high labor productivity decreases its risky asset share by 20%. Our results suggest that human capital risk is an important determinant of household portfolio holdings.

Professor Richard Stanton
Dissertation Committee Chair

Contents

List of Figures	iii
List of Tables	iv
1 Out of Balance: Housing and the Consumption Allocation of Households	1
1.1 Introduction	2
1.2 Setup of the Model	7
1.3 Derivation of the Model	12
1.3.1 Methodology	12
1.3.2 The Unconstrained Agent's Problem	12
1.3.3 The Homeowner's Problem	14
1.4 Calibration of the Model	21
1.4.1 Parameters	21
1.4.2 Results	25
1.4.3 Comparative Statics	29
1.5 Conclusion	32
2 Housing and the Consumption Allocation of Households: An Empirical Analysis	34
2.1 Introduction	35
2.2 Description of the Datasets	36
2.3 Identification of the Constrained Homeowners	42
2.4 Effect of Indivisibility on Housing Consumption Volatility	48
2.5 Conclusion	54
3 Hedging Labor Income Risk	56
3.1 Introduction	57
3.2 Theoretical Framework and Empirical Strategy	60
3.3 Description of the Datasets	63
3.4 Construction of Variables	65
3.4.1 Portfolios	65

3.4.2	Wage Volatility and Labor Productivity	69
3.5	Empirical Tests and Results	73
3.6	Conclusion	86
A	Derivation of the Portfolio Choice Model in Chapter 1	94
A.1	Format of the Solution	94
A.2	Solution for the Unconstrained Agent	96
A.3	Solution for the Homeowner	98
B	A Mean-Variance Approach to the Opportunity Cost of Housing Services in Chapter 1	101
B.1	The Homeowner's Problem	101
B.2	The Opportunity Cost of Housing Services	103
C	Information about the Dataset in Chapter 2	106
C.1	Defining a Household	106
C.2	Demographics	107
C.3	Assets	107
C.4	House Price Dynamics	109
C.5	Risk Aversion	109
C.6	Moves	110
C.7	Labor Occupation	111
C.8	Variable Definitions	111
D	Information about the Dataset in Chapter 3	115
D.1	Variable Definitions	115

List of Figures

1.1	Investment options in housing in the real world	10
1.2	Optimal consumption allocation for the homeowner	18
1.3	Interest rates and consumption allocation	20
3.1	Summary of model — Two sector example.	62
B.1	Mean-Standard Deviation Graph	104

List of Tables

1.1	Values of the baseline parameters	22
1.2	Portfolio allocation under the baseline parameters	25
1.3	Consumption allocation under the baseline parameters	27
1.4	Comparative statics	31
2.1	General summary statistics for households in 2005	38
2.2	Portfolio characteristics for households in 2005	39
2.3	Neighborhood characteristics for households in 2005	41
2.4	First-stage probit regression on the likelihood of being a landlord . .	47
2.5	Characteristics of adjustments in households' level of housing consumption	51
2.6	Effect of indivisibility on housing consumption volatility	53
3.1	Summary statistics from the 2001 wave of the LINDA data set. . . .	67
3.2	Participation rates and portfolio shares for households in 2001. . . .	68
3.3	Rankings of industries by their levels of wage volatility and labor productivity.	71
3.4	Pearson correlations between labor elasticity and wage measures. . .	74
3.5	Effect of wage volatility in 2002 on portfolio shares.	78
3.6	Effect of labor productivity in 2002 on portfolio shares.	79
3.7	Effect of a change in wage volatility between 1999 and 2002 on portfolio shares.	83
3.8	Effect of a change in labor productivity between 1999 and 2002 on portfolio shares.	84
C.1	Definition of age and cohort groups	108
C.2	Estimates of relative risk aversion for households in the PSID	110
C.3	Categories of moves in the PSID	111
C.4	Categories of Labor Occupations from the Census Bureau	112

Acknowledgments

Berkeley has had a profound influence on my intellectual development, and I am indebted to the Haas Ph.D. program for this opportunity. I am especially grateful to Richard Stanton, Martin Lettau, Johan Walden, and Nancy Wallace for their continuous guidance and support throughout the research stage. Their help in making my work more focused and meaningful was invaluable. I would also like to thank Jonathan Berk, Sara Holland, Dwight Jaffee, John O'Brien, Miguel Palacios, Christine Parlour, Adam Szeidl, and Carles Vergara-Alert for all their helpful comments and discussions. Kim Guilfoyle and June Wong went above and beyond to help me with the important and various administrative procedures. My fellow graduate students in Finance provided a wonderful working environment which was stimulating, supportive, and very fun!

I am also grateful to the Finance faculty at Carnegie Mellon for welcoming me during the last year of the program, and in particular to Christine Parlour and Burton Hollifield who made it possible.

I received generous financial support from the Institute for Business and Economic Research (IBER), the White Foundation, the Dean Witter foundation, and the Fisher Center for Real Estate and Urban Economics. My doctorate would simply not have been possible without this help.

Prior to Berkeley, I did not realize that I would choose an academic career path until I met Peter Lindert at U.C. Davis who introduced me to the world of research. I also draw inspiration from Marie-France Weber from the Lycée La Pérouse, who showed me how rewarding and influential teaching can be.

Katherine and Chloe deserve all the credit for providing the emotional stability required to complete this work, and most importantly, for giving me a purpose. Finally, I need to thank my parents, Philippe and Sylvie, who have provided unconditional support for all my decisions since the very beginning. This dissertation is dedicated to them.

Chapter 1

Out of Balance: Housing and the Consumption Allocation of Households

Unlike most other goods, housing provides both consumption and investment benefits. However, the structure of the housing market is such that most households cannot separate these benefits. In this essay, I propose a tractable theoretical framework to analyze the impact of this constraint on the consumption allocation of homeowners who would ideally like to own just a fraction of their home. I show that for these homeowners, the relative cost of living in their home is not just the imputed rental cost. It also includes an opportunity cost of having an unbalanced financial portfolio. This cost varies substantially over time, and it is especially high in “good” times, when available investment opportunities yield high returns and homeowners allocate a high fraction of their wealth to current consumption. As a result, this cost dampens variations in the level of their housing consumption, and it amplifies variations in both their level of non-housing consumption and the composition of their consumption baskets.

1.1 Introduction

Housing is a distinct non-standard good. On the one hand, it is arguably the single most important consumption good. Expenditures on housing consumption represent about 20-30% of total consumption expenditures.¹ On the other hand, as an extremely durable good in limited supply, it also provides investment benefits.² In addition, the current structure of the housing market is such that most households cannot separate these benefits. Households can either rent the home they live in or they can own all of it, but there is no easily available hybrid solution that allows them to own a share of their home and rent the remaining share from a co-owner.³ I refer to this friction as the *indivisibility* of the investment in housing. This indivisibility is important because, while housing provides multiple investment benefits ranging from portfolio diversification to a perfect hedge against rent risk, homeowners often have to invest a significant share of their wealth in housing. In 2005, the average ratio of housing wealth to net worth was 1.15, which is about three times the share of their wealth invested in stocks.⁴

The objective of this essay is to study the impact of this indivisibility friction on households' consumption and portfolio decisions. I derive a Merton (1971) portfolio choice model in a world with two consumption goods: a perishable good and a durable good (housing), and I model the decisions of a homeowner who would ideally like to own just a fraction of the home he lives in. In particular, I compare his decisions to what he would do in a hypothetical, frictionless world wherein he can own and consume separate amounts of housing.

The main result is that indivisibility makes the composition of the homeowner's

¹ Author's computations from the Panel Study of Income Dynamics (PSID), based on the dollar amount spent on rent and utilities by renters.

² Other goods like art or antique cars also share the same dual consumption-investment nature.

³ Homeowners can finance their home by borrowing, but in doing so they remain fully exposed to variations in the price of the home. They can also invest in assets like real estate investment trusts (REITs) or housing futures in order to hedge their exposure to house price risk, but these securities are not perfect hedges as they are often based on national or regional house price indices. De Jong et al. (2008) survey the literature on estimates of house price risk and find that more than 50% of it is idiosyncratic.

⁴ Author's computations from the Panel Study of Income Dynamics (PSID). See Table 2.2 for details.

consumption basket *persistently suboptimal* as well as *excessively volatile*. It is suboptimal relative to the frictionless case in that the homeowner always under-consumes housing services and over-consumes the perishable good. The reason is that for the homeowner, the effective cost of living in his home is not just the imputed rental cost. It also includes an opportunity cost of having a financial portfolio that is overweighted in housing. The composition of the homeowner's consumption basket is also too volatile, because this opportunity cost varies over time along with the available investment opportunities. This cost is especially high in "good" times, when the investment opportunities yield high returns, since this is precisely when the homeowner, for whom the housing investment-consumption ratio is already too high, would like to allocate a greater fraction of his wealth to current consumption.

I simulate the model in an economy where the interest rate is stochastic and estimate that under reasonable parameters, indivisibility increases the relative price of housing services by 33% on average. Variations in the opportunity cost component amplify variations in the composition of the homeowner's consumption basket by 41%. They also amplify variations in the growth rate of non-housing consumption by 27% and dampen variations in the growth rate of housing consumption by 13%.

The primary contribution of this essay is to show that indivisibility has a major effect on the consumption allocation of homeowners. This result adds a new dimension to a growing literature that has focused primarily on the effects of indivisibility on housing tenure decisions and portfolio allocations. Here, I abstract from the housing tenure decision and assume that the household is a homeowner, although he would ideally like to own just a fraction of his home. The extensive literature on the housing tenure decision suggests that indivisibility is binding for many households.⁵ Furthermore, Cauley, Pavlov and Schwartz (2007) propose a thought experiment where they allow a homeowner to sell a fraction of his home.⁶ They find that for realistic param-

⁵See Henderson and Ioannides (1983); Ortalo-Magné and Rady (2002); Sinai and Souleles (2005); Yao and Zhang (2005); Davidoff (2006); Corradin (2009); Van Hemert (2009); Sinai and Souleles (2009) among others.

⁶Kraft and Munk (2008) also model the decisions of an agent who can rent and own separate amounts of housing. They derive explicit solutions in the presence of borrowing constraints and transaction costs. In this setting, the composition of the consumption baskets of households can also be distorted. However, there is no indivisibility in their model, which is the object of my paper.

eter values, this alternative provides significant welfare benefits. In my model, I also compare the homeowner's decisions to what he would do if he could sell a fraction of his home. However, unlike Cauley et al. (2007), I focus on the homeowner's decisions to allocate his consumption between housing and the perishable good. I allow the homeowner to choose his level of housing consumption (at all times). His preferences are also non-separable between the two consumption goods in that he values the perishable good highly when his consumption of housing services is low.

The effect that indivisibility has on the composition of the consumption basket of a homeowner is related to Brueckner (1997), who models risky housing in a mean-variance setting and finds that indivisibility makes the homeowner's portfolio of stocks and housing mean-variance inefficient. I show that indivisibility also affects the trade-off between housing consumption and non-housing consumption in a related way. Since the homeowner cannot separate his investment demand from his consumption demand of housing, he must take into account the fact that his desired level of housing consumption may affect his financial portfolio. If the chosen level of housing consumption makes the homeowner's portfolio over-weighted in housing, then the relative price of housing services includes a positive opportunity cost, and as a result, the homeowner allocates a low fraction of his consumption budget to housing.

It is important to mention that this effect of indivisibility on consumption allocation does not require the presence of other frictions. My modeling approach relies on Damgaard, Fuglsbjerg and Munk (2003), who provide a tractable framework for solving the decisions of an agent who consumes a perishable good and a risky durable good (housing) in a world where the *only* friction is that the agent must own what he consumes of housing. I extend their model by adding a stochastic investment opportunity set.⁷ This simple approach allows me to explore in depth the effect of indivisibility on consumption allocation over time. I find that even in the absence of transaction costs in the housing market, the composition of the homeowner's consumption basket can

⁷If the investment opportunity set is constant, my results reduce to Damgaard, Fuglsbjerg and Munk (2003). However, I solve the model differently, by comparing the homeowner's decisions to those of the unconstrained agent. This methodology allows me to identify the effect of indivisibility on the homeowner's decisions.

be significantly distorted.⁸ The reason is that for the homeowner, the relative price of housing consumption varies substantially over time because of the opportunity cost component.

By contrast, the effect of indivisibility on portfolio holdings of stocks requires other frictions. If the correlation between the returns to individual house prices and aggregate stock prices is close to zero, which was documented in Flavin and Yamashita (2002), then indivisibility does not affect the portfolio holdings of stocks in my model. Flavin and Yamashita (2002), Cocco (2005), and Yao and Zhang (2005) show that by adding other frictions like a borrowing constraint or a cost to entering the stock market, indivisibility can explain the life-cycle pattern in the participation rates and portfolio allocations in stocks. Chetty and Szeidl (2009) argue that in the presence of consumption commitments, indivisibility can also affect households' demand for stock holdings.

The amplifying effect of indivisibility on the volatility of the composition of homeowners' consumption baskets has potential implications for consumption-based asset pricing. A recent strand of the asset pricing literature has looked at the effect of the role of housing as part of households' consumption baskets on asset prices. It builds on the idea that asset prices should be related not only to variations in the level of households' non-housing consumption as in the standard C-CAPM model,⁹ but also to variations in the composition of their consumption basket.¹⁰ As Piazzesi, Schneider and Tuzel (2007) put it, if housing and other goods are substitute goods, then recessions are perceived as particularly severe when the share of housing consumption is

⁸Grossmann and Laroque (1990), Cuoco and Liu (2000), Damgaard, Fuglsbjerg and Munk (2003) (in the second part of their paper), and Corradin, Fillat and Vergara-Alert (2009) derive optimal portfolio allocations with housing in the presence of transaction costs and find that it makes housing consumption very sticky.

⁹C-CAPM stands for Consumption Capital Asset Pricing Model. See Rubinstein (1976), Lucas, Jr. (1978), and Breeden (1979).

¹⁰For this result, households' preferences need to be non-separable between housing and the perishable good. Yogo (2006) and Pakos (2007) also derive consumption-based asset pricing models with non-separable preferences between non-durable goods and non-housing durable goods. Lustig and Van Nieuwerburgh (2005) and Favilukis, Ludvigson and Van Nieuwerburgh (2009) work with non-separable preferences between housing and non-housing but focus primarily on the effect of the role of housing as collateral for borrowing in a world where heterogeneous households face limited risk-sharing opportunities.

low. For asset pricing purposes, they refer to these variations in the share of housing consumption as “composition risk.” They model a representative-agent economy with a perfect rental market and find that movements in the observed aggregate expenditure share on housing services from the NIPA tables forecast excess stock returns and help to solve asset pricing puzzles. Fillat (2008) exploits the persistence in this expenditure share by adding recursive preferences and finds that the risk premium on stocks varies across different investment horizons. Siegel (2005) and Flavin and Nakagawa (2008) add transaction costs to housing and show that it further improves the standard C-CAPM at the household level.¹¹ In particular, it relaxes the troublesome constraint that the agent’s risk aversion and his intertemporal elasticity of substitution are inversely related.

Although the analysis in this paper is entirely partial equilibrium,¹² it suggests that if in equilibrium some stockholders are constrained homeowners,¹³ then the level of composition risk in the U.S. economy may be greater than previously estimated. The reason is threefold. First, the variations in the composition of households’ consumption baskets are particularly high for the homeowners who would ideally like to own just a fraction of their home. Second, these variations are important for asset pricing because homeowners represent the vast majority of stock-holders. In 2005, homeowners owned 94% of the aggregate wealth in stocks owned by U.S. households.¹⁴ Third, a large part of these variations come from variations in the opportunity cost component of the relative cost of housing services. This last result is important because it no longer ties composition risk to the volatility of rental prices. As Lustig and Van Nieuwerburgh (2005) point out, the way Piazzesi et al. (2007) measure composi-

¹¹Flavin and Nakagawa (2008) also show that if house prices are not correlated with stock prices then the C-CAPM at the aggregate level can still hold. Their model is an extension of Grossmann and Laroque (1990), who show that in a world with housing as the single consumption good, the C-CAPM does not hold in the presence of transaction costs.

¹²Ortalo-Magné and Prat (2009) show that in a general equilibrium setting with multiple locations, fractional home ownership can be the optimal outcome.

¹³In a world where there is only one house and one agent, the house price will adjust so that the agent is happy to own the entire home. However, in a world with heterogeneous agents and multiple homes that do not have perfectly correlated prices, this is no longer necessarily the case.

¹⁴Authors’s computations from the 2005 wave of the PSID. Stocks include U.S. and foreign stocks. See Table 2.2 for details.

tion risk in the presence of a perfect rental market can generate a volatility of rental prices that is an order of magnitude too high.¹⁵ Since the observed aggregate expenditure share of housing services does not vary much over time, it can only generate a volatile stochastic discount factor if the representative agent is inherently adverse to changing his expenditure share. This is the case when his intra-temporal elasticity of substitution is close to one. In this case, small changes in the expenditure shares can imply large changes in rental prices.

The rest of this essay is organized as follows. I set up the portfolio choice model in Section 1.2 and derive it in Section 1.3. I then calibrate the model in Section 1.4. Section 1.5 concludes. All the derivations are detailed in Appendix A. In Appendix B, I provide intuition on what the opportunity cost of housing services corresponds to in a simple mean-variance setting.

1.2 Setup of the Model

I provide a simple theoretical framework for understanding the effect of the indivisibility of the investment in housing on the consumption allocation of households. It extends the standard Merton (1971) model to incorporate the dual role of housing as a consumption good and an investment good.

Consider an agent who lives T years. At time t , he derives utility from a (non-separable) Cobb-Douglas consumption basket that is composed of two goods: a perishable non-housing good (C_t) and a durable housing good (K_t).¹⁶ The agent's utility is iso-elastic over that consumption basket and additively separable over time. The

¹⁵Sinai and Souleles (2005) use household-level data from the 1990-99 Current Population Surveys and estimate that the standard deviation of real rent growth between 1990 and 1998 at the MSA level is only 3.1% per year. I estimate from Piazzesi et al. (2007) that the implied annual volatility of rental price growth is about 85%. This calculation is based on an intra-temporal elasticity of substitution of 1.05, which is the value needed to match asset prices with reasonable coefficients of the relative risk aversion (5) and the discount factor (.99). Piazzesi et al. (2007) show that if we increase the risk aversion coefficient to 16 and the discount factor to 1.05, then the volatility of rental prices no longer needs to be that high.

¹⁶The agent consumes K_t units of housing. Here, a unit of housing is a one-dimensional representation of the size, the location, and the quality of the house. The housing market can be viewed as a continuum of houses of different units and whose prices are perfectly correlated. I also refer to housing consumption as consumption of housing services.

non-housing good is the numeraire. At time t , his lifetime expected utility takes the form

$$E \left[\int_t^T e^{-\delta s} U(C_s, K_s) ds \right],$$

where

$$U(C_s, K_s) = \frac{1}{1-\gamma} (C_s^\beta K_s^{1-\beta})^{1-\gamma}.$$

Here, γ is the coefficient of relative risk aversion over the entire consumption basket, and β measures the extent to which the agent values non-housing consumption relative to housing consumption. I assume that $\gamma > 1$, which also means that the two consumption goods are substitutes: living in a small home makes the non-housing good highly valuable.¹⁷

To consume one unit of housing services at time t , the agent pays a rental cost. I model this cost as a fraction ρ of $P_{H,t}$, which is the price of owning one unit of housing. The assumption that ρ is constant is constraining but necessary to solve the model in closed-form. I relax it in Section 1.4 when I simulate the model. If the agent also owns the home, then this rental cost can be viewed as the imputed rental cost, since the agent is essentially paying the rent to himself.

The agent can invest in a risky stock and a risk-free bond, which acts as a mortgage if the agent needs to borrow. The price $P_{B,t}$ of the risk-free asset is locally deterministic and the price of the stock $P_{S,t}$ follows a stochastic differential equation,

$$\begin{aligned} dP_{B,t} &= P_{B,t} \cdot r_t dt, \\ dP_{S,t} &= P_{S,t} \cdot (\mu_{S,t} dt + \sigma_{S,1,t} dZ_{1,t}). \end{aligned}$$

where $Z_{1,t}$ is a Wiener process, and r_t , $\mu_{S,t}$, and $\sigma_{S,1,t}$ are functions of both time and a state variable X_t .¹⁸ Hence, the investment opportunity set is stochastic and time-varying. I return to the dynamics and interpretation of X_t later on.

¹⁷As in Yogo (2006) the agent's utility is non-separable between housing services and the non-housing good when the inter-temporal elasticity of substitution (inter-ES) is different from the intra-temporal elasticity of substitution (intra-ES). The two goods are Edgeworth-Pareto substitutes (complements) when the intra-ES is greater (lower) than the inter-ES. In my setting, the intra-ES between the two goods is one because of the Cobb-Douglas specification of the consumption basket. The inter-ES over the entire consumption basket is $1/\gamma$, which means that when $\gamma > 1$ the two goods are substitutes.

¹⁸ X_t can be a vector of state variables. I argue later on that to solve the model in closed-form X_t must only depend on shocks $dZ_{1,t}$ and $dZ_{2,t}$.

As a durable good, housing also provides investment benefits. I assume that it does not depreciate over time, so that if the agent chooses not to trade, then $dK_t = 0$. Owning housing provides rental income as a dividend as well as risky capital gains,

$$dP_{H,t} = P_{H,t} (\mu_{H,t}dt + \sigma_{H,1,t}dZ_{1,t} + \sigma_{H,2,t}dZ_{2,t}),$$

where $\mu_{H,t}$ includes capital gains net of maintenance cost and depreciation. The moments $\mu_{H,t}$, $\sigma_{H,1,t}$, and $\sigma_{H,2,t}$ are functions of time and the state variable X_t . $Z_{2,t}$ is a Wiener process that is independent from $Z_{1,t}$. Since the stock is not exposed to this second source of uncertainty, $Z_{2,t}$, the agent cannot hedge away all the housing risk in the stock market.¹⁹

Despite the dual investment-consumption nature of housing, the structure of the housing market is such that the agent cannot separate these benefits. This is the indivisibility friction I explore in this model. Let H_t be the number of units of housing that the agent owns and $\phi_t = H_t/K_t$ his housing investment-consumption ratio. Figure 1.1 represents all of the agent's investment options in housing in the real world. The agent can rent his home ($\phi_t = 0$), own and occupy his home ($\phi_t = 1$), or even rent out a part of the home he owns ($\phi_t > 1$, e.g. landlords). However, there is no easily available hybrid solution that allows the agent to own a fraction of his home and rent the remaining fraction from the other co-owner ($0 < \phi_t < 1$).

This indivisibility friction can be viewed as a form of market incompleteness. Households do not have access to financial securities that provide a perfect hedge to their house price risk. Assets like real estate investment trusts (REITs) or housing futures only provide a partial hedge as they are often based on national or regional house price indices, whereas most of the house price risk is considered to be idiosyncratic. Recent studies have estimated the annual idiosyncratic house price volatility to range between 7% and 15%.²⁰

¹⁹This result is true unless either the correlation between shocks $dZ_{1,t}$ and $dZ_{2,t}$ is perfect or $\sigma_{H,2,t} = 0$. Note that in this model, the stock can be viewed as any portfolio of risky assets, including assets that are correlated with house prices like REITs.

²⁰See Case and Shiller (1987), Case and Shiller (1989), Goetzmann (1993), Flavin and Yamashita (2002), and De Jong, Driessen and Van Hemert (2008). Note again that in this model, if $\sigma_{H,2,t} = 0$ then the stock provides a perfect hedge against house price risk, so the indivisibility friction is no longer relevant.

Figure 1.1: Investment options in housing in the real world



I represent the household's options in terms of his housing investment-consumption ratio ϕ_t . The dotted line indicates the set of unavailable options.

In this essay, I abstract from the housing tenure decision and consider an agent who ideally would like to choose H_t so that he only owns a fraction of the home he lives in (e.g. the parameters of the model are such that the optimal ϕ_t is between 0 and 1), but who is not allowed to do so and instead *must* be a homeowner. He has to own what he consumes of housing ($\phi_t = 1$).²¹ Then, to identify the effect of indivisibility on the agent's decisions, I compare them to what he would do in a hypothetical world in which he could choose H_t freely, in the spirit of Cauley et al. (2007). Notation-wise, I refer to the agent as “the homeowner,” and I refer to the agent in the hypothetical case as “the unconstrained agent,” and denote his optimal decisions with a star superscript.

Indivisibility is the only friction in this model. Including other frictions would complicate the solution substantially and make it difficult to identify the effect of indivisibility. So I abstract from frictions such as transaction costs, borrowing con-

²¹Another way of phrasing it is that the agent prefers owning to renting for some reason outside of the model and that the cost of switching is high. This assumption is restrictive but it is necessary to obtain results in closed-form. In the empirical analysis in Chapter 2, I only focus on the consumption behavior of homeowners as a result. Note that this assumption makes policy implications regarding the value of the optimal ϕ difficult. In this model, an increase in the optimal ϕ makes the homeowner less constrained. However, once we take renting into account, an increase in the optimal ϕ incites some renters to switch and become homeowners, which does not necessarily make them less constrained.

straints, or the option to default. Except for indivisibility, markets are complete. I impose that the state variable X_t only depends on the shocks $[dZ_{1,t}, dZ_{2,t}]$,²²

$$dX_t = \mu_{X,t}dt + \sigma_{X,1,t}dZ_{1,t} + \sigma_{X,2,t}dZ_{2,t}$$

where $\mu_{X,t}$, $\sigma_{X,1,t}$, and $\sigma_{X,2,t}$ depend on time and X_t . So far, I have kept X_t as general and abstract as possible in order to show that this framework is flexible. In this general setting, all the results can be derived up to an ordinary differential equation. However, to solve this ODE, I need to impose further assumptions. I consider an affine setting in which the state variable X_t is the risk-free rate r_t . I assume it is mean-reverting and follows an Ornstein-Uhlenbeck process. That is, $\mu_X = \kappa(\mu_r - r_t)$, $\sigma_{X,t} = \sigma_r$, and $X_t = r_t$, where κ is the coefficient of mean-reversion. I assume that the excess expected returns and the volatility of returns of the home and the stock are constant.

Given this setup, the wealth process for both the homeowner and the unconstrained agent takes the following form,

$$\begin{aligned} dW_t = & \left[r_t W_t + \begin{pmatrix} \Theta_{S,t} & H_t P_{H,t} \end{pmatrix} \begin{pmatrix} \mu_{S,t} - r_t \\ \mu_{H,t} - r_t \end{pmatrix} - C_t - (K_t - H_t) \rho P_{H,t} \right] dt \\ & + \begin{pmatrix} \Theta_{S,t} & H_t P_{H,t} \end{pmatrix} \begin{pmatrix} \sigma_{S,1,t} & 0 \\ \sigma_{H,1,t} & \sigma_{H,2,t} \end{pmatrix} \begin{pmatrix} dZ_{1,t} \\ dZ_{2,t} \end{pmatrix} \end{aligned}$$

where $\Theta_{S,t}$ is the dollar amount spent on stocks, and $(K_t - H_t) \rho P_{H,t}$ is the net amount spent on the rent. Remember that for the homeowner, H_t is not a choice variable – it must be equal to K_t .

²²Wachter (2002) shows this assumption is necessary to solve the problem of an investor who has utility over intermediate consumption.

1.3 Derivation of the Model

1.3.1 Methodology

I use dynamic programming to provide analytical solutions for the optimal portfolio and consumption decisions of both the unconstrained agent and the homeowner. I extend the solution method of Damgaard et al. (2003) to an affine setting. All the details are provided in Appendix A. Both agents have a value function V (V^* for the unconstrained agent) that depends on four state variables: time t , Wealth W_t , the state variable X_t , and the house price $P_{H,t}$. The value functions V and V^* only differ by a constant,

$$V(t, W_t, X_t, P_{H,t}) \propto \frac{1}{1-\gamma} P_{H,t}^{-(1-\beta)(1-\gamma)} W_t^{1-\gamma} g(t, X_t)^\gamma,$$

$$g(t, X_t) = \epsilon \int_0^{T-t} e^{-\frac{\delta}{\gamma}\tau + \frac{1-\gamma}{\gamma}A_1(\tau) + \frac{1-\gamma}{\gamma}A_2(\tau)X_t} d\tau,$$

where ϵ is a constant, and the functions $A_1(\tau)$ and $A_2(\tau)$ solve a set of ordinary differential equations. The values of ϵ , $A_1(\tau)$ and $A_2(\tau)$ are given in Appendix A.

After a couple of variable changes, this problem can be reduced to the standard Merton (1971) problem. I begin by solving the unconstrained agent's problem, and then I turn to the homeowner and solve his decisions in terms of those of the unconstrained agent. In Appendix A, I show that it is possible to split his Hamilton-Jacobi-Bellman (HJB) equation into two components: the unconstrained agent's HJB equation, and a second component that only includes "distortion" terms.

1.3.2 The Unconstrained Agent's Problem

I begin by modeling the decisions of the unconstrained agent who can choose H_t freely. Let $\alpha_{H,t} = H_t P_{H,t}/W_t$, $\alpha_{S,t} = \Theta_{B,t}/W_t$, and $\alpha_{B,t} = \Theta_{B,t}/W_t$ be the agent's allocations of wealth in housing, the stock, and the risk-free asset respectively. It is optimal for the unconstrained agent to choose,

$$\alpha_{H,t}^* = \alpha_{H,t}^{MV} + \frac{g_X}{g} \frac{\sigma_{X,2,t}}{\sigma_{H,2,t}} + (1 - \beta) \left(1 - \frac{1}{\gamma} \right), \quad (1.1)$$

$$\alpha_{S,t}^* = \alpha_{S,t}^{MV} + \frac{g_X}{g} \frac{1}{\sigma_{S,t}} \left(\sigma_{X,1,t} - \frac{\sigma_{H,1,t}}{\sigma_{H,2,t}} \sigma_{X,2,t} \right), \quad (1.2)$$

$$\alpha_{B,t}^* = 1 - \alpha_{S,t}^* - \alpha_{H,t}^* \quad (1.3)$$

where $\alpha_{H,t}^{MV}$ and $\alpha_{S,t}^{MV}$ are the weights of the mean-variance efficient portfolio that includes housing and the stock.

Since the unconstrained agent can choose H_t and K_t separately, he is able to separate the investments benefits of housing from its consumption benefits. This is clear from equation (1.1). The optimal allocation of wealth invested in housing is independent of K_t . It only depends on three terms: the mean-variance efficient weight, a hedge term against changes in the value of the state variable X_t , and a hedge term against fluctuations in the price of the rental cost (Damgaard et al., 2003).

The unconstrained agent is also able to optimize his investments in the stock and the risk-free asset. The optimal allocation in the stock is also a function of the mean-variance efficient weight in the stock and a hedge term against changes in the value of the state variable X_t . It does not include a hedge term against fluctuations in rent. In this model, since the rental price is perfectly correlated with the house price, the unconstrained agent only uses his investment in housing for this hedge.

The separation of the investment and consumption benefits of housing is also reflected in the composition of the consumption basket of the unconstrained agent. The optimal ratio of non-housing consumption to housing consumption only depends on the relative preference for each good and the home's rental price,

$$\frac{C_t^*}{K_t^*} = \frac{\beta}{1 - \beta} \cdot \rho P_{H,t}. \quad (1.4)$$

1.3.3 The Homeowner's Problem

For the homeowner, H_t is not a choice variable. He must own what he consumes of housing, i.e. $H_t = K_t$. This friction affects all his trade-offs: the trade-off between housing consumption and non-housing consumption, the trade-off between investments in housing, the stock, and the risk-free asset, and the trade-off between overall investment and overall consumption. I begin by describing how indivisibility affects each trade-off and then solve the homeowner's decisions.

The Effect of Indivisibility of the Homeowner's Trade-offs

First and foremost, indivisibility affects the trade-off between housing consumption and non-housing consumption. The intra-temporal first-order condition is,

$$\frac{C_t}{K_t} = \frac{\beta}{1 - \beta} \cdot P_{H,t} \cdot \left[\rho + \gamma \sigma_{H,2,t}^2 (\alpha_{H,t} - \alpha_{H,t}^*) \right] \quad (1.5)$$

where $\alpha_{H,t}$ is the *implied* portfolio share in housing, i.e. $\alpha_{H,t} = \frac{K_t P_{H,t}}{W_t}$.

Like in the unconstrained case in (1.4), the ratio of non-housing consumption to housing consumption depends on the relative preference for each good and the home's rental cost. But here it also depends on how housing affects the agent's financial portfolio. The second term in the brackets on the RHS of (1.5) is part of the relative price of housing services. It is an opportunity cost.

I show in Appendix B that in a simple mean-variance setting, this opportunity cost corresponds to the extent to which an extra unit of housing affects the agent's marginal (indirect) utility over future wealth. Its sign depends on whether the implied portfolio share in housing $\alpha_{H,t}$ is greater or lower than the optimal share $\alpha_{H,t}^*$. If living in a nice home obliges the agent's portfolio to be under-weighted in housing ($\alpha_{H,t} < \alpha_{H,t}^*$), then the unit cost of housing services is lower than the unit rental cost, because an extra unit of housing actually improves the agent's portfolio by pushing $\alpha_{H,t}$ toward $\alpha_{H,t}^*$. On the other hand, if living in a nice home obliges the agent's portfolio to be over-weighted in housing ($\alpha_{H,t} > \alpha_{H,t}^*$), then the unit cost of housing services is higher than the unit rental cost, because an extra unit of housing makes

the homeowner's portfolio worse by pushing away $\alpha_{H,t}$ from $\alpha_{H,t}^*$. In this case, the agent would allocate a lower share of his consumption budget to housing than the unconstrained agent.

Indivisibility also affects the trade-off between investments in housing, the stock, and the risk-free bond. Equation (1.6) comes from the first-order condition with respect to the allocation of wealth in the stock, and (1.7) comes from the portfolio budget constraint,

$$\alpha_{S,t} = \alpha_{S,t}^* + \frac{\sigma_{H,1,t}}{\sigma_{S,1,t}} (\alpha_{H,t} - \alpha_{H,t}^*), \quad (1.6)$$

$$\alpha_{B,t} = 1 - \alpha_{S,t} - \alpha_{H,t}. \quad (1.7)$$

From (1.6), we see that the homeowner's allocation in the stock consists of two terms: the optimal allocation $\alpha_{S,t}^*$, and a hedge term against indivisibility. As with the first trade-off, the sign of this hedge term depends on whether the homeowner's implied portfolio share in housing is greater or lower than the optimal share $\alpha_{H,t}^*$. It also depends on the correlation between the returns to the prices of the stock and housing. If the correlation is high, then the agent can offset any suboptimal portfolio allocation in housing by adjusting the level of his stock-holdings. However, if the correlation is low, which has been documented by Flavin and Yamashita (2002), the agent cannot use the stock as a hedge, so he chooses to substitute housing with the risk-free asset. This result emphasizes the idea without other frictions like a borrowing constraint, an entry cost to the stock market, or a consumption commitment to housing, indivisibility does not affect the portfolio shares in stocks.²³

Finally, indivisibility affects the trade-off between overall investment and overall consumption, although only to a minor extent. Inserting all the first-order conditions back into the HJB equation, we get,

$$\begin{aligned} U(C_t, K_t) - U(C_t^*, K_t^*) - V_{W,t} \cdot \left[(C_t - C_t^*) + \rho P_{H,t} (K_t - K_t^*) \right] \\ = \frac{1}{2} \gamma \sigma_{H,2,t}^2 V_{WW,t} W_t^2 (\alpha_{H,t} - \alpha_{H,t}^*)^2, \end{aligned} \quad (1.8)$$

²³Studies such as Flavin and Yamashita (2002), Cocco (2005), Yao and Zhang (2005), and Chetty and Szeidl (2009) show that in the presence of these frictions, housing can crowd out the demand for stock holdings.

where $V_{W,t}$ and $V_{WW,t}$ are the first- and second-order derivatives of the homeowner's value function with respect to his wealth.

Equation (1.8) specifies how the homeowner plans to modify his consumption-investment allocation from his optimal one ($C_t = C_t^*$, $K_t = K_t^*$). Indivisibility affects this trade-off via the RHS of (1.8). Again, this term depends on whether the homeowner's implied portfolio share in housing is greater or lower than the optimal share $\alpha_{H,t}^*$. I show in appendix B that in a mean-variance setting this term corresponds to the loss in (indirect) utility over future wealth for any value of α_H . If the agent chooses K_t such that his portfolio is not well exposed to housing risk (e.g. $\alpha_H \neq \alpha_H^*$), then he is essentially poorer than the unconstrained agent, because he is not able to maintain an optimal financial portfolio.

The homeowner responds to this negative income effect by consuming a smaller fraction of his current wealth. Equation (1.8) can be rewritten as

$$\begin{aligned} \frac{1}{\beta(1-\gamma)} (C_t - C_t^*) - \left[(C_t - C_t^*) + \rho P_{H,t} (K_t - K_t^*) \right] \\ = \frac{1}{2} \gamma \sigma_{H,2,t}^2 W_t (\alpha_{H,t} - \alpha_{H,t}^*)^2. \end{aligned} \quad (1.9)$$

Since $\gamma > 1$, the negative income effect dominates the substitution effect toward consumption (Merton, 1971).²⁴ However, the net effect from this third trade-off is relatively small. The loss in indirect utility over future wealth is quadratic in the deviation of α_H from α_H^* , which means that it only matters when the deviation is large. This result is not new. Gains to diversification are large only if the agent's portfolio is quite undiversified. This means that unless indivisibility has an extremely strong effect on α_H , it will not have a strong direct effect on the homeowner's allocation of wealth to total consumption.

Result #1: The Composition of the Homeowner's Consumption Basket is Suboptimal

In this section, I merge all three trade-offs and solve for the homeowner's decisions. The first key result is that the composition of the homeowner's consumption basket is

²⁴The substitution effect toward consumption comes from the fact that indivisibility makes investment opportunities no longer as attractive as for the unconstrained agent.

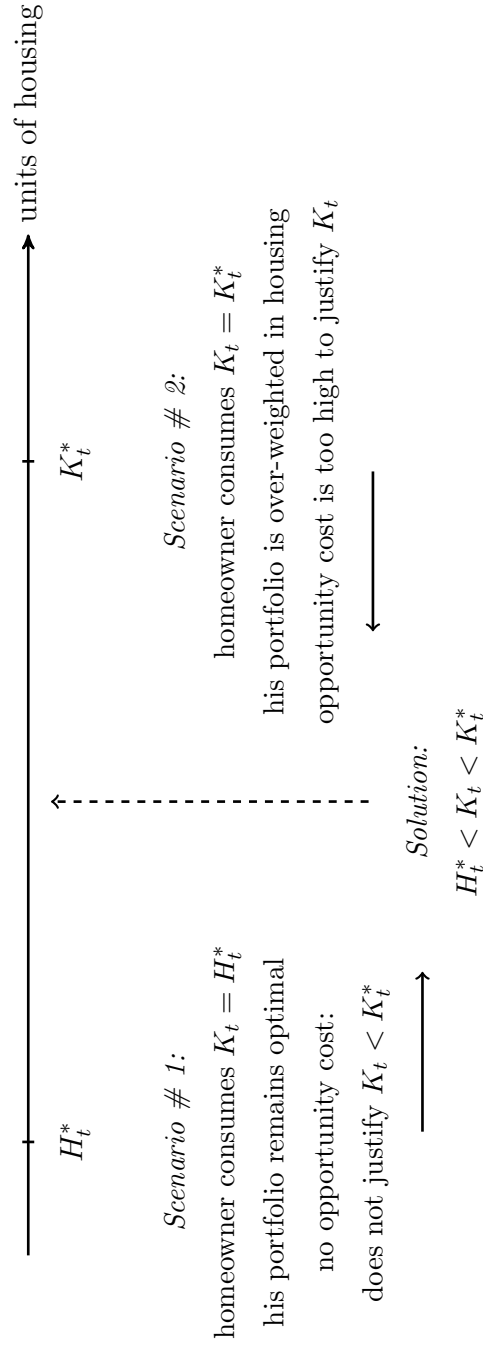
persistently suboptimal. Here, I provide intuition for this result and I then quantify it in Section 1.4. For a more formal proof, see Appendix A.

Figure 1.2 provides an easy way of visualizing why the composition of the homeowner's consumption basket is suboptimal by using a one-dimensional graph to plot the number of units consumed and invested in housing. I begin by plotting the optimal decisions of the unconstrained agent. Remember that in this model, I am considering the particular case in which the unconstrained agent chooses to be a partial homeowner, so $H_t^* < K_t^*$.

I show that for the homeowner, it is optimal to consume K_t such that $H_t^* < K_t < K_t^*$. It has to do primarily with the first trade-off between housing consumption and non-housing consumption. For the homeowner, the relative cost of consuming one unit of housing services is not just the rental cost. It also depends on how housing affects the homeowner's financial portfolio. The magnitude of this opportunity cost depends on the level of housing consumption K_t .

To see why the homeowner's choice of K_t is between H_t^* and K_t^* , consider what would happen if instead he chose either $K_t = H_t^*$ or $K_t = K_t^*$. If the homeowner were to maintain the same investment in housing as the unconstrained agent, i.e. $K_t = H_t^*$ (scenario #1 in Figure 1.2), then he would consume less housing than the unconstrained agent ($K_t < K_t^*$). This scenario is suboptimal because, as the homeowner's portfolio remains optimally invested in housing, the implied opportunity cost of housing services is zero and therefore too low to justify consuming $K_t = H_t^*$. Given this relative price of housing services, the homeowner would rather live in a bigger home and consume less of the perishable good.

Figure 1.2: Optimal consumption allocation for the homeowner



I plot the optimal level of housing consumption of the homeowner (K_t) in terms of the optimal levels of housing investment (H_t^*) and consumption (K_t^*) of the unconstrained agent.

On the other hand, if the homeowner were to consume the same amount of housing as the unconstrained agent, i.e. $K_t = K_t^*$ (scenario #2 in Figure 1.2), then his financial portfolio would be over-weighted in housing ($K_t > H_t^*$). This scenario is also suboptimal because the implied opportunity cost of housing services is positive and therefore too high to justify the decision of consuming as much housing as the unconstrained agent. Given the higher relative price of housing services, the homeowner would rather live in a smaller home and consume more of the perishable good.


The homeowner's optimal solution consists in choosing $H_t^* < K_t < K_t^*$. Essentially the homeowner is making a compromise between the investment and consumption benefits he gains from housing. The homeowner invests more in housing than the unconstrained agent ($H_t > H_t^*$), but not enough to stay in the same home ($K_t < K_t^*$). Since the homeowner's investment portfolio is over-exposed to housing risk, the opportunity cost of housing services is positive. Thus, the homeowner allocates a lower fraction of his consumption budget to housing consumption. I show in the next section that the exact "location" of K_t between H_t^* and K_t^* in Figure 1.2 depends on how far the unconstrained agent's optimal housing investment-consumption ratio ϕ_t^* is away from one. It also depends on the other parameters in equations (1.5) and (1.8) such as β and γ . I discuss how changes in some of these parameters affect this result in Section 1.4.3, once I calibrate the model.

Result #2: The Composition of the Homeowner's Consumption Basket is Excessively Volatile

The second key result is that the composition of the homeowner's consumption basket is also excessively volatile over time. This result comes from the fact that the opportunity cost of housing services varies with the available investment opportunities. Figure 1.3 provides intuition for this result by comparing the homeowner's optimal consumption in two environments: a low interest rate environment, and a high interest rate environment.

In the high interest rate environment, the allocation of wealth to current consumption (in housing and non-housing) is high for both the homeowner and the un-

Figure 1.3: Interest rates and consumption allocation



I plot the homeowner's allocation of wealth to housing consumption and non-housing consumption in high and low interest rate environments.

constrained agent. Both agents receive a positive income effect from the high interest rate. Since $\gamma > 1$, this income effect dominates the offsetting substitution effect that savings have become relatively more profitable. However, while the unconstrained agent can decrease his optimal housing investment-consumption ratio ϕ_t^* , the homeowner cannot. His over-investment in housing becomes more severe (i.e. $\alpha_{H,t} - \alpha_{H,t}^*$ increases), which means that in these “good times,” the opportunity cost of housing services is especially high. As a result, the homeowner decreases the share of his consumption budget to housing.

In the low interest rate environment, the allocation of wealth to current consumption is lower for both agents, because of the lower income effect. However, because the unconstrained agent is increasing his optimal housing investment-consumption ratio ϕ_t^* toward one, the homeowner is not as constrained by indivisibility. The opportunity cost of housing services decreases. As a result, his ratio of housing consumption to total consumption is higher than in the high interest rate environment.

1.4 Calibration of the Model

In Section 1.3, I have provided intuition on how indivisibility affects the consumption allocation of households. I now quantify these effects. I simulate the model in an economy where the interest rate is the state variable and follows an Ornstein-Uhlenbeck process, as in Vasicek (1977). I also relax the assumption that the rent is a constant fraction of the house price, which means that I have to solve the model numerically. I use the Crank-Nicholson finite difference method.

Note that as a word of caution, the purpose of this calibration exercise is not to make exact predictions on households' portfolio and consumption allocations, but rather to study the magnitude of the effect of indivisibility in a simple setting. Many relevant features such as stochastic labor income, borrowing constraints, or the option to default on one's mortgage will affect the homeowner's optimal housing investment-consumption ratio ϕ_t^* and hence the extent to which he is constrained by indivisibility. In Chapter 2, I test the effect of indivisibility on consumption allocation, controlling for these other effects.

1.4.1 Parameters

I report the parameters that I use in the baseline case in Table 1.1. All the asset returns are continuously-compounded and annualized. For stock returns, I use the value-weighted NYSE/AMEX/NASDAQ index from CRSP. The nominal return rate between 1930 and 2007 has a mean of 9.37% and a volatility of 19.24%. For interest rates, I use the 6-month commercial paper rate that is reported on Robert Shiller's website.²⁵ It is available until 2004, and to be consistent with the stock returns, I only select the data after 1930. The nominal rate has a mean of 4.5% and a volatility of 3.4%.

In order to deflate these returns, I cannot use the standard inflation index because it comes from an index of goods that includes housing services. In the model, the numeraire is the non-housing good. So I follow Piazzesi et al. (2007) and create a

²⁵After 1997, the 6-month Certificate of Deposit rate is used.

Table 1.1: Values of the baseline parameters

Description	Symbol	Parameter
<i>Asset properties</i>		
Interest rate (r)		
Mean	μ_r	.02
Mean-reversion	κ	.45
Exposure to shock Z_1	$\sigma_{X,1}$	0
Exposure to shock Z_2	$\sigma_{X,2}$	-.03
Stock (S)		
Premium	$\mu_S - \mu_r$.05
Exposure to shock Z_1	$\sigma_{S,1}$.16
Exposure to shock Z_2	$\sigma_{S,2}$	0
Home (H)		
Premium	$\mu_H - \mu_r$.02
Exposure to shock Z_1	$\sigma_{H,1}$	0
Exposure to shock Z_2	$\sigma_{H,2}$.1
Rental cost	ρ_0	.043
	ρ_1	-.148
	$\underline{\rho}, \bar{\rho}$.02, .06
<i>Utility parameters</i>		
Risk aversion	γ	6
Time discount factor	δ	.01
Preference for non-housing	β	.75
Number of years	T	30

I use the following parameters to simulate the model in Section 1.4. The results of the baseline case are reported in Tables 1.2 and 1.3.

new price index from the NIPA consumption tables, which excludes the following categories: housing services, durable goods, and cloth and shoes.²⁶ It behaves fairly similarly to the standard price index. The average inflation between 1930 and 2007 is about 3.5%. The real annual returns of the stock and the risk-free asset have means of 5% and 2% and volatilities of 16% and 3% respectively.

For housing returns, I base my estimates on Flavin and Yamashita (2002) who compute the moments of housing returns directly from household-level data from the Panel Study of Income Dynamics (PSID). The real annual return for homeowners from 1968 to 1992 has a mean of 6.6% and a volatility of 14.24%, and the correlation between stock and housing returns is zero. I choose a value for the average excess housing return of 2%, which is slightly smaller than their estimate. The reason is that in the model, I define $\mu_{H,t}$ as the expected return on housing excluding the rental “dividend” income. Since Flavin and Yamashita (2002) compute housing returns as capital gains plus risk-free interest as a proxy for the rent, I subtract the real interest rate twice to get the excess return. I also round down the value to 2% since the estimates in Flavin and Yamashita (2002) are higher than other empirical estimates at the household-level. As for the volatility, I choose a value of 10%, which is also slightly lower than their estimate.²⁷ This value remains above their estimates of volatility at the MSA-level and it accounts for potential noise in the measurements of housing returns. In the PSID, housing returns are computed from self-reported guesses of the market values of households’ homes.

The initial values of wealth W_0 , the house price $P_{H,0}$, and the stock price $P_{S,0}$ are not relevant for the results I am interested in. All the variables of interest, such as the opportunity cost of housing services as a fraction of the house price, the optimal housing investment-consumption ratio, the portfolio shares of wealth, and the percentage changes in wealth and consumption over time are independent of these initial values. While the irrelevance of the initial house price $P_{H,0}$ may appear puzzling at first, I show that what matters for the optimal housing investment-consumption ratio ϕ_t^* is ρ , the ratio of the rental cost to the price of housing.

²⁶Details are available upon request.

²⁷Yao and Zhang (2005) and Corradin (2009) also assume a volatility of 10%.

As I mentioned earlier, the assumption that ρ is constant is limiting because it would generate an unreasonably high volatility of rental prices. Sinai and Souleles (2005) use household-level data from the 1990-99 Current Population Surveys and estimate that the standard deviation of real rent growth between 1990 and 1998 at the MSA level is only 3.1% per year, which is less than one-third of the volatility of housing prices I am using.²⁸ In order to capture this low volatility while remaining within the structure of the model, I model ρ_t so that its variations partially offsets movements in the house price $P_{H,t}$. More specifically, I assume (i) that ρ_t varies with the state variable r_t and (ii) that the shocks to r_t come from the second source of uncertainty $Z_{2,t}$, which also affect $P_{H,t}$,

$$\rho_t = \rho_0 + \rho_1 \cdot r_t, \quad \rho_t \in [\underline{\rho}, \bar{\rho}],$$

where I choose the constants $\underline{\rho}$, $\bar{\rho}$, ρ_0 and ρ_1 to match an average ρ of 4% and an annual volatility of growth rates in rental prices of 3.1%.

The drawback of this modeling approach is that all the hedging against interest rate risk is achieved via the investment in housing. However, even though this hedging component may not correspond exactly to the true hedging role of housing, it provides in this simple setting a channel to generate a low investment demand for housing for the unconstrained agent. Therefore, it allows me to focus on the effect of indivisibility on the homeowner's consumption allocation, which is the main objective of this calibration exercise.

The parameter β measures how much the agent values the non-housing good relative to housing. I estimate it from the households' intra-temporal Euler equation. Given the Cobb-Douglas nature of their consumption baskets, the ratio of expenditures should be equal to $\beta/(1-\beta)$. It is difficult to back out this ratio for homeowners since their expenditures on housing services depend on the unobservable opportunity cost component. But for renters, the amount they spend on housing services is just the rent. According to the 2006 Consumer Expenditures Survey, rent constituted

²⁸This low volatility of rental prices is not due to rents being re-set only periodically. In Sinai and Souleles (2005), the annual volatility is computed using the within-MSA annual differences between the actual log rent and the calculated average growth rate over the prior nine years.

Table 1.2: Portfolio allocation under the baseline parameters

Variable	Statistic	UA _{mv}	UA _r	UA _{rent}	UA	HH
α_H	μ	.333	-.427	.208	.115	.334
α_S	μ	.325	0	0	.325	.325
α_B	μ				.56	.341
%ΔW	μ	-.024	-.01			
	σ	.055	.063			

I simulate the model 1,000 times and report the average values. UA refers to Unconstrained Agent and HH to Homeowner. α_H , α_B , and α_S correspond to the shares of wealth allocated to housing, the risk-free bond, and the stock. I show in equations (1.1) and (1.2) that for the unconstrained agent, α_H and α_B are both combinations of three portfolios: the tangency portfolio (mv), a hedging portfolio against interest rate risk (r), and a hedging portfolio against rent risk (rent). The columns UA_{mv}, UA_r, and UA_{rent} correspond to the weights of α_H and α_B in these three portfolios. %ΔW corresponds to the growth rate of wealth. μ and σ correspond to annual means and volatilities.

24.7% of the annual expenditures of renters, which implies a β of .75.²⁹

Finally, I use a risk aversion coefficient γ of 6 and a time discount factor δ of .01. I choose a time horizon T of 30 years to model the decisions of the average homeowner, who is about 50 years old in the PSID.³⁰ I simulate the model over a period of two years 1000 times and report the average values.

1.4.2 Results

Tables 1.2 and 1.3 report the simulation results for the baseline parameters. In this setting, the average optimal housing investment-consumption ratio ϕ_t^* for the unconstrained agent is 27%, which means that partial home ownership is optimal and the case I analyze in the model applies.

The effect of indivisibility on portfolio allocation is economically significant. The homeowner allocates on average 33% of his wealth to housing, which is about 20% higher than in the frictionless case. His portfolio share in the risk-free asset decreases

²⁹This value is close to the values that are used in related papers. Cocco (2005), Yao and Zhang (2005), and Corradin (2009), use values of β of .7, .8, and .9 respectively.

³⁰Author's computations. See Table 2.1 for details.

by the same amount, and his portfolio share of stock holdings does not change. As I mentioned earlier, the zero effect on the share of stock holdings comes from the zero correlation between stock and housing returns and the absence of frictions like borrowing constraints.

This portfolio distortion affects the homeowner's wealth dynamics. In comparison to the unconstrained agent, the homeowner's portfolio is over-weighted in housing and under-weighted in the risk-free asset. Since housing is risky and has a positive risk premium, the homeowner's wealth grows at a faster and more volatile rate than the wealth of the unconstrained agent. In the model, wealth is decreasing over time because the agents do not receive endowments past $t = 0$, but it is decreasing at a lower and more volatile rate for the homeowner.

This portfolio distortion is also reflected in the consumption allocation of the homeowner. Indivisibility increases the average relative price of housing services by about 33%. In addition to the rental cost of housing services, which is on average 4% of the house price, the homeowner is paying an opportunity cost, which corresponds to the second term in the brackets of the RHS of (1.5). It amounts to 1.3% of the house price on average.

This opportunity cost varies substantially over time. The annual volatility of its growth rate is 26%. These variations affect the dynamics of the homeowner's consumption basket. I report the unconditional volatilities of the growth rates of (i) non-housing consumption (C), housing consumption (K), and the ratio of non-housing consumption to housing consumption (C/K). I also report the expected values of these growth rates but the results are not very meaningful. By construction, the relative price of housing services is increasing over time as ρ is on average proportional to the house price (which has a positive drift). As a result, non-housing and housing consumption are increasing and decreasing over time respectively.

First, the ratio of non-housing consumption to housing consumption is 41% more volatile than in the frictionless case. This additional volatility comes from the variations in the opportunity cost, which are tied to the state of the economy. The opportunity cost of housing services is particularly high in the "good" times when the interest rate is high, and hence when the optimal housing investment-consumption

Table 1.3: Consumption allocation under the baseline parameters

Variable	Statistic	UA	HH	% Δ
opp. cost	μ	0	.013	
	σ	0	.268	
	$corr(., r)$	0	.984	
ϕ	μ	.27	1	
	σ	.055	0	
	$corr(., r)$	-.993	0	
% $\Delta C/K$	μ	.027	.045	
	σ	.059	.083	1.41
	$corr(., r)$.153	.23	1.503
% ΔC	μ	.002	.018	
	σ	.058	.0739	1.277
	$corr(., r)$.109	.124	1.138
% ΔK	μ	-.022	-.023	
	σ	.067	.058	.867
	$corr(., r)$.06	.013	.217

I simulate the model 1,000 times and report the average values. UA refers to Unconstrained Agent and HH to Homeowner. The column % Δ corresponds to the percentage change from UA to HH. “opp. cost” is the opportunity cost of housing services as a fraction of the house price. It corresponds to the second term in the brackets on the RHS of (1.5). ϕ is the housing investment-consumption ratio. % $\Delta C/K$, % ΔC , and % ΔK correspond to growth rates in the ratio of non-housing consumption to housing consumption, non-housing consumption, and housing consumption. μ and σ correspond to annual means and volatilities. $corr(., r)$ corresponds to the instantaneous volatility of the given variable with the interest rate. In the case of the opportunity cost variable, σ and $corr(., r)$ are computed from the growth rate of the opportunity cost of housing services (not scaled by the house price).

ratio ϕ_t^* is low. The instantaneous correlations between these two variables and the interest rate are almost perfect (.984 and $-.993$ respectively). As a result, the correlation between the ratio of non-housing consumption to housing consumption and the interest rate is 50% greater for the homeowner than for the unconstrained agent.

The second result is that indivisibility amplifies the volatility of non-housing consumption by 27%. The reason for this increase is twofold. For one, there is the composition effect that I just described. Since the relative price of housing services is high in times when the homeowner allocates a greater share of his wealth to current overall consumption, variations in level of non-housing consumption are amplified. The correlation between non-housing consumption growth and the interest rate is 14% higher for the homeowner than for the unconstrained agent. In addition, there is also direct wealth effect that I mentioned above as well. Since the homeowner's wealth is more volatile than in frictionless case, it also leads to an increase in the volatility of their non-housing consumption.

The third result is that indivisibility dampens the volatility of housing consumption by about 13%. Even though this effect is not as large as the effect on non-housing consumption volatility, it does not mean that the effect on housing consumption is less important. On the contrary, the fact that housing consumption volatility is lower for the homeowner indicates that the composition effect of indivisibility is quite strong. Since the relative price of housing services is high in times when the homeowner allocates a greater share of his wealth to current overall consumption, the variations in the level of housing consumption are dampened. The correlation between housing consumption volatility and the interest rate is 80% *lower* for the homeowner than for the unconstrained agent. This dampening effect offsets the amplifying effect from the homeowner's more volatile wealth. Hence, if housing consumption is less volatile for the homeowner, it is precisely because the composition effect of indivisibility is large enough to outweigh the increase in volatility from the homeowner's more volatile wealth.

1.4.3 Comparative Statics

I now explore how heterogeneity in the characteristics of homeowners affect the extent to which they are constrained by indivisibility. In Table 1.4, I present some comparative statics, by varying the coefficient of mean reversion of the interest rate κ as well as three characteristics of the homeowner: his horizon T ³¹, his level of risk aversion γ , and his relative preference for the perishable good β . I find that different values of these parameters have a large effect on (i) the magnitude of the opportunity cost of housing services and hence on (ii) the volatility of the homeowner's consumption of both goods. In the empirical analysis that I conduct in Chapter 2, I build on this result to test the effect of indivisibility on consumption volatility for various homeowners.

First, looking at different values of κ highlights the importance of the persistence of interest rates on the effect of indivisibility on consumption volatility. In the calibration, variations in the optimal housing investment-consumption ratio ϕ_t^* come from changes in the interest rate. The more persistent these changes, the greater the effect of indivisibility. For a low speed of mean-reversion ($\kappa = .2$), indivisibility amplifies variations in the composition of the homeowner's consumption basket by 76%, which is almost twice as important as in the baseline case.

In terms of household characteristics, homeowners with a short horizon T are relatively more constrained. A decrease in T from 50 years to 15 years leads to a increase in the opportunity cost from 1% to 2.1% of the house price. Consequently, variations in both the ratio of non-housing consumption to housing consumption and in the level of non-housing consumption are further amplified, and variations in the level of housing consumption are further dampened. For example, in the case of the ratio of the two goods, the amplification factor jumps from 1.35 to 1.52. These results come from the fact that the optimal investment demand for housing $\alpha_{H,t}^*$ increases with the horizon. In this setting, it is the hedging demand against interest rate risk that increases. Sinai and Souleles (2005) show that the hedging demand against rent risk can also increase. As a result, as T increases from 15 years to 50 years, the

³¹Varying the subjective discount factor δ provides similar results.

optimal housing investment-consumption ratio ϕ_t^* increases from 19% to 33%.

Homeowners who are highly risk averse are also relatively more constrained by indivisibility. In Table 1.4, an increase in γ from 4 to 8 leads to an increase in the opportunity cost from .02% to 2.3% of the house price. As in the previous case, variations in both the ratio of non-housing consumption to housing consumption and in the level of non-housing consumption are further amplified, and variations in the level of housing consumption are further dampened.

The positive effect of γ on the homeowner's opportunity cost comes from three channels. The primary channel is that a high value of γ leads to a low investment demand for housing. As γ increases from 4 to 8, the main decrease in the portfolio share in housing of the unconstrained agent comes from the weight in the tangency portfolio, which decreases from 50% to 25%. The hedging component against rent risk increases but only by 3%.

Second, a higher degree of risk aversion means that for any level of housing consumption, the portfolio distortion that comes from indivisibility is perceived as more severe. We see from (1.5) that a high value of γ scales up the opportunity cost component. However, this increase is small because at the same time it also makes the homeowner more reluctant to over-invest in housing (i.e. it pushes $\alpha_{H,t}$ toward $\alpha_{H,t}^*$).

Third, by the nature of the homeowner's utility function, a high value of γ also means that the two consumption goods are more substitutable in the sense that living in a small home make the perishable good more valuable. The substitution aspect of γ enters (1.9): if $\gamma > 1$ then a decrease in housing consumption (K_t) can be substituted by an increase in non-housing consumption (C_t). A higher degree of substitutability increases the opportunity cost of housing services because it makes the homeowner more willing to over-invest in housing (i.e. it pushes $\alpha_{H,t}$ away from $\alpha_{H,t}^*$). If the perishable good becomes more valuable to the homeowner for a given loss in housing consumption, then he becomes more willing to pay the high opportunity cost.

Table 1.4: Comparative statics

Param.	Value	opp. cost	ϕ^*	α_H				$\sigma(\% \Delta C / K)$				$\sigma(\% \Delta C)$				$\sigma(\% \Delta K)$			
				U _{A_{mv}}	U _{A_r}	U _{A_{rent}}	U _A	HH	U _A	HH	% Δ	U _A	HH	% Δ	U _A	HH	% Δ	U _A	HH
κ	.2	.0316	-.607	.333	-.816	.208	.208	-.275	.252	.051	.091	1.758	.062	.097	1.571	.075	.061	.812	
	.5	.013	.27	.333	-.427	.208	.208	.115	.334	.059	.083	1.41	.058	.074	1.277	.067	.058	.867	
	.8	.005	.7	.333	-.254	.208	.208	.287	.373	.066	.081	1.21	.057	.066	1.15	.066	.059	.9	
T	15	.021	.191	.333	-.408	.208	.208	.113	.482	.059	.09	1.52	.057	.081	1.43	.067	.056	.83	
	30	.013	.27	.333	-.427	.208	.208	.115	.334	.059	.083	1.41	.058	.074	1.277	.067	.058	.867	
	50	.01	.327	.333	-.433	.208	.208	.109	.275	.059	.079	1.35	.058	.071	1.21	.067	.06	.88	
γ	4	.002	.807	.5	-.345	.187	.187	.342	.405	.059	.076	1.295	.082	.089	1.09	.084	.076	.9	
	6	.013	.27	.333	-.427	.208	.208	.115	.334	.059	.083	1.41	.058	.074	1.277	.067	.058	.867	
	8	.023	-.003	.25	-.47	.219	.219	-.001	.283	.059	.087	1.48	.046	.066	1.431	.061	.052	.848	
β	.5	.0364	.2	.333	-.449	.417	.417	.301	.908	.059	.1	1.722	.065	.116	1.783	.067	.052	.782	
	.75	.013	.27	.333	-.427	.208	.208	.115	.334	.059	.083	1.41	.058	.074	1.277	.067	.058	.867	
	.9	.007	.027	.333	-.413	.083	.083	.004	.119	.059	.07	1.19	.053	.059	1.117	.068	.063	.937	

I simulate the model 1,000 times and report the average values. UA refers to Unconstrained Agent and HH to Homeowner. Each column % Δ corresponds to the percentage change from UA to HH. “opp. cost” is the opportunity cost of housing services as a fraction of the house price P_H . It corresponds to the second term in the brackets on the RHS of (1.5). ϕ^* is the housing investment-consumption ratio for the unconstrained agent. α_H corresponds to the share of wealth allocated to housing. For the unconstrained agent, it is a combination of three portfolios: the tangency portfolio (mv), a hedging portfolio against interest rate risk (r), and a hedging portfolio against rent risk (rent). The columns U_{A_{mv}}, U_{A_r}, and U_{A_{rent}} correspond to the weights of α_H in these three portfolios. $\sigma(\% \Delta C / K)$, $\sigma(\% \Delta C)$, and $\sigma(\% \Delta K)$ correspond to the volatilities of the growth rates in the ratio of non-housing consumption to housing consumption, non-housing consumption, and housing consumption respectively. κ is the speed of mean-reversion of the interest rate. T is the time horizon. γ is the coefficient of relative risk aversion. β is the relative preference for the non-housing good.

Finally, in terms of β , homeowners who value housing consumption highly (low β) are also relatively more constrained by indivisibility. A decrease in β from .9 to .5 is linked with an increase in the opportunity cost from .07% to 3.6% of the house price. As in the two previous cases, variations in both the ratio of non-housing consumption to housing consumption and in the level of non-housing consumption are further amplified, and variations in the level of housing consumption are further dampened. This result comes from the fact that if the homeowner cares more about housing consumption (low β), then he becomes more willing to pay a high opportunity cost to maintain the same level of housing consumption (i.e. $\alpha_{H,t} - \alpha_{H,t}^*$ increases).³²

1.5 Conclusion

In this essay I have proposed a tractable theoretical framework for understanding the effect of indivisibility on the consumption allocation of homeowners. I developed a Merton (1971) portfolio choice model that incorporates the dual consumption-investment nature of housing and studied the consumption and portfolio decisions of a homeowner who would ideally like to own just a fraction of his home. The main result is that indivisibility makes the composition of the homeowner's consumption basket persistently suboptimal as well as excessively volatile. The reason is that for the homeowner, the relative price of housing services includes an opportunity cost of having an unbalanced financial portfolio. This cost varies substantially over time, and it is especially high in the "good" times, when the homeowner allocates a greater fraction of his wealth to current consumption. As a result, this cost dampens variations in the level of his housing consumption, and it amplifies variations in both the level of his non-housing consumption and the composition of his consumption basket.

Although the analysis in this paper is entirely partial equilibrium, it suggests that the level of variations in the composition of homeowners' consumption baskets in the U.S. economy (i.e. "composition risk") may be greater than previously estimated.

³²This effect dominates two offsetting channels. First, as β decreases the investment demand for housing increases because of the higher hedging demand against rent risk. Second, for a given degree of substitution between the two goods, a low value of β means that the homeowner cares less about non-housing consumption. In (1.9) the sensitivity of non-housing consumption decreases.

These variations are likely to matter for asset prices since homeowners represent the vast majority of stock holders. Future research should test whether the dynamics of the composition of the consumption baskets of these homeowners can improve the current “Housing” C-CAPM model.

The results of this paper also raise the question of whether we should promote equity financing in the housing market. There may be significant welfare gains from the introduction of programs that would relax the indivisibility constraint. For example, Caplin, Chan, Freeman and Tracy (1997) motivate the creation of equity sharing programs that would allow owner-occupied homes to be partly owned by “limited partners.” I have argued in this paper that reduced exposure to housing risk may lead not only to a more optimal portfolio allocation but also to a consumption schedule that is more balanced and stable over time. It would be an interesting exercise to quantify how much more stable the U.S. economy would become in the presence of a liquid market for the fractionally owned real estate shares. Ortalo-Magné and Prat (2009) provide a first step to analyze the effects of indivisibility in a general equilibrium economy.

Nonetheless, despite these predictions for welfare, I have not addressed in this paper why equity sharing programs have not been more popular to this day. There may be other frictions like moral hazard between the various co-owners of the house, liquidity issues for the “limited partners,” or tax issues that make these programs currently unfeasible. Here again, further research is needed to assess whether we can design contracts that relax indivisibility in a way that does not lead to these other frictions.

Chapter 2

Housing and the Consumption Allocation of Households: An Empirical Analysis

In this essay, I test empirically the theory that I developed in Chapter 1 on the impact of the housing indivisibility constraint on the consumption dynamics of households. Using household-level data from the Panel Study of Income Dynamics (PSID), I test the hypothesis that homeowners who face a high opportunity cost choose *ceteris paribus* a low housing consumption volatility. I also develop a method to identify these constrained homeowners by comparing their characteristics to those of a subset of unconstrained homeowners: the landlords. The results are consistent with the predictions of the model. First, the characteristics of homeowners that determine how constrained they are in the model are strong predictors of those homeowners who choose to be landlords in the data. For example, homeowners with a low level of risk aversion, little value for housing consumption, and a long horizon are relatively more likely to be landlords. Second, I find evidence that the more constrained homeowners adjust their level of housing consumption much less over time.

2.1 Introduction

In Chapter 1, I argued that the dual consumption-investment nature of housing has strong implications for the way households allocate their consumption between housing and other goods over time. In particular, homeowners who would ideally like to own just a fraction of their home pay an opportunity cost to live in their home, in addition to the imputed rental cost. Since this opportunity cost is high in “good” times when homeowners allocate a high fraction of their wealth to current consumption, it amplifies variations in both their level of non-housing consumption and the ratio of non-housing consumption to housing consumption, and it dampens variations in their level of housing consumption.

In this essay, I use panel survey data from the Panel Study of Income Dynamics (PSID)¹ and exploit heterogeneity in the characteristics of homeowners to test empirically whether those who face a high opportunity cost choose *ceteris paribus* a low housing consumption volatility. I focus on housing consumption because, while the PSID also provides data on non-housing consumption, it is very limited. Only noisy estimates of food consumption are available, and, as Blundell, Pistaferri and Preston (2006) argue, the dynamics of food consumption differ in important ways from the dynamics of other non-durable consumption.²

This empirical analysis is challenging because the extent to which indivisibility affects most homeowners is unobservable. In the PSID, households report whether they rent or own their primary home, but they do not report whether they would rather be partial owners of their home. However, there is one subset of homeowners who are not constrained: the landlords. Unlike other homeowners, landlords have chosen to own more housing than what they consume. My identification strategy consists of exploiting information on the characteristics of landlords. I begin by

¹Some of the data used in this analysis are derived from Sensitive Data Files of the Panel Study of Income Dynamics, obtained under special contractual arrangements designed to protect the anonymity of respondents. These data are not available from the author. Persons interested in obtaining PSID Sensitive Data Files should contact through the internet at psidhelp@isr.umich.edu.

²Previous studies such as Zeldes (1989) and Mankiw and Zeldes (1991) have used food housing consumption from the PSID, by working with groups of households. It is more difficult to do so in my analysis. My identification strategy depends on multiple control variables, which have missing values, so I am left with a smaller sample size.

identifying which characteristics have led some homeowners to be landlords via a probit regression, and I then use the results to infer the extent to which the non-landlord homeowners are constrained based on how similar their characteristics are to those of landlords.

The empirical results are consistent with the predictions of the model. First, the characteristics that determine how constrained homeowners are in the model also predict those homeowners who have chosen to be landlords in the data. For example, homeowners with a low level of risk aversion, a low preference for housing consumption (relative to non-housing consumption), and a long horizon (e.g. low mobility) are relatively more likely to be landlords. Secondly, the decrease in housing consumption volatility that comes from indivisibility is quite significant. A homeowner who goes from being most constrained to least constrained increases *ceteris paribus* the extent to which he adjusts his level of housing consumption by 172% when he moves to a new house. He also increases the amount spent on making additions to his house by 47%. Furthermore, he remains as likely to move again or make new additions to his home, which rules out a potential transaction costs story.

The essay proceeds as follows. I provide a description of the datasets that I use in Section 2.2. I identify which homeowners are constrained by indivisibility in Section 2.3. Finally, I test the predictions of the model on the effect of indivisibility on the volatility of housing consumption in Section 2.4. Further details on how I constructed the data from the PSID are provided in Appendix C.

2.2 Description of the Datasets

My analysis relies on data from the Panel Study of Income Dynamics (PSID).³ It is a national survey of a representative sample of households living in the U.S. The survey has been conducted on an annual basis from 1968 to 1997 and then on a bi-annual basis, and it has tracked the same households over time. In 1968 there were 4800 households, and since then, the sample size has grown to more than 7000

³The PSID is conducted at the Survey Research Center, Institute for Social Research, University of Michigan.

households. During each interview, households were asked to report information such as their demographics, consumption behavior, earned income, and real estate holdings (house value, mortgage, tenure choice, etc) Furthermore, beginning in 1984, households were also asked to answer questions regarding their wealth, such as the net value invested in stocks or bonds. The wealth surveys are only conducted every 5 years from 1984 to 1999 and every two years after 1999. These surveys are key to my study, so I focus on the 1984-2005 period.

For this analysis, the PSID presents multiple advantages over other major surveys such as the Survey of Consumer Finances (SCF) or the Consumer of Expenditures Survey (CEX). For one, since it tracks households over time, it allows me to look at households' consumption over time instead of making inferences. Secondly, whereas the CEX and the SCF data sets gather information on either consumption or portfolio allocations, the PSID does both. While the focus of this study is on consumption, I find that wealth is one of the characteristics of homeowners that play a key role in determining the magnitude of their opportunity cost of housing services. Finally, for each interview the PSID provides detailed information on where all the survey participants live.⁴ This is also important for the opportunity cost, given the high level of heterogeneity in the housing market across various areas. The location data is available at the census tract level, so I use the U.S. Census Bureau to get neighborhood characteristics for each household. I access the census data via the Neighborhood Change Database (NCDB), which provides instant access to most of the census variables for years 1980, 1990, and 2000 while adjusting for changes in the boundaries of census tracts over the years.⁵ I use linear interpolation to infer values for all the other years.⁶

⁴This data is available only by a special request and a confidential data use contract.

⁵I thank Jesse Silva and Harrison Dekker from the UC Berkeley Library for having helped me to find and access the NCDB database.

⁶The census is conducted only at the beginning of each decade. For years 01, 03, and 05, I interpolate using census years 90 and 00. Note that not all the US was tracted in the 1980s, so there are rural areas for which data is unavailable. In these cases, I interpolate from the 1990 and 2000 census values.

Table 2.1: General summary statistics for households in 2005

Variable	All Households				Homeowners			
	Mean	Std Dev	Min	Max	Landlords		Non-landlords	
					Mean	Std Dev	Mean	Std Dev
<i>Demographics</i>								
age	50.73	15.06	21	94	53.18	14.55	51.97	14.2
family size	2.58	1.3	1	9	2.4	1.2	2.75	1.28
married	.69	.46	0	1	.73	.44	.79	.41
separated	.22	.41	0	1	.16	.37	.17	.37
single	.09	.29	0	1	.11	.31	.05	.21
high school diploma	.31	.46	0	1	.21	.41	.3	.46
college degree	.19	.39	0	1	.23	.42	.3	.31
risk aversion	8.8	6.14	.2	14.50	7.57	6.21	8.46	6.12
<i>Income, Wealth, and Housing</i>								
family income	87.15	154.59	1.2	5,500	106.58	75.08	97.54	177.99
net worth	378.84	867.12	0	20,715	1,003.91	2,055.92	405.08	733.19
homeowner	.79	.41	0	1	1	0	1	0
landlord	.07	.26	0	1	1	0	0	0
house value	247.40	199.27	0	1000.50	324.23	231.17	250.21	191.88
LTV	.36	.29	0	1.08	.34	.29	.36	.29
have a second mortgage	.11	.31	0	1	.12	.33	.14	.35
N	2,051	2,051	2,051	2,051	149	149	1,468	1,468

I report general summary statistics from three groups of the population from the 2005 wave of the PSID: the entire population, homeowners who also received rental income on real estate properties during the year prior to the interview (landlords), and other homeowners. I exclude households with negative net worth and only retain those who come from the Survey Research Center core sample. Reported variables are described in Section C.8.

Table 2.2: Portfolio characteristics for households in 2005

Variable	All Households		Homeowners			
			Landlords		Non-landlords	
	Part. rate	Value	Part. rate	Value	Part. rate	Value
<i>prop. of household's wealth in...</i>						
stocks	.31	.21	.46	.18	.35	.2
IRA-stocks	.45	.23	.66	.16	.51	.21
bonds	.22	.16	.26	.1	.24	.11
cash	.89	.14	.93	.05	.92	.09
primary house	.79	1.15	1	.81	1	1.18
other real estate	.2	.31	.7	.33	.19	.25
autos	.92	.25	.97	.06	.94	.12
business/farms	.12	.31	.28	.29	.13	.24
mortgage	.58	.79	.72	.54	.73	.82
other debts	.49	.22	.54	.1	.51	.51
<i>prop. of total PSID wealth in...</i>						
all stocks				.17		.77
all bonds				.1		.82
all real estate				.19		.63
business/farms				.3		.8
N	2,051	2,051	149	149	1,468	1,468

I report summary statistics on portfolio allocations from three groups of the population from the 2005 wave of the PSID: the entire population, homeowners who also received rental income on real estate properties during the year prior to the interview, and other homeowners. I exclude households with negative net worth and only retain those who come from the Survey Research Center core sample. Reported variables are described in Section C.8.

The observation unit is a “household-year.” Given the panel nature of the dataset, defining a household can be tricky because of possible changes in its composition over time. I follow the PSID conventions and use the household head member to represent his or her household. To abstract from major changes in the life of a household, I impose that any change in the civil status of the head couple results in the creation of a new household. Details of the criteria that I use are described in Appendix C.

I compute the households’ net worth for the years in which the wealth survey was conducted. The PSID asks the participants to report the net market value invested in stocks,⁷ bonds/insurance, businesses/farms, other real estate, and motor vehicles, as well as the amount of cash being held and the amount of debt that is not backed up by real estate or motor vehicles. In the regular questionnaires, households are also asked to provide an estimate of the value of their home and to report how much is left to repay on their mortgage. I use all these variables to compute net worth.

The second set of neighborhood characteristics consists of statistics on individual house prices, which I aggregate by Metropolitan Statistical Area (MSA). I compute the expected value and the volatility of returns of individual house prices, for which I have at least four consecutive observations between 1970 and 2005, and then I average these values (equal weights) across all houses in a given MSA. I also use the Wharton Land Regulation Index, which is provided by Joseph Gyourko, Albert Saiz, and Anita Summers online. The data is available for various cities. I aggregate them by MSA via a simple equal-weighted average.

I filter out outlier values for a series of variables (see Appendix C for the details). I also filter out households who have negative net worth or who have an income less than \$1,000 in year-2000 dollars, as well as those who live outside a Metropolitan Statistical Area (MSA). I require households to be present in the 1996 interview wave because of an important question on their risk aversion (see Appendix C). I present three tables of summary statistics for the remaining population in 2005: information on their demographics, income, and housing situation in Table 2.1, information

⁷Before 1999, this amount includes both retirement and non-retirement accounts. Afterwards, households were asked to report the values of both accounts separately. To be consistent with the pre-1999 years, I take the sum of these two accounts.

Table 2.3: Neighborhood characteristics for households in 2005

Variable	Mean	Std Dev
<i>Census variables</i>		
census tract pop	5,916.51	3,383.6
prop. White households	.82	.21
prop. Afr-Americans households	.1	.19
prop. Asian households	.04	.06
prop. Hispanic households	.08	.13
prop. foreign-born households	.08	.1
prop. farmers	.00	.01
prop. executives	.07	.04
prop. workers	.18	.06
prop. technicians	.11	.06
vacancy rate	.03	.03
recreational vacancy rate	.02	.06
prop. owner-occupied homes	.74	.19
prop. households in same house 5 yrs ago	.55	.14
unemployment rate	.05	.06
<i>MSA variables aggregated from the PSID</i>		
housing returns	.01	.16
Wharton land regulation index	.14	.72
N	1,763	1,763

I report two sets of neighborhood characteristics for households from the 2005 wave of the PSID. First, I get census-tract level statistics from the Census Bureau for all the census tracts in which households from the PSID resided. I filter out the households from the PSID with negative net worth or who were not members of the original Survey Research Center core sample. Since the U.S. Census is only conducted at the beginning of every decade, I interpolated 2005 data (linearly) from the 1990 and 2000 censuses. Reported variables are described in Section C.8.

on their portfolio allocations in Table 2.2, and information on their neighborhood characteristics from the US Census in Table 2.3. Because of the positive net worth restriction, the sample is slightly biased toward the older and richer households, as indicated by the high rate of homeowners (79%) and the high average net worth (\$378,840) and household income (\$87,150). About 38% of households have a negative net worth. It comes from the fact that we are missing some of the key elements of net worth like human capital.

2.3 Identification of the Constrained Homeowners

Since the objective of this empirical analysis is to understand the effect of indivisibility on consumption allocation, I need to begin by identifying which homeowners are constrained as well as how constrained they are. In the model, I defined a constrained homeowner as one whose optimal housing investment-consumption ratio ϕ_t^* is strictly between zero and one (see Figure 1.1). Unfortunately, for most homeowners, ϕ_t^* is unobservable to the econometrician. The PSID provides information on whether households own or rent their home, but there is no indication regarding whether these households would rather be partial home owners.

There is, however, one subset of homeowners who are not constrained. They are the landlords. Unlike other homeowners, these households have chosen to own more housing than what they consume, which suggests that the indivisibility constraint does not bind for them.⁸ In other words, it suggests that their observed investment-consumption ratio ϕ_t is greater than one and optimal ($\phi_{i,t} = \phi_{i,t}^*$) and that the opportunity cost of living in their home is zero.⁹

⁸In reality, landlords are not fully unconstrained. It may be difficult for them to achieve their exact value of ϕ_t^* , because of limited supply in the housing market. However, their options are much more varied for the homeowners whose optimal investment-consumption ratio ϕ_t^* is less than one. Landlords can choose to rent out various parts of the house they own, and they can choose to own additional homes nearby of various sizes.

⁹It could be that some landlords have consumption value for the house they rent out. For instance, they retain the option of moving into that house later on. I suspect however that this concerns only a minority of landlords.

My identification strategy consists of exploiting information on the characteristics of landlords. Since the opportunity cost of housing services depends on all of the characteristics of a household, such as β , γ , or T , I begin with the entire sample of homeowners and look for the characteristics that have led some homeowners to be landlords. I then focus on the sample of the non-landlord homeowners, and I sort them according to how similar their characteristics are to those of landlords. I assume that the households whose characteristics are *least similar* to those of landlords are *most constrained*. They pay the highest opportunity cost of housing services.

This strategy relies on there being heterogeneity in characteristics across all the homeowners. In Tables 2.1 and 2.2, I compare the characteristics of both groups.¹⁰ Only 7% of households have chosen to be landlords, and they are clearly different from the other homeowners. For example, the average landlord is much richer than the average non-landlord homeowner, with a difference in net worth of about \$600,000, and he is also less risk averse and more likely to invest in stocks, bonds, and private businesses.

In a related empirical analysis, Brueckner (1997) also identifies landlords as unconstrained homeowners and finds that a comparison of their portfolios to those of non-landlord homeowners, which he assumes are all constrained, yields mixed empirical results. In this analysis, I do not assume that all non-landlord homeowners are constrained. On the contrary, I consider the homeowners whose characteristics are extremely similar to those of landlords to be least constrained. I leave unanswered the question of why these homeowners have chosen not to be landlords. In reality, the transition from being a non-landlord homeowner to a landlord homeowner is not continuous, so the option to be a landlord may not be worthwhile for the homeowners whose optimal housing investment-consumption ratio ϕ^* is higher than but close to one. For example, becoming a landlord involves high transaction costs from having to buy another house to dealing with tenants. There are also tax differences in the treatment of owning rental housing and owner-occupied housing. Unlike regular homeowners, landlords are taxed on the rental income they earn, even though they

¹⁰A landlord is defined as a household who has received rent income in the year prior to her interview.

are allowed to depreciate the value of the house they rent out.

I begin by estimating a probit regression of $\text{Prob}[L_{i,t} = 1]$ on a series of determinants $Z_{i,t}$, where $L_{i,t}$ is a dummy variable that is one if the homeowner is a landlord,

$$L_{i,t} = \mathbb{1}\{a'Z_{i,t} + u_{i,t} > 0\}, \quad (2.1)$$

where $u_{i,t} \sim N(0, 1)$, $\mathbb{1}$ is an indicator function, and t refers to any of the years where net worth is available.

The vector $Z_{i,t}$ of independent variables is composed of characteristics that affect the magnitude of the opportunity cost of housing services. In the comparative statics analysis of the theoretical model in Section 1.4.3 I showed that this opportunity cost depends positively on homeowners' level of risk aversion γ , their horizon T , and their relative preference for housing $1 - \beta$. In the data I look for proxies of these parameters. I proxy for γ by computing an elicited risk tolerance measure from a set of questions in the PSID. In 1996, survey participants were asked to estimate how tolerant they were of various gambles on their lifetime income. I back out from their answers a coefficient of relative risk aversion (see Appendix C for the details).

I proxy for β by including family income and net worth. In the model, net worth does not play a role because of the homotheticity assumption. In reality however, housing is considered a basic good. Wealthier households are likely to value housing consumption relatively less.¹¹ I also use family size and the civil status of the household as proxies. Being married and having a large family may be indicative of a high consumption demand for housing.

For the horizon T , I rely on the Census Bureau, which provides information on the fraction of households in the same census tract who have been in the same house for the past five years. Homeowners who live in areas where there is low mobility are more likely to stay in their house for a long period of time. I also include fixed effects for various age groups (20-30, 30-40, ...).

I also include other variables that affect the homeowner's optimal housing investment-consumption ratio ϕ_t^* . Any variable that leads to a high optimal housing investment-consumption ratio ϕ_t^* should lead *ceteris paribus* to a high opportunity cost of housing

¹¹Wealthier households may also be more or less risk averse but I already control for γ .

services. For example, I include properties of the house prices. Households who live in houses with high expected returns and low volatility should have a high ϕ_t^* . I compute the expected value and the volatility of housing returns for each MSA from the entire PSID data (1968-2005). I also include the proportion of vacant houses that are for sale or for rent from the U.S. Census Bureau. Finally, I include a measure of the elasticity of housing supply from the Wharton Land Use Regulation Index from Gyourko, Saiz and Summers (2008).

Finally, I include other factors that may affect ϕ_t^* outside of the model. For example, I compute the proportions of households in the census tracts from different ethnic backgrounds (White, African-American, Hispanic, Asian, Foreign) to capture potential cultural effects. I control for any unobservable differences between renter-occupied housing and owner-occupied housing by adding the share of owner-occupied homes in the same census tract. I include dummies on whether the household head graduated from high school and college. Davidoff (2006) and Ortalo-Magné and Rady (2002) point out that the investment demand for housing decreases as the covariance between labor income and housing prices increases. I control for this effects by including the proportions of households in different industry sectors in the same census tract (see Appendix C). Finally, I add fixed effects for various cohorts, year, and states. Year fixed effects control for any changes in ϕ_t^* over time, such as the 1986 Tax Reform Act.¹² State dummies control for any specific state-level unobservable effect such as mortgages regulation and subsidies.

Note that since some of these variables from the wealth surveys are only available every five years before 1999 and then every two years, I only select the years during which they are available. Standard errors are robust and clustered by household. All dollar amounts are converted to 2000 real dollars using the consumer price index (CPI). Similarly, all returns are in real terms.

Table 2.4 presents the results of the probit estimation on the likelihood of being a landlord. There are 8,822 observations and the adjusted R-square is about 19.75%. The effects of the proxies for the homeowners' level of risk aversion γ , their relative

¹²The 1986 Tax Reform Act limited the extent to which many investors could deduct tax losses associated with their real estate investments against their gross income.

preference for housing β , and their horizon T are consistent with the predictions of the model. For instance, homeowners who report a high degree of risk aversion are less likely to be landlords. This effect is statistically significant at the 10% level.

In terms of the proxies for β , wealth has the most significant effect. Wealthier homeowners are much less likely to be constrained than other homeowners. Being married and having a large family are also negatively related to being a landlord, which confirms the hypothesis that these homeowners have a high consumption demand for housing.

Finally, the mobility variable, which proxies for the horizon T , is also a strong predictor of being a landlord. Homeowners in census tracts which have a high proportion of households who have been in the same house for the last five years are more likely to be landlords.

Other variables related to the investment value of the house matter as well. Homeowners who live in areas where housing returns are high and volatile on average are more likely to be landlords. Unreported state fixed effects have extremely strong effects on the likelihood of being a landlord. Similarly, unreported year fixed effects, which capture changes in the mortgage rate over time, are also quite significant.

College graduation has a negative effect on the probability of being a landlord. While this coefficient indicates that landlords may be less sophisticated investors than other homeowners, their portfolio characteristics in Table 2.2 suggest the contrary. Landlords are more likely to invest in stocks and bonds than other homeowners.

Given the coefficients a , I compute the predicted probability of being a landlord for all the non-landlord homeowners and use it as a proxy for how constrained they are. I denote this predicted probability as $\Phi(a'Z_{i,t})$, where $\Phi(x)$ corresponds to the cumulative normal probability of x . I assume that the homeowners who have similar characteristics to those of landlords are less constrained than the other homeowners and pay a smaller opportunity cost.

Table 2.4: First-stage probit regression on the likelihood of being a landlord

Variable	Prob[Landlord]	Variable	Prob[Landlord]
<i>Ind. characteristics</i>		<i>Census-Tract characteristics</i>	
family size	-.077*** (.024)	prop. White hh	-.478 (.864)
married	-.167** (.072)	prop. Afr-Amer. hh	-.382 (.868)
high school	-.009 (.066)	prop. Asian hh	-1.322 (1.272)
college	-.131* (.077)	prop. hispanic hh	-.552 (.659)
family income	.18*** (.05)	prop. foreign-born hh	-.718 (.784)
net worth	.356*** (.03)	prop. owner-occ. homes	-1.213*** (.216)
risk aversion	-.009* (.005)	prop. farmers	.36 (2.287)
		prop. executives	-1.664 (.247)
		prop. workers	-.229 (.749)
<i>MSA Characteristics</i>		prop. technicians	-.179 (.878)
avg. housing returns	4.461** (2.173)	vac. rate	.368 (1.293)
std. housing returns	1.315* (.783)	recr. vac. rate	.306 (.52)
Land reg index	-.036 (.52)	prop. hh same house 5 yrs ago	.503* (.279)
N	8,822		
Adj. R^2	.197		

I report the estimates of a probit regression of $L_{i,t}$, a dummy variable indicating whether household i has received rental income in year t . I regress $L_{i,t}$ on a vector of independent variables $Z_{i,t}$,

$$L_{i,t} = \mathbb{1}\{a'Z_{i,t} + u_{i,t} > 0\},$$

where $u_{i,t} \sim N(0,1)$, and $\mathbb{1}$ is an indicator function. I only select the years where net worth is available (1984, 1989, 1994, 1999, 2001, 2003). Standard errors are robust and clustered by households. I control for year, age, cohort, and US-state fixed effects. The sample is restricted to households with positive net worth, positive family income. Further information on the vector $Z_{i,t}$ of explanatory variables is provided in Section C.8.

2.4 Effect of Indivisibility on Housing Consumption Volatility

We are now ready to explore the effect of indivisibility on the optimal consumption allocation of homeowners. The model predicts that indivisibility amplifies variations in both the level of non-housing consumption and the composition of the consumption basket and dampens variations in their level of housing consumption. Here I focus on the volatility of housing consumption and find evidence that more constrained homeowners adjust their housing consumption less over time.

A more complete analysis would include a test of the effects of indivisibility on non-housing consumption and the composition of the homeowners' consumption baskets. However, while the PSID also provides data on non-housing consumption, it is very limited. Only noisy estimates of food consumption are available, and, as Blundell, Pistaferri and Preston (2006) argue, the dynamics of food consumption differ in important ways from the dynamics of other non-durable consumption.

Homeowners can adjust their housing consumption in two ways. They can either move to another house or choose to make improvements to their current house. Downing and Wallace (2001) and Davidoff (2006) provide evidence that improvements in housing represent a major component of these adjustments. Both options involve transaction costs, so I also consider the frequencies at which households move out of their home and make improvements to it. I have information on these four variables in the PSID.

Frequency of moves During each interview, homeowners were asked to report whether they have moved out of their home and explain why, whether it is job related, for consumption purposes (more or less space, better neighborhood), or for other reasons (being evicted, health reasons). I refer the reader to Appendix C for details. I create a dummy variable to identify the homeowners who moved for a consumption purpose in a given year, and I only retain observations where the homeowner was also a homeowner during the previous interview.

Size of the moves Homeowners were also asked to report or otherwise guess the current market value of their home. I select the homeowners who have moved for a consumption purpose and I compare the first available value of their new home to the last available value of their old home. Since the objective is to focus on the size of the adjustment, I compute the absolute value of the percentage change in housing consumption.¹³ While these house values are not necessarily the transaction prices, they are so close to the transaction period that the estimation error is likely to be small.¹⁴

Frequency and size of the improvements In the wealth surveys, homeowners were asked whether they have made major improvements above \$10,000 (since the last wealth survey) on any of their real estate properties. Homeowners were also asked to report the dollar amount that was spent on these improvements.¹⁵ Since these variables are backward looking, I take their forward lag. To avoid taking into account improvements that were done purely to increase the value of their house before selling it, I do not include observations where the households reported being likely to move within the next couple of years from the date of the interview. I also measure the amount that households spent on the improvements by scaling it to the current value of their real estate.

For each of these four measures, which I denote as $\sigma_{K,i,t}$, I set up the following regression equation:

$$\sigma_{K,i,t} = b'X_{i,t} + c \cdot \Phi(\gamma'Z_{i,t}) + \epsilon_{i,t} \quad (2.2)$$

where $\epsilon_{i,t} \sim N(0, \sigma_\epsilon^2)$, $X_{i,t}$ is a vector of independent variables, and $\Phi(a'Z_{i,t})$ is the predicted probability of being a landlord that I backed out from equation (2.1). I run probit regressions if $\sigma_{K,i,t}$ represents one of the two frequency measures. I also

¹³I filter out observations where the time lapse between the two houses values is greater than two years.

¹⁴Furthermore, Skinner (1994) finds that the reported values from the PSID are roughly similar to objective Commerce Department measures in the 1970s and 1980s.

¹⁵This amount includes work households have done themselves. It is not supposed to reflect general maintenance or upkeep.

exclude landlords from the sample at this point, in order to avoid capturing other unobservable effects that may have led the landlords to choose to be landlords.

Before I move forward with the results of the estimation, it is important to point out that this approach presents a couple of challenges. First, some of the low-wealth homeowners who are most constrained by indivisibility may also be liquidity constrained. In the model, I do not account for borrowing or default constraints that may lead these homeowners to choose a “corner” consumption allocation. I control for this effect by focusing on the sample of households who are older than 30 years, do not have a second mortgage and who do not come from the Survey of Economic Opportunities (SEO), which is a subsample from the PSID that includes primarily low-income families.

Second, since all four dependent variables depend on each other, one should not interpret the effect of indivisibility on each of them separately. For instance, homeowners who rarely move out of their home could still have a high housing consumption volatility as long as they improve their home a lot or move to a very different home when they do move. As Grossmann and Laroque (1990) initially pointed out, homeowners who face high transaction costs do not switch homes frequently, but they adjust their housing consumption significantly when they choose to move.¹⁶ In Table 2.5, I report summary statistics for these four variables. It is clear that transaction costs matter. Only 2% of homeowners move in a given year for consumption purposes,¹⁷ and when they move they adjust their housing consumption significantly (29% increase). Similarly, only 11% of homeowners choose to make major improvements to their house in a given period, and when they choose to do so they invest about 11% of the value of their real estate.

¹⁶Damgaard et al. (2003), who extend Grossmann and Laroque (1990) to a world with housing and non-housing consumption and stochastic house prices, show that transaction costs have the same effects on adjustments to housing consumption, although these effects are not as big quantitatively.

¹⁷The actual fraction of moves is slightly higher, but I only count observations where during the previous interview the homeowner was also a homeowner and reported the value of his house.

Table 2.5: Characteristics of adjustments in households' level of housing consumption

Variable	Main Sample				Full Sample			
	Frequency	Mean	Std Dev	N	Frequency	Mean	Std Dev	N
moves								
consumption	.039	.29	.51	8,307	.028	.41	.66	31,159
job	.008	.19	.31	8,307	.005	.19	.53	31,159
other	.004	.1	.4	8,307	.002	.15	.57	31,159
improvements (\$)	.11	31.5	38.07	5,702	.09	29.35	36.71	9,764
improvements / real estate value	.11	.16	.24	5,702	.09	.16	.22	9,764

I report summary statistics on changes in the level of housing consumption for two samples. The main sample includes observations from years 1984 to 2005. I exclude homeowners who (i) are younger than 30 years, (ii) who are not part of the Survey Research Center core sample, (iii) who have a second mortgage, and (iv) who also own rental housing. I also exclude households who have missing or negative net worth in the first wealth survey prior to the interview date. I also report results for the full sample, which covers the same years but does not impose any of the restrictions I used to construct the main sample. I report statistics for two types of variables: moves to new homes and improvements on owned real estate. These variables are described in Section C.8.

I control for these substitution effects by imposing a test that is simple and strict: indivisibility has a negative effect on housing consumption volatility if $c > 0$ for all four measures. A positive value of c implies a positive relationship between housing consumption volatility and the probability of being a landlord, which is negatively related with the extent to which households are constrained. In other words, if constrained homeowners (i) move less frequently, (ii) adjust their housing consumption less when they move, (iii) choose to make improvements to their home less frequently, and (iv) make smaller improvements, then I conclude that indivisibility leads to a decrease in the volatility of the level of their housing consumption. Since any substitution effect between these four measures would lead to coefficients c of different signs, any evidence of the same sign would strengthen my results.

The vector $X_{i,t}$ consists of some variables that I used in the first-stage estimation, such as the proxies for γ and β . I showed in the model that these variables can also affect the volatility of housing consumption independently from their effect on indivisibility. For example, in Table 1.4, an increase in the degree of risk aversion γ from 4 to 8 decreases the annual volatility of housing consumption of the unconstrained agent from 8.4% to 6.1%.

I also control for variables that may also affect the homeowners' level of housing consumption, such as their education, their ethnic background, their profession, or the unemployment rate at the census-tract level. In the regressions on the size of the moves, I include dummies on whether the homeowners changed MSAs or states. Finally, I add state, year, cohort, and age dummies.

I report the estimates of equation (2.2) for all four measures of housing consumption volatility in Table 2.6. As in the previous stage, I compute robust standard errors and cluster them by homeowner. It is worthwhile noting that the coefficients of this estimation are not as significant as those from the first-stage estimation, which suggests a higher degree of noise with the housing consumption data. This biases me against finding any significance in the effect of indivisibility.

For all four measures, the effect of indivisibility is positive. It is also similar for both moves and improvements. While the coefficients on the frequencies of moving and making improvements are positive but not statistically significant, the extent to

Table 2.6: Effect of indivisibility on housing consumption volatility

Variable	Move		Improvements	
	Frequency	Adjustment	Frequency	Adjustment
prob. landlord	.885 (.76)	1.722* (1)	.134 (.869)	.467* (.287)
family size	-.064** (.032)	-.01 (.034)	-.04 (.031)	-.001 (.475)
married	.259*** (.097)	-.041 (.11)	.21* (.115)	-.076 (.078)
risk aversion	-.001 (.004)	.002 (.003)	.001 (.003)	.001 (.005)
high school diploma	.064 (.079)	-.112 (.072)	-.008 (.081)	.06* (.033)
college degree	-.032 (.082)	.014 (.077)	.045 (.093)	.021 (.023)
family income	-.057 (.071)	.028 (.066)	.291*** (.07)	.02 (.03)
net worth	-.016 (.056)	-.12* (.069)	.074 (.061)	-.104*** (.025)
Δ MSA		.04 (.135)		
Δ State		.009 (.179)		
N	5,749	470	2,790	402
Adj. R^2	.1225	.287	.139	.334

I report the estimates of the second stage regressions of measures of housing consumption volatility on indivisibility. The sample includes observations from years 1984 to 2005. I consider four proxies for housing consumption volatility $\sigma_{K,i,t}$: the frequency of consumption-related moves (moves - frequency), the adjustment in the level of housing consumption during a move (moves - adjustment), the frequency of housing improvements (improvements - frequency), and the amount spent on improvements (improvements - adjustment). I regress each dependent variable on the households' predicted probabilities of being landlords $\Phi(a'Z_{i,t})$ that I backed out from equation (2.1),

$$\sigma_{K,i,t} = b'X_{i,t} + c \cdot \Phi(a'Z_{i,t}) + \epsilon_{i,t}$$

where $\epsilon_{i,t} \sim N(0, \sigma_\epsilon^2)$, and $X_{i,t}$ is a vector of independent variables. I run probit regressions if $\sigma_{K,i,t}$ represents one of the two frequency measures. Standard errors are robust and clustered within households. I control for year, age, cohort, and US-state fixed effects. I also control for the households' education, their ethnic background and professions but do not report these effects. Sample restrictions and reported variables are described in Section C.8.

which households adjust their housing consumption when they move or improve their house is quite significant. Both of these coefficients are statistically distinct from zero at the 10% level in the main sample, which means one-tailed tests on their being strictly positive are significant at the 5% level.

Economically, the effects of indivisibility on the size of the adjustments are both quite large. A homeowner who goes from facing the highest possible opportunity cost of housing services (e.g. a zero probability of being a landlord) to no cost (e.g. a full probability of being a landlord) implies a 172% increase in the extent to which he adjusts his level of housing consumption, conditional on moving. It also leads to a 47% increase in the amount (scaled by the current house value) spent on improvements in the house, conditional on reporting to be unlikely to move in the near future.

These results are consistent with the predictions of the model. Homeowners who pay a high opportunity cost of housing services adjust their level of housing consumption less than the “less constrained” homeowners. They do not necessarily move or improve their home less frequently than the less constrained homeowners, but when they do so they do make much smaller adjustments.

2.5 Conclusion

In this essay, I have provided a new empirical methodology for testing the effect of indivisibility of housing on the volatility of housing consumption. Using household-level data from the Panel Study of Income Dynamics (PSID), I developed a method to identify homeowners who are constrained by comparing their characteristics to those of a subset of unconstrained homeowners: the landlords. The results are consistent with the predictions of the model that I derived in Chapter 1. First, the characteristics of homeowners that determine how constrained they are in the model are strong predictors of those homeowners who choose to be landlords in the data. Second, I find evidence that the more constrained homeowners adjust their level of housing consumption much less over time.

These results, however, only provide evidence along one dimension of the model: the effect of indivisibility on housing consumption volatility. I have not explored

in this essay the effects on non-housing consumption because of data limitations. Further empirical research is needed to test and quantify how much this indivisibility constraint distorts the overall composition of households' consumption baskets.

Chapter 3

Hedging Labor Income Risk¹

We investigate the relationship between workers' labor income and their investment decisions. Using a detailed Swedish data set on employment and portfolio holdings we estimate wage volatility and labor productivity for Swedish industries and, motivated by theory, we show that highly labor productive industries are more likely to pay workers variable wages. We also find that both levels and changes in wage volatility are significant in explaining changes in household investment portfolios. A household going from an industry with low wage volatility to one with high volatility will *ceteris paribus* decrease its portfolio share of risky assets by 25%, i.e., 7,750 USD. Similarly, a household that switches from a low labor productivity industry to one with high labor productivity decreases its risky asset share by 20%. Our results suggest that human capital risk is an important determinant of household portfolio holdings.

¹This chapter was written in collaboration with Thomas Jansson, Christine Parlour, and Johan Walden.

3.1 Introduction

Labor income accounts for about two thirds of national income and, since the seminal work of Mayers (1973), it has been shown to play an important role in theoretical asset pricing. In studies such as Bodie, Merton and Samuelson (1992), Danthine and Donaldson (2002), Qin (2002), Santos and Veronesi (2006) and Parlour and Walden (2009), risky labor income affects the portfolio decisions made by investors, which in turn has general equilibrium asset pricing implications. However, the empirical evidence is mixed as to whether an aggregate labor factor can explain stock returns. Fama and Schwert (1977) find that adding a labor factor does not improve the performance of the unconditional CAPM. By contrast, Jagannathan and Wang (1996) find that an aggregate labor factor significantly improves the performance of a conditional CAPM in explaining the cross section of expected returns (see also, Palacios-Huerta (2003)).

Given a potentially incomplete market and noisy measurements, using aggregate labor income data to show the importance of human capital risk in investors' investment decisions, is a daunting task. So, we take a different approach: Since the effects of risky human capital on asset prices are driven by investors' portfolio decisions, we directly study their portfolio holdings. If there is no discernable relationship between agents' labor income and their investment decisions, then it is difficult to posit a plausible link between a labor factor and asset prices. We use panel data on employment and portfolio holdings of a large subset of the Swedish population; and we examine if there is a relationship between employees' labor productivity, wage structure (measured by wage level and volatility) and portfolio holdings.

We find that although there is only a weak link between the *levels* of employee labor productivity, wage structure and portfolio holdings, there is a strong link between *changes* in these variables. For example, households adjust their portfolios in response to job changes. In particular, for households where both adults switch industries in the same year, an increase in wage volatility by 1% will lead to a decrease in the share of risky assets by 1.07%. This effect is statistically significant at the 5% level. This means that a household going from the industry with the least variable

wage (recycling metal waste) to the industry with the most variable wage (fund management) will *ceteris paribus* decrease its share of risky assets by more than 25%, or 7,750 USD. Similarly, a household that switches from a low labor productivity industry to one with high labor productivity decreases its risky asset share by 20%. We also provide evidence on the link between wage volatility and the labor productivity of industries. We find that industries that require high levels of labor productivity also have wages that are (i) volatile and (ii) high on average.

Our results are consistent with a world in which households take human capital into account when making investment decisions, but in which other, offsetting, factors are also important, e.g., heterogeneity in risk-preferences, a familiarity bias, or heterogeneous information. If any of these other factors varies with the business cycle, then our results are consistent with a world in which a human capital factor is of little help in an unconditional CAPM (as argued in Fama and Schwert (1977)), but significantly improves the performance of a conditional CAPM (as argued in Jagannathan and Wang (1996)).

Our tests are based on the predictions in Parlour and Walden (2009). Briefly, the paper develops a general equilibrium model with multiple industry sectors in which workers accept employment contracts offered by firms and their effort is used as an input into production. Firms face a moral hazard problem in that they cannot observe the effort level of employees, so optimal wage contracts include risky compensation. The theory explicitly links the level of labor productivity in a sector to (i) both the level and the volatility of wages offered to employees, and (ii) the portfolios that these employees hold in equilibrium. Firms that require high labor productivity choose a highly variable wage structure that is linked to performance in order to induce effort from their employees. As a result, employees of the high-productivity firms choose to reduce their exposure to risky assets in their investment portfolio.

We use the LINDA database, which provides detailed income and wealth information for a large representative sample of about 3% of the Swedish population from 1999 to 2003. While we do not have information on their individual security holdings, we do know the share of the households' wealth invested in directly held stocks, mutual funds, and other financial assets such as derivative and capital insurance prod-

ucts. This information provides us with a measure of hedging of systematic risk. By definition, most firms bear a positive level of market risk. If we assume that the wages are on average positively correlated with the market then employees can hedge their labor income risk by holding a lower share of risky assets and mutual funds.

In addition to investigating the relationship between agents' portfolio composition and labor income, we also investigate individuals who change industries over the years. For these individuals, we look at their portfolio holdings one year before and one year after their industry switch, and we ask the following question: given their initial portfolio holdings, how does the industry switch affect the change in their portfolio holdings? In particular, do individuals who switch to sectors that are more productive and offer riskier income streams decrease their share of risky assets? Our measure of industry risk and volatility is estimated across all agents who work in the industry and therefore captures the *ex ante* uncertainty in an agent's human capital.

Our paper is related to a series of other empirical papers that use micro data to investigate the relationship between non-financial income risk and portfolio decisions. Malloy, Moskowitz and Vissing-Jorgensen (2005) find evidence that labor income risk, through a firing decision, can explain the value effect. Their focus is different from ours, however, since we are interested in the relationship between a firm's productivity, the wages it pays, its expected stock returns, and the portfolio holdings of investors.

Massa and Simonov (2006) also present a detailed study of the Swedish population. They look at individual stock holdings and find that households tend to hold stocks that are closely related to their labor income, which goes against the hypothesis of hedging. They argue this is because of a preference for familiar stocks, due to heterogeneous information. This is in line with our finding that the hedging motive does not appear in the levels of stock holding, but rather in the changes after a shock to human capital. In fact, this is consistent with one of Massa and Simonov's findings that investors' hedging demand is greater (or not as negative) for households who switch professions or locations or who experience an unemployment shock. They interpret this as familiarity shocks which prevent the investor from biasing his portfolio away from hedging. Our analysis thus differs from theirs in that we explicitly consider changes in employment but are agnostic about the determinants of portfolio

composition.

Our paper is also related to another series of papers that look at the relationship between wage volatility and labor productivity. Our results indicate that industries with high coefficients of labor elasticity also provide more volatile wages, which is consistent with our theory as well as with the results of other studies. Abowd, Kramarz and Margolis (1999) use a French longitudinal sample and find that firms with higher (total) wages are more productive. Furthermore, the proportion of executive compensation from high productivity firms is found to be higher than in low productivity firms. (See, for example Gaver and Gaver (1993), Gaver and Gaver (1995), Bizjak, Brickley and Coles (1993), and Smith and Watts (1992)). In the LINDA database, however, our workers are not necessarily executives.

The rest of the essay is organized as follows. In Section 3.2 we provide a brief review of the model in Parlour and Walden (2009) and describe the predictions on the relationship between firm productivity, wages, and portfolio decisions. In Section 3.3 we describe the data and the methodology, and in Section 3.5 we provide the empirical results. In Section 3.6 we offer some concluding remarks. Further information about the datasets is provided in Appendix D.

3.2 Theoretical Framework and Empirical Strategy

Our discussion in this section is aimed at providing an overview of, and intuition for, the predictions in Parlour and Walden (2009). For details, the reader is referred to the paper. The model is static and uses a CARA-normal framework. The economy is composed of N sectors, which for expositional purposes we will take to be two; each of which has a different level of labor productivity: sector 1 has high labor productivity and sector 2 has low labor productivity. Within each sector there are many firms and within each firm there are many workers. Workers need to exert effort to be productive, and since their effort level is not observable, firms face a moral hazard problem. As a result, firms choose to offer incentive contracts, which optimally consist

of a fixed part and a variable part. The variable part depends on the performance of the firm, e.g., its profits. For simplicity, firms are assumed to have unlimited liability.

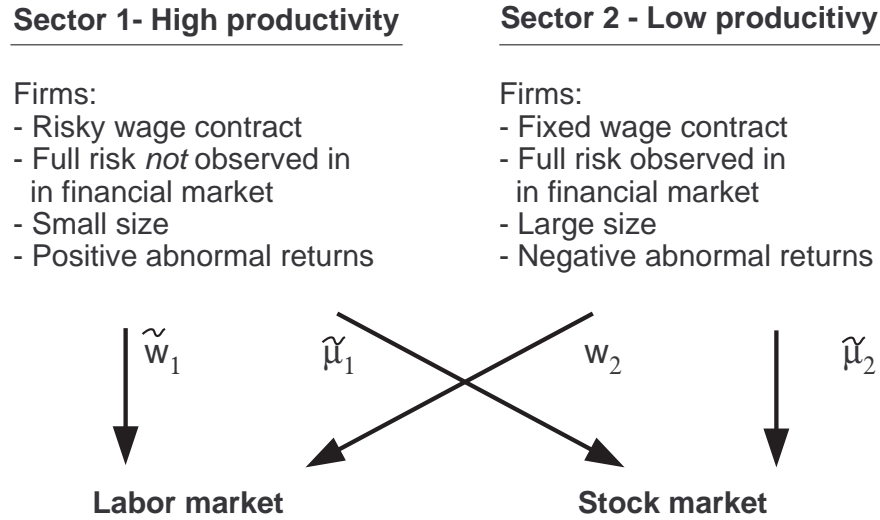
The central intuition of the paper is that an agent's stock portfolio does not accurately reflect his total exposure to systematic risk. Alternatively, in general equilibrium, a firm's equity also does not reflect all the systematic risk that it generates: firms payout risk through wages. Firms with high labor productivity find it optimal to pay most of their wage compensation as incentive wages, since it is relatively important for them to provide incentives to their workers. Thus, the compensation in the high productivity sector 1, \tilde{w}_1 , is risky. Low productivity firms, in sector 2, pay most of their wage compensation through the fixed part, so their compensation, w_2 , is essentially risk-free.

The model also provides implications for the cross-section of expected returns. For example, it is natural to obtain a size effect (and, under other additional assumptions, a value effect). In equilibrium, even though the total size of the high productivity sector is larger than that of the low productivity sector, high productivity firms are on average *smaller* than low productivity firms, because of the higher level of competition. Furthermore, since the high-productivity firms pay a *greater* fraction of their asset risk through wage compensation, their true risk is underestimated if one uses the stock market portfolio as a proxy for the true market portfolio. In other words, econometricians who use the stock market portfolio in their CAPM regressions should find that firms in sector 1 earn positive abnormal returns, $\tilde{\mu}_1$ in the stock market, whereas firms in sector 2 earn negative abnormal returns, $\tilde{\mu}_2$. The model is summarized in Figure 3.1.

While this framework generates several predictions about the relationship between the type of compensation (fixed versus variable) offered, the expected returns, and the type of firm that accords with existing empirical literature,² we focus on the novel implications that relate the productivity of the firm, the riskiness of the wage contract and the portfolio holdings of the workers.

²See for example, Abowd, Kramarz and Margolis (1999), Gaver and Gaver (1993), Gaver and Gaver (1995), Bizjak, Brickley and Coles (1993), Smith and Watts (1992), Kruze (2008), Mehran (1995), Frye (2004) and Kedia and Mozumdar (2002).

Figure 3.1: Summary of model — Two sector example.



In particular, two sorts of predictions arise. First, there are predictions on levels:

H1: The higher the labor productivity of the industry, the higher the wage volatility.

H2: Workers with more variable wages have lower exposure to the market through financial assets.

H3: Workers in higher labor productivity industries have lower exposure to the market through financial assets.

Second, there are predictions on changes. While there might be agent specific heterogeneity outside the model that affects portfolio holdings, if an employee moves to an industry that offers a different wage contract then he should rebalance his portfolio. For example, consider a worker who changes jobs from a low productivity sector to a high productivity sector. Through the labor market he has effectively increased his exposure to the market and therefore should decrease his exposure to risky assets in his investment portfolio.

H4: Workers who switch to a sector with higher wage volatility decrease their exposure to the market through financial assets.

H5: Workers who switch to higher labor-productive sectors decrease their exposure to the market through financial assets.

3.3 Description of the Datasets

To construct measures of portfolio holdings, we use a unique Swedish annual panel database called LINDA (standing for Longitudinal INdividual DAta for Sweden), a joint project between Uppsala University, The National Social Insurance Board,³ Statistics Sweden, and the Swedish Ministry of Finance. LINDA contains an annual cross-sectional sample of around 300,000 individuals, which is approximately 3% of the entire Swedish population. These individuals are tracked over the years. Family members of sampled individuals are also included; this allows us to examine household labor and investment decisions. The sampling procedure ensures that the panel is representative for the population as a whole, and each annual cohort is cross-sectionally representative.

The data are primarily based on filed tax reports (available on an annual basis from 1968) and include various measures of income, government transfers and taxes in addition to individual characteristics such as sex, marital status, education, municipality of residence, and country of birth. From 1999 onwards, the market values of financial and real assets (e.g. stocks, bonds, mutual funds, and owner-occupied homes) are estimated by Statistics Sweden and included in LINDA.

To investigate labor income and working conditions, we rely on Statistics Sweden to obtain two more data sets. The first provides information on industry characteristics, and we use it to compute a measure of labor productivity for each industry. Every year, Statistics Sweden collects firm data such as total sales, the number of employees, and value added. Data from the 558 largest companies is collected through complete surveys. Information about the remaining number of companies is provided by the Swedish Tax Authorities. The coverage rate in 2006 was around 85%. However, the percentage of missing companies as shares of the total number of employees

³The National Social Insurance Board, “Förskringskassan,” manages the Swedish social security system, see, <http://www.fk.se/sprak/eng>.

or net income was only around 3%. The data are reported by industry, which are classified according to 5-digit SNI codes. These are equivalent to the NAICS/SIC codes in the USA. We have access to industries at the 3-digit SNI level from 1997 to 2005.

In LINDA, any working individual is assigned a 5-digit SNI code each year, depending on the industry he or she works in.⁴ Using the SNI codes, we can therefore merge the industry-level data with the household-level data from LINDA. We do this at the three digit SNI code level, which provides sufficient granularity. In total, there are 223 3-digit codes, however we only use a subset of these because the classification changed in 2002, and the mapping between the old codes (1992 classification) and the new codes (2002 classification) is not one-to-one. This classification change matters for our study because it occurs in the time period we are studying. To avoid any potential bias, we only use the subset of SNI codes that remain the same. In addition to other filters, we still end up with 104 SNI codes.

Finally, to control for agent heterogeneity, we use a Statistics Sweden demographic data set. Since LINDA provides information on the region where individuals live, we can also merge this one with the LINDA database and use population density as a control in our regressions on portfolio holdings. This data set groups regions into 6 different categories, based on the population composition at the end of year 2002.

We exclude observations in which information on the wage volatility or the level of labor productivity is missing, and households (defined below) whose financial wealth, net wealth or family income are extreme: less than 3000 SEK, 1000 SEK, and 1000 SEK respectively, and those with negative net holdings of risky assets. As we are interested in labor market participation, we also exclude households in which the largest income goes to someone younger than 18 years or older than 65 years, and households whose family income in 2000 ranges in the top 0.1% of the remaining sample.

⁴In the event the individual has had two jobs during a year, the reported SNI code corresponds to the sector in which she generated most of her income during that year.

3.4 Construction of Variables

Our tests require a measure of portfolio holdings in addition to agents' employment (the source of their returns to human capital). To understand the relationship between returns to human capital and portfolio returns, we also require a measure of wage volatility. Finally, to relate wage characteristics to industry characteristics, we need to estimate an aggregate industry production function.

3.4.1 Portfolios

Since portfolio decisions are typically made at the household level, we use these as our units of observation. However, we also keep track of the individuals within each household as each may work in a different industry. While aggregating household financial holdings is straightforward, imputing wage volatility or labor productivity to a household is less so.

Our sample includes information on wealth from 1999 to 2003, we take 2001 as the base year in order to maximize the sample size. In 2001, we select the two adults within each household who generate the greatest levels of income. We then sort these two individuals by income, and adopt the convention that Individual #1 (Ind1) generates the highest income in 2001 and Individual #2 (Ind2) is the other adult.⁵ We then retain and keep track of these two individuals over the years.

We define a “switcher” as a household where at least one of the two adults changed sectors between 1999 and 2002. More precisely, in order to take into account the fact that investors may not adjust their portfolios immediately before or after a job change, we only look at the adults who switched industries between 2000 and 2001.⁶ A change in industries is recorded as a 3-digit SNI code change. We eliminate switcher households that undergo a major change in their civil status, such as marriage or divorce and those that have increased or decreased their portfolio holdings of either

⁵If the two individuals have the same income, we adopt the convention that Individual #1 is the oldest individual.

⁶In other words, adults households did not switch industries between 1999 and 2000 and between 2001 and 2002.

risky assets, mutual funds, or stocks, by more than 100% between 1999 and 2002.⁷

Summary statistics for the overall population as well as for the switchers are displayed in Table 3.1 for the reference year 2001. A first glance at the table indicates that the sample of switchers is fairly representative of the overall population, which is reassuring because it indicates that on the whole switchers and other households have similar characteristics. However, the switchers tend to be slightly wealthier. In addition, a greater fraction of them are homeowners, and they are more likely to be married and to have a college degree.

For each household (h), we look at its non-retirement portfolio⁸ of risky assets (ra), which contains directly-held stocks and risky mutual funds. We do not consider other risky financial assets, such as capital insurance products, as we do not have any information on the composition of their investments. Calvet, Campbell and Sodini (2007) find that including capital insurance products do not change their results on the level of diversification of household's portfolios. Risky mutual funds include pure-equity funds as well as funds that invest only a positive fraction of their assets in stocks. Ideally we would like to separate these two types of mutual funds but unfortunately this information is not available after 1999 in LINDA. From the 1999 data, however, it seems that the vast majority of these funds is pure-equity. We also decompose the portfolio of risky assets and study in detail the portfolios of directly-held stocks (s) and risky mutual funds (mf). At the end of each year t , we define $w_{h,t}^i$ as the share of household h 's holdings of portfolio i over its financial wealth, which is the sum of cash (checking and savings accounts, money-market funds), bond-only mutual funds, stocks, and risky mutual funds. So, $w_{12,2003}^s$ refers to household #12's share of directly held stocks in its financial wealth at the end of the year 2003.

We report summary statistics on portfolio shares of the overall population as well as those of switchers in 2001 in Panel A of Table 3.2. Again, the switchers are fairly representative of the population, even though they are slightly more likely to invest in

⁷These are absolute values.

⁸Retirement portfolios are not available in LINDA. Until 1998, Sweden had a low-risk defined benefit system, "Allmän Tjänst Pension," ATP, which was then replaced by a defined contribution system (see Sundén (2006)). Since no changes were made retroactively, a large part of the Swedish pension capital was therefore low-risk in our studied time period.

Table 3.1: Summary statistics from the 2001 wave of the LINDA data set.

Variable	All Households				Switchers	
	Mean	Std Dev	Min	Max	Mean	Std Dev
<i>Demographics</i>						
age	44.9	10.14	18	65	44.62	9.11
nordic	.97	.17	0	1	.98	.13
number of children	1.1	1.14	0	13	1.35	1.16
<i>Civil Status</i>						
married	.57	.49	0	1	.69	.46
partnered	.15	.36	0	1	.16	.37
single	.19	.4	0	1	.11	.31
<i>Education</i>						
student	.06	.23	0	1	.06	.23
college degree	.47	.5	0	1	.53	.50
business degree	.15	.35	0	1	.21	.40
<i>Population Density</i>						
high	.35	.48	0	1	.40	.49
medium	.54	.50	0	1	.51	.50
low	.11	.32	0	1	.09	.29
<i>Labor income</i>						
family income	366.07	161.32	1.07	1243.31	409.13	157.70
is unemployed	.13	.34	0	1	.17	.37
is retired	.13	.34	0	1	.13	.34
<i>Housing and Wealth</i>						
homeowner	.88	.33	0	1	.93	.25
net worth	1.09	1.81	0	154.18	1.3	1.71

All monetary values are defined in Swedish kronor (SEK). The average SEK/USD exchange rate on December 28th, 2001 was 10.67. There are 102,049 observations and 6,428 switchers. Reported variables are described in Section D.1.

Table 3.2: Participation rates and portfolio shares for households in 2001.

Variable	All Households			Switchers		
	Mean	Std Dev	Part.	Mean	Std Dev	Part.
<i>Panel A: LINDA</i>						
risky assets	.58	.33	.91	.57	.31	.95
stocks	.22	.26	.56	.22	.24	.63
mutual funds	.48	.32	.84	.46	.3	.88
<i>Panel B: SCF I</i>						
stocks	.40	.31	.41			
mutual funds	.30	.26	.30			
<i>Panel C: SCF II</i>						
stocks	.29	.26	.41			
mutual funds	.19	.19	.30			

Panel A refers to observations from the LINDA dataset. The data set has 102,049 observations overall and 6,429 observations for the switchers. Panels B and C refer to observations from the Survey of Consumer Finances (SCF). In Panel B, we adjust the SCF portfolios so that they are comparable to the ones computed in LINDA. In particular, we exclude retirement assets and we sum up the holdings of pure-equity and mixed mutual funds. Panel C reflects more closely the true risky portfolio shares in the USA. The holdings of mixed mutual funds are halved in order to reflect the fact that they are not fully invested in stocks, and the retirement assets are included.

the stock market. To benchmark, we compare the Swedish participation rates in risky assets and their portfolio shares with those from the USA, which we glean from the 2001 Survey of Consumer Finances (SCF). Since the information on household wealth is more precise in SCF, we present two Tables. In Panel 3.2B, we adjust the SCF portfolios so that they are comparable to the ones computed in LINDA. In particular, we exclude retirement assets and we sum up the holdings of pure-equity and mixed mutual funds. Panel 3.2C reflects more closely the true risky portfolio shares in the USA. The holdings of mixed mutual funds are halved in order to reflect the fact that they are not fully invested in stocks, and the retirement assets are included.

Comparing panels A and B of Table 3.2, it is evident that the participation rate in risky assets is much higher in Sweden than in the USA. Part of this is mechanical: In

LINDA bank accounts for which the annual interest earned is under 100 SEK do not have to be reported. Since we impose a minimum wealth of 3000 SEK, we eliminate all the households who do not make the threshold because of their missing bank accounts and who do not participate in the stock market. The SCF, which is a survey and not a report from a tax authority, does not exclude such observations. However, these missing bank accounts do not completely explain the difference in participation rates. Indeed, if we relax the minimum financial wealth threshold, participation rates in stocks and mutual funds are still about 75% and 69% respectively, which is still considerably higher than in the USA. This result indicates that the selection bias in stock market participation in Sweden is not as important as in the USA. The widespread stock market participation among the Swedish population is well known, and has, e.g., been explained with a high degree of trust and sociability in the Swedish society (see Georgarakos and Pasini (2009)). Second, Swedish households tend to invest much more of their risky assets in mutual funds than American households. This may be due to the introduction in the late 1970's of highly accessible, mutual funds (so-called "Allemansfonder"), which offered high tax-incentives. The tendency towards well-diversified investments is consistent with our empirical analysis; since our measure of hedging is the share of financial assets invested in risky assets. As we do not know how Swedish households compose their portfolio of direct stock holdings, observing a high portfolio share in mutual funds indicates that these households are likely to be mostly invested in the overall stock market. As a result, if these households hedge their labor income risk, they are likely to do so by leveraging up or down their holdings of mutual funds.

3.4.2 Wage Volatility and Labor Productivity

Given that our focus is on households who have switched jobs, computing a measure of annual wage volatility that comes directly from their total household income is difficult, because we only have data for two years after their switch in 2001. We proceed as follows. We begin by computing industry-averages of wage volatility, given the large LINDA sample from 1993 to 2003, and then we attribute these values to

all individuals given the industry they work in each year.⁹ Finally, we aggregate by household each year.

We proceed similarly for our measure of labor productivity, using the industry characteristics data from Statistics Sweden. While using industry-averages may not necessarily reflect an agent's exact wage volatility or labor productivity, it is not unreasonable to view them as *ex ante* measures of both productivity and wage volatility, given that agents are unaware of how their particular careers will evolve.

In the large LINDA sample from 1993 to 2003, we select all the individuals who work in the same industry for at least 5 consecutive years.¹⁰ (Data on wages is also available from the Statistics Sweden output files, but we only have access to the aggregate wage per industry, which provides less information than the micro data from LINDA.) We calculate the wage growth volatility of each individual, which we then aggregate by industry sector. Then, we compute the volatility of the annual growth rate of their real disposable income during these years,¹¹ and average this volatility across all the households within the same 3-digit sector. We only select industries for which we have more than 30 observations, and in doing so we have a measure of wage volatility for 191 industries. This measure takes into account unemployment risk. If a worker is let go during a year, he will still be assigned his former SNI code as long as he was employed during part of the year.

⁹In some cases individuals have worked in two or more industries during the same year. We unfortunately do not have access to this information, but the SNI code that is reported is the one for which the individual earned most of his annual combined salary.

¹⁰We restrict these individuals to have the same 5-digit SNI code in order to make sure they do not switch jobs. We also exclude individuals who are receiving student aid and new job training (if they are unemployed), in order to exclude part-time jobs. Finally, we exclude individuals who are either self-employed or who are owners (or who are a close relative to an owner) of a closely held company, e.g. "3:12" firms, because these individuals are more likely to report their income in a non-conventional way. We choose a period of 5 consecutive years in order to maximize the sample size but results are robust to different specifications.

¹¹We work with disposable income because it is more reliable than pre-taxed income. One weakness of using disposable income is that we may be picking up tax effects that are not related to the individuals' labor income situation. On the other hand, it allows us to capture all the tax effects that are related to their labor income situation. Disposable income is available at the individual-level because in Sweden individuals do not file their taxes jointly.

Table 3.3: Rankings of industries by their levels of wage volatility and labor productivity.

Wave Volatility Rankings			Labor Productivity Rankings		
SNI	Description	Est.	SNI	Description	Est.
Bottom 10					
371	Recycling of metal waste and scrap	.07	552	Youth hostels, camping sites	.02
271	Manufacturing of iron and steel	.08	519	Other wholesale	.02
131	Mining of iron and ores	.08	504	Sales and maintenance of motorcycles	.03
173	Finishing of textile	.09	371	Recycling of metal waste and scrap	.04
272	Manufacturing and casting of iron tubes	.09	202	Manufacturing of fibreboard, plywood, etc...	.04
172	Weaving of cotton	.09	364	Manufacturing of sports goods	.05
365	Manufacturing of games and toys	.09	911	Activities of business organizations	.05
274	Production of precious metals, copper	.10	263	Manufacturing of construction products	.06
403	Steam and hot water supply	.10	205	Manufacturing of other products of wood	.06
175	Manufacturing of ribbons, curtains	.10	245	Manufacturing of perfumes and toilet preparations	.06
Top 10					
21	Renting of household goods	.21	453	Installation wiring, insulation, plumbing	.32
13	Mixed farming	.21	741	Legal representation activities	.32
722	Publishing of software	.22	742	Architecture activities	.32
741	Legal representation activities	.23	211	Manufacturing of pulp and newsprint	.32
672	Other finance activities	.24	601	Urban transport via railways	.33
744	Advertising	.24	341	Manufacturing of motor vehicles	.33
924	Other Entertainment	.25	452	General construction	.33
553	Restaurants	.26	244	Pharmaceutical operations	.34
921	Motion picture and video production	.26	722	Post activities	.35
671	Finance administration, fund management	.30	401	Publishing of software	.35

Wage Volatility is defined as the average volatility of annual returns to real disposable income across all individuals within a 3-digit SNI code who have stayed in the same 5-digit SNI code for at least 5 consecutive years between 1993 and 2003. The level of labor of productivity is defined as the elasticity of output with respect to labor. It is estimated via a random coefficients panel regression on the Output tables from Statistics Sweden. Rankings of the wage volatility and labor productivity measures are based on 191 and 104 observations, respectively.

Table 3.3 reports the top and bottom ten industries ranked by wage volatility. It is not surprising to find that industries such as “fund management,” “legal representation activities,” and “motion picture and video production” have high wage volatility whereas industries such as “recycling of metal waste and scrap” and “mining of iron and ores” have low wage volatility.

In Parlour and Walden (2009), the agents from the highly productive industries who receive volatile wages also receive higher wages on average, in order to be compensated for the high level of labor income risk. It is easy to test this relationship using data from LINDA. We select the same individuals as those from our measure of wage volatility and compute the average annual level of real disposable income for each 3-digit SNI code.

Once we have computed these measures on the volatility and level of wages for each 3-digit industry, we assign them to each individual-year given their SNI code. Finally, we aggregate these measures by household, weighting each individual by the amount of disposable income he or she earned during that year. In other words, if the household is composed of two working individuals, then the household labor income volatility measure is a weighted average of the individuals’ volatility. In reality, the household labor volatility should also include the covariance between both individuals’ labor income. However, given that we are working with industry-level estimates for their labor income, estimating precisely this covariance is difficult. In our regression we try to correct for this by creating a dummy to catch whether both individuals work in the same 3-digit SNI code.

According to Parlour and Walden (2009), the volatility of wages should reflect the level of labor productivity for each industry sector. As a robustness check, we construct a measure of labor productivity from the Statistics Sweden Output tables that does not come directly from wages. This specification allows us to test hypotheses (1), (3) and (5). We look at the elasticity of labor under the assumption that the industry’s production function is Cobb-Douglas,

$$\log(Y_{j,t}) = \log(A_j) + a_j * \log(L_{j,t}) + b_j * \log(K_{j,t}) + \epsilon_{j,t}, \quad (3.1)$$

where indices j and t refer to the 3-digit SNI code j and year t , and where Y is the

aggregate value added in real terms, L is the number of employees, and K is the real value of the industry's assets. We filter out a few SNI codes where data was missing or that had very few firms.¹²

We estimate the elasticity of labor, a_j via a random coefficients panel regression, where a_j , b_j , and $\log(A_j)$ are treated as random effects. We also add year fixed effects, and we impose an AR(1) structure on the errors within each industry j to allow for potential serial correlation over time. The results conform with standard intuition. Summary statistics of a include a mean of .21, a standard deviation of .09, a minimum of .02, and a maximum of .35. In Table 3.3 we also report the top and bottom ten industries ranked by their level of labor productivity. Industries such as “manufacturing of construction products” and “recycling of metal waste and scrap” have low labor productivity, whereas industries such as “legal representation activities,” “architecture,” and “publishing of software” have high labor productivity. We have data on labor elasticity for 104 industries. As with our measures of labor income risk, once we have computed a measure of productivity for each industry, we assign it to each individual-year and aggregate these values by household.

3.5 Empirical Tests and Results

We are now in a position to test hypotheses H1-H5 in Section 3.2. For convenience, we repeat the hypotheses below.

H1: The higher the labor productivity of the industry, the higher the wage volatility.

One of the first conclusions of the optimal contracting approach is that in industries in which labor productivity is high, employers have a stronger incentive to elicit high effort and so expose workers to more risk in order to motivate them. Furthermore, if agents are risk averse then in order to induce them to accept more volatile wages, they must be paid a higher wage. Therefore, there should be a positive correlation between wage levels and wage volatility. We report the correlations between

¹²Details are available upon request.

Table 3.4: Pearson correlations between labor elasticity and wage measures.

Variable	Labor Elasticity	Wage Level Mean	Wage Growth Vol
Labor Elasticity	1		
Wage Level Mean	.26***	1	
Wage Growth Vol	.20**	.189*	1

There are 104 observations. Labor Elasticity is computed from the Statistics Sweden Output tables. Wage Level Mean is the average level of log real wages per industry, and Wage Growth Vol is the average volatility of annual growth rate of real wages per industry. Both measures are computed from the LINDA data set for individuals who worked in a given sector for at least 5 consecutive years. Test statistics indicate the probability of observing the empirical correlation under the null hypothesis that the correlation is zero. Statistical significance is represented by *** for 1%, ** for 5%, and * for 10%.

the average level and the volatility of wages and the elasticity of labor in Table 3.4. The data suggest that the higher the labor elasticity (or productivity), the higher the mean level of wages. This is consistent with a payment that compensates a risk averse agent for wage volatility. In addition, there is a positive correlation between elasticity and wage growth volatility which is consistent with an optimal contracting framework.¹³

Having established the positive correlation between labor productivity and wage volatility, we address the effect of both on portfolio levels and then, on portfolio changes. Here are the hypotheses on the levels.

H2: Workers with more variable wages have lower exposure to the market through financial assets.

H3: Workers in higher labor productivity industries have lower exposure to the market through financial assets.

¹³There can be other effects going on. For example, some industries may be productive but also have an easier time monitoring their employees. Alternatively, incentive to shirk may be lower at the more productive industries because work is more enjoyable there. Both effects would push the correlation to be negative, which strengthens our results. However, note that this is only a simple correlation analysis which does not allow us to fully identify whether the cause of this positive correlation is a moral hazard story.

First, consider H2. If human capital is an asset that generates a cash flow stream, then those working in high-productivity sectors, which *ceteris paribus* have riskier income streams, should have a lower share of risky assets and mutual funds. Of course, both employment and human capital are potentially endogenous variables.

As in Vissing-Jorgensen (2002) and Massa and Simonov (2006), we assume that the investment decision takes place in two steps: first, the investor decides whether to enter the stock market, and then he selects his portfolio holdings. In order to account for the first stage participation decision, we use a two-step estimation procedure following Heckman (1979). We model the decision to enter the stock market by estimating $p_{h,t}^{ra}$, the observed probability of participation in the portfolio of risky assets, with the probit regression,

$$p_{h,t}^{ra} = \alpha_{1,t}^{ra} + \beta_{1,t}^{ra'} \cdot \Phi_{h,t} + \gamma_{1,t}^{ra'} \cdot X_{h,t} + \epsilon_{1,h,t}^{ra}, \quad (3.2)$$

where $X_{h,t}$ is a vector of explanatory variables for household h in year t , and $\Phi_{h,t}$ includes either wage volatility or labor productivity along with interaction variables for households where both individuals work in the same industry.

In this and the subsequent regressions, the choice of control variables in the vector $X_{h,t}$ is critical because of the potential endogeneity issues. We control for each household's composition, where it is located, the sources and composition of household wealth and financial sophistication.

To control for differences in household composition, we include the age of the head of the household, as well as age squared, dummies that indicate the civil status of the head (married, partnered but not married, single parent, or single household), the number of children who are minors in the household, a dummy for whether at least one of the adults was born in a Nordic country, and dummies for the number of individuals who used to be part of the household but who are deceased or have emigrated.

Location may affect portfolio decisions and so we use dummies for the population density of the area in which the household lives (high, medium, low). A high density region indicates one of the three metropolitan areas in Sweden: the Stockholm region, the Gothenbourg region, or the Malmo/Lund/Trelleborg regions. A medium

density region is one in which the household lives in an other (less) urban area, which consists of municipalities with (i) more than 27,000 inhabitants, (ii) less than 90,000 inhabitants within 30 km (19 miles) of the municipality center, and (iii) more than 300,000 inhabitants within 100 km (62 miles) of the municipality center. Finally, a low density region represents all the other regions of Sweden.

Measures of labor income and employment include the logarithm of family disposable income, a dummy on whether at least one of the adults is receiving unemployment insurance, a dummy on whether at least one of the adults is receiving a retirement pension, and the ratio of debts to family income. Measures of real estate include a dummy on whether the household owns real estate and the ratio of house value to net worth.

Measures of education include dummy variables on whether at least one of the adults has a college degree and studied business after high school. We also add a dummy variable on whether at least one of the adults is receiving student aid. We avoid controlling for portfolio shares in previous years, because as we will see in the next section, portfolio shares are extremely predictable over time, which means that including them would capture most of the information from the other variables, including $\Phi_{h,t}$. We also avoid net wealth and financial wealth for the same reasons.

Then, in the second stage, we regress the portfolio shares $w_{h,t}^i$ on $\Phi_{h,t}$, our vector of proxies for either wage volatility (for H2) or labor productivity (for H3). Our main focus is on the portfolio share of risky assets ($i = ra$), but we also repeat the exercise for the portfolio shares of stocks and mutual funds. We also include the same vector $X_{h,t}$ of control variables and Heckman's lambda variable ($\lambda_{h,t}$), which controls for possible selection at the first stage. The equation is as follows,

$$w_{h,t}^i = \alpha_{2,t}^i + \beta_{2,t}^{i'} \cdot \Phi_{h,t} + \gamma_{2,t}^{i'} \cdot X_{h,t} + \theta_{2,t}^i \cdot \lambda_{2,h,t} + \epsilon_{2,h,t}^i, \quad (3.3)$$

where i refers to the asset class (risky assets, stocks, and mutual funds). Households hedge their labor income risk if $\beta_{2,t}^i < 0$.

The results of the second stage regressions are reported in Table 3.5 for wage volatility and Table 3.6 for labor productivity. We only report the results for the year 2002, but the results are almost identical across the years. The coefficients of

the control variables are similar across Table 3.5 and Table 3.6. The table with wage volatility has 102,049 observations. The table with labor productivity only has 38,403 observations, the reason being that there are fewer industries for which we were able to compute a measure of labor productivity. In these cross-sectional regressions, both switcher and non-switcher households are included.

Multiple variables are strong predictors of portfolio shares. This is not surprising, and it is consistent with the results in Vissing-Jorgensen (2002), Massa and Simonov (2006), and Calvet, Campbell and Sodini (2007). The richer and more educated households tend to tilt their portfolio toward stocks. This is especially the case for the households for which at least one adult has a business degree.

In terms of real estate and location, we find that conditional on owning real estate, a high ratio of house value to net worth does crowd out participation in the stock market, in line with Cocco (2005). Furthermore, while living in a small urban area does lead to an increase in the share of risky assets and mutual funds, relative to living in a rural area, living in one of Sweden's three metropolitan areas leads to a decrease in the share of risky assets. This may be due to the crowding out effect of the higher home prices in these areas.

In terms of other household characteristics, households who come from Scandinavian countries tend to invest more in mutual funds, which suggests a cultural effect that is consistent with the summary statistics presented earlier. Married, partnered, and single parent households tend to be more invested in risky assets than single households, but less invested in stocks. The coefficient on the number of children is similar, which suggests a potential risk aversion story.

The coefficient on λ_t confirms the selectivity among market participants, despite the high overall participation rate in risky assets. We report the bootstrapped standard errors of the estimates, and for both the shares of risky assets and mutual funds, θ_t is significantly different from 0.

Clearly, from Table 3.5, controlling for selection bias, the effect of the wage volatility variable is weakly consistent with H2. An increase in wage volatility does lead to a decrease in the portfolio shares of risky assets that is significant at the 5% level. However, it is not necessarily significant from an economic perspective. A 1% increase

Table 3.5: Effect of wage volatility in 2002 on portfolio shares.

Variable	Risky Assets		Stocks		Mutual Funds	
	Est.	Std. Err	Est.	Std. Err	Est.	Std. Err
Intercept	.912**	(8.98)	-.614***	(.157)	1.53***	(.232)
wage vol.	-.081**	(.041)	.536***	(.034)	-.616***	(.048)
wage vol. same ind.	.046*	(.026)	.083***	(.022)	-.036	(.027)
age	-.007***	(.001)	-.003***	(.001)	-.004**	(.002)
age ²	.06***	(.01)	.034***	(.001)	.03*	(.001)
nordic	.077***	(.014)	-.009	(.019)	.086***	(.028)
has deceased	-.14***	(.041)	-.003	(.031)	-.137***	(.035)
has emigrated	-.037***	(.014)	.001	(.019)	-.037***	(.014)
no. children	.032***	(.001)	-.006***	(.001)	.038***	(.002)
single parent	.039***	(.006)	-.032***	(.006)	.071***	(.013)
partnered	.008	(.007)	-.027***	(.008)	.035**	(.016)
married	-.002	(.006)	-.038***	(.006)	.036***	(.011)
student	.024***	(.006)	.016***	(.005)	.009	(.008)
college degree	.023**	(.004)	.029***	(.004)	-.007	(.008)
business major	.01***	(.003)	.025***	(.003)	-.014***	(.004)
high pop. density	-.031***	(.005)	.025***	(.006)	-.056***	(.007)
medium pop. density	.023***	(.004)	.005*	(.005)	.019***	(.003)
family income	-.029***	(.007)	.054***	(.053)	-.082***	(.016)
is unemployed	-.003	(.004)	-.003	(.003)	-.001	(.004)
is retired	-.002	(.003)	.004	(.003)	-.007*	(.004)
homeowner	.033***	(.006)	.046***	(.006)	-.013	(.09)
house / networth	.022***	(.003)	-.017***	(.006)	.038***	(.008)
debt / income	-.001	(.001)	.001	(.003)	-.001	(.004)
lambda	.269***	(.056)	.206**	(.086)	.063	(.136)
No. Obs	102,049		102,049		102,049	
F	3,209***		4,495***		17,291***	

We report second-stage estimates of portfolio holdings as a percentage of financial assets in 2002. The sample is restricted to households with positive holdings only. Three separate OLS regressions are run. The dependent variables are the shares of risky assets (stocks and mutual funds) over financial wealth, the share of directly-held stocks over financial wealth, and the share of risky mutual funds (equity and mixed) over financial wealth. Financial wealth is defined as the sum of cash (checking and savings accounts, money-market funds), bond-only mutual funds, stocks, and risky mutual funds. Lambda is the inverse mills ratio from the first stage estimation of equation (3.2). We report the bootstrapped standard errors. The superscripts ***, **, and * refer to coefficients statistically distinct from 0 at the 1, 5, and 10% level respectively. F refers to the Wald goodness-of-fit test. Other explanatory variables are described in Section D.1.

Table 3.6: Effect of labor productivity in 2002 on portfolio shares.

Variable	Risky Assets		Stocks		Mutual Funds	
	Est.	Std. Err	Est.	Std. Err	Est.	Std. Err
Intercept	1.325***	(.138)	-.555***	(.203)	1.88***	(.318)
labor prod.	.064	(.047)	.111***	(.034)	-.048	(.042)
labor prod. same ind	.038**	(.017)	.07***	(.017)	-.032	(.021)
age	-.004***	(.018)	-.001	(.001)	-.002	(.002)
age ²	.03*	(.02)	.01	(.01)	.02	(.02)
nordic	.04*	(.021)	-.007	(.026)	.046	(.555)
has deceased	-.094*	(.056)	-.036	(.048)	-.057	(.068)
has emigrated	-.02	(.023)	-.003	(.014)	-.017	(.022)
no. of children	.031***	(.003)	-.011*	(.003)	.042***	(.003)
single parent	.033***	(.01)	-.018*	(.003)	.051***	(.016)
partnered	-.007	(.012)	-.018	(.013)	.011	(.02)
married	-.004	(.009)	-.03***	(.009)	.026*	(.013)
student	.02**	(.008)	.018**	(.007)	.001	(.01)
college degree	.012**	(.006)	.037***	(.007)	-.024**	(.012)
business major	.001	(.009)	.027***	(.004)	-.026***	(.007)
high pop. density	-.007	(.006)	.039***	(.005)	-.046***	(.009)
medium pop. density	.035***	(.004)	.012***	(.004)	.023***	(.007)
family income	-.063***	(.009)	.048***	(.013)	-.111***	(.02)
is unemployed	-.003	(.005)	-.004	(.004)	-.001	(.006)
is retired	-.005	(.005)	.006	(.004)	-.012*	(.006)
homeowner	.14	(.009)	.044***	(.008)	-.03**	(.014)
house / networth	.033**	(.005)	-.017***	(.007)	.05***	(.01)
debt / income	-.001	(.001)	.002	(.002)	-.003	(.003)
lambda	.111	(.08)	.256**	(.111)	-.145	(.173)
No. Obs	38,403		38,403		38,403	
F	2,059***		1,571***		3,485***	

We report second-stage estimates of portfolio holdings as a percentage of financial assets in 2002. The sample is restricted to households with positive holdings only. Three separate OLS regressions are run. The dependent variables are the shares of risky assets (stocks and mutual funds) over financial wealth, the share of directly-held stocks over financial wealth, and the share of risky mutual funds (equity and mixed) over financial wealth. Financial wealth is defined as the sum of cash (checking and savings accounts, money-market funds), bond-only mutual funds, stocks, and risky mutual funds. Lambda is the inverse mills ratio from the first stage estimation of equation (3.2). We report the bootstrapped standard errors. The superscripts ***, **, and * refer to coefficients statistically distinct from 0 at the 1, 5, and 10% level respectively. F refers to the Wald goodness-of-fit test. Other explanatory variables are described in Section D.1.

in wage volatility only leads to a .08% decrease in share of risky assets. Furthermore, from Table 3.6, the effect of an increase in the labor productivity variable actually leads to an increase, though not significant, in the share of risky assets.

The decomposition of risky assets into directly-held stocks and mutual funds provides some extra insight. For one, there is a clear substitution effect between stocks and risky mutual funds. While an increase in wage volatility leads to a significant increase in the share of stocks, it also leads to a similar decrease in the share of mutual funds. This result is consistent with Massa and Simonov (2006), who look at the levels of individual stock holdings and find that households' investments in stocks also come from factors other than hedging, such as a preference toward stocks they are more familiar with, for information reasons. Indeed, they argue that less-informed agents choose to invest more in stocks closely related to their labor income because they are more familiar with these stocks, via either location or professional proximity.

One can conjecture other sources of heterogeneity correlated with labor income that affect portfolio selection. For example, households in high productivity industries could be more financially-educated and choose to invest more in individual stocks and less in mutual funds. They might also be of a different type and have separate investment policies. For example, it may be that the less risk averse agents choose to work in riskier industries and invest more in the stock market. Since our cross-section analysis cannot control for these issues, we turn to our main estimation strategy and look instead at changes in the portfolio shares of the switchers, conditional on their initial portfolio shares. This analysis allows us to abstract from the potential heterogeneity in the levels of portfolio shares, and to test H4 and H5 and consider how household investment behavior changes with employment changes.

H4: Workers who switch to a sector with higher wage volatility decrease their exposure to the market through financial assets.

H5: Workers who switch to higher labor-productive sectors decrease their exposure to the market through financial assets.

As with the levels analysis, we implement a two-stage analysis where we begin by controlling for the possibility of a selection bias among market participants, with 1999

as the base year. Equation (3.4) is similar to equation (3.2), except that t now refers to year 1999. Then, in the second-stage, we retain the switchers who participated in the stock market in 1999 and study the effect of a change in $\Phi_{h,t}$ between 1999 and 2002 on their portfolio holdings (recall that $\Phi_{h,t}$ contains either wage volatility or labor productivity along with interaction variables for households where both individuals work in the same industry). The equations take the form

$$p_{h,t}^{ra} = \alpha_{3,t}^{ra} + \beta_{4,t}^{ra'} \cdot \Phi_{h,t} + \gamma_{3,t}^{ra'} \cdot X_{h,t} + \epsilon_{3,h,t}^{ra}, \quad (3.4)$$

$$\begin{aligned} \Delta w_{h,t}^i &= \alpha_{4,t}^i + \beta_{4,t}^{i'} \cdot \Delta \Phi_{h,t} + \gamma_{4,t}^{i'} \cdot X_{h,t} + \varphi_t^{i'} \cdot Y_{h,t} + \kappa_t^{i'} \cdot Z_{h,t} \\ &\quad + \theta_{4,t}^i \cdot \lambda_{4,h,t} + \epsilon_{4,h,t}^i. \end{aligned} \quad (3.5)$$

where t refers to year 1999, h indexes switchers, $X_{h,t}$, $Y_{h,t}$, and $Z_{h,t}$ are vectors of control variables, and $\Delta X_{h,t}$ refers to a change in X_h from year t (1999) to year $t+3$ (2002). As in the previous section, we expect households to hedge their labor income risk if $\beta_{4,t}^{ra} < 0$.

To control for different possible explanations, we decompose switchers into three groups: first, a group in which individual #1 switches industries, a second group in which individual #2 switches industries, and a third in which both individuals switch industries. For each group of households we also add an interaction variable that captures change in either their wage volatility or their labor productivity. These groups are not mutually exclusive. For example, the first group includes individuals #2 who are also switching. The idea behind this decomposition is to see whether (i) the hedging effect is strongest for the third group where both individuals switch during the same year, and also whether (ii) the hedging effect is stronger when (the rich) individual #1 switches than when (the less rich) individual #2 switches. We also add two interaction variables for the households who switch in such a way that they end up in the same industry in 2002 and for those who switch in such a way that they are no longer in the same industry in 2002.

In the second stage, we include the vector of controls $X_{h,t}$ (described for the previous hypotheses) as well as two other sets of control variables, which we denote as $Y_{h,t}$ and $Z_{h,t}$. These include key variables such as the initial level of net worth and the initial portfolio shares of stocks and risky mutual funds, which captures all the

information on the individuals' types under the assumption that types do not vary over time. $Y_{h,t}$ is defined as the vector of these extra controls.

In addition to employment, other household characteristics may have changed during 1999-2002. $Z_{h,t}$ is defined as the vector of these changes. These variables include a dummy on whether the household moved from a rural area to a metropolitan area, a dummy on whether at least one member of the household has died or emigrated, and a variable that computes the change in the number of children. We also look at the change in family disposable income, the change in the Debt-to-Income ratio and create dummies on whether at least one of the adults found a job, lost a job, or retired from the job market during the time period. In terms of real estate, we include two dummies on whether households started or stopped owning real estate as well as a variable that captures the change in the ratio of house value to net worth. In terms of education, we include a dummy on whether at least one of the households has graduated.¹⁴ We avoid controlling for changes in wealth during the time period since some of it comes from the proceeds of the household's portfolio holdings.

The results of equation 3.5 are reported in Table 3.7 for wage volatility and Table 3.8 for labor productivity. For parsimony we do not report the coefficients of the $X_{h,t}$ control variables in 1999. The tables with wage volatility and labor productivity have 6,428 and 1,580 switchers respectively. As expected, the effects of the 1999 levels of portfolio shares are extremely significant, which confirms the high degree of predictability in portfolio shares.

Here we find strong evidence that switchers are hedging their labor income risk. Beginning with Table 3.7, the effect of a change in the level of wage volatility is significant for the switchers, both economically and statistically. In particular, for the double-switchers, an increase in wage volatility by 1% will lead on average to a decrease in the share of risky assets by 1.07%, in absolute terms. This effect is statistically significant at the 5% level.¹⁵ We stress that this percentage is of financial wealth, which in 2002 was around 310,000 SEK (approximately 31,000 USD). This

¹⁴We define graduation as a stop in the individual's student aid.

¹⁵Although the reported p-value for this coefficient in Table 3.7 is about 6%, our test of hedging is a one-tail test, and so the relevant p-value is about 3%.

Table 3.7: Effect of a change in wage volatility between 1999 and 2002 on portfolio shares.

Variable	Δ Risky Assets		Δ Stocks		Δ Mutual Funds	
	Est.	Std. Err	Est.	Std. Err	Est.	Std. Err
Intercept	-.44	(.39)	-.97***	(.221)	.527**	(.29)
Δ wage vol.						
ind#1 switchers	-.126	(.157)	-.01	(.094)	-.115	(.149)
ind#2 switchers	.346	(.292)	.305	(.194)	.041	(.249)
double-switchers	-.1.073*	(.569)	-.322	(.38)	-.751*	(.409)
to same industry	.218	(.536)	-.177	(.297)	.395	(.495)
from same industry	-.09	(.497)	.021	(.257)	-.111	(.409)
Δ household size	.022***	(.008)	-.006	(.004)	.028***	(.007)
has graduated	-.017	(.023)	-.009	(.014)	-.008	(.02)
low to high pop. density	-.029	(.031)	.015	(.021)	-.044*	(.023)
Δ family income	-.154***	(.018)	-.013	(.01)	-.141***	(.013)
found a job	-.028*	(.016)	-.01	(.009)	-.018	(.015)
lost a job	.012	(.014)	-.001	(.007)	.012	(.012)
has retired	.008	(.016)	.0158	(.01)	-.007	(.012)
Δ debt / income	-.017***	(.003)	-.001	(.002)	-.016***	(.003)
bought house	.019	(.035)	.014	(.016)	.005	(.023)
sold house	-.086**	(.038)	-.01	(.02)	-.075**	(.031)
Δ house / networth	.054***	(.007)	-.001	(.005)	.054***	(.007)
net worth	.006*	(.004)	.006**	(.003)	.001	(.003)
stocks	-.51***	(.017)	-.464***	(.014)	-.042***	(.013)
mutual funds	-.512***	(.012)	-.01***	(.007)	-.502***	(.013)
lambda	.634***	(.164)	.37***	(.101)	.257**	(.153)
No. Obs	6,428		6,428		6,428	
F	8,907***		2,813***		5,096***	

Second-stage estimates of changes in the shares of portfolio holdings between 1999 and 2002. Three separate OLS regressions are run. The sample is restricted to households with positive holdings of risky assets in 1999. The dependent variables are the change in the share of risky assets (stocks and mutual funds) over financial wealth, the change in the share of directly owned stocks over financial wealth, and the change in the share of mutual funds (equity and mixed) over financial wealth. Financial wealth is defined as the sum of cash (checking and savings accounts, money-market funds), bond-only mutual funds, stocks, and risky mutual funds. Lambda is the inverse mills ratio from the first stage estimation of equation (3.4). We report the bootstrapped standard errors. The superscripts ***, **, and * refer to coefficients statistically distinct from 0 at the 1, 5, and 10% level respectively. F refers to the Wald goodness-of-fit test. Other explanatory variables are described in Section D.1.

Table 3.8: Effect of a change in labor productivity between 1999 and 2002 on portfolio shares.

Variable	Δ Risky Assets		Δ Stocks		Δ Mutual Funds	
	Est.	Std. Err	Est.	Std. Err	Est.	Std. Err
Intercept	.001	(.845)	-.88**	(.414)	.889*	(.507)
Δ labor prod.						
ind#1 switchers	-.32	(.205)	-.227**	(.105)	-.092	(.192)
ind#2 switchers	-.616**	(.278)	-.221	(.165)	-.395	(.284)
double-switchers	-.037	(.559)	-.159	(.43)	.121	(.548)
to same industry	.661	(.488)	.306	(.291)	.355	(.624)
from same industry	.35	(.532)	-.019	(.326)	.369	(.402)
Δ household size	.053***	(.015)	.003	(.009)	.05***	(.017)
has graduated	-.018	(.061)	.038	(.034)	-.056	(.056)
low to high pop. density	-.05	(.061)	.037	(.04)	-.087*	(.052)
Δ family income	-.193***	(.038)	.017	(.02)	-.209***	(.02)
found a job	-.006	(.031)	.025	(.017)	-.032	(.027)
lost a job	.05*	(.026)	.025*	(.014)	.026	(.031)
has retired	.006	(.03)	.01	(.018)	-.004	(.022)
Δ debt / income	-.018***	(.005)	.001	(.003)	-.018***	(.004)
bought house	-.004	(.061)	.01	(.023)	-.014	(.049)
sold house	.0126	(.078)	.032	(.037)	-.0196	(.075)
Δ house / networth	.05***	(.015)	-.013	(.011)	.063***	(.013)
net worth	.007	(.009)	.001	(.005)	.006	(.009)
stocks	-.513***	(.034)	-.44**	(.029)	-.073***	(.028)
mutual funds	-.52***	(.023)	-.002	(.012)	-.518***	(.024)
lambda	.51	(.42)	.178	(.228)	.33	(.251)
No. Obs	1,580		1,580		1,580	
F	2,280***		1,345***		1,357***	

Second-stage estimates of changes in the shares of portfolio holdings between 1999 and 2002. Three separate OLS regressions are run. The sample is restricted to households with positive holdings of risky assets in 1999. The dependent variables are the change in the share of risky assets (stocks and mutual funds) over financial wealth, the change in the share of directly owned stocks over financial wealth, and the change in the share of mutual funds (equity and mixed) over financial wealth. Financial wealth is defined as the sum of cash (checking and savings accounts, money-market funds), bond-only mutual funds, stocks, and risky mutual funds. Lambda is the inverse mills ratio from the first stage estimation of equation (3.4). We report the bootstrapped standard errors. The superscripts ***, **, and * refer to coefficients statistically distinct from 0 at the 1, 5, and 10% level respectively. F refers to the Wald goodness-of-fit test. Other explanatory variables are described in Section D.1.

means that a household going from the industry with the least variable wage (recycling metal waste) to the industry with the most variable wage (fund management) will *ceteris paribus* decrease its share of risky assets by almost 25%, or 7,750 USD. The decomposition of risky assets into stocks and mutual funds indicates that the decrease in risky assets is fairly balanced among the two asset classes.

The hedging effect is not as strong but still there for the households where individuals #1 switched. For example, an increase in wage volatility by 1% leads the individual #1 switcher to reduce its share of risky assets by almost .12%.

Table 3.8 presents similar results with labor productivity. An increase in the coefficient of labor elasticity by 1% leads switchers to decrease their share of risky assets by .32% to .61%. Again, from an economic perspective, it means that households going from the least productive industry to the most productive industry would re-balance their share of risky assets by up to almost 20%. These effects are statistically significant at the 10% and 5% levels, respectively (one-tailed t-tests). We note, however, that the change in labor productivity has little effect on the portfolios of double-switchers. This is not surprising, since the sample size is much smaller with labor productivity. There are only 45 double-switchers, instead of 208 in Table 3.7. There is also little effect for the households where the individuals switch either to or from the same industry in both tables.

One alternative potential explanation for the fact that the coefficients of the changes in wage volatility and labor productivity are negative is if wage volatility is correlated with wealth. If so, a change in wage volatility could be associated with a change in wealth, which could be the real reason for portfolio changes. We control for this potential factor by looking at the change in household income between 1999 and 2002. Supposedly, households who switch to an industry where they obtain a wage increase have become wealthier. The addition of this variable acts not only as a control but it also indicates the effect of an increase in wealth on the portfolio share of risky assets. In both Tables 3.7 and 3.8, we find that an increase in household income leads to a significant decrease in the share of risky assets. This result suggests that this other potential explanation goes the other way, hence strengthening our results.

3.6 Conclusion

The literature on labor income risk and the levels of portfolio holdings has led to mixed results. On the one hand, there is evidence that agents hedge human capital risk (Guiso, Jappelli and Terlizzese (1996), Vissing-Jorgensen (2002)). On the other hand, at the individual stock holdings level, households tend to own stocks that are closely related to their labor income (Massa and Simonov, 2006).

In this paper we take advantage of a unique Swedish panel dataset and provide a new approach to this issue by focusing on the households who switched industries between 1999 and 2002. We study the effect of their industry change – in particular the effect of changes in their wage volatility and labor productivity – on their portfolio holdings of risky assets. Focusing on changes in portfolio holdings for households who switch industries, we find that households do hedge their labor income risk, although the effects do now show in the cross section of levels of portfolio holdings. The effect is economically significant — A household that moves from the lowest to the highest productivity industry decreases its exposure to risky assets by risky by about 25%.

Our results are therefore in line with the findings of Guiso, Jappelli and Terlizzese (1996) and Vissing-Jorgensen (2002), as well as with those of Massa and Simonov (2006), and suggest that both hedging and other, offsetting, effects are important in households' portfolio decisions. If the strength of these two offsetting effects vary with the business cycle, it is not suprising that the unconditional CAPM with human capital fails (as documented by Fama and Schwert (1977)) but the conditional CAPM with human capital is successful in explaining the cross section of stock returns (as documented by Jagannathan and Wang (1996)).

Bibliography

- Abowd, J., Kramarz, F. and Margolis, D. (1999), ‘High wage workers and high wage firms’, *Econometrica* **67**(2), 251–333.
- Bizjak, J., Brickley, J. and Coles, J. (1993), ‘Stock-based incentive compensation and investment behavior’, *Journal of Accounting and Economics* **16**, 349–372.
- Blundell, R., Pistaferri, L. and Preston, I. (2006), Imputing consumption in the PSID using food demand estimates from the cex. Working paper.
- Bodie, Z., Merton, R. and Samuelson, W. (1992), ‘Labor supply flexibility and portfolio choice in a life cycle model’, *Journal of Economic Dynamics and Control* **16**, 427–449.
- Breeden, D. (1979), ‘An intertemporal asset pricing model with stochastic consumption and investment opportunities’, *Journal of Financial Economics* **7**(3), 265–296.
- Brueckner, J. (1997), ‘Consumption and investment motives and the portfolio choices of homeowners’, *The Journal of Real Estate Finance and Economics* **15**(2), 159–180.
- Calvet, L., Campbell, J. and Sodini, P. (2007), ‘Down or out: Assessing the welfare costs of household investment mistakes’, *Journal of Political Economy* **115**(5), 707–747.
- Caplin, A., Chan, S., Freeman, C. and Tracy, J. (1997), *Housing Partnerships: A New Approach to a Market at a CrossRoads*, MIT Press, Cambridge and London.

- Case, K. and Shiller, R. (1987), ‘Prices of single-family homes since 1970: New indexes for four cities’, *New England Economic Review* **79**(1), 46–56.
- Case, K. and Shiller, R. (1989), ‘The efficiency of the market for single-family homes’, *American Economic Review* **79**(1), 125–137.
- Cauley, S., Pavlov, A. and Schwartz, E. (2007), ‘Homeownership as a constraint on portfolio allocation’, *Journal of Real Estate Finance and Economics* **34**, 283–311.
- Chetty, R. and Szeidl, A. (2009), The effect of housing on portfolio choice. Working Paper.
- Cocco, J. (2005), ‘Portfolio choice in the presence of housing’, *Review of Financial Studies* **18**(2), 535–567.
- Corradin, S. (2009), Household leverage. Working paper.
- Corradin, S., Fillat, J. and Vergara-Alert, C. (2009), Optimal portfolio choice with predictability in house prices and transaction costs. Working Paper.
- Cuoco, D. and Liu, H. (2000), ‘Optimal consumption of a divisible durable good’, *Journal of Economic Dynamics and Control* **24**, 561–613.
- Damgaard, A., Fuglsbjerg, B. and Munk, C. (2003), ‘Optimal consumption and investment strategies with a perishable and an indivisible durable good’, *Journal of Economic Dynamics and Control* **28**(2), 209–253.
- Danthine, J.-P. and Donaldson, J. (2002), ‘Labour relations and asset returns’, *The Review of Economic Studies* **69**(1), 41–64.
- Davidoff, T. (2006), ‘Labor income, housing prices and homeownership’, *Journal of Urban Economics* **59**(2), 209–235.
- De Jong, F., Driessen, J. and Van Hemert, O. (2008), Hedging house price risk: Portfolio choice with housing futures. Working paper.

- Downing, C. and Wallace, N. (2001), A real options approach to housing investment. Working paper.
- Fama, E. and Schwert, W. (1977), 'Human capital and capital market equilibrium', *Journal of Financial Economics* **4**(1), 95–125.
- Favilukis, J., Ludvigson, S. and Van Nieuwerburgh, S. (2009), The macroeconomic effects of housing wealth, housing finance, and limited risk-sharing in general equilibrium. Working paper.
- Fillat, J. (2008), Housing as a measure for long-run risk in asset pricing. Working Paper.
- Flavin, M. and Nakagawa, S. (2008), 'A model of housing in the presence of adjustment costs: A structural interpretation of habit persistence', *The American Economic Review* **98**(1), 474–495.
- Flavin, M. and Yamashita, T. (2002), 'Owner-occupied housing and the composition of the household portfolio', *The American Economic Review* **92**(1), 345–362.
- Frye, M. (2004), 'Equity-based compensation for employees: Firm performance and determinants', *Journal of Financial Research* **27**(1), 31–54.
- Gaver, J. and Gaver, K. (1993), 'Additional evidence on the association between the investment opportunity set and corporate finance, dividend, and compensation policies', *Journal of Accounting and Economics* **16**, 125–160.
- Gaver, J. and Gaver, K. (1995), 'Compensation policy and the investment opportunity set', *Financial Management* **24**(1), 19–35.
- Georgarakos, D. and Pasini, G. (2009), Trust, sociability and stock market participation. Working paper NETSPAR No.04/2009-015.
- Goetzmann, W. (1993), 'The single family home in the investment portfolio', *Journal of Real Estate Finance and Economics* **6**(3), 201–222.

- Grossmann, S. and Laroque, G. (1990), ‘Asset pricing and optimal portfolio choice in the presence of illiquid durable consumption goods’, *Econometrica* **58**, 25–51.
- Guiso, L., Jappelli, T. and Terlizzese, D. (1996), ‘Income risk, borrowing constraints, and portfolio choice’, *The American Economic Review* **86**(1), 158–172.
- Gyourko, J., Saiz, A. and Summers, A. (2008), ‘A new measure of the local regulatory environment for housing markets’, *Urban Studies* **45**, 693–729.
- Heckman, J. (1979), ‘Sample selection bias as a specification error’, *Econometrica* **47**, 153–161.
- Henderson, J. and Ioannides, Y. (1983), ‘A model of housing tenure choice’, *The American Economic Review* .
- Jagannathan, R. and Wang, Z. (1996), ‘The conditional capm and the cross-section of expected returns’, *The Journal of Finance* **51**, 3–53.
- Kedia, S. and Mozumdar, A. (2002), Performance impact of employee stock options. Working Paper.
- Kraft, H. and Munk, C. (2008), Optimal housing, consumption, and investment decisions over the life-cycle. Working Paper.
- Kruze, D. (2008), ‘Profit sharing and productivity: Microeconomic evidence from the united states’, *The Economic Journal* **102**(410), 24–26.
- Liu, J. (2007), ‘Portfolio selection in stochastic environments’, *Review of Financial Studies* **20**(1), 1–39.
- Lucas, Jr., R. (1978), ‘Asset prices in an exchange economy’, *Econometrica* **46**(6), 1429–1445.
- Lustig, H. and Van Nieuwerburgh, S. (2005), ‘Housing collateral, consumption insurance, and risk premia: An empirical perspective’, *The Journal of Finance* **60**(3), 1167–1219.

- Malloy, C., Moskowitz, T. and Vissing-Jorgensen, A. (2005), Job risk and asset returns. Working Paper.
- Mankiw, N. and Zeldes, S. (1991), ‘The consumption of stockholders and nonstockholders’, *Journal of Financial Economics* **29**, 97–112.
- Massa, M. and Simonov, A. (2006), ‘Hedging, familiarity, and portfolio choice’, *The Review of Financial Studies* **19**(2), 633–685.
- Mayers, D. (1973), ‘Nonmarketable assets and the determination of capital asset prices in the absence of a riskless asset’, *The Journal of Business* **46**(2), 258–267.
- Mehran, H. (1995), ‘Executive compensation structure, ownership, and firm performance’, *The Journal of Financial Economics* **38**, 163–184.
- Merton, R. (1971), ‘Optimum consumption and portfolio rules in a continuous-time model’, *Journal of Economic Theory* **3**, 373–413.
- Ortalo-Magné, F. and Prat, A. (2009), Spatial asset pricing: A first step. Working Paper.
- Ortalo-Magné, F. and Rady, S. (2002), ‘Tenure choice and the riskiness of non-housing consumption’, *Journal of Housing Economics* **11**, 266–279.
- Pakos, M. (2007), Asset pricing with durable goods: Potential resolution of some asset pricing puzzles. Working Paper.
- Palacios-Huerta, I. (2003), ‘The robustness of the conditional capm with human capital’, *The Journal of Financial Econometrics* **1**(2), 272–289.
- Parlour, C. and Walden, J. (2009), Capital, contracts, and the cross-section of stock returns. Working Paper.
- Piazzesi, M., Schneider, M. and Tuzel, S. (2007), ‘Housing, consumption, and asset pricing’, *Journal of Financial Economics* **83**, 531–569.

- Qin, J. (2002), ‘Human-capital-adjusted capital asset pricing model’, *The Japanese Economic Review* **53**, 182–198.
- Rubinstein, M. (1976), ‘The valuation of uncertain income streams and the pricing of options’, *The Bell Journal of Economics* **7**(2), 407–425.
- Santos, T. and Veronesi, P. (2006), ‘Labor income and predictable stock returns’, *The Review of Financial Studies* **19**, 1–44.
- Siegel, S. (2005), Consumption-based asset pricing: Durable goods, adjustment costs, and aggregation. Working Paper.
- Sinai, T. and Souleles, N. (2005), ‘Owner-occupied housing as a hedge against rent risk’, *The Quarterly Journal of Economics* **120**(2), 763–789.
- Sinai, T. and Souleles, N. (2009), Can owning a home hedge the risk of moving. Working paper.
- Skinner, J. (1994), Housing and saving in the united states, *in* Y. Noguchi and J. Poterba, eds, ‘Housing Markets in the United States and Japan’, University of Chicago Press, Chicago, pp. 191–213.
- Smith, C. and Watts, R. (1992), ‘The investment opportunity set and corporate financing, dividend, and compensation policies’, *Journal of Financial Economics* **32**, 263–292.
- Sunden, A. (2006), ‘The swedish experience with pension reform’, *Oxford Review of Economic Policy* **22**(1), 133–148.
- Van Hemert, O. (2009), Household interest rate risk management. Working Paper.
- Vasicek, O. (1977), ‘An equilibrium characterization of the term structure’, *Journal of Financial Economics* **5**, 177–188.
- Vissing-Jorgensen, A. (2002), Toward an explanation of household portfolio choice heterogeneity: Nonfinancial income and participation cost structures. Working Paper.

- Wachter, J. (2002), 'Portfolio and consumption decisions under mean-reverting returns: An exact solution for complete markets', *Journal of Financial and Quantitative Analysis* **37**(1), 63–91.
- Yao, R. and Zhang, H. (2005), 'Optimal consumption and portfolio choices with risky housing and borrowing constraints', *Review of Financial Studies* **18**(1), 197–239.
- Yogo, M. (2006), 'A consumption-based explanation of expected returns', *The Journal of Finance* **61**(2), 539–580.
- Zeldes, S. (1989), 'Consumption and liquidity constraints: An empirical investigation', *Journal of Political Economy* **97**(2), 305–346.

Appendix A

Derivation of the Portfolio Choice Model in Chapter 1

A.1 Format of the Solution

For notational convenience, I eliminate the time subscript in this appendix. At time t , the value function of the homeowner V is given as,

$$V(t, W, X, P_H) = \sup_{\Gamma} \int_0^{T-t} e^{-\delta\tau} U(C_\tau, K_\tau) d\tau,$$

with boundary condition

$$V(T, W, X, P_H) = 0.$$

$\Gamma = \{\Theta, K, C\}$ is the set of admissible controls for the homeowner. Since the value function of the unconstrained agent (V^*) takes the same form, (with $\Gamma^* = \{\Theta, K, C, H\}$), I proceed through the rest of this section with the homeowner. The Hamilton-Jacobi-Bellman (HJB) equation of this problem can be written as

$$\delta V(t, W, X, P_H) = \sup_{\Gamma} \left[U(C, K) + E[dV(t, W, X, P_H)] \right]. \quad (\text{A-1})$$

I begin by expanding (A-1),

$$\begin{aligned} \delta V = \sup_{\Gamma} \left[U(C, K) + \frac{\partial V}{\partial t} + V_W E(dW) + \frac{1}{2} V_{WW} E(dW)^2 \right. \\ \left. + V_X E(dX) + \frac{1}{2} V_{XX} E(dX)^2 + V_{P_H} E(dP_H) + \frac{1}{2} V_{P_H P_H} E(dP_H)^2 \right. \\ \left. + V_{WX} E(dW dX) + V_{WP_H} E(dW dP_H) + V_{XP_H} E(dX dP_H) \right]. \quad (\text{A-2}) \end{aligned}$$

Then I plug in the processes for dW , dP_H , and dX from Section 1.2,

$$\begin{aligned} \delta V = \sup_{\Gamma} \left[U(C, K) + \frac{\partial V}{\partial t} + V_W \left(rW + \Theta'(\mu - r\mathbf{1}) - C - (K - H)\rho P_H \right) + \frac{1}{2} V_{WW} \Theta' \sigma \sigma' \Theta \right. \\ \left. + V_X \mu_X + V_{XX} \frac{1}{2} \|\sigma_X\|^2 + V_{P_H} \mu_H P_H + \frac{1}{2} V_{P_H P_H} \|\sigma_H\|^2 P_H^2 \right. \\ \left. + V_{WX} \Theta' \sigma \sigma_X + V_{WP_H} \Theta' \sigma \sigma_H + V_{XP_H} \sigma'_H \sigma_X P_H \right] \quad (\text{A-3}) \end{aligned}$$

I guess that the value function has the form

$$V(t, W, X, P_H) = \frac{1}{1 - \gamma} g(t, X)^\gamma P_H^{-(1-\beta)(1-\gamma)} W^{1-\gamma}. \quad (\text{A-4})$$

Let $\tilde{W} = W/P_H$. Since V is homogeneous, it can be reduced to $V(t, W, X, P_H) = P_H^{\beta(1-\gamma)} v(t, \tilde{W}, X)$, where $v(t, \tilde{W}, X) = \frac{1}{1-\gamma} g(t, X)^\gamma \tilde{W}^{1-\gamma}$. I re-express all the derivatives of $V(t, W, X, P_H)$ in terms of $v(t, \tilde{W}, X)$,

$$\begin{aligned} \frac{\partial V}{\partial t} &= P_H^{\beta(1-\gamma)} \frac{\partial v}{\partial t}, & V_W &= P_H^{\beta(1-\gamma)-1} v_{\tilde{W}}, \\ V_{WW} &= P_H^{\beta(1-\gamma)-2} v_{\tilde{W}\tilde{W}}, & V_X &= P_H^{\beta(1-\gamma)} v_X, \\ V_{XX} &= P_H^{\beta(1-\gamma)} v_{XX}, & V_{P_H} &= P_H^{\beta(1-\gamma)-1} v_{P_H}, \\ V_{P_H P_H} &= P_H^{\beta(1-\gamma)-2} v_{P_H P_H}, & V_{WX} &= P_H^{\beta(1-\gamma)-1} v_{\tilde{W}X}, \\ V_{WP_H} &= P_H^{\beta(1-\gamma)-1} v_{\tilde{W}P_H}, & V_{XP_H} &= P_H^{\beta(1-\gamma)-1} v_{XP_H}, \end{aligned}$$

where, with a slight abuse of notation, the derivatives of $v(t, \tilde{W}, X)$ with respect to P_H are really just functions of \tilde{W} and X ,

$$\begin{aligned} v_{P_H} &= \beta(1 - \gamma)v - \tilde{W}v_{\tilde{W}}, \\ v_{P_H P_H} &= (\beta(1 - \gamma) - 1)\beta(1 - \gamma)v - 2(\beta(1 - \gamma) - 1)\tilde{W}v_{\tilde{W}} + \tilde{W}^2 v_{\tilde{W}\tilde{W}}, \\ v_{\tilde{W}P_H} &= (\beta(1 - \gamma) - 1)v_{\tilde{W}} - \tilde{W}v_{\tilde{W}\tilde{W}}, \\ v_{XP_H} &= (\beta - 1)(1 - \gamma)v_X. \end{aligned}$$

Let $\tilde{C} = C/P_H$ and $\tilde{\Theta}_S = \Theta_S/P_H$. So $\tilde{\Theta} = (\tilde{\Theta}_S, H_t)'$. Once I insert these changes into (A-3), the term $P_H^{\beta(1-\gamma)}$ cancels out and we are left with

$$\begin{aligned} \delta v = \sup_{\Gamma} & \left[U(\tilde{C}, K) + \frac{\partial v}{\partial t} + v_{\tilde{W}} \left(r\tilde{W} + \tilde{\Theta}'(\mu - r\mathbf{1}) - \tilde{C} - (K - H)\rho \right) + \frac{1}{2}v_{\tilde{W}\tilde{W}}\tilde{\Theta}'\sigma\sigma'\tilde{\Theta} \right. \\ & + v_X\mu_X + v_{XX}\frac{1}{2}\|\sigma_X\|^2 + v_{P_H}\mu_H + \frac{1}{2}v_{P_HP_H}\|\sigma_H\|^2 \\ & \left. + v_{\tilde{W}P_H}\tilde{\Theta}'\sigma\sigma_H + v_{\tilde{W}X}\tilde{\Theta}'\sigma\sigma_X + v_{XP_H}\sigma'_H\sigma_X \right], \end{aligned} \quad (\text{A-5})$$

with boundary condition $v(T, \tilde{W}, X) = 0$. The modified HJB equation (A-5) is now independent of P_H . Before I solve for the optimal controls for both agents, I apply one additional change of variable. Let $\alpha_C = C/W$, $\alpha_K = KP_H/W$, $\alpha_H = HP_H/W$, $\alpha_{\Theta} = \Theta/W$ and $\alpha_Q = \alpha_C/\alpha_K$. It follows that $\alpha_C = \tilde{C}/\tilde{W}$, $\alpha_K = K/\tilde{W}$, $\alpha_H = H/\tilde{W}$, and $\alpha_{\Theta} = \tilde{\Theta}/\tilde{W}$. Then, (A-5) becomes

$$\begin{aligned} \delta v = \sup_{\Gamma} & \left[\tilde{W}^{1-\gamma}U(\alpha_C, \alpha_K) + \frac{\partial v}{\partial t} - v_{\tilde{W}}\tilde{W}(\alpha_C + \alpha_K\rho) \right. \\ & + v_{\tilde{W}}\tilde{W}r + v_{\tilde{W}}\tilde{W}\alpha'_{\Theta}(\bar{\mu} - r\mathbf{1}) + \frac{1}{2}v_{\tilde{W}\tilde{W}}\tilde{W}^2\alpha'_{\Theta}\sigma\sigma'\alpha_{\Theta} \\ & + v_X\mu_X + v_{XX}\frac{1}{2}\|\sigma_X\|^2 + v_{P_H}\mu_H + \frac{1}{2}v_{P_HP_H}\|\sigma_H\|^2 \\ & \left. + v_{\tilde{W}X}\tilde{W}\alpha'_{\Theta}\sigma\sigma_X + v_{\tilde{W}P_H}\tilde{W}\alpha'_{\Theta}\sigma\sigma_H + v_{P_HX}\sigma'_H\sigma_X \right], \end{aligned} \quad (\text{A-6})$$

where

$$\bar{\mu} = \begin{pmatrix} \mu_S \\ \mu_H + \rho \end{pmatrix}.$$

A.2 Solution for the Unconstrained Agent

I begin by deriving the optimal portfolio and consumption decisions of the unconstrained agent. The optimal weights and the value function v are denoted with a star superscript. The first order conditions with respect to α_C , α_K and α_{Θ} are

$$U_{\alpha_C^*} = v_{\tilde{W}}^*\tilde{W}^{\gamma}, \quad (\text{A-7})$$

$$U_{\alpha_K^*} = \rho v_{\tilde{W}}^*\tilde{W}^{\gamma}, \quad (\text{A-8})$$

$$\alpha_{\Theta}^* = -\frac{v_{\tilde{W}}^*}{v_{\tilde{W}\tilde{W}}^*\tilde{W}}(\sigma\sigma')^{-1}(\bar{\mu} - r\mathbf{1}) - \frac{v_{\tilde{W}X}^*}{v_{\tilde{W}\tilde{W}}^*\tilde{W}}(\sigma')^{-1}\sigma_X - \frac{v_{\tilde{W}P_H}^*}{v_{\tilde{W}\tilde{W}}^*\tilde{W}}(\sigma')^{-1}\sigma_H. \quad (\text{A-9})$$

where $U_{\alpha_C^*}$ and $U_{\alpha_K^*}$ are the marginal utilities of the agent with respect to the shares of his wealth allocated to non-housing consumption and housing consumption respectively. (A-9) is the vector-form expression for (1.1) and (1.2). Merging (A-7) and (A-8), and defining $\alpha_Q^* = \alpha_C^*/\alpha_K^*$, I get that α_Q^* is constant,

$$\alpha_Q^* = \frac{\rho\beta}{1-\beta}.$$

Hence, I can re-express α_C^* and α_K^* as

$$\begin{aligned}\alpha_C^* &= \left(v_{\tilde{W}}^* \tilde{W}^\gamma \beta^{-1} (\alpha_Q^*)^{(1-\beta)(1-\gamma)}\right)^{-\frac{1}{\gamma}}, \\ \alpha_K^* &= \left(v_{\tilde{W}}^* \tilde{W}^\gamma \beta^{-1} (\alpha_Q^*)^{1-\beta(1-\gamma)}\right)^{-\frac{1}{\gamma}}.\end{aligned}$$

I insert the optimal controls into (A-6). After tedious computations, it becomes

$$\begin{aligned}0 &= -v^* \left(\delta - r - \mu_H \bar{\beta} - \frac{1}{2} \|\sigma_H\|^2 \bar{\beta} (\bar{\beta} - 1) \right) \\ &\quad + \frac{\gamma}{1-\gamma} (v_{\tilde{W}}^*)^{\frac{\gamma-1}{\gamma}} \epsilon + \frac{\partial v^*}{\partial t} + v_{\tilde{W}}^* \tilde{W} (r - \mu_H) \\ &\quad + v_X^* (\mu_X + \bar{\beta} \sigma_{H_1} \sigma_{X_1} + v_{XX}^* \frac{1}{2} \|\sigma_X\|^2 \\ &\quad - \frac{1}{2} \frac{(v_{\tilde{W}}^*)^2}{v_{\tilde{W}\tilde{W}}^*} \|\lambda\|^2 - \frac{1}{2} \frac{(v_{\tilde{W}X}^*)^2}{v_{\tilde{W}\tilde{W}}^*} \|\sigma_X\|^2 - \frac{v_{\tilde{W}X}^* v_{\tilde{W}}^*}{v_{\tilde{W}\tilde{W}}^*} \lambda' \sigma_X),\end{aligned}\tag{A-10}$$

where $\epsilon = \beta^{\frac{1-\gamma}{\gamma}} (\alpha_Q^*)^{-\frac{(1-\beta)(1-\gamma)}{\gamma}}$, $\bar{\beta} = \beta(1-\gamma)$ and λ is the 2×1 vector of the market prices of risk for the shocks $[Z_1, Z_2]$.

I insert v^* , and divide (A-10) by $\tilde{W}^{1-\gamma}$, $g(t, X)^{\gamma-1}$, and $\frac{\gamma}{1-\gamma}$. The \tilde{W} term cancels out and we are left with the classic second-order, non-homogeneous, linear ordinary differential equation

$$\begin{aligned}0 &= \epsilon + \frac{\partial g}{\partial t} - g \left(\frac{\delta}{\gamma} - \frac{1-\gamma}{\gamma} \bar{r} - \frac{1-\gamma}{2\gamma^2} \|\bar{\lambda}\|^2 \right) \\ &\quad + g_X \left(\bar{\mu}_X + \frac{1-\gamma}{\gamma} \sigma_X' \bar{\lambda} \right) + \frac{1}{2} g_{XX} \|\sigma_X\|^2,\end{aligned}\tag{A-11}$$

with the boundary condition $g(T, X) = 0$ and the coefficients

$$\begin{aligned}\bar{r} &= r - \mu_H(1-\beta) - \frac{1}{2} \|\sigma_H\|^2 \beta (\bar{\beta} - 1), \\ \bar{\mu}_X &= \mu_X + \bar{\beta} \sigma_{H_1} \sigma_{X_1}, \\ \bar{\lambda} &= \lambda + (\bar{\beta} - 1) \sigma_H.\end{aligned}$$

Since r is affine in X and \bar{r} is affine in r , \bar{r} is also affine in X . Hence, the set-up of (A-11) is a special case of Liu (2007), so $g(t, X)$ has the form

$$g(t, X) = \epsilon \int_0^{T-t} e^{-\frac{\delta}{\gamma}\tau + \frac{1-\gamma}{\gamma}A_1(\tau) + \frac{1-\gamma}{\gamma}A_2(\tau)X} d\tau.$$

The coefficients $A_1(\tau)$ and $A_2(\tau)$ are the solutions to a set of ordinary differential equations,

$$\begin{aligned} A_1(\tau) &= \left(\frac{\|\bar{\lambda}\|^2}{2\gamma} \tau + \left(\bar{r} + \frac{1-\gamma}{2\gamma\kappa^2} (\sigma_r^2 - 2\kappa\sigma_r\bar{\lambda}) \right) (\tau - A_2(\tau)) - \frac{1-\gamma}{4\gamma\kappa} \sigma_r^2 A_2^2(\tau) \right), \\ A_2(\tau) &= \frac{1}{\kappa} (1 - e^{-\kappa\tau}). \end{aligned}$$

A.3 Solution for the Homeowner

Unlike the unconstrained agent, the homeowner must own what he consumes of housing. In this appendix, I formulate this constraint by assuming that the homeowner must own a fraction ϕ of the home he lives in, where $\phi = 1$. I re-express (A-6) in terms of α_H . The optimal weights for the homeowner are denoted without any superscript. The first order condition with respect to α_C , α_H and α_S are

$$U_{\alpha_C} = v_{\tilde{W}} \tilde{W}^\gamma, \tag{A-12}$$

$$\alpha_H = \alpha_H^* - \frac{1}{v_{\tilde{W}\tilde{W}} \tilde{W}^2 \sigma_{H_2}^2} (U_{\alpha_K} - v_{\tilde{W}} \tilde{W} \rho), \tag{A-13}$$

$$\begin{aligned} \alpha_S &= \alpha_S^* + \frac{\sigma_{H_1}}{\sigma_S} \frac{1}{v_{\tilde{W}\tilde{W}} \tilde{W}^2 \sigma_{H_2}^2} (U_{\alpha_K} - v_{\tilde{W}} \tilde{W} \rho), \\ &= \alpha_S^* - \frac{\sigma_{H_1}}{\sigma_S} (\alpha_H - \alpha_H^*). \end{aligned} \tag{A-14}$$

To get α_Q , I merge (A-12) and (A-13),

$$\alpha_Q = \frac{\rho\beta}{1-\beta} - \frac{\beta}{1-\beta} \frac{v_{\tilde{W}\tilde{W}} \tilde{W}}{v_{\tilde{W}}} \sigma_{H_2}^2 (\alpha_H - \alpha_H^*). \tag{A-15}$$

Instead of inserting these optimal controls back into (A-6) and deriving another ODE, I show that the homeowner's HJB can be re-expressed into the HJB of the uncon-

strained agent. To see why, I begin by re-expressing (A-6) in terms of α_Θ^* ,

$$\begin{aligned} \delta v = & \tilde{W}^{1-\gamma} U(\alpha_C, \alpha_K) + \frac{\partial v}{\partial t} - v_{\tilde{W}} \hat{W} (\alpha_C + \alpha_K \rho) \\ & + v_{\tilde{W}} \tilde{W} r - \frac{1}{2} v_{\tilde{W}\tilde{W}} \tilde{W}^2 \alpha_\Theta^{*'} \sigma \sigma' \alpha_\Theta^* + \frac{1}{2} v_{\tilde{W}\tilde{W}} \tilde{W}^2 (\alpha_\Theta - \alpha_\Theta^*)' \sigma \sigma' (\alpha_\Theta - \alpha_\Theta^*) \\ & + v_{P_H} \mu_H + \frac{1}{2} v_{P_H P_H} \|\sigma_H\|^2 + v_X \mu_X + v_{X X} \frac{1}{2} \|\sigma_X\|^2 + v_{P_H X} \sigma_H' \sigma_X. \end{aligned} \quad (\text{A-16})$$

Then I split (A-16) into two components: one component T_1 that just depends on the optimal weights of the unconstrained agent, and another component T_2 that includes all the remaining terms,

$$\delta v = T_1 + T_2,$$

where

$$\begin{aligned} T_1 = & \tilde{W}^{1-\gamma} U(\alpha_C^*, \alpha_K^*) - v_{\tilde{W}} \tilde{W} (\alpha_C^* + \alpha_K^* \rho) + \frac{\partial v}{\partial t} + v_{\tilde{W}} \tilde{W} r - \frac{1}{2} v_{\tilde{W}\tilde{W}} \tilde{W}^2 \alpha_\Theta^{*'} \sigma \sigma' \alpha_\Theta^* \\ & + v_{P_H} \mu_H + \frac{1}{2} v_{P_H P_H} \|\sigma_H\|^2 + v_X \mu_X + v_{X X} \frac{1}{2} \|\sigma_X\|^2 + v_{X P_H} \sigma_H' \sigma_X, \end{aligned} \quad (\text{A-17})$$

and

$$\begin{aligned} T_2 = & \tilde{W}^{1-\gamma} [U(\alpha_C, \alpha_K) - U(\alpha_C^*, \alpha_K^*)] \\ & - v_{\tilde{W}} \tilde{W} [(\alpha_C + \alpha_K \rho) - (\alpha_C^* + \alpha_K^* \rho)] + \frac{1}{2} v_{\tilde{W}\tilde{W}} \tilde{W}^2 (\alpha_\Theta - \alpha_\Theta^*)' \sigma \sigma' (\alpha_\Theta - \alpha_\Theta^*). \end{aligned} \quad (\text{A-18})$$

If the value functions of the homeowner and the unconstrained agent are proportional, i.e. $v = \xi v^*$, then it is easy to see that (A-17) is the HJB equation (A-6) for the unconstrained agent, i.e. $\delta v = T_1$. This implies that $T_2 = 0$. This second component only includes distortion terms, which I denote from now on with a hat. For instance, $\hat{\alpha}_H = \alpha_H - \alpha_H^*$.

I insert (A-14) and (A-12) into (A-18), and I divide all the terms by $W^{1-\gamma}$, $g(X)^\gamma$. It becomes

$$0 = \left(\frac{1}{\beta(1-\gamma)} - 1 \right) \hat{\alpha}_C - \rho \hat{\alpha}_K - \frac{1}{2} \sigma_{H_2}^2 \hat{\alpha}_H^2. \quad (\text{A-19})$$

The proportionality between v and v^* also implies that (A-15) becomes

$$\hat{\alpha}_Q = \frac{\beta}{1-\beta} \gamma \sigma_{H_2}^2 \hat{\alpha}_H. \quad (\text{A-20})$$

Equations (A-20) and (A-19) form a system of equations. I substitute for $\hat{\alpha}_H$ and show that the solution is one of the roots of a quadratic equation,

$$\hat{\alpha}_H = \frac{-\Lambda_1 - \sqrt{\Lambda_1^2 - 4\Lambda_0\Lambda_2}}{2\phi\Lambda_2} - \alpha_H^*. \quad (\text{A-21})$$

where the coefficients Λ_0 , Λ_1 , and Λ_2 are given as,

$$\begin{aligned} \Lambda_0 &= -\alpha_C^* \frac{\gamma}{\beta(1-\gamma)} - \frac{1}{2} \sigma_{H_2}^2 \gamma (\alpha_H^*)^2, \\ \Lambda_1 &= \frac{\gamma}{(1-\gamma)(1-\beta)} (\rho - \gamma \sigma_{H_2}^2 \phi \alpha_H^*), \\ \Lambda_2 &= \gamma \sigma_{H_2}^2 \phi^2 \left(\frac{\gamma}{(1-\gamma)(1-\beta)} + \frac{1}{2} \right). \end{aligned}$$

The existence of a solution relies on $\Lambda_1^2 - 4\Lambda_0\Lambda_2 \geq 0$ and $\hat{\alpha}_C > -\alpha_C^*$. In section 1.4, I find that these assumptions hold for most reasonable parameter values. It can be shown that only one of these roots lead to positive values of α_C and α_K .

Finally, the proportionality constant ξ is just the ratio of the marginal utilities over non-housing consumption between the homeowner and the unconstrained agent,

$$\xi = \frac{v_{\tilde{W}}}{v_{\tilde{W}}^*} = \left(\frac{\alpha_C}{\alpha_C^*} \right)^{\beta(1-\gamma)-1} \left(\frac{\alpha_K}{\alpha_K^*} \right)^{(1-\beta)(1-\gamma)}$$

Appendix B

A Mean-Variance Approach to the Opportunity Cost of Housing Services in Chapter 1

In order to provide intuition on what the opportunity cost of housing services means for homeowners, I consider a two-period mean-variance setting where the returns to the stock and the home are uncorrelated. I model the portfolio allocation of an agent who only consumes his wealth in the second period and who is constrained to hold a fixed value of α_H , in the spirit of Brueckner (1997) and Flavin and Yamashita (2002).¹ To remain consistent with the paper, I refer to this agent as a homeowner. The notation also remains the same. The intuition behind the results of this toy model extends to the continuous time model.

B.1 The Homeowner's Problem

Let U^{MV} be the homeowner's mean-variance utility,

$$U^{MV}(\mu_W, \sigma_W) = \mu_W - \frac{\gamma}{2}\sigma_W^2,$$

¹Parts of this appendix were derived in collaboration with Carles Vergara-Alert.

where γ is the “risk-aversion” constant, μ_W is the expected future wealth, and σ_W is the standard deviation of future wealth.

In this toy model, I assume that the returns of the stock and the home are uncorrelated (i.e., $\sigma_{H,1} = 0$) and α_H is fixed. Then, μ_W and σ_W^2 are given as

$$\begin{aligned}\mu_W &= r + \alpha_S \cdot (\mu_S - r) + \alpha_H \cdot (\mu_H + \rho - r), \\ \sigma_W^2 &= \alpha_H^2 \cdot \sigma_H^2 + \alpha_S^2 \sigma_S^2.\end{aligned}$$

To derive the budget constraint, substitute for α_S and express μ_W as a function of σ_W ,

$$\mu_W = r + \lambda_S \cdot \sqrt{\sigma_W^2 - \sigma_{H,2}^2 \alpha_H^2} + \alpha_H \cdot (\mu_H + \rho - r), \quad (\text{B-1})$$

where α_H is exogenous and λ_S is the Sharpe ratio of asset S . (B-1) is the budget constraint for the homeowner and a given value of α_H . The graph in Figure B.1 displays two budget constraints. The loosely dashed line is the budget constraint for $\alpha_H = \alpha_H^* = 0.6$ and the loosely dotted line is for $\alpha_H = 2$. The budget constraint for the unconstrained agent is the thick solid line. It is the Capital Markets Line (CML). As we can see, any budget constraint for the homeowner lies below the CML. Furthermore, for each value of α_H the homeowner’s budget constraint has a unique tangency point to the CML. It is only tangent to the CML at the optimal portfolio (σ_W^*, μ_W^*) if $\alpha_H = \alpha_H^*$.

Given the budget constraint (B-1), the homeowner maximizes his expected utility. Solving the first order conditions yields

$$\sigma_W^2 = \gamma^{-2} \lambda_S^2 + \sigma_H^2 \cdot \alpha_H^2, \quad (\text{B-2})$$

$$\mu_W = r + \gamma^{-1} \lambda_S^2 + \alpha_H \cdot (\mu_H + \rho - r). \quad (\text{B-3})$$

To-express μ_W in terms of σ_W (i.e. in our mean-variance space), combine (B-2) and (B-3),

$$\mu_W = r + \gamma^{-1} \lambda_S^2 + \lambda_H \sqrt{\sigma_W^2 - \gamma^{-2} \lambda_S^2}. \quad (\text{B-4})$$

Equation (B-4) is depicted in Figure B.1 as the densely dashed line. It is the locus of optimal portfolios for the homeowner and different values of α_H . It is tangent to the CML only at (σ_W^*, μ_W^*) . When α_H increases, the optimal values of σ_W^2 and μ_W both go up. It means that in Figure (B.1), an increase in α_H leads to a movement along (B-4) toward the North-East.

To understand why, note that given the zero correlation between the returns of the stock and the home, the share in wealth that is invested in the stock does not depend on α_H ,

$$\alpha_S = \alpha_S^* = \frac{\mu_S - r}{\gamma \sigma_S^2}.$$

This means that any movement along the locus of optimal portfolios comes from a change in the allocations to the risk-free asset and the housing. As α_H increases, the homeowner decreases the level of risk-free holdings, so his portfolio becomes riskier and more exposed to housing risk.

B.2 The Opportunity Cost of Housing Services

As I show in the paper, the opportunity cost of housing services depends on how housing affects the homeowner's portfolio. In this section, I show that it depends more specifically on how housing affects his marginal indirect utility over future wealth.


To see why, I begin by comparing the homeowner's indirect utility V^{MV} to that of the unconstrained agent V^{MV*} . V^{MV} can be derived by plugging (B-2) and (B-3) into the homeowner's utility function,

$$V^{MV} = r + \gamma^{-1} \cdot \lambda_S^2 + \alpha_H(\mu_H + \rho - r) - \frac{\gamma}{2} (\gamma^{-2} \lambda_S^2 + \sigma_H^2 \cdot \alpha_H^2).$$

The marginal utility of V^{MV} with respect to α_H is

$$V_{\alpha_H}^{MV} = (\mu_H + \rho - r) - \gamma \sigma_H^2 \cdot \alpha_H. \quad (\text{B-5})$$

Figure B.1: Mean-Standard Deviation Graph



The parameter values are $r = .02$, $\lambda_S = .4$, $\lambda_H = .3$, $\sigma_H = .2$, $\mu_H + \rho - r = .6$, and $\gamma = .25$. μ_W and σ_W are the expected value and volatility over the homeowner's future wealth. The two thin solid lines correspond to the homeowner's indifference curves. The thick solid line is the Capital Markets Line. The loosely dotted and dashed lines are the budget constraints given values of α_H of .6 and 2. The densely dashed line is the set of optimal allocations for all values of α_H for the homeowner.

Remember that the unconstrained agent can choose H_t . It is easy to show that his optimal share of wealth in housing α_H^* is

$$\alpha_H^* = \frac{\mu_H + \rho - r}{\gamma \sigma_H^2}.$$

Hence, we can re-express V^{MV} and $V_{\alpha_H}^{MV}$ in terms of whether the homeowner's portfolio is over- or under-weighted in housing,

$$V^{MV} = V^{MV*} - \frac{1}{2} \gamma \sigma_H^2 (\alpha_H - \alpha_H^*)^2, \quad (\text{B-6})$$

$$V_{\alpha_H}^{MV} = -\gamma \sigma_H^2 (\alpha_H - \alpha_H^*). \quad (\text{B-7})$$

Here, there are two interesting results. First, the difference in indirect utilities between the homeowner and the unconstrained agent in (B-6) is equivalent to the RHS of (1.8) in the paper. It says that the extent to which the homeowner is worse off

from the unconstrained agent is proportional to the square of the difference in their shares of wealth invested in housing. In the paper, this loss in indirect utility matters for the homeowner's trade-off between overall consumption and overall investment. However, it is negligible for small deviations of α_H from α_H^* .

Second, the homeowner's marginal indirect utility with respect to α_H in (B-7) is equivalent to the second term in the brackets of the RHS of (1.5). It corresponds to the homeowner's opportunity cost of housing services. In equilibrium, the relative price of housing consumption is equal to the marginal rate of substitution between non-housing consumption and housing consumption. Since the homeowner cannot separate his consumption demand from his investment demand for housing, an increase in housing consumption does not provide marginal utility just today. It also affects the homeowner's marginal indirect utility over future wealth.

The homeowner's marginal indirect utility with respect to α_H is decreasing over α_H . Furthermore, it is positive when $\alpha_H < \alpha_H^*$ and negative when $\alpha_H > \alpha_H^*$. To understand why, it is useful to look at what happens for various values of α_H . If α_H is very low, we are on the South-West part of the densely dashed line. The optimal portfolio of the homeowner is suboptimal relative to the portfolio of the unconstrained agent in that it has a lower Sharpe ratio and a lower risk exposure ($\mu_W < \mu_W^*$ and $\sigma_W < \sigma_W^*$). The marginal indirect utility with respect to α_H is positive because an extra unit of housing would push the homeowner's portfolio toward the optimal portfolio. It is also very high.

As α_H increases up to α_H^* , the homeowner's portfolio is becoming better diversified (i.e. higher Sharpe ratio) and has a greater risk exposure. $V_{\alpha_H}^{MV}$ is still positive but not as high. At $\alpha_H = \alpha_H^*$, the homeowner is at his optimal portfolio. $V_{\alpha_H}^{MV}$ is zero. From there, it becomes negative. As α_H increases past α_H^* , the Sharpe ratio falls again as the portfolio is over-weighted in housing, and the homeowner is forced to take on undesired risk exposure. An extra unit of housing pushes the portfolio away from his optimal portfolio. This effect becomes worse as α_H keeps increasing.

Appendix C

Information about the Dataset in Chapter 2

In this appendix I provide details on how I construct the PSID dataset and define variables for the empirical analysis.

C.1 Defining a Household

I formatted the PSID data so that the observation unit is a “household-year.” Given the panel nature of the dataset, defining a household can be tricky because of possible changes in its composition over time. Here are the criteria that I used:

- A household is represented by a head member. Each year, the PSID automatically selects one member of an interviewed family unit to be the head and it ranks all the other members in terms of their relationship to that person (partner, child, sibling, ...), so I follow their definition. For multiple-member units, the convention is that the head must be older than 16 years old and the person with the most financial responsibility. If that person is female and she has a partner (husband, boyfriend living in the same unit, or civil partner), then he is designated as the head, unless he is incapacitated. I add the restriction that the household head member must be older than 18 years and younger than 100 years.

- Over the years, some households survive, some no longer exist, and new households are created. Here are the criteria:
 - A household no longer exists when there is a change in the marital status of the head couple. A change can come from either divorce, separation, death, or a new partnership in the case of a head member who used to be single. In the event any of the other household members keeps being interviewed by the PSID afterwards, I consider him/her as a member of a new household. I do so to avoid capturing changes in housing consumption that come from these major life changes.
 - A new household is created when a member of an existing household who is not the head or his partner (e.g. child, sibling...) leaves and creates his or her own household.
 - A new household is also created when the PSID extends the sample to new families and interviews them for the first time.

C.2 Demographics

I define cohort and age dummies in Table C.1. I filter out households who only appear in the survey once and whose real total taxable family income is lower than \$1,000 (in 2000 dollars). For the landlords, I filter out the top .1% values of rental income earned. Values for total taxable family income were unavailable between 1994 and 1996 and in 2001, so I use linear interpolation from the first available surrounding years. Similarly, civil status was not available between 1994 and 1997.

C.3 Assets

In the years 1984, 1989, 1994, 1999, 2001, 2003, and 2005, the PSID conducted a wealth survey. Households were asked to report the market value of their assets, via the following type of question: *if you sold all [the amount in asset x that you or*

Table C.1: Definition of age and cohort groups

Birth date	Cohort group	Age	Age group
1940-	1	18-30	1
1941-1950	2	31-40	2
1951-1960	3	41-50	3
1961-1970	4	51-60	4
1971-1980	5	61-70	5
1981+	6	70+	6

anyone in your family own] and paid off anything you owed on it, how much would you have? I compute the net worth of a household as the sum of its assets minus the sum of its debts. The sum of assets includes the net value invested in stocks (including mutual funds and retirement accounts), the net value invested in bonds (including the cash value in life insurance policies), the amount of cash (including checking and savings accounts, CDs, government savings bonds or Treasury bills), the value of the household's primary house, the net value of his other real estate properties, the net value invested in farms and private businesses, and the net value invested in cars. The sum of debts include the amount remaining on non-collateralized debt (such as credit card charges, student loans, medical bills, or loans from relatives) and the mortgage on the households' primary house. For each asset, I compute its share of net worth.

Any asset with a missing value resulted in a missing value for net worth. Some assets were particularly prone to outliers, so I filter out the top .1% values of net worth as well as the .1% shares of net worth in bonds, cash, autos, and house value. I also restrict the share of net worth in total real estate to be between -4 and 4, the share of net worth in debt and cash to be less than 1, and the loan-to-value ratio on the household's primary home to be less than 1.2.

C.4 House Price Dynamics

I compute the expected value and the volatility of returns of individual house prices and then average these two moments by MSA (equal weights). Returns only include capital gains. Here, to maximize the sample size, I use the large PSID sample from 1970 to 2005. I require at least four consecutive observations for each house. I also eliminate observations where the absolute value of an individual yearly return is greater than 100%. While this computation does not take into account factors such as the property tax, maintenance costs, or depreciation, I show that heterogeneity in capital gains can already predict those households who choose to be landlords.

I also use the Wharton Land Regulation Index, which is provided by Joseph Gyourko, Albert Saiz, and Anita Summers online, as a measure of the elasticity of supply for housing. The data is available for various cities in the U.S. I aggregate them by MSA via a simple equal-weighted average.

C.5 Risk Aversion

In 1996, the PSID asked participants a series of questions to elicit their level of relative risk aversion. The first question is,

Suppose you had a job that guaranteed you income for life equal to your current, total income. And that job was [your/your family's] only source of income. Then you are given the opportunity to take a new, and equally good, job with a 50-50 chance that it will double your income and spending power. But there is a 50-50 chance that it will cut your income and spending power by a third. Would you take the new job?

Households can respond “yes,” “no,” “don’t know,” or “not-available.” Then, depending on the household’s response, the PSID asks a set of follow-up questions, where it changes the extent to which the income is cut in the bad state of the risky gamble (50%, 75%, 20%, 10%). I only retain households who participated in the 1996

interview and who responded yes or no to these questions. Given their responses, I back out for each of them a coefficient of relative risk aversion and report the results in Table C.2

Table C.2: Estimates of relative risk aversion for households in the PSID

Accept	Reject	γ
NA	10%	14.5
10%	20%	5.1
20%	33%	2.7
33%	50%	1.4
50%	75%	.6
75%	NA	.2

Note: Accept corresponds to the biggest loss in lifetime income that the household is willing to gamble. Reject corresponds to the smallest loss in lifetime income that the household is reluctant to gamble. γ is the elicited measure of relative risk aversion that I backed out.

C.6 Moves

I classify moves as being either job-related, consumption-related, or for some other reason. The descriptions of these categories are reported in Table C.3.

Table C.3: Categories of moves in the PSID

Category	Reason for moving (from the PSID)
Job	to take another job; transfer; stopped going to school
Job	to get nearer to work
Consumption	expansion of housing: more space; more rent; better place
Consumption	other house-related: want to own home; got married
Consumption	neighborhood-related: better neighborhood; go to school; to be closer to friends and/or relatives
Consumption	contraction of housing: less space; less rent
Consumption	to save money; all my old neighbors moved away; retiring (NA why)
Other	HU coming down; being evicted; armed services, etc.; health reasons; divorce; retiring because of health

C.7 Labor Occupation

The Census Bureau categorizes individuals by their labor occupation and reports the number of 16+ years old employed in various occupations for each census tract. I group some of their categories as described in Table C.4. I then divide the number of individuals in each category by the total population of the census tract.

C.8 Variable Definitions

Table 2.1: Reported demographics include the age of the household head, the family size, dummy variables on whether the household head is married, separated, or single and on whether he/she has received a high school diploma and a college degree, and his/her elicited risk aversion (see Section C.5). Reported income, wealth,

Table C.4: Categories of Labor Occupations from the Census Bureau

Category	Description from the Census Bureau
Farmers	farm workers or in forestry and fishing
Executives	executives, managers, and administrators (excl. farms)
Workers	sales workers; administrative support and clerical workers; precision production, craft, and repair workers; operators, assemblers, transportation, and material moving workers; nonfarm laborers; service workers
Technicians	professional and technical occupations

and housing variables include total taxable family income earned during the year prior to the interview and net worth (see Section C.3) in thousands of dollars, dummy variables on whether the household owns his/her primary home and whether he/she is a landlord, the reported market value of the household's primary home (in thousands of dollars) and its loan-to-value ratio (LTV), conditional on being a homeowner, as well as a dummy variable indicating whether the household has an outstanding second mortgage.

Table 2.2: I report shares of net worth invested in (i) non-retirement stocks (stocks) and (ii) retirement stocks (IRA-stocks), which include mutual funds, (iii) bonds, including the cash value of insurance policies, (iv) cash, which consists of checking, savings, and CD accounts and government savings bonds, (v) the household's primary house, (vi) other real estate owned, (vii) automobiles, and (viii) farms and private businesses. All categories except for the primary house are net of any related collateralized debt outstanding. I also report shares of net worth owed in the

form of (ix) a mortgage on the households' primary house and (x) non-collateralized debt, such as credit card charges, student loans, medical bills, or loans from relatives. Then, for landlord homeowners and non-landlord homeowners, I group some asset categories (all stocks: retirement and non-retirement holdings, all bonds: cash and bonds, all real estate: primary house and other real estate, and businesses and farms. For each category, I compute the aggregate value of the holdings of all the participants in the 2005 wave of the PSID, and then I compute the fraction of this value that is owned by landlord homeowners and non-landlord homeowners.

Table 2.3: Reported statistics include the census tract population, the proportion of White, African-American, Asian, Hispanic, and foreign-born individuals, the proportion of individuals who are farmers, executives, workers, and technicians (see Section C.7 for these definitions), the rate of non-occupied homes that are for rent or for sale (vacancy rate), the rate of non-occupied homes for seasonal, recreational, or occasional use (recreational vacancy rate), the fraction of homes that are owner-occupied, the fraction of households who resided in the same home five years prior to the census date, and the proportion of persons 16+ years old who are in the civilian work force and unemployed (unemployment rate).

Table 2.5: Moves are decomposed by categories (see Section C.6): consumption-related moves, job-related moves, and other moves. I only retain moves for which households have reported the value of their home both prior to and after the move. For each sample, I report the frequency of moves, as well as the expected value and the standard deviation of the difference (in % terms) between the first available real value of the new home and the last available real value of the old home. Information on housing improvements is only available during the years in which the wealth surveys were conducted. It measures the real amount spent on additions (above \$10,000, excluding general maintenance and upkeep) on all of the households' properties since the previous wealth survey. I also compute the amount spent on additions over the current market value of all the real estate owned (improvements / real estate value). For each sample, I report the frequency of these additions, as well as the expected value and the standard deviation of the amounts spent (conditional on having improved one's properties).

Table 2.6: Δ MSA and Δ State are dummy variables that equal one if a household moved to a new Metropolitan Statistical Area (MSA) or U.S. State. The other variables in vector $X_{i,t}$ are described in Table 2.1. I exclude homeowners who (i) are younger than 30 years, (ii) who are not part of the Survey Research Center core sample, (iii) who have a second mortgage, and (iv) who also own rental housing. I also exclude households who have missing or negative net worth in the first wealth survey prior to the interview date.

Appendix D

Information about the Dataset in Chapter 3

In this appendix I provide details on how I define variables from the LINDA dataset for the empirical analysis.

D.1 Variable Definitions

Table 3.1: Reported variables include the age of the household head (age), the number of children, the debt-to-income ratio, the house value-to-net worth ratio, household disposable income, in thousands of SEK, industry averages of wage growth volatility, the average level of wages, and unemployment rate, and household net wealth in millions of SEK (which does not include the value of real assets such as yachts etc. unless the household is subject to wealth tax. Further, net wealth does not include any retirement – tax-deductible – assets, human capital, and the values of private businesses and bank accounts for which less than SEK 100 is earned annually. All debt is included). We also report the following dummy variables which are 1 if at least one adult satisfies the criterion: unemployed, Nordic, college education, business degree, married, partnered, single parent, single, deceased, emigrated, student, lives in a high population density area (Stockholm, Gothenburg or Malmo/Lund/Trelleborg), medium population density (not a high density area but with more than 27,000 in-

habitants and more than 300,000 within 100 km), low population density (not a high or medium density area), retired, homeowner.

Tables 3.5 and 3.6: In addition to the variables described in Table 3.1, “age²” is the squared value of age (scaled by 1000), “house / networth” is the ratio of housing wealth over net worth, and “debt-to-income” corresponds to the ratio of debts to household disposable income. Both family income and net worth are in log terms. “wage vol.” is defined as the average volatility of annual returns to real disposable income across all individuals within a 3-digit SNI code who have stayed in the same 5-digit SNI code for at least 5 consecutive years between 1993 and 2003. “wage vol. same ind.” is an interaction variable that is equal to wage volatility if the two adults in the household work in the same 1-digit SNI code. “labor prod.” is the elasticity of output with respect to labor estimated via a random coefficients panel regression from the Output tables from Statistics Sweden. “labor prod. same ind” is an interaction variable that is equal to labor productivity if the two adults in the household work in the same 1-digit SNI code.

Tables 3.7 and 3.8: Explanatory variables are changes to family disposable income in logs (family income), changes to house-to-net wealth-ratio (house / networth), changes in the debt-to-income ratio (debt / income), and changes in wage volatility or labor productivity (Δ wage vol. or Δ labor prod.) for various groups: “individual#1 (2)” switcher consists of households where individual #1 (2) has switched industries between 2000 and 2001 and stayed in the same industry between 2001 and 2002, “double-switchers” consists of households where both individual #1 and individual #2 switched industries. We include interaction variables, “to (from) the same industry” consists of households where individuals switched industries in a way that they are both (no longer) in the same 1-digit SNI code in 2002. Furthermore, we include dummy variables that equal 1 if at least one in the household satisfies the criteria: moved from a low population density to a high one (low to high), stopped receiving student aid between 1999 and 2002 (has graduated), retired between 1999 and 2002 (has retired), unemployed in 1999 but not in 2002 (found job), employed in 1999 but not in 2002 (lost job), owned no real estate in 1999 but owned real estate in 2002 (bought house), and owned real estate in 1999 but owned no real estate in 2002 (sold

house). We also control for 1999 levels of net worth (logs) and shares of stocks and mutual funds.