Structural Estimation and Policy Analysis Applications in Housing and Education

by

Brant Abbott

B.A., Simon Fraser University, 2006
M.A., The University of Victoria, 2008

A THESIS SUBMITTED IN PARTIAL FULFILLMENT OF
THE REQUIREMENTS FOR THE DEGREE OF

DOCTOR OF PHILOSOPHY

in

The Faculty of Graduate Studies

(Economics)

THE UNIVERSITY OF BRITISH COLUMBIA

(Vancouver)

June 2013

Abstract

The first chapter studies the role of mortgage constraints in life-cycle housing decisions. I argue that observed Loan-to-Value (LTV) ratios contradict the view that large down-payments limit home ownership among young households. I estimate a model of life-cycle housing decisions that includes high LTV mortgages using a method of simulated moments. The model closely replicates home ownership rates over the life-cycle, but restricted models with maximum-LTV constraints cannot. Differences in estimated parameters lead to differences in the importance of credit-frictions, family composition and income risk in shaping ownership decisions.

The second chapter studies the relationships among housing consumption inequality, credit market frictions and the price of housing in the context of the recent U.S. housing boom. Loosening lending standards, falling interest rates and speculation have all been cited as potential causes of the increase in the average price of housing. I identify the relative importance of each of these causes through a structural model that is estimated using housing consumption micro data. The price of housing is an endogenous feature of the model, which explains 61.8% of actual house price growth, of which 24.8% is due to the falling real interest rate, 20.1% is due to investor speculation and the remainder is due to a loosening debt-to-income ratio constraint. The estimated model replicates the increase in housing consumption inequality observed over the time period.

The third chapter compares partial and general equilibrium effects of alternative financial aid policies intended to promote college participation. We build an overlapping generations life-cycle, heterogeneous-agent, incomplete-markets model with education, labor supply, and consumption/saving decisions. Altruistic parents make inter vivos transfers. Student labor supply and government grants and loans complement parental transfers as sources of college funding. We find that the current U.S. financial aid system improves welfare, and removing it would reduce GDP by two percentage points. Relaxation of government-sponsored loan limits would have no salient effects. The short-run partial equilibrium effects of expanding tuition grants are sizeable. However, long-run general equilibrium effects are 3-4 times smaller. Every additional dollar of government grants crowds out 20-30 cents of parental transfers.

Preface

- Chapter three of this thesis, titled ”Education Policy and Intergenerational Transfers in Equilibrium,” is co-authored work with Giovanni Gallipoli (UBC), Giovanni L. Violante (NYU) and Costas Meghir (Yale). Our contributions to this work were equal.

Table of Contents

Appendices

List of Tables

List of Figures

Acknowledgements

Thank you to my supervisors Paul Beaudry, Giovanni Gallipoli and Henry Siu, the external examiner Martin Gervais, and co-authors Costas Meghir and Giovanni L. Violante.

Chapter 1

Housing Consumption Inequality, Credit Market Frictions, and the U.S. Housing Boom

1.1 Introduction

The average price of housing in the United States increased substantially between the late 1990s and mid 2000s. For example, the Federal Housing Finance Administration's national index (normalized by CPI) increased by 53% between 1996 and 2006. The housing market instability that followed the boom has led to a great deal of interest and research into the causes of the boom. Hypotheses concerning these causes generally involve an increased demand for housing resulting from falling interest rates, loosening credit constraints or speculation about future capital gains.

The current paper develops a framework that accommodates a unified analysis of the previously disconnected hypotheses about the causes of the housing boom. The framework accommodates measurement of the relative contributions of the various causes of the housing boom, as well as an investigation of complementarities between the causes. One of the key components of the analysis is an empirical strategy to identify the importance of each cause of the housing boom. The strategy I employ involves matching the model's predictions to changes in the joint distribution of housing consumption, mortgage debt and related variables.

Relative to the literature, the current paper analyzes the possible causes of the housing boom within a unified framework, rather than individually. No work to date (that I am aware of) has analyzed a model that contains channels for interest rate reduction, credit constraint loosening and speculation within a single quantitative framework. This allows me to assess the relative contributions of these channels, and to study possible complementarity between causes.

The growing literature on the housing boom (discussed below) has provided valuable insights into the causes of house price growth, but no work to date has considered evidence from housing consumption micro data. Micro data reveal that average housing consumption (measured in square-feet) within the top quintile of household income increased by 3.6%

between 1999 and 2007, whereas it *decreased* by 5.8% within the lowest quintile of household income. The overlapping generations model of housing market equilibrium estimated in this paper endogenously generates an increase in housing consumption inequality, as well as an increase in the price of housing, when credit market parameters change.

The real interest rate is a commonly discussed cause of fluctuations in house prices. The World Bank estimates that the U.S. real prime lending rate (prime rate adjusted by GDP deflator) fell from 6.85% in 1998 to 1.49% in 2004. Glaeser, Gottlieb, and Gyourko (2010) estimate a smaller rate reduction, nearly two percent, based on ten-year treasury bill rates. Other research has connected a decrease in the mortgage risk premium to rising house prices (Favilukis, Ludvigson, and Van Nieuwerburgh, 2012). The real interest rate is the primary determinant of the user-cost (annualized price) of housing, which is the relevant price for consumption decisions. When the real interest rate falls so does the user-cost, thus a household with standard preferences would wish to increase housing consumption.

Where housing supply is relatively inelastic, increased housing demand will bid up the price of housing substantially. However, the falling real interest rate does not explain the entire rise in the average price of housing. In fact, according to analysis by Glaeser, Gottlieb, and Gyourko (2010) the real interest rate explains only 20% of house price growth between 1996 and 2006. Another set of hypotheses suggest that relaxation of credit constraints generated greater housing demand and pushed up the price of housing. Two features of the boom period related to credit constraints have been considered in the literature: innovative mortgage products and loosening lending standards.

Innovation in the types of mortgages available was a prominent feature of the period in question. Previously exotic contracts, such as interest-only and graduated payment mortgages, became very popular. For example, Barlevy and Fisher (2011) report market shares of interest only mortgages reaching 50-60% in some cities. A generalization of these types of innovations is that they allowed households with fixed incomes to take on larger loans by lowering the necessary payments. Another type of innovative mortgage, referred to as a piggy-back mortgage, allowed households to make relatively smaller down-payments when buying a home (Chambers, Garriga, and Schlagenhauf, 2009a,b).

On the topic of lending standards, attention has often focused on the so-called subprime lending boom. Subprime loans provided mortgage credit to those who would not have qualified previously due to insufficient income and/or collateral, or a poor credit score (Mian and Sufi, 2009). The market share of such loans grew very quickly, from less than 5% in the late 1990s to about 30% in 2005-06 (Coleman IV, LaCour-Little, and Vandell, 2008). Subprime lending was often localized in poor neighborhoods, and there is evidence that lax lending standards supported price growth particularly at the lower end of the real estate market (Landvoight, Piazzesi, and Schneider, 2012). One study finds that loosening credit standards can explain 53% of house price variation between 1992 and 2010 (Favilukis, Kohn, Ludvigson, and Van

Nieuwerburgh, Forthcoming).

Another common hypothesis regarding the cause of the housing boom is speculation about future capital gains. Evidence of a substantial increase in housing market activity among investors during the housing boom period is provided by Haughwout, Lee, Tracey, and van der Klaauw (2011). Models that formalize beliefs about future capital gains are provided by Barlevy and Fisher (2011) and Burnside, Eichenbaum, and Rebelo (2011).

Consideration of the rich patterns that exist in housing consumption may help validate or reject theories of the housing boom. This is because it is difficult for a model to replicate the opposite changes in housing consumption experienced by high income households compared to low income households. To illustrate these disparate changes, Figure 1.1 plots average housing consumption among the four quartiles of household income from 1993 to 2007. Simple models with standard preferences are at odds with this data because they predict that all income groups would change their consumption in the same direction in response to a price change. Reconciliation of this apparent contradiction requires careful detailing of housing consumption dynamics and the borrowing constraints that shape them. In the current paper, interest rate and borrowing constraint changes generate an increase in the price of housing, and the interaction of these elements with idiosyncratic income risk generates increased housing consumption inequality.

To understand the mechanism by which a housing boom leads to greater housing consumption inequality, it is useful to first clarify why the changes in housing consumption observed in the data may be puzzling. In a simple model, if the cost of housing rises or falls then the housing demands of all households would move in the same direction. Thus, disparity in the resulting consumption changes would have to be the result of high income households increasing their demands proportionally more than low income households.[1] However, with standard preferences (e.g. CES or Cobb-Douglas) the proportional change in demand in response to a price movement is independent of income. Thus, a simple model will not be sufficient to understand the observed disparity in housing consumption changes. A more elaborate model with credit market frictions is needed in order to understand this phenomenon.

An obvious possibility is that credit constraints prevented low income households from demanding more housing. However, simple loan-to-value or debt-to-income type constraints may not be enough. If all households are able to borrow up to a common multiple of their income, and all households spend the same fraction of their income on housing (e.g. they have CES preferences), then there is no reason low income households would be constrained and high income households would not. For example, if all households borrow the entire value of their home, and they all spend equal proportions of their income on housing, then they would all have the same debt-to-income ratio.

[1] Another possibility is that relative incomes changed; however, the distribution of income was generally quite stable over the short period of the housing boom.

However, consideration of income dynamics reveals that low income households might naturally have larger debt-to-income ratios than high income households. To understand how this could occur we must distinguish between lifetime income and current income. In a dynamic model all households would choose housing consumption to be the same fraction of their lifetime income, rather than their current income. Thus, if two households have the same lifetime income, but different current incomes, their debt-to-income ratios (based on current income) may be different. For example, households that experience an unexpected decline in their current income usually borrow in order to smooth consumption over these periods. This would imply that they have higher than normal current debt and lower than normal current income, hence their debt-to-income ratios are higher than normal. Higher than normal debt-to-income ratios make them prone to being borrowing constrained, which would force them to reduce consumption.

To generate lower average housing consumption among low income households at the end of the boom than the beginning, it has to be the case that more households are affected by the maximum debt-to-income constraint. This occurs because of a higher price of housing and a lower real interest rate, which cause households to borrow much more in order to finance the purchase of their homes. It is important to note that the relevant price for housing consumption decisions is the user-cost of housing, which is approximately the real interest rate multiplied by the price of a house. If the real interest rate falls by the same percentage that the price of housing rises, then the user cost stays constant.[2] This implies that an unconstrained household would not alter their housing consumption in response to these changes. However, because the price of housing has risen, households would need larger mortgages in order to finance the same amount of housing consumption. These larger mortgages imply that all households will naturally have larger debt-to-income ratios than before the price of housing increased. Hence, when facing an unexpectedly low current income, households are more likely to hit the maximum debt-to-income ratio constraint and be forced to downsize their housing consumption.

The theory that the reduction in average home size among low income households was driven by those who experienced adverse idiosyncratic income shocks provides a valuable restriction to help identify the estimated credit constraint parameters. Credit constraints regulate how frequently households that experience unexpected income losses hit the borrowing limit, and how much they must reduce their housing consumption when this occurs. In order for the model to realistically capture the observed disparity in housing consumption, the borrowing limits will have to be set quite precisely.

In addition to credit constraints, life-cycle features can also help identify changes in mortgage interest rates. A second feature of housing consumption data that is documented in this paper is the negative relationship between the age of homeowners and growth in home

[2]Ignoring taxes, maintenance, etc.

sizes between 1999 and 2007. Housing consumption among young households whose head is between 21 and 30 years old increased by 1.9%, while housing consumption among older households whose head is between 61 and 70 years old decreased by 1.0%. This tilting of the consumption age-profile is a natural result of falling real interest rates. Thus, the estimated fall in the real interest rate is in part identified by the observed tilting of the life-cycle consumption profile.

I employ the following method: first, I solve the steady-state of an overlapping generations model of housing market equilibrium. The inelastic supply of housing in this initial equilibrium is set so that the price per square-foot of housing equals that observed in 1999. Second, I solve a new equilibrium after a set of credit constraint and interest rate changes have been applied. The price of housing in the new equilibrium is endogenously generated by the model. Estimated preferences are held constant across the two equilibria, thus changes in the joint data density, such as the consumption changes discussed above, identify the changes in credit constraints and interest rates.

The empirical approach involves utilization of the entire joint distribution of housing consumption, housing tenure (own vs. rent), age, income, mortgage debt and loan-to-value ratio. I employ the Indirect Inference estimator of Gouriéroux, Monfort, and Renault (1993), in which I match a likelihood based approximation of the joint data density to a model generated counterpart. This involves estimating the deep structural parameters by minimizing the weighted squared deviation between auxiliary model parameters estimated from the data and those estimated from structural model simulated data. The Indirect Inference approach has been used in other literature to estimate dynamic models when maximum likelihood approaches are intractable (e.g. Johnson (2010) in the education literature).

Parameter estimates indicate that the real interest rate earned by savers fell by 1.01 percentage points, the mortgage interest premium fell by 0.24 percentage points, and the rental housing premium fell by 1.21 percentage points.[3] In terms of borrowing constraints, I estimate that the maximum debt-to-income ratio increased substantially, from 2.18 to 3.23, while that maximum loan-to-value ratio increased only by a trivial amount. The model replicates the joint distribution of variables used in estimation reasonably well. In particular, the estimated model replicates gradients of housing consumption change on income and age.

I conduct counterfactual experiments to determine the effect of each credit market change in isolation, and to assess for complementarities between various changes. These experiments indicate that the decreases in the real interest rate and mortgage premium explain 24.8% of actual (observed) house price growth. The fall in the rental premium, which may reflect speculation by housing investors, accounts for 20.1% of actual house price growth. The re-

[3]The rental housing premium is the difference between the gross return on rental housing and the real interest rate earned by savers. This premium reflects the costs of holding rental property, net of expected capital gains.

mainder of the model's explanatory power is due to loosening of the maximum debt-to-income ratio constraint. This has some explanatory power in isolation, but its larger role is through amplification of the effect of interest rate changes. The increase in housing consumption inequality is larger when interest rates fall in isolation, and smaller when the debt-to-income ratio constraint is loosened in isolation.

The remainder of this paper explores important features of housing boom data and analyzes them within a structural model. Section two documents the changes in housing consumption and mortgage variables observed over the housing boom period. Section three presents the structural model and highlights the important mechanisms for house price growth and increased housing consumption inequality. Section four describes the estimation procedure and some important computational innovations that make it feasible. Section five presents parameter estimates, goodness-of-fit measures and counterfactual experiments. Section six draws conclusion from the preceding evidence.

1.2 Housing and Mortgage Micro Data

Three types of data are presented in this section: housing consumption, mortgage finance and non-housing expenditure. Housing consumption and mortgage finance patterns are the key sources of variation used in estimation. The non-housing expenditure data are not used in estimation, rather they used as out-of-sample variation to assess the model's predictions.

Data are from the 1999 and 2007 waves of the American Housing Survey (AHS). The AHS is a biennial survey conducted by the United States Census Bureau to study the characteristics of housing units and their occupants. Sampling weights are used. Data reported in nominal dollars were normalized using the Consumer Price Index with base years 1982-84. The 1999 sample includes 34125 complete observations and the 2007 data includes 32822 complete observations. Observations with missing data on the variables of interest (incomplete observations) are assumed to occur randomly, and are excluded.

1.2.1 Housing Consumption

Housing consumption is measured by the square-footage of a household's dwelling. By this measure average housing consumption increased by a modest 22 square-feet, or 1.4%, between 1999 and 2007. This magnitude is consistent with estimates of the aggregate housing stock, which indicate an increase of 1.15% per capita between 2002 and 2007 (United States Census Bureau, 2012). Despite this small average change, some segments of the population experienced rather large changes in consumption. These changes can be summarized by two patterns: (1) the housing consumption of young households increased, while that of older households fell, and (2) the housing consumption of low income households decreased, while that of high income households increased.

Table 1.1 documents changes in housing consumption by real-income quintile.[4] Home sizes increased within the top four income quintiles, but fell by a substantial 5.8% for the lowest income quintile. While the percentage increases within the second to fourth quintiles are closer to the average increase, the increase within the fifth income quintile is more than double the average increase. This pattern indicates an important reallocation of housing from low to high income households over the housing boom period.

Table 1.2 documents changes in housing consumption by age groups, which are blocked into ten year increments from 21 to 70. While the level change is non-monotonic in age, the percentage change of housing consumption steadily falls over the life-cycle. Households with a head in their twenties consumed 1.9% more housing on average in 2007 than in 1999. This declines to a 1.0% decrease for households whose head is in their sixties.

1.2.2 Historical Housing Consumption

An important question one might ask is whether the differences in housing consumption between 1999 and 2007 are truly associated with the boom period, or are part of a long-term trend. To address this question, Figure 1.1 plots average home size within each income quartile relative to housing consumption in 1993, along with the (real) FHFA national house price index. The deviation of each housing consumption path represents the percentage by which average home size within that income quartile is above or below the 1993 level at each point in time. From 1993 to 1999 average home size grew steadily and there was regularity across income groups in the growth rate. However, after 1999 growth rates of housing consumption begin to show great disparity. Housing consumption growth accelerates for the highest income households, stagnates (on average) for middle income households, and becomes negative for low income households. Comparison of relative housing consumption paths with housing prices indicates an association between the acceleration of house price growth after 1999 with increasing housing consumption inequality. This association is consistent with the mechanism of the structural model, in which rising house prices lead to more frequent financial distress among households that experience periods of unexpectedly low income.

Table 1.3 provides historical trends in the levels of housing consumption. It shows that from the early to late 1990s housing consumption generally increased, and only after 1999 did it begin to fall for low income households. Both low income and older households increased their housing consumption by modest amounts between 1993 and 1999. The only groups for whom housing consumption fell slightly from 1993 to 1999 are middle income and middle age households.

[4] Boundaries are computed using combined 1999 and 2007 data, and thus are held constant across samples. If quintile boundaries had been computed within sample years instead, they would have been nearly identical.

Income Quintile	1999	2007	Difference	SE(Difference)	Percent Change
1	1257	1184	-73	(15.0)	-5.8%
2	1383	1379	-4	(13.6)	-0.3%
3	1523	1541	18	(12.9)	1.2%
4	1750	1762	12	(12.7)	0.7%
5	2124	2200	76	(14.1)	3.6%

Notes: Table 1.1 documents changes in housing consumption by income group over the housing boom period. These data show that the changes associated with the housing boom caused low income households to substantially decrease their consumption of housing, and also caused high income households to increase their consumption at a historically high rate.

Table 1.1: Housing consumption (in square-feet) by income

Age Group	1999	2007	Difference	SE(Difference)	Percent Change
21-30	1201	1224	22.9	(13.1)	1.9%
31-40	1585	1597	11.3	(13.4)	0.7%
41-50	1726	1729	3.5	(14.0)	0.2%
51-60	1768	1755	-12.9	(15.3)	-0.7%
61-70	1735	1717	-17.9	(18.2)	-1.0%

Notes: Table 1.2 documents changes in housing consumption by age group over the housing boom period. These data show that the changes associated with the housing boom caused young households to consume more housing than previously, and also caused older households to have lower housing consumption than previously.

Table 1.2: Housing consumption (in square-feet) by age

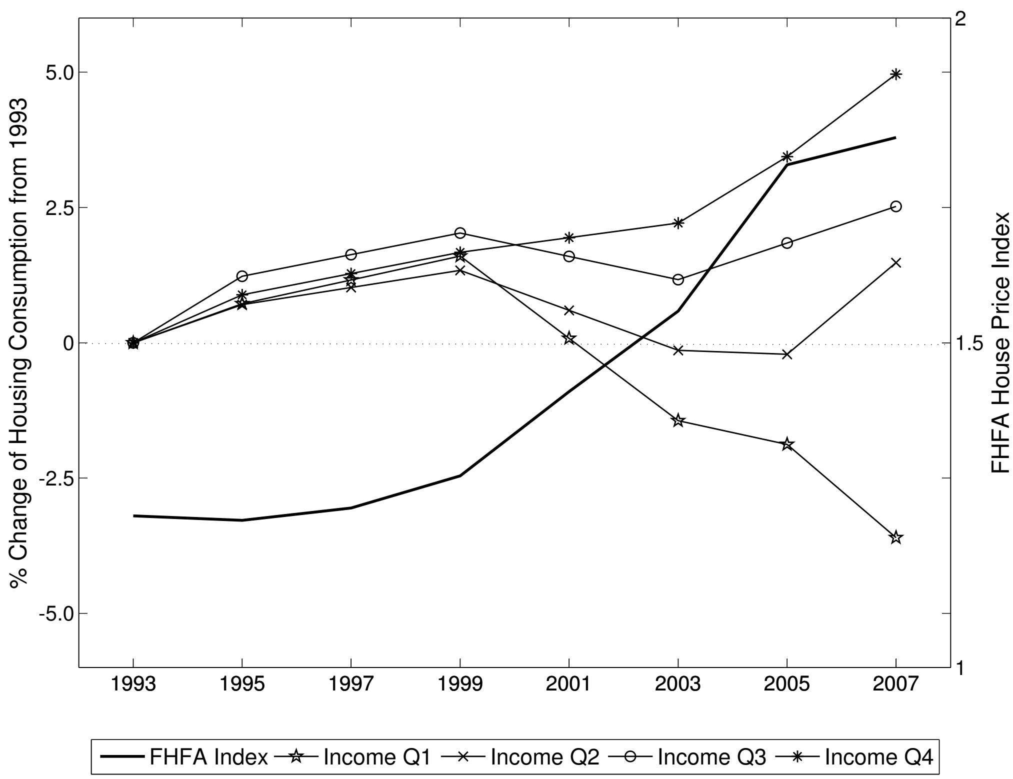

Notes: Figure 1.1 plots the changes in housing consumption compared to 1993 for the four quartiles of household income. Also plotted is the path of the real price of housing, as measured by the Federal Housing Finance Administration's Index (normalized by CPI). From 1993 to 1999 housing consumption grows steadily for all income groups, and house prices grow modestly. After 1999 growth in the price of housing accelerates, and disparity in housing consumption growth begins to appear. For high income households, housing consumption growth accelerates over the boom period, whereas for low income households, housing consumption shrinks.

Figure 1.1: Historical housing consumption inequality and house prices

	1993	1999	2003	2007
Income Quintile				
1	1229	1257	1226	1184
2	1355	1383	1348	1379
3	1547	1523	1512	1541
4	1707	1750	1739	1762
5	2077	2124	2137	2200
Age Group				
21-30	1199	1201	1212	1224
31-40	1563	1585	1586	1597
41-50	1763	1726	1724	1729
51-60	1749	1768	1749	1755
61-70	1669	1735	1720	1717

Notes: Table 1.3 provides historical details on housing consumption by age and income. For most income and age groups housing consumption increased steadily from 1993 to 1999. After 1999 increasing inequality in housing consumption begins to grow. Low income households begin to decrease their housing consumption, and high income households accelerate the growth of their housing consumption.

Table 1.3: Housing consumption (in square-feet) trends by income and age

1.2.3 Mortgages

American Housing Survey respondents are asked about the origination characteristics of their mortgage, and whether the mortgage origination coincided with the purchase of their home. Therefore, the data include the mortgage principal of a home buyer, which can be used to compute the loan-to-value ratio of a buyer when combined with data on the purchase price. The sample is restricted to purchases made within two years of the survey date, thus the 1999 and 2007 samples consider originations from 1997-98 and 2005-06, respectively. Furthermore, only mortgages associated with the purchase of a home are used, which encompasses 96% of observations.

Average initial mortgage principal increased substantially for all age groups, as evidenced by Table 1.4. Younger households appear to have increased their borrowing by more than relatively older households, but the pattern is not entirely clear. In contrast, the pattern of borrowing percentage changes by income quintile is very clear. The average increase in initial mortgage principal for the lowest income quintile is nearly double the average increase for highest income quintile.

Data on LTV ratios by income and age groups are presented in Tables 1.5 and 1.6. These data add very little information to the details presented by Glaeser et al. (2010). The reason presenting trends by age and income does not add new information is that it is generally very difficult to conclude that any changes occurred in this variable. The median appears to have risen significantly between 1999 and 2007, which is consistent with Glaeser et al.; however, their more detailed year-by-year analysis indicates that this may be noise rather than a trend.

1.2.4 Supplementary Consumer Expenditure Data

Although estimation is based on American Housing Survey data, and housing consumption is the main focus, the estimated model also has implications for consumption of non-housing goods. These predictions will be compared against supplementary Consumer Expenditure Survey (CEX) data as an out-of-sample test of the model. Furthermore, expenditure on housing in the CEX can be used to verify the housing consumption patterns observed in the AHS. Expenditure data, taken directly from the CEX annual reports, are normalized by the Consumer Price Index with base years 1982-84.

Table 1.7 reports average housing and non-housing consumption expenditure by income quintile. The striking feature of this data is the broad shifting of consumption expenditure towards housing. Indeed, non-housing expenditure fell for all but the highest income quintile. The differential amounts by which housing expenditure increased across income groups is consistent with the observed changes in actual housing consumption: within the lowest income quintile average housing expenditure increased by the smallest percentage, only 6.4%, while in the top income quintile housing expenditure increased by 12.1%.

	1999	2007	Difference	SE(Difference)	% Change
Income Quintile					
1	34295	55871	21567	(3635)	62.9%
2	39739	60028	20289	(2165)	51.1%
3	50704	69801	19097	(2122)	37.7%
4	63196	85437	22241	(2120)	35.2%
5	84514	112706	28192	(2235)	33.4%
Age Group					
21-30	52280	77198	24918	(2055)	47.7%
31-40	64898	90192	25294	(2054)	39.0%
41-50	63994	90994	27000	(2591)	42.2%
51-60	63995	77935	13940	(3224)	21.8%
61-70	52690	72567	19877	(5080)	37.7%

Notes: Table 1.4 shows how mortgage debt grew during the housing boom period by income and age. We see that low income households increased their mortgage debt more than high income households, but there is no clear pattern across age groups.

Table 1.4: Original mortgage principal by age

Income Quintile	1999	2007	Difference	se(Diff)
Averages				
1	0.831	0.850	0.019	(0.026)
2	0.862	0.849	-0.012	(0.015)
3	0.852	0.851	-0.001	(0.013)
4	0.855	0.864	0.009	(0.010)
5	0.830	0.844	0.015	(0.009)
Medians				
1	0.958	0.971	0.013	(0.012)
2	0.956	0.984	0.028	(0.006)
3	0.935	0.975	0.040	(0.007)
4	0.922	0.950	0.028	(0.006)
5	0.883	0.904	0.021	(0.008)
90th Percentiles				
1	1.009	1.026	0.017	(0.008)
2	1.000	1.019	0.019	(0.010)
3	1.000	1.000	0.000	(0.007)
4	1.000	1.000	0.000	(0.006)
5	0.984	1.000	0.016	(0.028)

Notes: Table 1.5 presents details of the distribution of loan-to-value ratios across income groups. In general, loan-to-value ratios changed very little over the boom period.

Table 1.5: Loan-to-Value ratio by income

Age Group	1999	2007	Difference	se(Diff)
Averages				
21-30	0.833	0.837	0.004	(0.017)
31-40	0.860	0.845	-0.015	(0.012)
41-50	0.865	0.864	-0.001	(0.008)
51-60	0.856	0.863	0.007	(0.006)
61-70	0.818	0.835	0.017	(0.006)
Medians				
21-30	0.958	0.989	0.031	(0.004)
31-40	0.927	0.965	0.038	(0.005)
41-50	0.905	0.935	0.030	(0.008)
51-60	0.891	0.894	0.003	(0.013)
61-70	0.871	0.892	0.021	(0.020)
90th Percentiles				
21-30	1.008	1.015	0.007	(0.006)
31-40	1.000	1.007	0.007	(0.012)
41-50	1.000	1.011	0.011	(0.008)
51-60	1.000	1.009	0.009	(0.007)
61-70	1.000	1.000	0.000	(0.004)

Notes: Table 1.6 presents details of the distribution of loan-to-value ratios across age groups. In general, loan-to-value ratios changed very little over the boom period.

Table 1.6: Loan-to-Value ratio by age

An important aspect of the housing expenditure part of the CEX data is that the pattern of expenditure changes by income mirrors the changes in housing consumption by income in the AHS. That is, housing expenditure increased much more for high income households than low income households. To see how this could validate the AHS data, suppose that an increase in housing expenditure of 9% was required in order for a household to keep their housing consumption constant over the boom. Households in the bottom quintile of income increased their expenditure by less than 9%, thus their housing consumption would have fallen; however, high income households increased their consumption by more than 9%, thus their housing consumption would have increased.

1.3 Structural Model

1.3.1 Demographics and Preferences

Life Cycle

The economy is populated by J overlapping generations. Each cohort consists of a continuum of agents. Let $j = 1, \ldots, J$ denote age. These agents receive endowment income for J^R periods, and then a pension for the next $J - J^R$ years of retirement. Mortality risk is present during retirement: age j agents survive for one more year with probability ψ_{j+1}. Each cohort initially consists of a unit mass of agents, but because of mortality the mass of each retired cohort is only Ψ_j. If agents reach age J they will die with probability one at the end of the period.

Commodities and Prices

There are two types of goods: a non-durable good, which is the numeraire, and a housing good, the relative price of which is p. Non-durable consumption is denoted c and housing consumption, measured in square-feet, is denoted h. Because the housing good is durable it has a second role as an asset whenever it is owned, as opposed to rented. The stock of housing owned by a particular household is denoted s. Finally, b denotes a household's quantity of financial assets, where a positive quantity represents ownership of savings bonds, and a negative quantity represents mortgage debt.

Endowments

Over the course of working life each household i receives a stream of endowments of the non-durable good, $\{y_{ij}\}_{j=1}^{J-J^R}$, which is the product of a deterministic age-specific part, $\overline{y}_j$, and a stochastic idiosyncratic part. The logarithm of the idiosyncratic part, z_{ij}, is modeled as in

Storesletten, Telmer, and Yaron (2004) as follows:

$$z_{ij} = \alpha_i + \eta_{ij} + \epsilon_{ij}, \tag{1.1}$$

where

$$\eta_{ij} = \rho\eta_{ij-1} + \nu_{ij}. \tag{1.2}$$

The α term is a permanent fixed effect component, where $\alpha_i \sim N(0, \sigma_\alpha^2)$. The time-variant shocks are also normally distributed draws: $\epsilon_{ij} \sim N(0, \sigma_\epsilon^2)$ and $\nu_{ij} \sim N(0, \sigma_\nu^2)$.

Annuity Markets and Initial Assets

Perfect annuity markets are assumed during retirement periods, and no bequests or intergenerational transfers are given. As a result, agents begin their lives with zero assets. The housing wealth of expiring households is sold and transferred to surviving households as financial assets. Non-convex adjustment costs do not apply to these sales.

Preferences

Preferences are composed of the present discounted value of utility from non-durable and housing consumption:

$$U = \sum_{j=1}^{J} \beta^j \Psi_j \frac{1}{1-\sigma} \left(\left[\theta c_j^\tau + (1-\theta) h_j^\tau \right]^{\frac{1}{\tau}} \right)^{1-\sigma}. \tag{1.3}$$

As usual, $\beta \in (0, 1)$ is the discount factor. Households also discount by cumulative survival probabilities Ψ_j, which are less than unity during the retirement stage. Inter-temporal risk aversion is given by the CRRA parameter σ, while the intra-temporal elasticity of substitution is controlled by τ. The weights on consumption goods are given by θ.

1.3.2 Markets

Financial Market

Households have access to one-period savings bonds, which pay a net return equal to the real interest rate, r. Households can borrow through mortgages, for which they pay the real interest rate plus a premium, r^m. Because of the mortgage premium, households always extinguish mortgage debt before buying savings bonds. Mortgage borrowing is subject to two credit constraints: the loan-to-value ratio cannot exceed ξ, and the debt-to-income ratio cannot exceed ζ. Financial market equilibrium is not modeled, thus r, r^m, ξ and ζ are exogenous parameters to be estimated.[5]

Housing Market

Demand in the housing market is due to the consumption decisions of both owners and renters. Aggregate housing supply is inelastic at the amount $\overline{h}$. At the equilibrium price of housing, p, the sum of demands from renters and owners equals the supply. A non-convex adjustment cost applies to all housing sales. The investors who own rental housing are part of the capital market. These investors require a gross return on rental housing that is the real interest rate plus a premium, r^h. Thus, a renter pays $(r + r^h)ph$ to rent h square-feet of housing. The fraction of the housing stock that is rented is an endogenous feature of the model. Given the exogenous credit market parameters, rental premium and price of housing, households optimally choose whether to rent or own. The only equilibrium restriction is that aggregate housing demand, which includes renters and owners, equals aggregate supply. The estimation routine matches the home-ownership rate in the model to the home-ownership rate in the data, mostly through adjustment of the rental premium.

Home sellers incur an adjustment cost ϕ percent of the value of the sold property, which can be interpreted as a real estate commission, although it also encompasses other costs. The non-convex nature of these costs results in lumpy housing adjustments with households increasing their housing consumption infrequently and by large amounts.

Changes in the parameter r^h can be interpreted as reflecting changes in investors beliefs about future capital gains. Suppose landlords require a return $r+\tilde{r}$ to invest in rental housing, and some of that return comes from dividend payments from renters and the remainder from capital gains. The expected capital gain is simply $\tilde{r} - r^h$, thus if r^h falls this indicates that expected capital gains have increased.

1.3.3 Household Decision Problems

A household's decision problem consists of two parts: a discrete choice between renting and owning, and continuous choices for non-durable and housing consumptions.

At any age j the value functions for the discrete choices are $V_j^{rent}(s_{j-1}, b_{j-1}, y_j)$ for renting, and $V_j^{own}(s_{j-1}, b_{j-1}, y_j)$ for owning. It is important to note that renting and owning reflect the current choice, not the previous. Whether the household rented or owned in the previous period affects the current value functions only through the amount of housing wealth the

[5] Fluctuations in interest rates and lending standards over the boom period are often considered to be the result of exogenous fluctuations, rather than an equilibrium fluctuation in a closed U.S. economy. International capital flows contributed to interest rate changes, while regulation and securitization are often cited as the reason for relaxed lending standards (Favilukis, Ludvigson, and Van Nieuwerburgh, 2012). Thus, taking these parameters as given and estimating their values is both tractable and reasonable.

household entered the period with (s_{j-1}). The overall value function for a household is:

$$V_j(s_{j-1}, b_{j-1}, y_j) = \max\left\{V_j^{rent}(s_{j-1}, b_{j-1}, y_j), V_j^{own}(s_{j-1}, b_{j-1}, y_j)\right\}. \tag{1.4}$$

The decision rule for the housing tenure (rent or own) decision is denoted $d_j(s_{j-1}, b_{j-1}, y_j)$, the value of which is unity for an owner and zero for a renter. The continuous part of a household's optimization problem is presented recursively. I distinguish the problem of those who rent from those who are home-owners.

Value of Owning:

Households who will own their current dwelling solve the following problem:

$$V_j^{own}(s_{j-1}, b_{j-1}, y_j) = \max_{c_j, h_j}\left\{u(c_j, h_j) + \psi_{j+1}\beta E_{y_{j+1}}\left[V_{j+1}(s_j, b_j, y_{j+1})\right]\right\}, \tag{1.5}$$

$$s.t.$$

$$y_j + \frac{1}{\psi_j}\left(b_{j-1}(1+\hat{r}) + ps_{j-1}\right) = c_j + ph_j + b_j + \mathbf{1}_{(h_j \neq s_{j-1})}\phi p s_{j-1}$$

$$\frac{-b_j}{ph_j} \leq \xi$$

$$\frac{-b_j}{y_j} \leq \zeta$$

$$s_j = h_j$$

$$\hat{r} = \begin{cases} r, & b_{j-1} \geq 0 \\ r + r^m, & b_{j-1} < 0 \end{cases}.$$

The first equation is the standard Bellman Equation representation of the optimization problem. The future value is discounted by ψ_{j+1}, which will equal one during working ages and be less than one during retirement. Mortality risk also enters the second equation, which is the budget constraint, because of perfect annuity markets during retirement. During retirement each household's assets are inflated by ψ_j^{-1} to account for redistribution of the wealth of expired households. During working ages this term disappears because $\psi_j = 1$. The right hand side of the budget constraint potentially includes adjustment costs. The indicator function $\mathbf{1}_{(h_j \neq s_{j-1})}$ is activated if this household has moved to a new home, and thus they pay ϕ percent of the value of their previous home. If they were previously renters then $s_{j-1} = 0$ and there is no adjustment cost.

The second and third constraints are the credit constraints. The second limits households' loan-to-value ratios to be smaller than ξ, and the third limits households' debt-to-income ratios to be smaller than ζ. The fourth constraint indicates that housing consumption will equal housing wealth because this is a homeowner. The last constraint indicates that the interest rate will vary depending on whether the financial position is positive or negative.

Value of Renting

Households who rent solve the following problem:

$$V_j^{rent}(s_{j-1}, b_{j-1}, y_j) = \max_{c_j, h_j} \left\{ u(c_j, h_j) + \psi_{j+1} \beta E_{y_{j+1}} \left[V_{j+1}(s_j, b_j, y_{j+1}) \right] \right\}, \tag{1.6}$$

$$s.t.$$

$$y_j + \frac{1}{\psi_j} \left(b_{j-1}(1+r) + p s_{j-1} \right) = c_j + (r + r^h) p h_j + b_j + \phi p s_{j-1}$$

$$b_j \geq 0$$

$$s_j = 0.$$

Mortality risk enters the Bellman Equation and budget constraint of a renter's problem in the same way it enters an owner's. Non-convex adjustment costs enter the budget constraint of a renter without an indicator function. This is because anyone who owned their home the previous period must now sell it to become a renter. For anyone who rented in the previous period $s_{j-1} = 0$, so adjustment costs do not apply. For renters the credit constraints simplify to a non-negative asset constraint because they have no collateral to allow them to borrow. The last constraint indicates that the rented home is not an asset, rather it only provides a consumption flow, hence $s_j = 0$.

1.3.4 Equilibrium

To simplify the definition of equilibrium it is useful to denote the state vector (s_{j-1}, b_{j-1}, y_j) by $\mathbf{x}_j$, the state space at each age by X_j, and the overall state space by $\mathbf{X}$.

A stationary housing market equilibrium is a collection of household decision rules for consumption, housing consumption, assets, and housing tenure $\{c_j(\mathbf{x}_j), h_j(\mathbf{x}_j), b_j(\mathbf{x}_j), d_j(\mathbf{x}_j)\}_{j=1}^{J}$, a price of housing p, and a set of age-specific measures μ_j such that:

1. Household decision rules $\{c_j(\mathbf{x}_j), h_j(\mathbf{x}_j), b_j(\mathbf{x}_j), d_j(\mathbf{x}_j)\}_{j=1}^{J}$ solve the respective household optimization problems (5) and (6) given the price of housing p,

2. The housing market clears at the price p,

$$\sum_{j=1}^{J} \int_{X_j} h_j(\mathbf{x}_j) d\mu_j = \overline{h}, \tag{1.7}$$

3. The distribution of households over the state space is invariant: $\mu(\mathbf{X}) = Q(\mathbf{X}, \mu)$, where $\mu = (\mu_1, \ldots, \mu_J)$, and Q is a transition function generated by individual decision rules and survival probabilities $\{\psi_j\}$.

1.3.5 Price and Inequality Mechanisms

Because of non-convex adjustment costs and credit constraints, simple Euler equations describing housing consumption cannot be derived in general. However, abstracting from those frictions and considering the optimality conditions that would hold in a simplified model provides intuition for the price adjustment and inequality generating mechanisms.

First consider the decision problem of a household who knows with certainty their future endowments and faces no adjustment costs or credit frictions. Representing an individual i's (known) lifetime income by M_i, one can easily show that their housing demand would be

$$h_i = \left(\frac{1}{rp}\right)^{\frac{1}{1-\tau}} \alpha M_i, \tag{1.8}$$

where α is a constant composed of preference parameters. Notice that the relevant price is the user-cost of housing, rp, rather than the full price. To see how the price of housing is affected by the interest rate it is useful to transform equation (8) into an aggregate equilibrium condition by summing over all households. In this case the left hand side of (8) can be replaced by $\overline{h}$, and the right hand side would include aggregate lifetime income. The critical feature of the aggregate version is that the multiplier, which includes the price and interest rate, is the same as in equation (8). It is straightforward then to see that any percentage decrease in the interest rate will necessitate an equal percentage increase in the price of housing to maintain equilibrium.

Why is the relationship between the interest rate and the price of housing non-trivial in the full model? In the full model aggregation of demand functions is not straightforward. Firstly, there is heterogeneity in the user-cost of housing because the relevant rate might be $r + r^m$, $r + r^h$, or simply r, which depends on owner or renter status, and whether or not an owner has a mortgage. If it turns out that $r^h > r^m$, then a fall of the real interest rate implies a larger percentage decrease in the user cost of housing for an owner than a renter, and hence a reallocation towards owners will occur, as well as a housing price adjustment.

The relationship between interest rates and the price of housing is further complicated if credit frictions and uncertainty are reintroduced. In this case there will be a fraction of the population for whom credit constraints are binding, and as a consequence housing demand is independent of the interest rate. In the maximization problem (5), binding borrowing constraints provide expressions that pin down future debt. These expressions imply a demand for housing that is independent of the interest rate. The simplest case is when both borrowing constraints bind, in which case housing demand is

$$h_i = \frac{\zeta}{\xi}\frac{y_i}{p}. \tag{1.9}$$

This expression is easily derived by making both of the credit constraints in (5) strict equalities. If the interest rate falls there will be no direct effect on such a household's demand. However, if the effect on unconstrained households causes the price of housing to increase, this will indirectly reduce constrained households' demands. Thus a reallocation from constrained to unconstrained households will occur. Returning to equation (8), this reallocation will imply that the price of housing will rise by a smaller percentage than the interest rate fell in order to maintain equality. Importantly, the elasticity of substitution between housing and non-housing consumptions regulates how much the user-cost of housing should ultimately adjust when the consumption of housing rises.

The reallocations of housing that make price adjustments non-trivial are also the source of changes in housing consumption inequality. The largest source turns out to be reallocations from households who have demands similar to equation (9), to households who have demands similar to (8). Connecting this back to the data, households who have demands similar to (9) are predominately those who have experienced particularly bad income shocks, and hence tend to be near the bottom of the income distribution.

1.3.6 Parameterization

Exogenous Parameters

The number of cohorts is set to $J = 60$, where the final year corresponds to age 80, and the first to age 21. Retirement begins at age $J^R = 46$. Several structural parameters are pre-specified, rather than estimated. The preference parameters β and σ are set to 0.965 and 2, respectively. The survival probabilities ψ_j are set such that a measure of size 1/15 dies each retirement period, which results in death of the last surviving agents after the 15th period of retirement. This parameterization also captures the increasing rates of mortality with age observed in actual survival data. Lastly, the aggregate housing supply, $\overline{h}$, is set in the initial equilibrium so that the equilibrium price of housing is \$43.69 per square-foot, which is the average (real) price in the AHS for a house purchased in 1999. For the second equilibrium $\overline{h}$ is set to be 1.4% larger than the initial equilibrium to account for growth in average housing consumption over the period.

Income Process

Recall that log income is modeled as the sum of a deterministic age profile, $\overline{y}_j$, and a random component z_{ij}. The parameters of the random part are set using the estimates of Storesletten, Temler and Yaron (2000). Specifically, the over-identified estimates from Panel D of Table 1 in that paper are used: $\rho = 0.977$, $\sigma_\alpha = 0.244$ and $\sigma_\nu = 0.024$. The age profile is estimated by fitting a cubic to real income data from the AHS sample being used:

$$\overline{y}_j = 5.56 + 0.254 \times (j + 20) - 0.00424 \times (j + 20)^2 + 0.0000233 \times (j + 20)^3. \tag{1.10}$$

The white noise component of idiosyncratic shocks is assumed to be measurement error, and thus is ignored.

1.4 Indirect Inference Structural Parameter Estimation

1.4.1 Technical Details

The structural parameters are estimated by Indirect Inference (see Gouriéroux, Monfort, and Renault (1993), and also McFadden (1989)). This involves estimation of an auxiliary model using both actual data and model simulated data. The parameters of the auxiliary model are a k-vector $\mathbf{b}$. The estimated auxiliary model parameters estimated from the T data observations are denoted $\hat{\mathbf{b}}_T$, and those estimated from H model simulations of T observations are denoted $\tilde{\mathbf{b}}_{HT}$. The important detail is that the auxiliary parameters estimated from model simulated data are functions of the vector of deep structural parameters, γ, i.e. $\tilde{\mathbf{b}}_{HT}(\gamma)$. The estimator then minimizes a quadratic form as follows:

$$\widehat{\gamma} = \arg\min_{\gamma} \left[\hat{\mathbf{b}}_T - \tilde{\mathbf{b}}_{HT}(\gamma)\right]' W \left[\hat{\mathbf{b}}_T - \tilde{\mathbf{b}}_{HT}(\gamma)\right], \tag{1.11}$$

where W is a positive definite weighting matrix. For simplicity W is chosen to be an identity matrix. A consistent estimator for the standard errors is provided in Gouriéroux et al. (1993).

1.4.2 Data and Auxiliary Model

The intention of the auxiliary model is to represent the joint distribution of housing consumption, housing tenure, income, age, mortgage principal and loan-to-value ratio. To accommodate this in a tractable fashion the data were categorized.[6] For income, house size and mortgage principal, intervals based on quintile boundaries were used to generate five categories. For age, evenly spaced twelve year increments were used to make five groups. For loan-to-value ratio to following groupings were adopted: (0,65],(65,80],(80,91],(91,100],(100,∞). Quintiles could not be used for loan-to-value ratio because of data clustering: over 15% of the sample has a loan-to-value ratio of exactly 100%, and thus this value spans roughly the 70th to the 85th percentile. Housing tenure is naturally a discrete variable.

[6]Housing tenure refers to ownership status and moving characteristics. Here three possibilities are considered: renter, owner who bought within the past year, and owner who bought previous to the last year.

A second tractability issue relates to sample size for the mortgage principal and loan-to-value ratio variables, which are only observed for home buyers who borrow to finance their purchase. There are 2360 and 2042 usable observations of these variables in the 1999 and 2007 data waves, respectively. However, using the full joint distribution of all six variables would result in $3 \times 5^5 = 9375$ categories per year, thus the data would be insufficient. To overcome this, two marginal distributions are fit for each year, rather than the full data density. The first is the joint distribution of housing consumption, housing tenure, age and income, and the second is the joint distribution of mortgage principal and loan-to-value ratio. For the first marginal distribution there are 34125 and 32822 observations to identify 375 category frequencies in 1999 and 2007, respectively, and for the second there are 2360 and 2042 observations to identify 25 category frequencies in 1999 and 2007, respectively. The overall likelihood is the product of the four marginal likelihoods.

A natural and flexible way to characterize the resulting categorical marginal densities is by a multinomial distribution. The parameters of the likelihood function for this auxiliary model are the probabilities of randomly drawn observations falling into each of the categories. The likelihood maximizing estimates of these parameters are simply the category frequencies. This implies that $\hat{\mathbf{b}}_T$ and $\tilde{\mathbf{b}}_{HT}(\gamma)$ are vectors of category frequencies for the data and model simulated data, respectively, and that the problem in (9) is minimizing squared category frequency deviations.

1.4.3 Identification

It is customary to provide intuition for parameter identification in instances when a formal proof is infeasible. This is provided in Table 1.8. Note that although the identification arguments are provided at an intuitive level, they have been confirmed through inspection of the gradient matrix of category frequencies on the structural parameters, which is computed as part of the standard error estimation.

Subtle differences across age and income groups in the behavior of housing consumption provide much of the identification. Firstly, the utility weight on non-housing consumption, θ, regulates the fraction of lifetime income that is spent on housing, and hence the general level of housing consumption. Given the general level in 1999, the overall increase in housing consumption in 2007 depends on households' reactions to the estimated change in user-cost, which is regulated by τ. Non-convex adjustment costs, ϕ, pin down the frequency with which home-owners adjust their consumption. The slope of the consumption profile is pinned down largely by the real interest rate. The real interest rate also affects many other moments, but the slope of the consumption profiles seems to be uniquely determined by the interest rate, whereas those other moments are usually generated by the real interest rate in combination with other parameters. The desire to own a home depends on the financial gain from ownership; hence

	Non-Housing			Housing		
	1999	2007	% Change	1999	2007	% Change
Income Quintile						
1	6334	5878	-7.2%	3720	3958	6.4%
2	9840	9532	-3.1%	5070	5494	8.4%
3	13568	13535	-0.2%	6257	6941	10.9%
4	19271	18460	-4.2%	8339	9174	10.0%
5	31468	31481	0.1%	13559	15192	12.1%

Notes: Table 1.7 shows how housing and non-housing expenditure in the Consumer Expenditure Survey changed over the boom period by income group. Non-housing expenditure decreased by a large amount for the lowest income group, and was roughly constant for the highest income group. Housing expenditure increased for all income groups, but by a larger percentage for higher income groups. High income households increased housing expenditure by almost double the amount that low income households did, which is consistent with the changes in home size observed in the AHS.

Table 1.7: Housing and non-housing consumption expenditure by income

Parameter	Description	Strongest Identifying Variation
θ	CES Utility Weight	Overall level of housing consumption
τ	Substitutability Parameter	Relative Overall Housing consumption 2007 vs. 1999
ϕ	Non-Convex Adjustment Cost	Percentage of Owners Buying
r_{99}	Real Interest Rate 1999	Housing Consumption Age Profile 1999
r_{07}	Real Interest Rate 2007	Housing Consumption Age Profile 2007
r^h_{99}	Rental Premium 1999	Percentage of Home-Owners 1999
r^h_{07}	Rental Premium 2007	Percentage of Home-Owners 2007
r^m_{99}	Mortgage Premium 1999	Mortgage Principals 1999
r^m_{07}	Mortgage Premium 2007	Mortgage Principals 2007
ζ_{99}	Max Debt-to-Income Ratio 1999	Housing Consumption Income Profile 1999
ζ_{07}	Max Debt-to-Income Ratio 2007	Housing Consumption Income Profile 2007
ξ_{99}	Max Loan-to-Value Ratio 1999	Loan-to-Value Ratios 1999
ξ_{07}	Max Loan-to-Value Ratio 2007	Loan-to-Value Ratios 2007

Notes: Table 1.8 describes the main variation that identifies each of the structural parameters. I examined the gradient matrix containing the slopes of the category frequencies on the structural parameters to confirm these sources of identification.

Table 1.8: Identification

the home ownership rate is determined by the rental premium. The amount of consumption that a home owner wishes to finance with a mortgage depends on the cost of borrowing, i.e. the mortgage premium. For the maximum debt-to-income ratio, differences in housing consumption across income groups provides identification, as discussed in the introduction of the paper. Lastly, the maximum loan-to-value ratio influences the fraction of mortgagors who choose loan-to-value ratios greater than 1.0, hence the fraction of borrowers in that category identifies the maximum loan-to-value ratio.

1.4.4 Computational Details

Computational cost is reduced in two ways. First, a variant of mathematical programming subject equilibrium constraints (MPEC) is employed (Su and Judd, 2012). Second, the problem is highly parallelized.

To understand the first source of computational simplification it is useful the recall the usual nested fixed point approach to estimating structural equilibrium models. In this approach the household optimization problems are solved and the model is simulated repeatedly until the equilibrium housing price is attained. At that point the objective function of the estimator would be evaluated, the vector of parameters to be estimated would be updated, and the process would be repeated until the estimator objective arrives at a minimum. The partial implementation of MPEC in the current work involves treating housing market equilibrium as a constraint on the problem and including the price of housing in the vector of estimated parameters. By imposing a constraint that the excess demand for housing is zero in the optimization problem, the solution for the price in the minimization routine will necessarily be the equilibrium price. The reduction in computational cost arises from the fact that the housing market does not need to be in equilibrium for every evaluation of the estimator objective function. Rather, the housing market only needs to be in equilibrium at the final evaluation.

Parallel programming was implemented by breaking the state space into 192 groups. The number of state space grid points was large (200 asset $\times$ 192 housing $\times$ 25 income process $\times$ 60 ages $=$ 57.6 million), thus reduction by a factor of 192 was very efficient. The estimation routine had an efficiency of just over 96%, taking roughly five hours to complete.[7] This implies that if scaled to a single processor the routine would have required $192 \times 5 \times 0.96 = 921.6$ hours, or just over 38 days.

[7]Inefficiency arises from processors that are idle while other processors in the team complete their tasks. Thus, in my implementation, processors spend 4% of their time idle.

1.5 Results

1.5.1 Parameter Estimates

Table 1.9 presents the estimated parameters and associated standard errors. The first two estimates are the utility function parameters. As discussed in section 1.4.3 above, the elasticity of substitution between housing and non-housing consumption has a great deal of impact on how much housing demand changes in response to a change in the user-cost. The estimated τ implies an elasticity of substitution of 0.36, and thus housing consumption will respond very little when the user-cost of housing changes. It is noteworthy that a very similar elasticity of substitution (0.32) was estimated by Li, Liu, and Yao (2009). The percentage cost of selling a home, ϕ, is estimated to be 5.1% of the home's value.

The remaining parameters are estimates of credit market conditions, which are presented in pairs of estimates of the same parameter across years. First is the real interest rate, which is estimated to have fallen by 1.01% between 1999 and 2007. Next is the mortgage premium, which is estimated to have fallen by 24 basis points. This can be interpreted as a reflection of securitization and other innovations that reduced the risk associated with mortgage lending. The reduction of the rental premium is estimated to be substantially larger than the reduction of the mortgage premium, falling from 2.94% to 1.73%. This 1.21% reduction would be due in part to the reduced mortgage premium, which lowers landlords' costs, but almost a full percentage point is not explained in this way. Some of the unexplained reduction in the rental premium is likely due to speculation by housing investors, whose reservation rental rate fell due to beliefs about future capital gains.

The last two pairs of estimates are of the credit constraint parameters. While the maximum debt-to-income constraint is estimated to have increased by almost 50%, the maximum loan-to-value constraint is estimated to have increased very little. These changes mirror the mortgage data presented in section 2.4 above, which show that while mortgage debt increased very substantially, loan-to-value ratios changed relatively little. One important question is whether the estimated changes in the maximum debt-to-income ratio and mortgage interest rate allow for larger allowable mortgage payments. By multiplying these estimates together it can be seen that the maximum mortgage interest payment increased from 9.4% of income to 10.1% of income between 1999 and 2007.

1.5.2 Goodness of Fit

There is no perfect measure of the fit of a structural model like there is for an OLS regression. A common suggestion is the squared-correlation between the data and model category frequencies, which is 0.58 for the current exercise. This number can be interpreted in a similar manner as an OLS regression R^2. However, it is often more meaningful to directly examine

Parameter	Description	Estimate	(Standard Error)
θ	CES Utility Weight	0.187	(0.0007)
τ	Substitutability Parameter	-1.787	(0.0017)
ϕ	Non-Convex Adjustment Cost	0.0510	(0.0020)
r_{99}	Real Interest Rate 1999	0.0275	(0.0008)
r_{07}	Real Interest Rate 2007	0.0174	(0.0031)
r^m_{99}	Mortgage Premium 1999	0.0160	(0.0010)
r^m_{07}	Mortgage Premium 2007	0.0136	(0.0030)
r^h_{99}	Rental Premium 1999	0.0294	(0.0004)
r^h_{07}	Rental Premium 2007	0.0173	(0.0015)
ζ_{99}	Max Debt-to-Income Ratio 1999	2.178	(0.016)
ζ_{07}	Max Debt-to-Income Ratio 2007	3.229	(0.086)
ξ_{99}	Max Loan-to-Value Ratio 1999	1.098	(0.045)
ξ_{07}	Max Loan-to-Value Ratio 2007	1.101	(0.042)

Notes: Table 1.9 provides estimates of the structural parameters and their standard errors. The first three estimates, which are preference and adjustment cost parameters, are reasonable and consistent with other literature. Credit market parameters are presented in pairs to highlight the differences across years. The real interest rate, the mortgage premium and the rental premium all fall. The maximum debt-to-income ratio loosens considerably, but the maximum loan-to-value ratio is relatively constant.

Table 1.9: Parameter estimates

how well the model replicates the data patterns that identify the parameters, and to check the model's predictions against out-of-sample data. To this end, the fitted patterns of housing consumption change by income and age, as well as the changes in mortgage borrowing, are presented. The model's predictions for non-housing consumption are compared against the out-of-sample CEX data presented in section 2.5.

Table 1.16 repeats the presentation of housing consumption percentage changes by age and income group, along with the fitted model counterparts. The model generates decreasing housing consumption for very low income households and positive growth for higher income households; however, the pattern is not monotonic. The greatest increases in housing consumption in the model are among middle income households, whereas the biggest increases in the data are among the highest income households. The gains in the middle of the income distribution generated by the model are a reflection of the decreased consumption among borrowing constrained households. Constrained households are not able to borrow against future income when a bad shock is realized, hence their consumption will grow as their income state improves. In terms of replicating the patterns of housing consumption growth by age the model does reasonably well. The greatest increases in housing consumption in the model are for 31-40 year old households, rather than 21-30 year old households. However, the numbers are very similar.

Table 1.11 provides the data and model counterparts for variables that were argued to be important for identification in section 4.3 above. Overall, the model replicates these basic features well. One source of error is the tendency of the model to overshoot home-buyer rates for 1999 and undershoot them for 2007. This is a product of the adjustment costs being assumed constant across time periods. Another source of error in the fit is the growth of average initial mortgage principal, which overshoots the data average by almost \$3000. The median loan-to-value ratio generated by the model is slightly lower than the data for both years, but still relatively close. The home ownership rate is slightly high for 1999, but only by 1.3%.

Lastly, consider the out-of-sample predictions for non-housing consumption expenditure presented in Table 1.12. The model predicts that non-housing expenditure will fall for all income groups, but by a smaller amount in percentage terms for higher income groups. This gradient is present in the CEX data as well, although the magnitudes across income groups differ from the model. The largest non-housing expenditure decrease in the data is in the bottom income quintile, whereas in the model the largest decrease occurs in the second income quintile. There are also differences at the top of the income distribution. The model predicts small decreases in non-housing consumption for the fourth and fifth income quintiles, but in the data the fourth income quintile has a substantial non-housing consumption decrease and the top income quintile actually increases non-housing consumption slightly.

One important aspect of the model is that aggregate non-housing consumption differs

By Income Quintile			By Age Group		
Quintile	Model	Data	Age Group	Model	Data
Income Quintile					
1	-4.6%	-5.8%	21-30	1.6%	1.5%
2	2.4%	-0.3%	31-40	1.7%	0.7%
3	4.5%	1.2%	41-50	-0.01%	0.2%
4	2.1%	0.7%	51-60	-0.5%	-0.7%
5	0.7%	3.6%	61-70	-0.6%	-1.0%

Notes: Table 1.16 describes how well the model replicates the changes in housing consumption observed in the AHS data between 1999 and 2007. The pattern of decreased housing consumption among low income households and increased housing consumption among high income households is replicated. By age group, the pattern of increased housing consumption among young people and decreased housing consumption among older people is replicated.

Table 1.10: Fitted housing consumption changes by age and income

	1999		2007	
Moment	Model	Data	Model	Data
Percent Home-Owners	67.4%	66.1%	68.5%	68.2%
Percent Home-Buyers	5.28%	4.76%	3.93%	4.33%
Average Mortgage Principal	59300	60600	86400	83700
Median Loan-to-Value	0.915	0.920	0.927	0.947

Notes: Table 1.11 shows that the model replicates features of the data above and beyond housing consumption. A good fit of the model to these data features is important because they provide variation that identifies some key structural parameters.

Table 1.11: Various fitted moments

from the aggregate endowment of non-durable goods only by the difference between interest income and rent/interest payments. Hence, when aggregate non-durable consumption falls in the model it is because a greater part of the aggregate endowment is flowing away from consumers to the capital market. In other words, non-housing consumption in the model falls only if expenditure on housing rises. If the beneficiaries of this increased flow were included in the model we would see their non-housing consumption rise, and they would likely be in the top quintile of income, thus rationalizing the CEX data for the top income quintile.

1.5.3 Counterfactual Experiments

I consider two sets of counterfactual experiments. The first explores the relative contributions of real interest rate, credit constraint and speculation changes to house price growth. The second set of experiments explores the roles of these shocks in generating the observed increase in housing consumption inequality.

Contributions to House Price Growth

Experiments in which shocks are applied to isolated subsets of parameters are documented in Table 1.13. The first column lists the shocks that are being considered in the particular experiment, and the second column lists the percentage of actual growth in the average price of housing that is explained. Actual growth is measured as the change of the average price per square foot of recently purchased homes in the 1999 and 2007 AHS samples. The percentage change measured in this way is 50.6%, which is similar to the change in the real FHFA index between 1998 and 2006.

The first row shows that if all of the estimated shocks are applied then 61.8% of actual price growth is explained. In all of the experiments the aggregate supply of housing grows by 1.4%, as was modeled in the estimation routine. To set a benchmark for comparison, the second row of Table 1.13 indicates how the price would adjust if no changes other than the increase in the supply of housing occurred. Because of very inelastic preferences, the increase in housing supply in isolation would have caused the price of housing to fall by 7.3%.

The first shock added to the supply shock is the estimated change in the debt-to-income ratio constraint. Recall from Table 1.9 that this constraint loosened by nearly 50%. This shock results in the price of housing rising 6.1% above the supply shock only case, but it is not enough to generate an increase in the price of housing above the initial price.

The shocks to the real interest rate and the mortgage premium are applied together. In combination these provide much of the model's explanatory power, causing the price of housing to rise by 24.8%. As explained in section 1.4.3, these shocks work through the mechanism of consumers bidding up the price of housing until the user-cost is roughly constant. The other aspect of user-cost that applies to renters is the rental premium, which likely reflects housing investors' beliefs about future capital gains. The fall in the rental premium causes the price of housing to rise 20.1% above the initial level.

It is noteworthy that the estimated contribution of the falling interest rate to house price growth is larger than the 20% estimated by Glaeser, Gottlieb, and Gyourko (2010). To understand this difference we can return to equation (8). According to that equation movements in the real interest rate cause the price of housing to move by an equal amount in percentage terms. That is, the relationship between house prices and interest rates has a log-log form. However, the analysis by Glaeser, Gottlieb and Gyourko is based on a log-linear (constant semi-elasticity) relationship. Thus, in the model presented here the price of housing is much more sensitive to real interest rate fluctuations when the real interest rate is low.

The last row of table 1.13 applies the changes in mortgage interest rates and the maximum debt-to-income ratio together. One very important point is that when these changes are applied together the explanatory power is greater than the sum of the isolated effects of these changes. In other words, there is complementarity between the mortgage interest rate and the maximum debt-to-income constraint. This complementarity can be understood in terms of the demand equations in section 3.5. According to equation (9), the housing demand of a constrained household does not depend on the real interest rate, thus a constrained household does not respond directly to the interest rate shock. Similarly, the demand of an unconstrained household does not respond directly to a borrowing limit shock. However, when these shocks are applied in combination there may be two effects on constrained households: first, the loosening borrowing constraint might change them into an unconstrained household, which is associated with increased demand, and second, these unconstrained households experience a further increase in demand due to the interest rate shock. Thus, the complementarity arises because the mass of households to whom the interest rate shock directly applies is larger after the borrowing constraint loosens.

Contributions to Housing Consumption Inequality

This section explores the source of the increased housing consumption inequality that occurred between 1999 and 2007. The main findings are that relaxation of the maximum debt-to-income ratio constraint worked to reduce housing consumption inequality, whereas falling interest rates worked to increase housing consumption inequality. Because the interest rate effects are larger inequality increased overall.

The first two columns of Table 1.14 present the benchmark changes in housing consumption, and the changes that result from a pure supply shock ($\overline{h}$). Because the price of housing falls when the supply increases, debt levels fall and the borrowing constraints of lower income households are somewhat relieved. Hence, all households are able to increase their consumption in response to a supply shock, although low income households increase home size by relatively less.

Extending the maximum debt-to-income ratio has an important consequence in terms of the reallocation of housing. In this case low income households increase their housing consumption more than three times as much as when housing supply increases in isolation. The rest of the bottom of the income distribution also increases housing consumption by more than the pure supply shock case, although the differences are less severe. At the top of the income distribution housing consumption increases by less than when supply increases in isolation because the vast majority of households in the top of the income distribution are not constrained.

One result of interest rates falling in isolation is that many more households hit the borrowing limit than in the benchmark. Hence, there is a reallocation of housing towards high income households who are able to increase their borrowing in order to consume more. The lower user-cost that results from falling interest rates is what drives them to do so. Obviously, when the borrowing limit also loosens, as in the benchmark, this offsets some of the interest rate induced reallocation resulting in the benchmark patterns of housing consumption change.

1.6 Extension - Pervasive Speculation

An important component of the main analysis was the discipline imposed by rational household decision making. While speculation was incorporated through the exogenous "investors" who supply rental housing, the decisions made by the consumers of housing, who ultimately determine the price, were rational. However, the estimated model does not fully explain the house price boom, nor does it fully account for the rise in housing consumption inequality. This extension explores the possibility that irrational speculation among home owners can explain some of this residual variation.

It is not clear how one might identify the irrational beliefs of home-owners. Here I discipline the estimated home-owner speculation to be equal to the estimated investor speculation, which is identified by the behavior of renters. Thus, estimated speculation among home-owners does not simply absorb residual price variation. Speculation provides and additional force to increase the price of housing, but must be consistent across owners and investors.

An important aspect of this extension is that the beliefs of home-owners are irrational. The price of housing will be a stable equilibrium price, however, home-owners will believe the price is not stable. Home-owners will always believe that the price of housing is going to be g percent higher next year, even though the price will turn out to not have grown at all next year. This inflates the continuation value of a home-owner so that equation 1.5 becomes:

$$V_j^{own}(s_{j-1}, b_{j-1}, y_j; p) = \max_{c_j, h_j} \left\{ u(c_j, h_j) + \psi_{j+1} \beta E_{y_{j+1}} \left[\widetilde{V}_{j+1}(s_j, b_j, y_{j+1}; (1+g)p) \right] \right\}. \quad (1.12)$$

Here $\widetilde{V}_{j+1}$ is the counterfactual future value that households would attain if their beliefs about housing appreciation turn out to be true. I interpret the reduction in the rental premium net of the reduction in the mortgage premium as the expected capital gains of investors. Hence, g is identified by the restriction $g = (r^h_{99} - r^h_{07}) - (r^m_{99} - r^m_{07})$. To implement this I first solve for the counterfactual value functions assuming that the speculative beliefs are correct. Then I solve for the optimal decision rules at the current price and use these to simulate the model. In equilibrium the price will clear the market, and the price will turn out to be stable because no further shocks occur.

Table 1.15 provides the parameter estimates for this extended model, as well as the price of housing in each equilibrium. There are subtle differences between the parameter estimates in the extended model and the main estimates in Table 1.9. The estimated reductions in interest rates and the rental premium are slightly smaller, as the mechanisms are relatively less important in generating the increased borrowing etc. observed in the data. The estimated belief about housing appreciation is $g = 0.0153$, thus households believe housing will appreciate at a real rate of 1.53% per year. This is relatively modest compared to the actual appreciation between 1999 and 2007, but is enough to make the current price of housing substantially higher.

The extended model generates a 39.8% increase in the price of housing, which explains 78.7% of the observed increase in the price of housing between 1999 and 2007. This is relatively more than the 61.8% that the main (fully rational) model generates. Thus, speculation among home-owners explains between 15% and 20% of the housing boom. Combining this with the 20.1% that investor speculation explained in the main model, speculation accounts for more than 1/3 of the housing boom.

Adding speculation among home-owners also improves the model fit to the observed changes in housing consumption. In particular, the largest increase in housing consumption occurs at the top of the income distribution, as in the data, rather than in the middle of the income distribution, as in the main model. Unfortunately, these gains in explanatory power are somewhat offset by a reduction in the model's fit to the age profile of housing consumption changes. The increases in housing consumption among young households, and decreases among older households, are far larger than in the basic model and the data. This is because young households believe they have the most to gain from investing in an appreciating asset. On the other hand, older households may be reducing their consumption relative to the initial steady state because their housing investments have not yielded the returns they were expecting.

1.7 Conclusion

This paper has explored the relationships between the U.S. housing boom, increasing housing consumption inequality, and credit market frictions. This paper documented a substantial increase in housing consumption inequality using data from the American Housing Survey. Particularly stark was the fact that housing consumption fell by 5.8% within the bottom quintile of household income, but increased for all other income groups, especially high income households. It was argued that these patterns of housing consumption change could be due to downsizing by borrowing constrained households who have experienced negative income shocks. A life-cycle model with idiosyncratic earnings risk and credit constraints was developed and estimated in order to explore this possibility.

The shocks to interest rates and credit constraints that generate the housing boom in the model were identified by changes in the joint distribution of housing consumption, mortgage finance and demographic variables. These shocks generate an increase in the price of housing that is 61.8% as large as the actual increase. Furthermore, the increased housing consumption among higher income households, and reduced housing consumption among low income households, is replicated by the fitted model. In the model, housing consumption falls by 4.6% within the lowest income quintile, and increases for all other income groups.

The largest cause of increasing house prices is the falling real interest rate, which accounts for 24.8% of the observed increase in house prices. Relaxation of the debt-to-income ratio constraint played a measurable role as well. In isolation, a relaxed debt-to-income constraint has some explanatory power, but the larger effect is through amplification of the effect of interest rates. Together these changes account for 39.9% of the observed housing appreciation. A falling rental housing premium, which largely reflects speculation by housing investors, accounts for 20.1% of the actual rise in the average price of housing.

Similarly, increased housing consumption inequality is primarily driven by falling real interest rates, which generate a reallocation of housing from constrained to unconstrained households. The loosening debt-to-income ratio constraint offsets some of this reallocation, and hence reduces the extent of increased inequality.

In an extension I considered the potential role of irrational speculation among home-owners. I found that this extension increased the models explanatory power such that the model accounts for 78.7% of the observed increase in house prices. Thus, total speculation among investors and home-owners can account for more than 1/3 of the observed house price boom.

Income Quintile	Model	Data
Income Quintile		
1	-3.2%	-7.2%
2	-7.7%	-3.1%
3	-0.5%	-0.2%
4	-0.1%	-4.2%
5	-0.7%	0.1%

Notes: Table 1.12 describes how well the model's predictions for non-housing consumption compare to out-of-sample data from the Consumer Expenditure Survey. The model performs well, replicating the greater decrease in non-housing consumption among low income households than high income households.

Table 1.12: Fitted non-housing consumption expenditure changes

Shocks	Explained Percent of Actual Price Growth
All	61.8%
$\overline{h}$	-7.3%
$\overline{h}, \zeta$	-2.2%
$\overline{h}, r, r^m$	24.8%
$\overline{h}, r^h$	20.1%
$\overline{h}, r, r^m, \zeta$	39.9%

Notes: Table 1.13 describes the results of counterfactual experiments that investigate the causes of the increased price of housing. The first row of the table reports the effects of all of the estimated parameter changes combined, and the remaining rows report the results when certain estimated parameters are changed in isolation. For the first two experiments the price falls, rather than rises, thus the percentage of the actual change in house prices that is explained is negative. The first experiment increases the supply of housing only, and all remaining experiments include this change. The second experiment increases the supply of housing and loosens the debt-to-income limit. The third experiment applies the reductions in interest rates, the fourth applies the fall in the rental premium, and the last applies the changes in interest rates and the debt-to-income constraint together.

Table 1.13: Price causation experiments

shocks	all	$\overline{h}$	$(\overline{h}, \zeta)$	$(\overline{h}, r, r^m, r^h)$
Income Quintile				
1	-4.6%	0.4%	1.3%	-5.9%
2	2.4%	1.7%	2.7%	1.3%
3	4.8%	1.5%	1.6%	4.1%
4	2.1%	1.3%	0.7%	2.2%
5	0.7%	0.9%	0.3%	1.3%

Notes: Table 1.14 shows the differential effects of debt-to-income constraint relaxation and changes in user-cost on housing consumption inequality. The first column repeats the changes in housing consumption by income group due to all estimated parameter changes combined. The second column shows the effect of the increased supply of housing alone. The third column shows that relaxing the debt-to-income constraint in isolation reduces housing consumption inequality. The last column shows that if the interest rate, mortgage premium and rental premium fell, but the debt-to-income constraint was held constant, housing consumption inequality would have increased by more than it did in the benchmark.

Table 1.14: Inequality effects of isolated shocks

Parameter	Description	Estimate
θ	CES Utility Weight	0.196
τ	Substitutability Parameter	-1.734
ϕ	Non-Convex Adjustment Cost	0.0816
r_{99}	Real Interest Rate 1999	0.0258
r_{07}	Real Interest Rate 2007	0.0181
r^m_{99}	Mortgage Premium 1999	0.0154
r^m_{07}	Mortgage Premium 2007	0.0140
r^h_{99}	Rental Premium 1999	0.0316
r^h_{07}	Rental Premium 2007	0.0149
ζ_{99}	Max Debt-to-Income Ratio 1999	2.251
ζ_{07}	Max Debt-to-Income Ratio 2007	3.137
ξ_{99}	Max Loan-to-Value Ratio 1999	1.103
ξ_{07}	Max Loan-to-Value Ratio 2007	1.119
g	Speculative Belief 2007	0.0153
p_{99}	Housing Price 1999	43.69
p_{07}	Housing Price 2007	61.07

Notes: In comparison to the main estimates presented in table 1.9, these results show that almost 20% more of the observed housing appreciation is explained when speculation by home-owners is included. The estimated belief is that housing will appreciate at a real rate of 1.53% per year. Also note that the estimated changes in credit constraints and interest rates are slightly smaller than in table 1.9.

Table 1.15: Parameter estimates with pervasive speculation

Quintile	Model	Data	Age Group	Model	Data
By Income Quintile			By Age Group		
Income Quintile					
1	-3.7%	-5.8%	21-30	5.6%	1.5%
2	0.4%	-0.3%	31-40	4.2%	0.7%
3	0.5%	1.2%	41-50	5.1%	0.2%
4	2.1%	0.7%	51-60	-5.7%	-0.7%
5	3.7%	3.6%	61-70	-8.2%	-1.0%

Notes: Table 1.16 describes how well the model replicates the changes in housing consumption observed in the AHS data between 1999 and 2007. The pattern of decreased housing consumption changes by income is better replicated in this case, but the quality of the fit to the patterns by age is reduced. l

Table 1.16: Fitted Housing Consumption Changes by Age and Income with Pervasive Speculation

Chapter 2

Loan-to-Value Ratios and Life-Cycle Housing Decisions

2.1 Introduction

Increased home-ownership has been a longstanding goal for public policy. This is because policy makers believe home-ownership "exerts a positive effect on local community development" largely because owners are more interactive with their neighbors and have a greater propensity to participate in local politics (Megboluge and Linneman, 1993). This has led to substantial research into policy and trends in home ownership. One of the striking patterns in the data is the increasing home-ownership rate over the life-cycle, with the vast majority of young households being renters and the vast majority of older households being owners. Some potential explanations for the low home-ownership rate of young households are down-payment constraints, income risk and anticipated changes in household composition. The effectiveness of policies aimed at increasing home-ownership depends on the significance of each explanation.

This paper investigates the relative importance of the factors that shape life-cycle housing decisions, with an emphasis on early life home-ownership. I flexibly accommodate the role of credit constraints by introducing a schedule of mortgage interest rates that depends on the loan-to-value ratio (LTV), and which nests the canonical maximum LTV model. I assess the extent to which this formulation improves our ability to replicate life-cycle housing patterns, and how it influences the predicted effects of policies that promote home ownership.

Credit frictions are an important determinant of home-ownership. Early empirical evidence of this is found in Linneman and Wachter (1989), and more recent evidence is provided by Calem, Firestone, and Wachter (2010). In the structural life-cycle housing literature credit frictions have usually been incorporated through maximum LTV constraints, or equivalently minimum down-payment constraints (examples include Yang, 2009; Li and Yao, 2007; Li, Liu, and Yao, 2009). It has commonly been assumed that mortgage LTVs cannot exceed a value in the neighborhood of eighty percent. For example, Bajari, Chan, Krueger, and Miller (2013) estimate and analyze a life-cycle housing model, subject to the assumption that a 20% down-payment is required. This is entirely at odds with U.S. micro-data, which show that more than 70% of mortgages initially have LTVs greater than 0.8, and one-quarter initially have LTVs of 1.0 or greater.

One hypothesis is that high costs associated with low or negative equity borrowing may be a form of credit constraint, leading low wealth families to rent when they would prefer to be home owners. That is, the true credit friction may be the high cost of high LTV loans, rather than the impossibility of high LTV loans. One of the contributions of this paper is to examine if and how predicted life-cycle housing consumption decisions are altered when high LTV mortgages are allowed for, albeit at a high cost. In particular, this paper shows that allowing for high LTV mortgages substantially improves the model fit to early life home-ownership data.

The observed distribution of purchase LTVs provides evidence of borrowing costs that increase with the LTV. As illustrated in Figure 2.1, there is substantial clustering of purchase LTVs just below 0.8 and 1.0, which indicate changes in borrowing costs at those LTV levels. Increased costs above 0.8 are likely due to mortgage insurance requirements, while increased costs above 1.0 are likely due to the need to borrow from an unconventional lender at those levels. In the model employed in this paper, non-convex jumps in the mortgage interest rate capture increased costs at these LTV levels. The size of the jumps is identified by the extent of the observed clustering just below these levels.

Income related credit constraints may also limit home-ownership among young households (Attanasio, Bottazzi, Low, Neisham, and Wakefield, 2012). Young households incomes grow quickly and they may wish to borrow against higher future income, but maximum debt-to-income or payment-to-income constraints prevent them. As incomes grow households could potentially borrow greater amounts in order to increase their housing consumption, but this would involve repeatedly incurring large adjustment costs. Instead, young households may prefer to remain renters so that their housing consumption can grow with their income without repeatedly paying real estate commissions and other adjustment costs. This theory is consistent with the steep age-profile of housing consumption up to age 35, where home size grows by approximately 60% (see Yang (2009), as well as the estimates below). If these young households can eventually buy without making a large down-payment it makes this strategy more desirable because less early life consumption needs to be forgone in order to save.

Earnings risk can also cause young households to rent (Diaz-Serrano, 2005). A large literature has shown that life-cycle income risk is substantial and that household consumption is only partially insured against these shocks (for example Pistaferri and Meghir, 2004; Heathcote, Perri, and Violante, 2010a; Krueger and Perri, 2011; Blundell, Pistaferri, and Preston, 2008a). If realized lifetime earnings differ substantially from ex-ante expected earnings then optimal home-size will also differ. If a home was purchased early in life then such a realization would result in either payment of adjustment costs or a suboptimal mix of housing and non-housing consumption. The availability of high LTV mortgages might play a role in this channel, as a household that has experienced a temporary income shock can borrow to smooth consumption, rather than be forced to sell.

Young households might also prefer to rent because of the rapid increase in the average number of household members during early adulthood. People may delay home purchases because they anticipate marriage and the arrival of children, which correspond to an increased optimal home size. As with the effects of rapidly increasing income during early life, the preference for renting because of family size growth is due to non-convex adjustment costs. The importance of this effect is highlighted by Fisher and Gervais (2011), who show that changes in the timing of changes in family composition can account for longer-run changes in the home-ownership rate of young households.

The relative importance of these factors in shaping the housing demand of young adults will influence the effects of policies intended to promote home ownership. By allowing for high LTV mortgages a model will rely more heavily on family composition and earnings risk in matching the low ownership rates of young adults. This affects estimates of structural parameters, which in turn affect the response to policy changes.

One counterfactual policy experiment I analyze in this paper is an increase in the mortgage interest subsidy (tax deduction). When applied to the estimated benchmark model with high LTV mortgages available, an increase in the subsidy substantially increases the home ownership rate of young households, whereas when applied to a model with a 0.8 maximum LTV constraint the increase in the subsidy has little effect. This contrast is due to differences in the margins that shape ownership decisions. In the benchmark model the main tension preventing ownership is that future adjustment costs outweigh the benefits of early ownership, whereas in the max-LTV model the main tension is insufficient resources to overcome credit constraints. Increasing the subsidy improves the relative value of ownership, but does not provide wealth to overcome down-payment requirements.

This paper is related to contributions by Chambers, Garriga, and Schlagenhauf (2009a,b) who explore how changes in the flexibility of mortgage arrangements and down-payment requirements affect home ownership over the life-cycle. Their results illustrate the relationships well, but they do not consider the degree to which low equity mortgages are utilized in the data. The current paper considers the life-cycle ownership pattern in greater detail, and emphasizes the relative importance of the factors that influence the rent-own decision. One paper that does consider the distribution of LTVs in detail is Korteweg and Sorensen (2011), but their questions are very different than the current paper.

The remainder of this paper is as follows. Section two summarizes the AHS data on LTVs and housing consumption decisions. Section three presents the structural model and highlights the key mechanisms. Section four describes the estimation procedure and presents results. Section five discusses these results and their implications for our understanding of life-cycle housing consumption. Section six concludes.

2.2 Mortgage and Housing Data

The primary source of data is the American Housing Survey (AHS). This survey is produced by the U.S. Census Bureau, and is "the most comprehensive national housing survey in the United States." The primary purpose of the study is to assess housing units, but occupants' characteristics are also assessed. The primary advantage of using this data is that construction of LTVs is possible through purchase price and purchase financing data. This data is a vital source of variation to discipline the nature of down-payment constraints and their effects over the life-cycle. Another advantage of using this data is that it allows an assessment of housing consumption through the characteristics of units, rather than the value of property.

The AHS is a biennial survey conducted in the spring and summer of the survey year. Sampling weights allow one to construct a nationally representative sample. I focus on survey years 1997-2009, which are readily available on the AHS web-site. Useable sample sizes range from 24,744 to 41,073 per wave, where useable observations have no missing data in the key variables.

2.2.1 Loan-to-Value Ratios

Home-owner survey respondents are asked about the purchase price of their home, as well as any mortgages outstanding. Among the questions asked of mortgagors are what the initial mortgage principal was, and whether the mortgage origination coincided with the purchase of the mortgaged home. An initial LTV can be constructed for any observation where the mortgage origination did indeed coincide with the purchase of the home.

Some survivorship bias occurs in that mortgages that have been re-financed or paid-off will not appear in the data. This error is minimized by only using data from homes purchased in the year prior to the survey (e.g. purchased in 2000 and reported in the 2001 survey). Only 2.9% of purchases one-year prior to the survey are refinanced by the survey date.

It is useful to begin with an overview of the distribution of LTVs. A histogram of purchase LTVs covering all age groups is provided in Figure 2.1. More than 62% of the probability mass is between 0.80 and 1.0, with another 9% for LTVs greater than 1.0. Somewhat hidden in the histogram is the fact that 19.4% of mortgages have and original LTV of exactly 1.0. Together these imply that more than 70% of mortgages begin with an LTV above 0.80.

A striking feature of Figure 2.1 is that it is bimodal with sharp drops in frequency above the modes. The number of LTVs in the 0.75-0.80 range is nearly double the number in the 0.80-0.85 range, and even more extreme, the number of LTVs in the 0.95-1.0 range is many times greater than the number just over 1.0. These jumps suggest non-convexities in the cost of borrowing at these thresholds. For example, mortgage insurance requirements apply when the LTV exceeds 0.80, which causes a cluster of home buyers who borrow just under this threshold. Similarly, a loan from a conventional source may not be available for those wishing

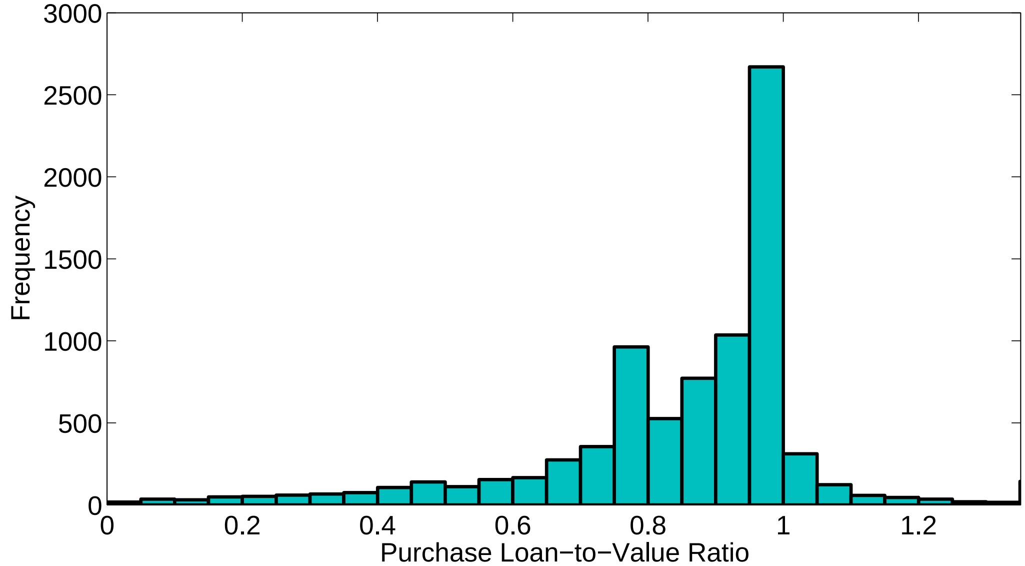

Notes: This figure illustrates the bimodal nature of the distribution of LTVs, as well as the jumps in frequency just above the modes. I hypothesize that these jumps reflect non-convexities in the borrowing cost function. At LTV=0.80 mortgage insurance premiums begin to apply, and above LTV=1.0 more expensive non-traditional lending sources are usually required.

Figure 2.1: Histogram of Loan-to-Value ratios

to borrow in excess of the value of their home (LTV>1.0), thus a secondary source of funds must be attained. It is likely that switching to a secondary lender is associated with a large increase in borrowing costs, causing a cluster of LTVs at or just below 1.0.

These findings are not a special feature of the AHS data. The summary of DataQuick LTVs reported by Glaeser, Gottlieb, and Gyourko (2010) is consistent with the AHS. While the AHS data are self-reported by owners, the DataQuick data are reported by lenders.

A natural consequence of life-cycle wealth accumulation is that the average purchase LTV becomes smaller with age; however, the progression is only subtle as households continue to utilize high LTV mortgages throughout the life-cycle. Figure 2.2 illustrates the rates at which LTVs fall into various categories for four age groups. In the right tail of the distribution we observe that the utilization rate for negative equity mortgages declines with age, but only falls by about 1/3 from the 21-30 age group to the 50+ age group. For LTVs at or just below 1.0 the decline with age is greater, with the utilization rate being nearly half as large for those over 50 compared to those in their 20s. For intermediate LTVs the patterns by age are less clear, as the age-gradient changes direction. For LTVs less than 0.75 the utilization rate is clearly increasing with age, as many more 50+ households choose mortgages with a small LTV than households in their 20s.

2.2.2 Housing Consumption

The aim of this subsection is to summarize the empirical facts that are the focus of the life cycle housing consumption literature. These data, along with the distribution of LTVs, provide identifying variation for the model estimation. As such, they are also the variation upon which the relative performance of competing models is based.

While these facts have been presented in previous work (e.g. Yang (2009) and Li and Yao (2007)), there is one aspect of measurement that is novel to this paper. I measure housing consumption as dwelling size, the units of which is square-feet, whereas previous work has used value as a measure of housing consumption. Each of these measures has advantages and disadvantages, but the similarity between them limits the need to debate their relative merits.

The disadvantage of using property values is the required normalization by a house-price index. As city level price indices show, there is substantial geographic variability in housing price changes. Even at the within city level there can be substantial variation in housing appreciation (e.g. Landvoigt, Piazzesi, and Schneider (2012)). Thus, it is difficult to differentiate a home that yields a large consumption flow from a home in an expensive city. Using size as a measure of consumption eliminates this to some degree because, all else equal, a 1500 sq-ft home will yield a greater consumption flow than a 1000 sq-ft home. The disadvantage of using size to measure consumption is that aspects of quality will be missed.

Figure 2.3 presents average home-size and the home-ownership rate from age 21 through

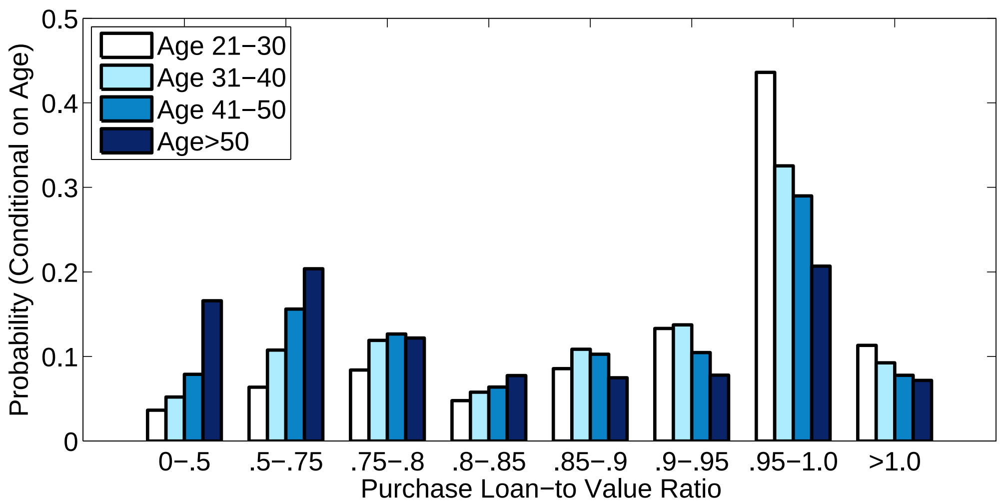

Notes: Each bar represents the fraction of mortgages that has an original Loan-to-Value ration within the corresponding interval for the corresponding age group. Interval boundaries are closed at the upper limit and open at the lower limit. This figure illustrates the decreasing (but continual) reliance on high LTV mortgages over the life-cycle. Low LTV mortgages become more prominent with age, but more than one-quarter of households over 50 continue to buy with 5% equity or less.

Figure 2.2: Loan-to-Value ratio utilization rates by age

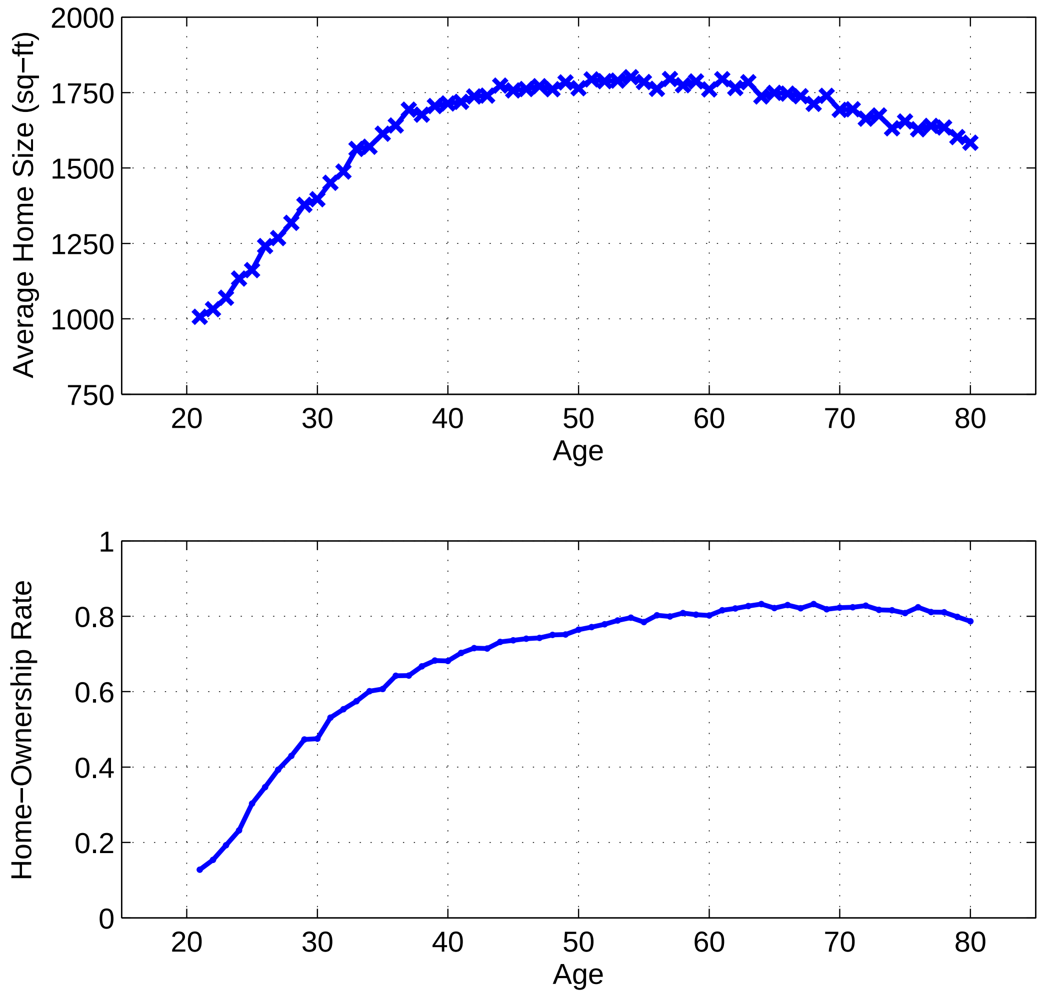

Notes: Home size and the ownership rate both grow quickly through the early stages of adult life. Home size peaks and plateaus around age 40, and declines through retirement. Home ownership does not peak until around retirement age, although the rate of increase in home ownership slows after age 40. The ownership rate does not decline as much during retirement as home size.

Figure 2.3: Average Home size and the home-ownership rate by age

80. Housing consumption grows quickly over the first twenty years of adult life, is flat for the remainder of working life, and then declines slowly during retirement. This pattern is consistent with that presented in Li and Yao (2007), but differs from that presented in Yang (2009), both of whom measure consumption as home value reported in the SCF. Here I have measured housing consumption as raw home size, whereas Yang (2009) measures housing consumption per adult-equivalent household member. For completeness, I experimented with measuring housing consumption as square-feet per household member, in which case consumption rises monotonically in age, just as Yang (2009) reports. For my empirical exercise raw home size is the appropriate measure because adjustments in housing consumption due to variation in family composition are an important determinant of early life decisions. Family size will enter the utility function such that housing consumption per person is valued, rather than raw consumption, but households must adjust their home size (potentially incurring adjustment costs) in order to attain optimal housing per person.

The second panel of Figure 2.3 illustrates the home-ownership rate by age. Few very young households own their home, but the ownership rate increases quickly. After age 40 the increase in the home-ownership rate slows, and the peak does not occur until about the age of retirement. Home ownership peaks later than housing consumption, and does not decline as much during retirement.

2.3 Structural Model

The structural model combines elements of incomplete markets heterogeneous agent macroeconomic models with elements of empirical housing consumption models. Like heterogeneous agent models, the driver of variation in life-cycle outcomes is an idiosyncratic earnings process. Variability in realized idiosyncratic earnings shocks and initial endowments leads to non-degenerate distributions of the empirically relevant state variables. Crucially, the model also includes incomplete markets so that earnings risk is not insurable, except (somewhat) through saving/borrowing.

The model adds a second consumption good (housing) to the usual heterogeneous agent framework. This necessitates the inclusion of several model features, in addition to adding the second good to the utility function and budget constraints. Housing consumption can be attained either through ownership or rental, and costs may differ between these modes of consumption. There are pervasive frictions in the market for financing of owner occupied housing. Adjustments to consumption of owner-occupied housing are subject to substantial non-convex adjustment costs.

To flexibly accommodate these important features of housing consumption a partial equilibrium approach is adopted. Most general equilibrium housing models adopt an ad-hoc no arbitrage condition to price rental versus owner occupied housing (e.g. Iacoviello and Pavan

(2013)). These assumptions are useful in maintaining the tractability of general equilibrium models, but, as Díaz and Luengo-Prado (2008) have demonstrated, they can introduce bias. Explaining ownership rates over the life-cycle is an important component of this paper, thus rental and ownership costs are treated as exogenous parameters to be estimated in order to avoid any bias.

2.3.1 Demographics and Preferences

Life Cycle

The economy is populated by J generations of households. Each cohort consists of a continuum of households. Let $j = 1, \ldots, J$ denote age. These households receive endowment income for J^R periods, and then a pension for the next $J - J^R$ years of retirement. Mortality risk is present during retirement: age j households survive for one more year with probability ψ_{j+1}. Each cohort initially consists of a unit mass of households, but because of mortality the mass of each retired cohort is only Ψ_j. The exogenous number of members of each household at age j is given by n_j.

Commodities and Prices

There are two types of goods: a non-durable good, which is the numeraire, and a housing good, the relative price of which is p. Non-durable consumption is denoted c and housing consumption, measured in square-feet, is denoted h. Because the housing good is durable it has a second role as an asset whenever it is owned, as opposed to rented. The stock of housing owned by a particular household is denoted s. Finally, b denotes a household's quantity of financial assets, where a positive quantity represents ownership of savings bonds, and a negative quantity represents mortgage debt.

Endowments

Over the course of working life each household i receives a stream of endowments of the non-durable good, $\{y_{ij}\}_{j=1}^{J-J^R}$, which is the product of a deterministic age-specific part, $\overline{y}_j$, and a stochastic idiosyncratic part. The logarithm of the idiosyncratic part, z_{ij}, is modeled as in Storesletten, Telmer, and Yaron (2004)[8] as follows:

$$z_{ij} = \alpha_i + \eta_{ij}, \tag{2.1}$$

where

$$\eta_{ij} = \rho\eta_{ij-1} + \nu_{ij}. \tag{2.2}$$

[8]For tractability the transitory component of z_{ij} is assumed to be measurement error. This component is only a very small part of the variance of earnings.

The α term is a permanent fixed effect component, where $\alpha_i \sim N(0, \sigma_\alpha^2)$. The time-variant shocks are also normally distributed draws: $\nu_{ij} \sim N(0, \sigma_\nu^2)$.

Initial assets, b_0, are drawn from a mixture distribution. That is, $b_0 = \tilde{b}_{0,1}\tilde{b}_{0,2}$, where $\tilde{b}_{0,1}$ and $\tilde{b}_{0,2}$ are underlying random variables. The first part, $\tilde{b}_{0,1}$, is distributed binomial, such that the probability a newborn household receives a positive endowment is $\bar{b}_{0,1}$. The second part, $\tilde{b}_{0,2}$, is distributed log-normal: $\ln(\tilde{b}_{0,2}) \sim N(\bar{b}_{0,2}, \sigma^2_{b_{0,2}})$. The motivation for this structure is that only a fraction of young households have financial assets (captured by the binomial draw), but those who do have initial wealth seem to have received substantial endowments (captured by the log-normal draw).

During retirement agents receive a pension equal to 45% of economy wide average earnings (see Mitchell and Phillips, 2006).

Preferences

Preferences are composed of the present discounted value of utility from non-durable and housing consumption:

$$U = \sum_{j=1}^{J} \beta^j \Psi_j \frac{1}{1-\sigma} \left([\theta(c_j/n_j)^\tau + (1-\theta)(h_j/n_j)^\tau]^{\frac{1}{\tau}} \right)^{1-\sigma}. \tag{2.3}$$

Recall that n_j is the exogenous size of the household at age j, thus the utility function values consumption per family member.

As usual, $\beta \in (0,1)$ is the discount factor. Households also discount by cumulative survival probabilities Ψ_j, which are less than unity during the retirement stage. Inter-temporal risk aversion is given by the CRRA parameter σ, while the intra-temporal elasticity of substitution is controlled by τ. The weights on consumption goods are given by θ.

2.3.2 Markets

Financial Market

Households have access to one-period savings bonds, which pay a net return equal to the real interest rate, r. Households can borrow through mortgages, for which they pay a premium, r^m. The rental premium is constructed so as to capture features of the cost of borrowing that are indicated by the distribution of LTVs. In particular,

$$r^m(LTV) = \begin{cases} r_0^m & \text{if} \quad LTV \leq 0.80 \\ r_0^m + r_{80}^m & \text{if} \quad 0.8 < LTV \leq 1.0 \\ r_0^m + r_{80}^m + r_{100}^m & \text{if} \quad LTV > 1.0 \end{cases}. \tag{2.4}$$

The mortgage premium is a piecewise constant function with non-convex jumps. The non-convexities and changes in slope are assumed to occur at LTVs of 0.8 and 1.0, which are the points in the LTV distribution where jumps in the density appear to occur (see section 2.1 above). The r^m_x parameters capture the size of non-convexities at point $LTV = x$. The intention of this highly stylized approach is to provide an alternative to the even more stylized maximum-LTV model. The max-LTV=0.8 model is nested within the current model, and would be represented by $r^m_{80} = \infty$. Similarly, the max-LTV=1.0 model is nested within the current model, and would be represented by $r^m_{100} = \infty$. Thus, this setup provides a relatively flexible alternative that performs well empirically. I have also considered piecewise linear functions (rather than piecewise constant functions), but the estimation results indicated that the slope terms were poorly identified. Because of the mortgage premium, households always extinguish mortgage debt before buying savings bonds. Mortgage borrowing is also subject to a borrowing constraint: a household's debt-to-income ratio cannot exceed ζ.

Housing Market

Households may either purchase housing at price p per square foot, or rent it a cost equal to a rental premium, r^h, times the price per square foot. That is, the annual cost of renting is $r^h p$ per square foot. The premium reflects the fact that the gross return earned by landlords needs to compensate them for effort costs associated with owning rental property over and above those associated with owning a savings bond. A non-convex adjustment cost applies to all housing sales. The seller incurs a cost of ϕ percent of the value of the sold property, which can be interpreted as a real estate commission, although it also encompasses other costs. Given the exogenous credit market parameters, rental premium and price of housing, households optimally choose whether to rent or own. The home-ownership rate that arises in the model is disciplined in the estimation routine by explicitly targeting the home-ownership rate by age.

2.3.3 Household Decision Problems

A household's decision problem consists of two parts: a discrete choice between renting and owning, and continuous choices for non-durable and housing consumptions.

The value functions at age j for the discrete choices are $V_j^{rent}(s_{j-1}, b_{j-1}, y_j)$ for renting, and $V_j^{own}(s_{j-1}, b_{j-1}, y_j)$ for owning. Importantly, these renting and owning values reflect the current choice, not the previous. Whether the household owned or rented in the previous year affects the current value functions only through the amount of housing wealth they enter the period with (s_{j-1}). The overall value function for a household is:

$$V_j(s_{j-1}, b_{j-1}, y_j) = \max\left\{V_j^{rent}(s_{j-1}, b_{j-1}, y_j), V_j^{own}(s_{j-1}, b_{j-1}, y_j)\right\}. \tag{2.5}$$

The continuous part of a household's optimization problem is presented recursively. I distinguish the problem of those who rent from those who are home-owners.

Value of Owning:

Households who will own their current dwelling solve the following problem:

$$
\begin{array}{rcl}
 & & V_j^{own}(s_{j-1}, b_{j-1}, y_j) = \max_{c_j, h_j} \left\{ u(c_j, h_j) + \psi_{j+1} \beta E_{y_{j+1}} \left[V_{j+1}(s_j, b_j, y_{j+1}) \right] \right\}, \\
 & \text{s.t.} & \\
y_j - \tau_w(y_j - \mathbf{1}_{(b_{j-1}<0)} \hat{r} b_{j-1}) + (1+\hat{r}) b_{j-1} + p s_{j-1} & = & c_j + p h_j + b_j + \mathbf{1}_{(h_j \neq s_{j-1})} \phi p s_{j-1} \\
-b_j \zeta & \leq & y_j \\
s_j & = & h_j \\
\hat{r} & = & \left\{ \begin{array}{ccc} r & \text{if} & b_{j-1} \geq 0 \\ r^m(LTV) & \text{if} & b_{j-1} < 0 \end{array} \right. .
\end{array}
\tag{2.6}
$$

The first equation is the standard Bellman Equation representation of the optimization problem. The future value is discounted by ψ_{j+1}, which will equal one during working ages and be less than one during retirement. The left hand side of the budget constraint reflects available resources, which depend on income, wealth and taxes. Mortgage interest payments are treated as a deduction for taxation purposes. The right hand side of the budget constraint potentially includes adjustment costs: the indicator function $\mathbf{1}_{(h_j \neq s_{j-1})}$ is activated if this household has moved to a new home, and thus they pay ϕ percent of the value of their previous home. If they were previously renters then $s_{j-1} = 0$ and there is no adjustment cost. The second constraint simply states that future housing wealth will equal current housing consumption as the household owns this home. The third constraint is the maximum debt-to-income ratio constraint. Income must be at least ζ times the mortgage principal. The final constraint indicates that the interest rate will vary depending on whether the financial position is positive or negative, and if it is negative then the mortgage interest rate includes a premium defined in equation 2.4.

Value of Renting

Working age households who rent solve the following problem:

$$V_j^{rent}(s_{j-1}, b_{j-1}, y_j) = \max_{c_j, h_j} \left\{ u(c_j, h_j) + \psi_{j+1}\beta E_{y_{j+1}} \left[V_{j+1}(s_j, b_j, y_{j+1}) \right] \right\},$$

$$\begin{array}{rcl} & \text{s.t.} & \\ y_j - \tau_w(y_j - \mathbf{1}_{(b_{j-1}<0)}\hat{r}b_{j-1}) + (1+\hat{r})b_{j-1} + ps_{j-1} & = & c_j + r^h ph_j + b_j + \phi ps_{j-1} \\ b_j & \geq & 0 \\ s_j & = & 0 \\ \hat{r} & = & \begin{cases} r & \text{if} \quad b_{j-1} \geq 0 \\ r^m(LTV) & \text{if} \quad b_{j-1} < 0 \end{cases}. \end{array} \tag{2.7}$$

Mortality risk enters the Bellman Equation of a renter's problem in the same way it enters an owner's. Non-convex adjustment costs enter the budget constraint of a renter without an indicator function. This is because anyone who owned their home the previous period must now sell it to become a renter. For anyone who rented in the previous period $s_{j-1} = 0$, so adjustment costs do not apply. For renters the credit constraints simplify to a non-negative asset constraint because they have no collateral to allow them to borrow. The last constraint indicates that the rented home is not an asset, rather it only provides a consumption flow, hence $s_j = 0$.

Terminal Condition

The model of mortgage costs does not impose any requirement that households are able to repay their debt. A terminal condition, which requires that agents cannot die with negative net worth, is thus added to the model. At age j this condition is imposed by adding $(1-\psi_{j+1})V_{j+1}^{1-\psi}$ to the Bellman Equation, where

$$V_{j+1}^{1-\psi} = \begin{cases} 0 & \text{if} \quad b_j + (1-\phi)ps_j \geq 0 \\ -\infty & \text{if} \quad b_j + (1-\phi)ps_j < 0 \end{cases}. \tag{2.8}$$

This part of the problem has been repressed in the decision problems above because all retired households ensure that their future net worth is always non-negative, and hence $V_{j+1}^{1-\psi} = 0$. There is no income risk during the retirement period, which allows these agents to be certain of their future net worth, and hence the mortality term drops out. This terminal condition imposes a natural borrowing limit on the model as young agents anticipate the requirement that they have non-negative wealth it in the future.

Parameter	Value	Description	Source
J	70	Maximum Model Age	
J^R	45	Retirement Age	
n_j	varies	Household Sizes	AHS: Cubic in age fit to AHS data
$\bar{b}_{0,1}$	0.127	Fraction with $b_0 \geq 0$	PSID: Frac. of 19-21 year olds with assets
$\bar{b}_{0,2}$	8.35	Average log of b_0 draw	PSID: Average ln assets of 19-21 year olds.
$\sigma^2_{b_{0,2}}$	1.91	SD of log of b_0 draw	PSID: SD of ln assets of 19-21 year olds.
$\bar{y}_j$	varies	Income Age Profile	PSID: Cubic in age fit to data
τ_w	0.27	Labor tax rate	Domeij and Heathcote (2003)
σ_α	0.25	SD of earnings fixed effect	Storesletten et al. (2004)
σ_η	0.0225	SD of earnings shocks	Storesletten et al. (2004)
ψ_j	varies	Retirement Survival Rate	U.S. Life Tables
p	44	Price per sq-ft of housing	AHS: Average Real Price per sq-ft in 2000

Table 2.1: Auxiliary parameter values

2.3.4 Parameterization

While the key parameters are estimated formally, the remaining auxiliary parameters are specified either based on other work or exogenous estimation proceedures. These auxiliary parameters are reported in Table 2.1.

The maximum age in the model is set to $J = 70$, which corresponds to real world age 90. Retirement age is set to $J^R = 45$, which corresponds to real world age 65. Survival probabilities are set based on the U.S. life tables. There is a small residual survival probability in the life tables at age 90, but I assume all surviving households die at that age. The number of people in the household is reported in the AHS. A cubic polynomial in the age of the head of the household was fit to form estimates of n_j by age.

Initial endowments are estimated using asset data that has been included in the PSID since 1999. I include households between 19 and 21 in the sample, and all reported assets including bonds, stocks, annuities, businesses and real estate. All variables are net of any debt owed against them, and real estate includes both the households own home (if owned) and other real estate investments. The first step is to estimate $\bar{b}_{0,1}$ as the fraction of agents who report positive assets. The second step is to estimate the mean and variance of the logarithm of initial assets, conditional on those assets being positive. The age-profile of earnings is captured by fitting a cubic polynomial in age (of the head) to household labor income reported in the PSID. The variances of income fixed effects and persistent shocks, as well as the persistence of endowments shocks, are taken from the over-identified cases in Table 1 of Storesletten, Telmer, and Yaron (2000).

Lastly, the real price per square-foot of housing is set to that observed in the AHS in

2000. To attain \$44 per square-foot I first normalize observed purchase prices by CPI with base year 1983-84 (the same normalization performed for income), and then normalize again by the size of the purchased unit. This delivers real dollars paid per square-foot, the average of which is \$44.

2.4 Estimation

2.4.1 Details

The remaining parameters are estimated through a method of simulated moments approach (see McFadden, 1989; Pakes and Pollard, 1989). The identifying moments naturally group into three parts. The first two are life-cycle profiles of average home-size and the home ownership rate, which are the objects I wish to explain. These moments are described in section 2.2 above. The third group of identifying moments reflect important features of the distribution of purchase LTVs, which are vital in identifying the features of the mortgage premium schedule. In particular, the category frequencies illustrated in Figure 2.2 are included as moments, but I aggregate these category frequencies over the life-cycle. I drop the first category ($LTV \leq 0.5$) because the category frequencies must sum to unity, thus one is redundant.

Let $\mathbf{m}_n$ be a vector containing the identifying moments estimated using n data observations. Furthermore, let $\tilde{\mathbf{m}}_{n,s}(\gamma)$ be the model counterpart computed using the sth model simulation of n observations, where γ is the vector of structural parameters to be estimated. Then the estimated parameter vector solves

$$\hat{\gamma} = \arg\min_{\gamma} \ (\mathbf{m}_n - \frac{1}{S}\sum_{s=1}^{S}\tilde{\mathbf{m}}_{n,s}(\gamma))' W_*^{-1} (\mathbf{m}_n - \frac{1}{S}\sum_{s=1}^{S}\tilde{\mathbf{m}}_{n,s}(\gamma)). \tag{2.9}$$

The positive definite weight matrix W_*^{-1} is the inverse of an estimate of the optimal weighting matrix. This was attained using a two-step estimation procedure as described in Adda and Cooper (2003). The total sample size (n) was 234,600, which I simulated 50 times (S).

Two restricted models are estimated in addition to the benchmark model. The first restricted model, referred to as the max-LTV=0.8 model, replaces the benchmark's mortgage premium schedule with a simpler one that consists of a standard mortgage premium up to $LTV = 0.8$ and a maximum-LTV constraint at that point. The second restricted model, referred to as the max-LTV=1.0 model, allows a non-convex jump in the premium at $LTV = 0.8$ and imposes a maximum-LTV constraint at $LTV = 1.0$. In estimation of the restricted models the parameter vector γ is replaced with a restricted vector γ_R. Importantly, the optimal weighting matrix W_*^{-1} remains that computed for estimation of the unrestricted benchmark model. Because the restricted models are nested in the unrestricted benchmark model, this allows formal Wald-type tests of the validity of the restrictions.

Parameter	Description	Model Function
θ	CES Utility Weight	General home size
τ	Intra-temporal Substitution	Utility cost of intra-temporal distortion
σ	Inter-temporal Substitution	Utility Cost of inter-temporal distortion
ϕ	Non-Convex Adjustment Cost	Resource cost of re-optimizing housing
β	Discount Factor	Life-cycle consumption growth rate
r^h	Rental Premium	Resource cost of renting
r	Real Interest Rate	Benefit of positive savings
r_0^m	Basic Mortgage Premium	Basic cost of borrowing (or ownership)
r_{80}^m	Mort. Prem. Jump at $LTV = 0.8$	LTV dist. jump at 0.80
r_{100}^m	Mort. Prem. Jump at $LTV = 1.0$	LTV dist. jump at 1.00
ζ	Max Debt-to-Income Ratio	Hard borrowing limit

Table 2.2: Description of estimated parameters

2.4.2 Parameter Estimates

Table 2.2 provides a summary of the estimated parameters and a description of their roles in the model, while Table 2.3 provides the actual estimates. Estimates from both the benchmark and restricted models are presented. First I compare the benchmark estimates to those reported in previous studies, and then I contrast them with the estimates from the restricted models.

The estimate of θ is highly dependent on the scale of housing consumption. Because one model unit of housing consumption is ten square-feet, the estimate of 0.214 is difficult to compare to previous estimates.

The estimated elasticity of substitution between housing and non-housing consumption implied by the estimate of τ is approximately 0.36. This number is only slightly larger than the estimate by Li, Liu, and Yao (2009), which is consistent with micro-empirical estimates of this parameter (Hanushek and Quigley, 1980; Seigal, 2008; Flavin and Nakagawa, 2008). However, this is substantially smaller than the estimate by Bajari, Chan, Krueger, and Miller (2013), which is slightly larger than unity.

The non-convex adjustment cost of selling a home, which can be interpreted as a real-estate commission, is estimated to be just over 9.3%. This is somewhat larger than common calibration values, but is smaller than the estimate by Li, Liu, and Yao (2009).

The preference parameters β and σ are not commonly estimated in life-cycle housing studies, but rather are calibrated to commonly accepted values. The estimated values here are well within the range of those commonly accepted values. Because much of the emphasis

Parameter	Benchmark	Max-LTV=1.0	Max-LTV=0.8
θ	0.202 (0.046)	0.210 (0.064)	0.681 (0.149)
τ	-1.690 (0.294)	-1.705 (0.171)	-0.468 (0.029)
ϕ	0.0945 (0.0151)	0.1021 (0.061)	0.0346 (0.097)
β	0.960 (0.019)	0.963 (0.021)	0.955 (0.024)
σ	2.46 (1.21)	2.51 (0.345)	2.91 (0.694)
r^h	0.0356 (0.0022)	0.0354 (0.0027)	0.0563 (0.0201)
r	0.0342 (0.0061)	0.0340 (0.0063)	0.0378 (0.0400)
r_0^m	0.0471 (0.0028)	0.0474 (0.0033)	0.0356 (0.0171)
r_{80}^m	0.00179 (0.0009)	0.00152 (0.0012)	∞*
r_{100}^m	0.0205 (0.0054)	∞*	–
ζ	-0.186 (0.299)	-0.013 (68.4)	-0.559 (0.232)
prob(χ^2-stat)	–	$p < .01$	$p < .01$

Notes: For descriptions of parameters see Table 2.2. Standard errors are in parentheses below estimates. Parameters marked with an * are restricted values, rather than estimates.

Table 2.3: Parameter estimates

of this paper is on inter-temporal substitution (relative to family composition and income risk), it is important to estimate these parameters in order to allow flexibility in the desire to smooth over time.

The real interest rate estimate of 3.4% is reasonable for the time period of the data. The rental rate is slightly larger than the base mortgage interest rate (once the tax deduction is accounted for), which generates a primary motivation to own housing, but the mortgage premia for high LTV mortgages exceeds the rental premium. Then why to households choose high LTV mortgages rather than renting? The benefit of a high LTV mortgage is access to credit for life-cycle smoothing purposes. By owning a home households can borrow against future income in order to finance housing *and non-housing* consumption.

The non-convex jump in the mortgage premium at $LTV = 0.8$ is approximately 16 basis points. In reality, mortgage insurance premiums are normally paid upfront, adding 1.75% to the initial mortgage principle.[9] For a conventional 30 year mortgage with a 7% interest rate, this would increase payments by roughly the same amount that a 20 basis point increase in the mortgage interest rate would, thus 16 basis points is a reasonable estimate. The non-convex jump at $LTV = 1.0$ is 2.06%, which is very substantive and explains the reluctance of borrowers to exceed this level unless the life-cycle smoothing motive is very strong.

In comparing estimates across models I focus on the substantive differences between the benchmark and max-LTV=0.8 models. One of the striking differences between the benchmark and max-LTV=0.8 estimates is the much lower substitutability of consumption across both time and goods in the max-LTV=0.8 model. In essence these lower elasticities of substitution imply a greater utility punishment for deviations from a first-best (unconstrained) allocation. One consequence of this is that households are much more concerned about the frictions that prevent them from consuming optimally across the life-cycle. Another consequence is that households become more concerned about the impact of income risk on the intra-temporal optimality of their consumption bundle. In this regard we can understand why the estimated adjustment costs are substantially reduced when we move from the benchmark to the max-LTV=0.80 model. To induce a desire for home-ownership, households must be able to adjust with a lower real cost because the utility cost of remaining at a suboptimal bundle is much higher.

In terms of costs, the basic mortgage premium and the rental premium are estimated to be larger in the max-LTV=0.8 model than the benchmark model, but the real interest rate is estimated to be lower. This drives a larger wedge between the costs of renting and the opportunity cost of owning for those with sufficient capital to buy without a mortgage. This greater cost differential increases the desire to be a home-owner in the restricted model. Lastly, the maximum debt-to-income ratio is much smaller in the max-LTV=0.8 model than the benchmark model. The income constraint plays a much larger role in determining who

[9]This is the current FHA premium

will be a home-owner.

For completeness, the results of Wald specification tests are reported. For the max-LTV=0.8 model there are two parameter restrictions, and for the max-LTV=1.0 model there is a single parameter restriction. For both tests the benchmark model dominates; however, I place less emphasis on these results, and more on the visual comparisons presented in the next subsection. This is because the restricted models cannot possibly replicate the LTV distribution moments, and the sample size is so large that even small reductions in fit should result in rejection of the restricted model.

2.4.3 Goodness-of-Fit

The benchmark model fits the targeted moments well. Figure 2.4 illustrates the fit of the benchmark model to empirical moments of the LTV distribution. Utilization of high LTV mortgages is slightly greater in the model than in the data, and utilization of low LTV mortgages is slightly less than in the data. Because the restrictions in the alternative models prevent a reasonable fit of the model LTV distribution to the data, I do not present those results.

Home ownership is the dimension in which model performance is clearly differentiated. Figure 2.5 plots the home-ownership rate across the life-cycle for each of the estimated models and the data values, as well as model deviations from the data in the lower panel. Because the benchmark model allows for negative equity mortgages, this model generates growth in the home-ownership rate among young households that closely matches the data. Some households anticipate rapid enough income growth that they are willing to pay a large premium in order to smooth both aspects of consumption over time, and this incentive to own housing allows the model to replicate the data well. Moving to the max-LTV=1.0 and max-LTV=0.8 models results in progressively poorer fits to life-cycle home ownership rates, despite re-estimation of the other parameters. The max-LTV=1.0 model is closest to the benchmark in matching the data, but fails to attain a fast enough growth rate for ownership. A zero equity mortgage allows perfect smoothing of housing consumption, but not non-housing consumption, hence the incentives to own housing are not as high in the max-LTV=1.0 model compared to the benchmark.

Figure 2.6 plots average home size over the life-cycle for the models and data, as well as the deviation of each model from the data. There is little differentiation between the benchmark and max-LTV=1.0 models in terms of average home size, as both models perform quite well. However, the max-LTV=0.8 model generates home sizes that are too large for middle-aged households and too small late in life.

Lastly, to illustrate the fit to a non-targeted aspect of the data, Figure 2.7 plots average mortgage debt over the life-cycle. The benchmark model outperforms the max-LTV models

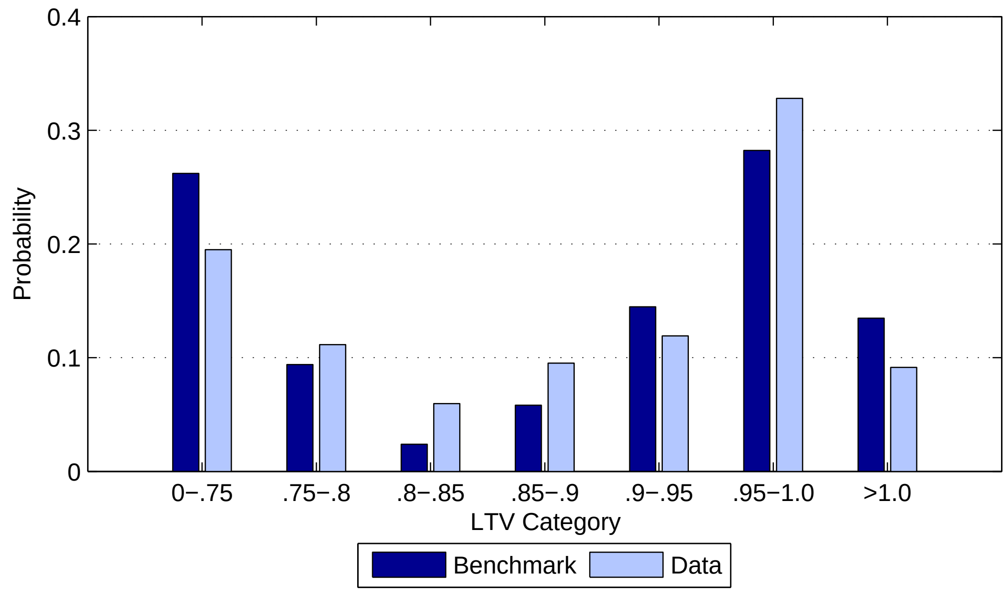

Notes: Fitted LTV moments are provided for the benchmark model only. The restricted models cannot replicate the data distribution by assumption. The benchmark model fits the empirical moments well, but there is slightly higher utilization of high LTV mortgages in the model than the data.

Figure 2.4: Fit of benchmark model to LTV data

in this dimension, particularly the max-LTV=0.8 model. The max-LTV=0.8 model does not generate enough mortgage borrowing over any stage of life. The benchmark and max-LTV=1.0 models both generate mortgage debt that is similar to the data, particularly during the early part of adulthood. However, mortgage debt declines too quickly in these models, before increasing at the end of life, which is not a feature of the data.

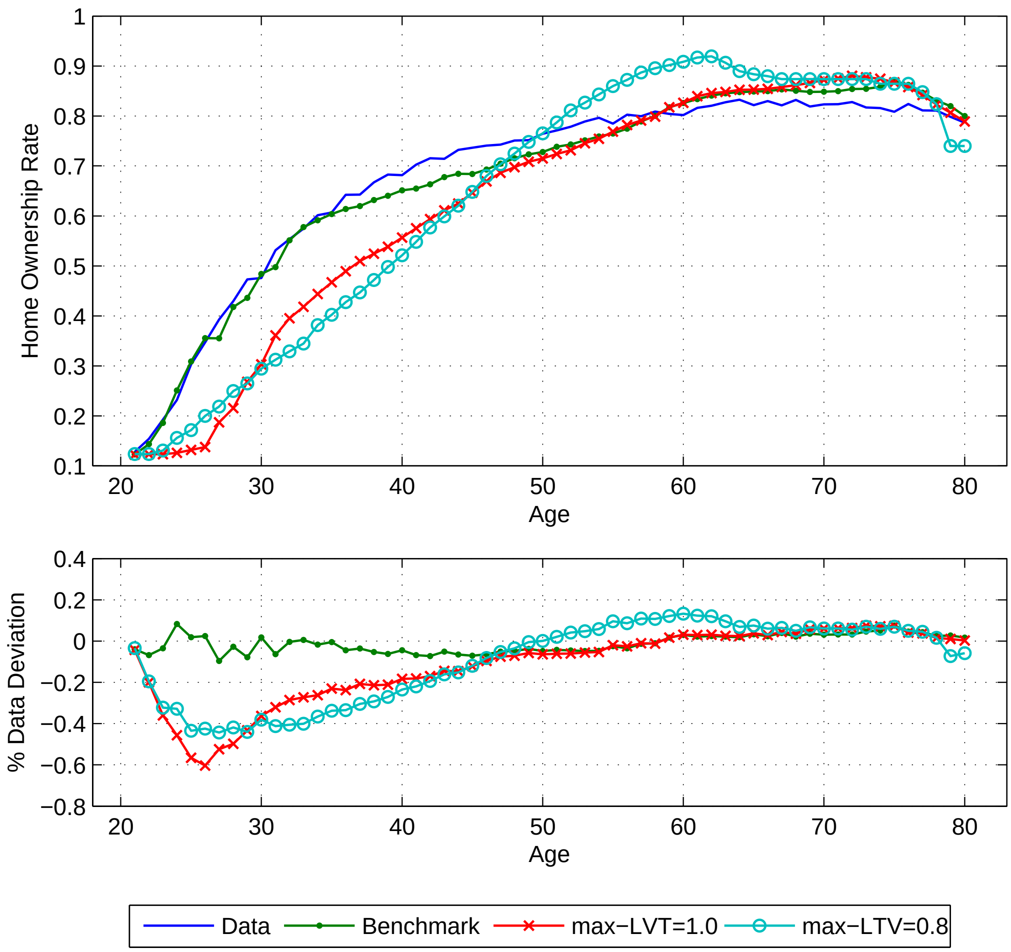

Notes: The top panel plots the home-ownership rate by age for each of the models and the AHS data. The bottom panel plots percentage deviations of the model fit from the data. The benchmark model dominates the max-LTV models in fitting home-ownership, particularly for young households.

Figure 2.5: Fit to home-ownership data

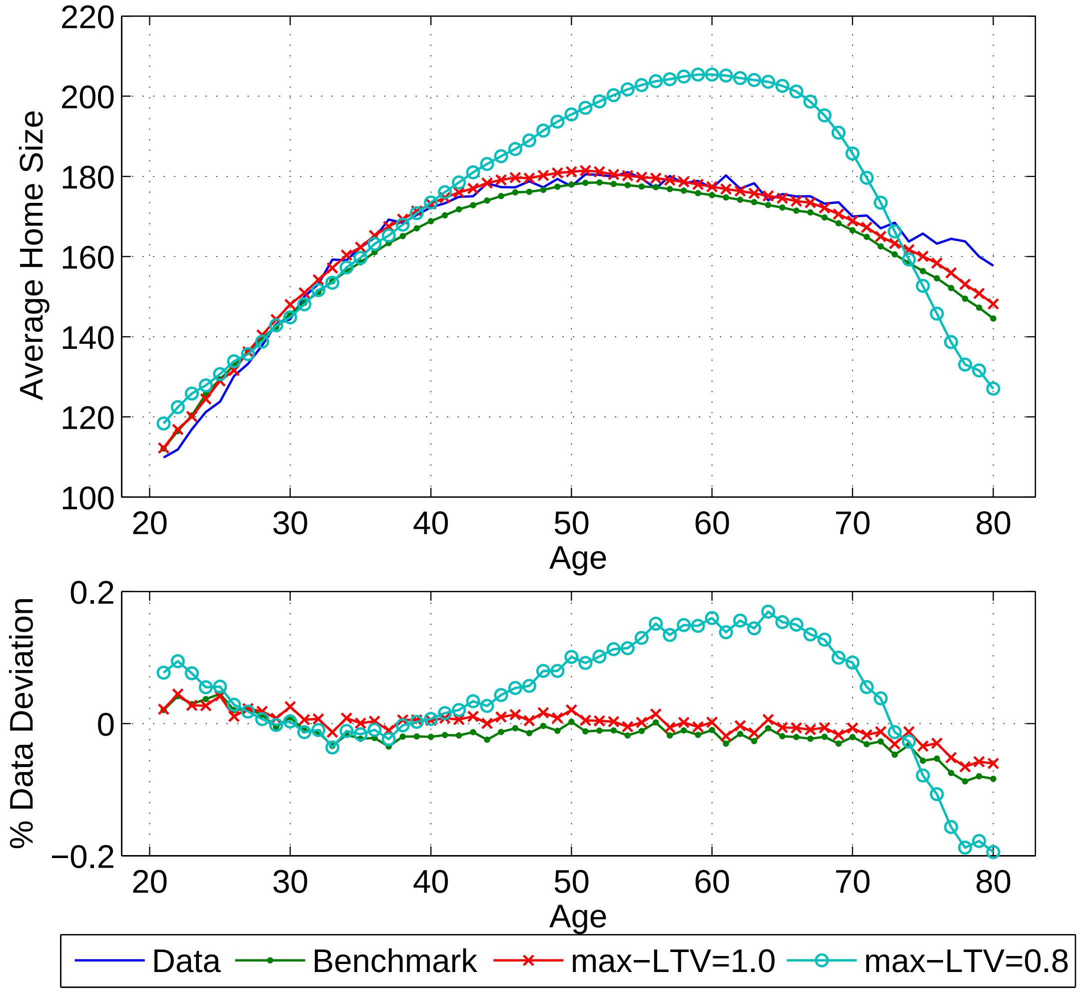

Notes: The top panel plots average home size by age for each of the three models and the AHS data. The bottom panel plots percentage deviations of of each model from the data. The benchmark and max-LTV=1.0 models are similar in fit, but the max-LTV=0.8 model fits the home size poorly for the later stages of life.

Figure 2.6: Fit to home-size data

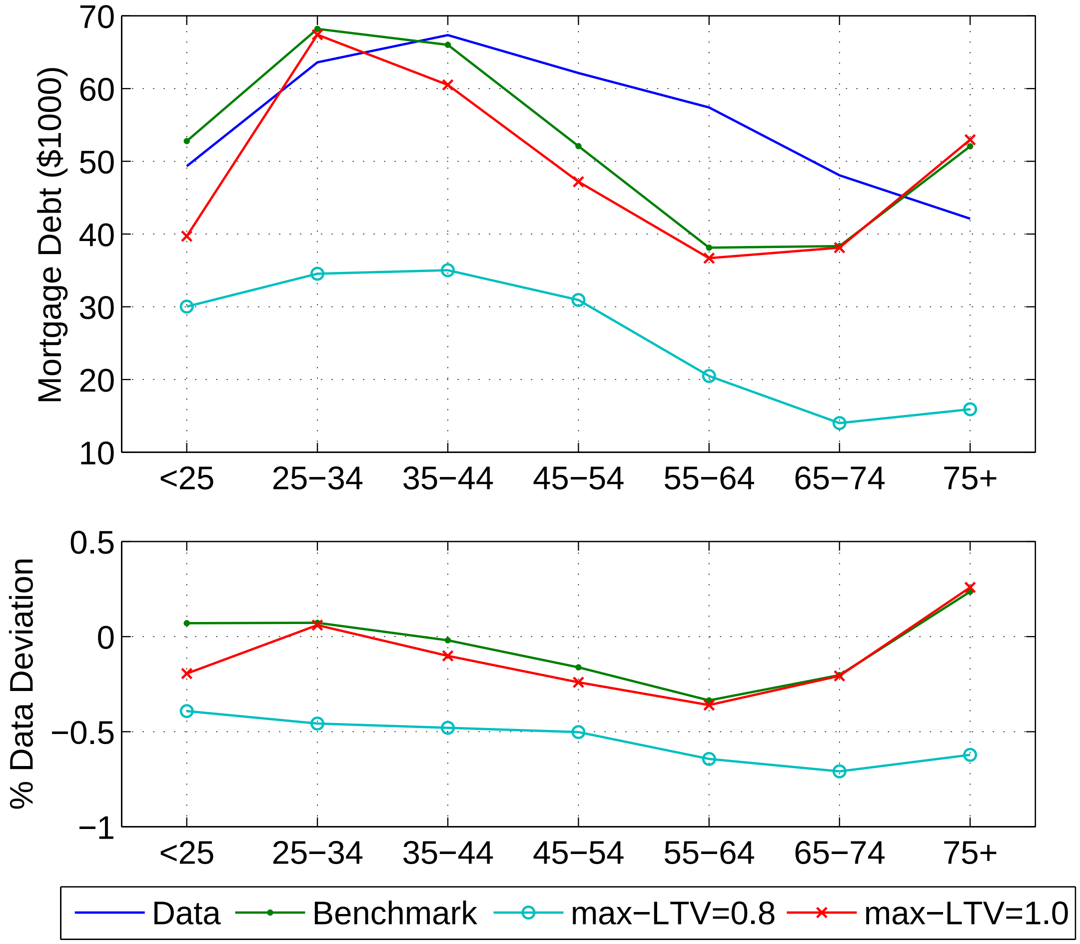

Notes: Average mortgage debt for various age groups is plotted for each model and the AHS data. The bottom panel indicates percentage deviations of each model from the data. These moments were not targeted in the estimation. The benchmark model performs only slightly better than the max-LTV=1.0 model, and the max-LTV=0.8 model does not generate nearly enough mortgage debt.

Figure 2.7: Fit to non-targeted mortgage debt data

2.5 Policy Experiments

This section reports the results of three policy experiments carried out on the benchmark and max-LTV=0.8 models. The experiments include increasing the mortgage interest subsidy, increasing the threshold LTV at which mortgage insurance is required, and relaxing the maximum debt-to-income constraint. The substantial differences in the estimated structural parameters between the benchmark and max-LTV=0.8 models leads to very different policy effects. This highlights the importance of understanding the ways in which credit constraints affect home ownership decisions when making policy decisions.

2.5.1 Mortgage Interest Subsidy

The effective interest rate paid on mortgage debt is $(1 - \tau_w)r^m$, where r^m depends on the LTV. In this experiment the subsidization rate is increased by three percentage points, from 0.27 to 0.30 percentage points. Importantly, the labor income tax rate is held constant at 0.27, which implies that mortgagors are now able to deduct slightly more than their interest payments for tax purposes.

Figure 2.8 plots home ownership across the life-cycle with and without the increased subsidy, where the top panel applies to the benchmark model and the bottom panel applies to the max-LTV=0.8 model. The increased subsidy has a large effect on the home-ownership decisions of agents in the benchmark model, but has very little effect in the max-LTV=0.8 model. Focusing on households 40 or younger, an increase in the subsidy causes the home-ownership rate to increase by almost 15 percentage points in the benchmark model, but by only 0.5 percentage points in the max-LTV=0.8 model.

The disparity in the effects of the policy change across models is due to differences in the forces that shape home-ownership decisions. In the benchmark model the relative cost of owning versus renting is the main tension. When the interest subsidy rises (holding the rental rate constant) owning becomes the lower cost option for a substantial number of households. However, in the max-LTV=0.8 model relative costs are less important because credit constraints are the main force preventing home-ownership.

2.5.2 Mortgage Insurance Threshold

This experiment considers the effect of moving the mortgage premium non-convexity from LTV=0.8 to LTV=0.9, which can be interpreted as relaxing the threshold for insurance requirements. This is a natural interpretation for the benchmark model where the interest jump at LTV=0.8 reflects increasing costs due to insurance premiums. For the max-LTV=0.8 model this interpretation is less obvious; however, the main motivation for setting maximum LTV constraints at 0.8 in the literature has been the requirement to pay mortgage insurance

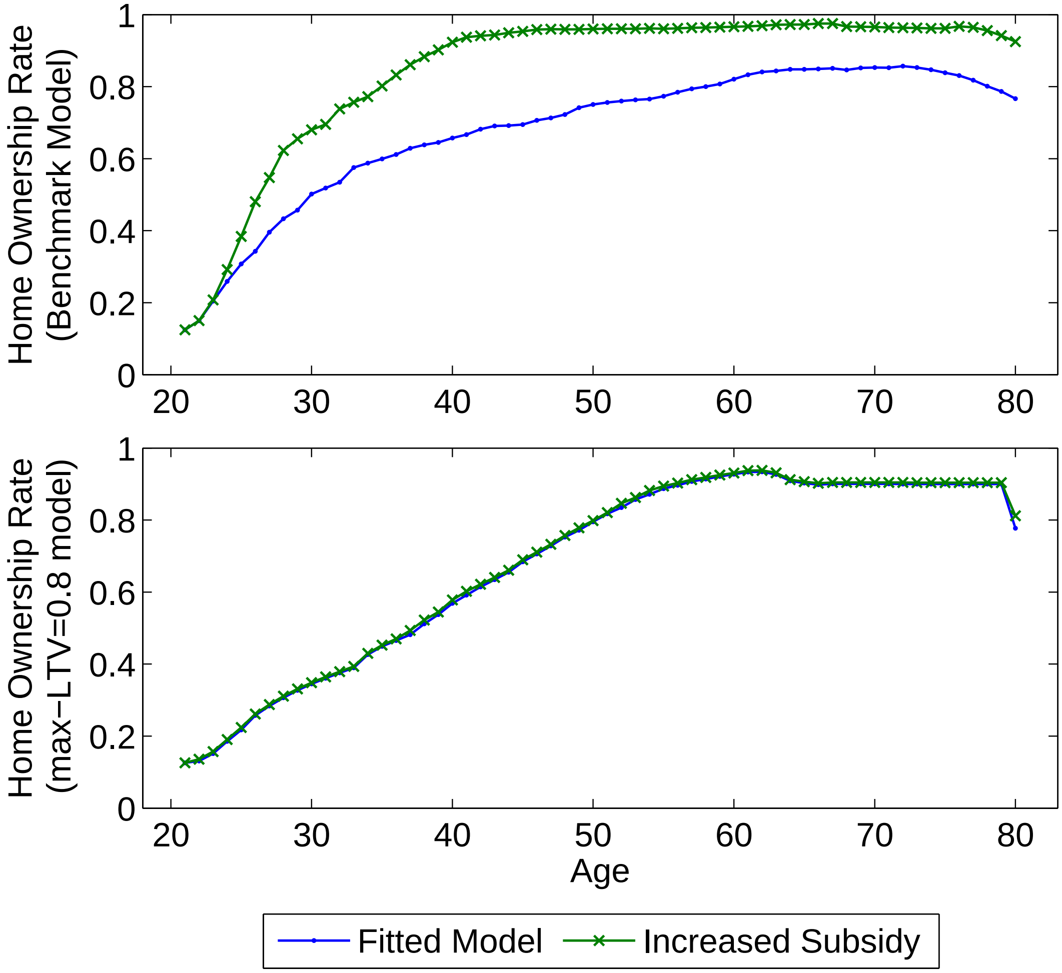

Notes: Only the benchmark model responds to a mortgage interest subsidy because the main tension preventing ownership in that model is relative cost, whereas resource availability is the main tension in the max-LTV=0.8 model.

Figure 2.8: Home-ownership effect of increased mortgage interest subsidy

premiums beyond this level.

Figure 2.9 plots the effects of increasing the LTV threshold from 0.8 to 0.9. For both models there are modest effects on home ownership. The home-ownership rate for those 40 or younger rises by 3.5 percentage points in the benchmark model, and by 4 percentage points in the max-LTV=0.8 model.

As with the first experiment, the reasons for the changes are different. In the benchmark model the cost of owning a home with an LTV between 0.8 and 0.9 falls in the experiment, making home ownership relatively cheaper. In the max-LTV=0.8 model LTVs in this range become possible, whereas they were not previously.

2.5.3 Maximum Debt-to-Income Ratio

In this experiment the parameter ζ is increased by 0.05 percentage points, which represents a loosening of the debt-to-income ratio constraint. The interpretation of this change is that a households income is allowed to be 5 percentage points smaller relative to their debt. Note that under the benchmark model estimates households can borrow about five times their income, whereas under the max-LTV=0.8 model estimates they can only borrow about twice their income. In comparison to the above experiment in which the insurance threshold LTV is relaxed, this experiment is a pure relaxation of credit constraints without lowering borrowing costs in the benchmark model.

Figure 2.10 illustrates the effects of relaxing the debt-to-income ratio constraint. For the benchmark model the only discernable effects are late in life, when households have low (pension) incomes and are attempting to borrow against their housing wealth. For the max-LTV=0.8 model there is a moderate increase in home-ownership across the life-cycle. In the max-LTV=0.8 model credit constraints are the main tension that shapes ownership decisions, and the income based constraint is an important component of that.

2.5.4 Discussion

A common theme in comparing the results of the counterfactual experiments across models is that differences are generated by the differing tensions that shape housing tenure decisions. In the canonical max-LTV=0.8 model credit frictions entirely prevent low wealth or low (relative) income households from buying homes, whereas in the less restrictive benchmark these characteristics only increase the relative cost of owning a home.

From the mortgage insurance threshold experiment we can conclude that the costs associated with the frictions measured in the benchmark model estimation can have meaningful effects on ownership decisions. Thus, although the credit frictions are more subtle in the benchmark model, they are still a very important determinant of life-cycle housing consumption patterns.

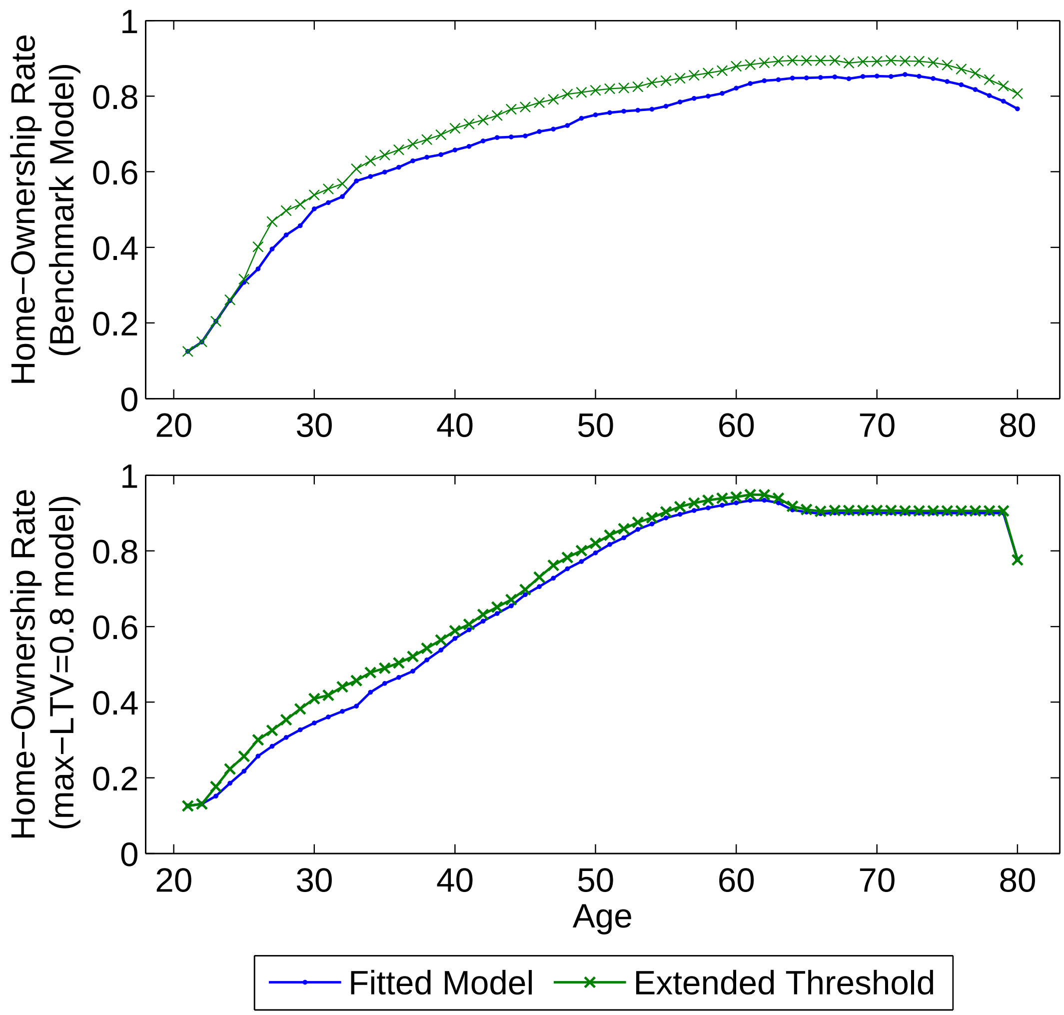

Notes: Increasing the threshold at which insurance is required from 0.8 to 0.9 increases ownership in both models. In the benchmark model this is due to a lower relative cost of ownership in that range, and in the max-LTV=0.8 model it is because fewer resources are required in order to buy a home.

Figure 2.9: Home-ownership effect of extended insurance requirement threshold

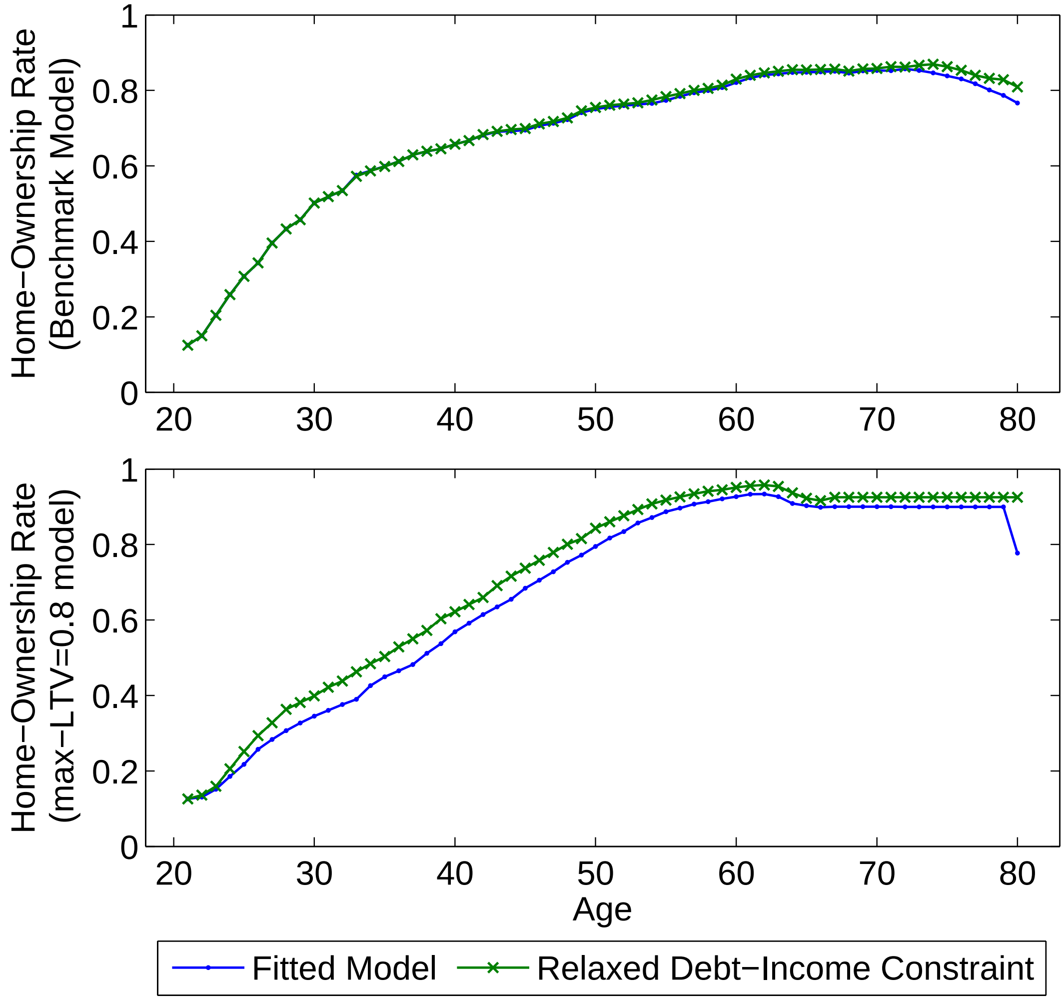

Notes: This experiment involves a pure reduction in the resources required to buy a home. this causes increased in ownership in the max-LTV=0.8 model, but has little effect in the benchmark model.

Figure 2.10: Home-ownership effect of relaxed Debt-to-Income constraint

Though credit market frictions are important in both models, the way that they affect home-ownership decisions is very different, which leads to very different policy implications. Because the relative cost of rental versus owner-occupied housing is the important determinant of tenure in the benchmark model, mortgage interest subsidies have large effects on ownership rates. However, binding credit constraints are less important in the benchmark, which means that relaxed lending standards have less effect. In the max-LTV=0.8 model binding credit constraints are the main tension, thus relaxed lending standards have meaningful effects in that model, whereas reduced mortgage insurance premiums do not.

2.6 Conclusion

The paper has explored the importance of allowing for high LTV mortgages in models of life-cycle housing decisions. I found that a model with high LTV mortgages and a mortgage interest rate schedule that is increasing in the LTV fits observed data much better than models with maximum LTV constraints. Furthermore, the effects of policies intended to increase home-ownership can be far different in these type of models.

I presented data on the distribution of LTVs, and how this distribution varies over the life-cycle. There are jumps in the density of LTVs at 0.8 and 1.0, which indicate non-convexities in the schedule of borrowing costs. The estimated size of the non-convex jump at LTV=0.8 is consistent with mortgage insurance premiums. The non-convex jump at LTV=1.0 is approximately two percentage points, which is substantial and reflective of the secondary markets for these types of loans.

I showed that both the benchmark model and the model with a 1.0 maximum LTV do reasonably well in replicating home-size by age, but the model with a 0.8 maximum LTV does not. The biggest gains in explaining the data that arise when high LTV mortgages are allowed for are in early life home-ownership patterns. Young households use negative equity mortgages to smooth non-housing and housing consumption, which generates enough incentive to own homes to replicate the rapid growth of ownership observed in early adulthood. The restricted models perform successively worse in this dimension as the maximum LTV constraint becomes tighter.

A very important difference between the estimated benchmark and max-LTV=0.8 models is the primary margin that dictates ownership decisions. In the benchmark model tenure decisions reflect the relative costs of owning and renting, whereas in the max-LTV=0.8 model renters are those who do not have sufficient resources to overcome strict borrowing constraints. These differences lead to differences in the effects of policy changes. Mortgage interest subsidies, which affect the relative cost of ownership, only improve ownership rates in the benchmark model. The ownership rate in the max-LTV=0.8 model is only responsive to changes in lending standards. This highlights the importance of understanding the mechanisms by which credit frictions work when formulating public policy.

Chapter 3

Education Policy and Intergenerational Transfers in Equilibrium

3.1 Introduction

Investment in human capital is a key source of aggregate productivity growth and a powerful vehicle for social mobility. Motivated by these considerations, governments promote the acquisition of education through a variety of interventions. Financial aid for college students is a pillar of education policy in many countries. For example, the US Federal government spent roughly 150 billion dollars on loans and grants for college students in 2012 (source: Trends in Student Aid, College Board, 2012). Given the vast resources spent by these programs it is paramount to accurately quantify the effects of policies intended to advance college enrollment.

In this paper, we address this question by providing an empirical and quantitative analysis of the impact of financial aid policies on college attainment and the aggregate economy. Measuring how government-sponsored loans and tuition grants affect college enrollment decisions is extremely challenging because of the interplay between four crucial economic factors: (i) individual heterogeneity and uncertainty in the returns to college education, (ii) imperfections in financial markets, (iii) private sources of funding, and (iv) general equilibrium feedback effects.

As recognized by the micro-econometric literature, there is extensive heterogeneity in the return to education. Higher ability individuals have higher expected pecuniary returns from higher education and self-select into college.[10] There is also substantial heterogeneity in non-pecuniary (psychic) costs of college attendance. Modelling these psychic costs is necessary because pecuniary returns can only account for a part of the observed college attendance patterns by ability (see Cunha, Heckman, and Navarro, 2005; Heckman, Lochner, and Todd, 2006). This vast heterogeneity means that unless policies are precisely targeted towards

[10] The first studies linking human capital investment to life cycle earnings (Mincer, 1958; Becker, 1964; Ben-Porath, 1967) sidestepped the important issue of self-selection into education, as described in the seminal contributions of Rosen (1977) and Willis and Rosen (1979).

specific groups most students are infra-marginal and largely unaffected by an intervention.

Labor economists have convincingly documented that earnings risk during working life is substantial and only partially insurable through borrowing, saving, labor supply adjustments, and family transfers (Blundell, Pistaferri, and Preston, 2008b; Low, Meghir, and Pistaferri, 2010; Heathcote, Storesletten, and Violante, 2012; Gallipoli and Turner, 2011). As Levhari and Weiss (1974) originally emphasized, college education is a multi-period investment that requires an ex-ante commitment of resources and time, and hence uncertainty in its return is a key determinant of education decisions. Therefore students may be unwilling to finance college using loans when uncertainty about their future earnings and ability to repay is high. In this respect grants may be a better instrument to promote enrollment (e.g., Johnson, 2011).

At least since Becker (1964) it has been well understood that the amount of college education in an economy may not be optimal because young individuals own only small amounts of pledgeable assets and their future human capital cannot be used as collateral in imperfect financial markets. The key question is: how pervasive are credit constraints? Studies using the 1979 cohort of the National Longitudinal Survey of Youth (NLSY79) concluded that in the 1980s family income played a small role in college-attendance decisions, after controlling for child ability and several other family background characteristics (Cameron and Heckman, 1998; Keane and Wolpin, 2001; Carneiro and Heckman, 2002; Cameron and Taber, 2004). This result implies a stark policy recommendation: rather than easing credit constraints at the age of college enrollment decisions, government interventions should be directed to earlier stages of life in order to ameliorate college preparedness and reduce "psychic costs" of schooling (e.g., by improving the quality of the public school system). However, more recent studies based on the 1997 cohort of the NLSY (NLSY97) have reached a different conclusion. For example, Belley and Lochner (2007) find that parental financial resources matter much more for college attendance in the 2000s (they estimate roughly twice the effect observed in the 1980s). These new contributions reopened the debate on the role of financial aid for college students.

The extent to which credit constraints distort college attendance depends on the availability of other private sources of funding, such as parental transfers and in-school part-time work. Garriga and Keightley (2009) show that omitting the labor supply margin of college students may lead to large overestimates in the effects of tuition subsidies. Gale and Scholz (1994) show that inter vivos transfers (IVTs) for education are sizeable, thus they should be incorporated into models of education acquisition. Keane and Wolpin (2001) and Johnson (2011) estimate parental IVTs as a function of observable characteristics from the NLSY79. Brown, Scholz, and Seshadri (2012) show that while parental contributions are assumed and expected in financial aid packages they are not legally enforceable nor universally given, implying substantial heterogeneity in access to resources for students with observationally similar families. Using this insight Winter (2012) argues that ignoring parental transfers may lead to wrong inference about the extent of credit constraints. We build on this body of evidence and account for some of its criticisms by developing a structural model where the parental transfer decision rule determines the distribution of initial wealth in equilibrium. A crucial implication of our framework is that more generous education policies may crowd out parental IVTs, and hence displace students' use of private financial resources to attend college.

The evaluation of large-scale financial aid policies, such as Stafford loans or Pell grants, targeted to the wider population rather than specific groups requires a general equilibrium model. Heckman, Lochner, and Taber (1998b,c) have led the way in advocating an approach to evaluation of education policies based on structural models that do not omit equilibrium feedbacks between aggregate quantities of labor of different education types and their relative prices. This approach has been also followed by Lee (2005), Lee and Wolpin (2006), Garriga and Keightley (2009), Bohacek and Kapicka (2012), Johnson and Keane (2013), and Krueger and Ludwig (2013).

Our framework of analysis combines these elements. We build a life-cycle, heterogeneous-agent model with incomplete markets of the type popularized by Huggett (1993) and Ríos-Rull (1995) and with inter-generational links in the tradition of Laitner (1992). Young individuals make education decisions based on productive ability (inherited from parents), psychic costs of schooling, and wealth. Imperfectly altruistic parents make inter vivos transfers that determine their children's initial wealth. Labor supply during college, government grants and loans, and private education loans, complement parental transfers as sources of funding for college education. During working life individuals make labor supply and consumption/saving decisions, and repay student loans. They face uninsurable random fluctuations in their labor productivity throughout the work stage of life. Labor inputs of different education levels are imperfect substitutes in the aggregate production function. Wages and the real interest rate clear input markets.

We obtain parameter estimates for our model by estimating key components such as the education specific stochastic wage processes and the aggregate production function by combining data from the PSID, the NLSY, the CPS and the macroeconomy. For some parameters where microeconometric estimation is not feasible or practical we use calibration to replicate key moments from the data. The US federal system of grants and loans is modelled in some detail in order to ensure that we capture the main sources of public funding for education allowing us to obtain a better estimate of the incidence of liquidity constraints and their impact. To lend additional credibility to our economic model we establish that simulated data are consistent with empirical data along a number of crucial dimensions that are not targeted in the parameterization. For example, cross-sectional life-cycle profiles of the mean and dispersion of hours worked, earnings, consumption, and wealth are consistent with their empirical counterparts (e.g. Guvenen, 2009; Heathcote, Storesletten, and Violante, 2012; Kaplan, 2011). The intergenerational correlation of income between parents and children is around 0.4, as documented by Solon (1999) for the US. Our modeling choices for federal financial aid imply marginal effects of parental wealth on college enrollment, controlling for child's ability, that are similar to those estimated by Belley and Lochner (2007) from the NLSY97. Moreover, when we use the model to simulate a randomized experiment in which a (treated) group of high-school graduates receives an additional $1,000 in yearly tuition grants and another (control) group does not, we estimate a treatment effect on college enrollment that is slightly larger than 3 percentage points. This estimate is consistent with the effects of quasi-randomized policy shifts surveyed by Kane (2003), and Deming and Dynarski (2009).

We conduct a number of different policy experiments in which we change the size and nature (need-based/merit-based) of the federal grant program and government-sponsored loan limits. These experiments yield five main results. First, existing federal grant programs improve welfare in our model and, by promoting skill acquisition through college education, add about 1.5% to aggregate output. Second, the expansion of tuition subsidies induces nontrivial crowding out in both parental IVTs and in the labor supply of college students. We estimate that a $1,000 reduction in tuition fees lowers annual hours worked by college students by 4%. Further, we estimate that every additional dollar of government grants crowds out 20-30 cents of parental IVTs on average. There is considerable variation in crowding out. One important difference is that richer parents' transfers are crowded out considerably more than poorer parents' transfers. This introduces a regressive element to increases in grants. Third, expansions of the grant program may have sizeable effects on enrollment in the short run, but long-run general equilibrium effects are 3-4 times smaller. Besides the attenuating role played by the crowding out of parental IVTs and student's labor supply, the central force at work is the relative price response. As college enrollment rises in response to more generous grants, the relative price of college educated labor falls and offsets the direct partial equilibrium impact of the policy on the quantity of college graduates. We verify that the equilibrium adjustment of the interest rate and the tax rate has only minor consequences. Fourth, government-sponsored loans are an especially valuable source of college financing for high-ability children whose parents cannot afford to fund the entire cost of college. These children recognize the high return from college education for their ability type, and are willing to borrow, heavily at times, to acquire a tertiary degree. As a result, the current government-sponsored loan program is welfare improving in the model, and we find that removing it would reduce aggregate output by 1.5 percentage points. Indeed, we estimate that the combined system of federal aid to college students (grants and loans) is worth 2.5 percent of GDP. Fifth, it is more effective to condition aid on family means than student ability. This result follows from the fact that, given the existing institutional and credit environment, most high ability individuals would choose to attend college regardless of the additional transfers, meaning that ability-testing targets many individuals who are infra-marginal to the education choice.

Our calculations suggest that borrowing limits for college students are already quite generous: expanding them would have negligible consequences on enrollment, output and welfare in the long run. However, we also find that parental wealth is a significant determinant of college attendance, as is the case in the NLSY97 data (Belley and Lochner, 2007). We reconcile these findings by noting that estimated psychic costs of college attendance decrease in family wealth. Therefore, relaxing credit limits for able children of poor households does not augment college enrollment much because these high-school students have large disutility of further education.[11] We interpret this correlation as capturing the fact that the quality of local public schools is highly linked to the neighborhood housing price and wealth level of its residents (as in, e.g., Fernandez and Rogerson, 1998). This interpretation is not inconsistent with the relevance of earlier credit constraints, as discussed in Restuccia and Urrutia (2004) and Caucutt and Lochner (2012), or with the substantial amount of heterogeneity in students' preferences for college documented by Fu (2012) who shows that a large fraction of students, mainly low-ability ones, may prefer the outside option over any college option.

The remainder of the paper is organized as follows. Section 3.2 outlines the model economy and defines equilibrium. Section 3.3 describes the model's parameterization. Section 3.4 further validates the model by assessing its behavior along several key dimensions not explicitly targeted in the calibration. Section 3.5 presents all the policy experiments. Section 3.6 provides a general discussion of the main findings and some context. Section 3.7 concludes. The Appendix contains additional details on the parametrization and on the results of the policy experiments.

3.2 Model

3.2.1 Brief Overview

We specify an overlapping generations general equilibrium model. Individuals start by making sequential education decisions: whether to attend and complete high school and, following that, whether to attend and complete college. When education is complete, they choose their work hours until retirement. Idiosyncratic uninsurable labor market risk makes individual earnings, and hence the return to education, stochastic. Throughout their life individuals choose their consumption expenditures. They can borrow only up to a limit, and save through a non state-contingent asset. The life cycle has a maximum length, but individuals may die earlier. There is no aggregate risk.

The alternative to education is either work or leisure: these define the opportunity cost

[11] Johnson (2011) also finds small effects of relaxing borrowing constraints in spite of estimating his model on the NLSY97. The reason is that students in his model are very reluctant to take up education loans in the first place due to the uncertainty about dropping out of college and future earnings risk. Lochner and Monje-Naranjo (2011) point out that one problem with this result is that students actually borrow much more in the data than in the model. Our findings are not subject to this critique since in calibrating the model we match numerous statistics on student access to credit, including their average cumulative loans at graduation.

of education. In addition to this opportunity cost, individuals face tuition costs and psychic costs of education that depend both on their ability and family background. While in college students can support themselves by choosing to work, running down savings, and obtaining grants and government or private loans; access to these external sources of funding depends on parental wealth.

Importantly, savings originate from parental transfers that are assumed to occur just before the first education decision is made. Preferences are altruistically linked so that the parent values the child's utility at some rate equal to, or below, her own. Because of altruism, the optimal size of a parental transfer depends on the marginal value of wealth of their child. This, in turn, depends on the child's ability (which the parent knows at that point) and, as a result, on the level of education that the child is expected to achieve given any level of transfers. Ability itself is assumed exogenous, but correlated to that of their parents through an estimated intergenerational transition matrix.

The economy includes a production sector with an aggregate production function whose inputs are physical capital and three different types of human capital, corresponding to the three levels of education (statutory, high school and college). The human capital inputs are the efficiency units supplied at the going prices and depend on the number of individuals working, their education, their ability, and their stochastic individual productivity. Since the various types of human capital are imperfect substitutes (based on our estimates and others before us) the returns to education are endogenous and will depend on the relative supplies of each type. All prices, including the return to physical capital, are determined by market clearing.

The government sector consumes, runs a social security system, raises taxes on income and consumption, and funds college education either by direct subsidies or by offering loans at a subsidized rate to lower income individuals and at a market rate to those from middle income families whose parental guarantees are not sufficient to obtain loans from the private sector. In addition, there is a rudimentary private banking sector that intermediates loans at an exogenously given rate.

3.2.2 Demographics and the Life Cycle

Demographics: The economy is populated by $J+1$ overlapping generations. Let $j = 0, \ldots, J$ denote age. The probability of surviving from age $j-1$ to age j is denoted by ζ_j. We let $\zeta_j = 1$ as long as the individual is in school or at work $\left(j \leq j^{WK}\right)$, but $\zeta_j < 1$ during retirement, from $j = j^{WK} + 1$ to J. Conditional on reaching age J, death is certain at the end of the period $(\zeta_{J+1} = 0)$. We set the size of the newborn cohort so that total population is normalized to 1.

Life Cycle: The life cycle of an individual has three distinct stages. In the first stage, the individual goes to school and acquires education. There are three levels of educational attainment: Less than High-School, High-School degree and College degree, which are denoted by $e \in \{LH, HS, CL\}$, respectively. Let j^e denote the last period of the school cycle e, with the convention that $j^{LH} = -1$ is the last period of compulsory high-school education. Until that age individuals are dependent "children". Starting from the following period ($j = 0$) individuals begin making independent decisions.

At age $j = 0$ individuals immediately choose whether to drop out of high school or continue. This decision, denoted $d^{HS} \in \{0, 1\}$, entails commitment to enroll in school until age j^{HS}. At age $j^{HS} + 1$ the agent decides whether to enroll in college, a choice which we denote by $d^{CL} \in \{0, 1\}$. This decision requires full commitment to completing college at age j^{CL} because drop-outs are not modeled.[12] During schooling, students choose their level of consumption/saving. Labor supply in college is flexible, but the time endowment available for work is reduced by $\bar{t}$ units to reflect the time required for learning. High school students do not work and their leisure is exogenously fixed at $\bar{l}$.

Agents begin the work stage of their lives at age 0, $j^{HS} + 1$ or $j^{CL} + 1$ depending on their education decision. During this stage, which lasts until mandatory retirement age j^{WK}, agents choose labor supply and consumption/saving. Retirement starts at age $j^{WK} + 1$. During retirement, individuals do not work ($l = 1$), receive a pension from the government, and allocate consumption/saving over their remaining lifetime of uncertain length.

3.2.3 Preferences and Intergenerational Links

Preferences: The period utility of workers and retirees $u(c_j, l_j)$ is strictly increasing and strictly concave in consumption $c \geq 0$ and leisure $l \in [0, 1]$, continuously differentiable, and satisfies Inada conditions. Utility in school has an additional separable component, $\kappa^e(\theta, q)$, $e \in \{HS, CL\}$, which is a function of fixed individual innate "ability" $\theta \in [\theta_{\min}, \theta_{\max}]$ and parental wealth group $q \in \{1, 2, 3\}$ (explained below). The function $\kappa^e(\theta, q)$ reflects psychic costs of schooling in terms of effort, preparedness, or taste for education (see, e.g., Heckman, Lochner, and Todd, 2006).

Intergenerational Links: Individuals are partially altruistic towards their offspring. Their child's expected lifetime utility enters their own value function with weight $\omega \in [0, 1]$. This one-sided altruism manifests itself as a monetary transfer once in the lifetime. At age j^{TR} (during the work stage) each individual (now a parent) has the opportunity to choose a non-negative amount of resources to transfer to their child. The parental transfer fully determines

[12] To avoid further complexity, we abstract from modelling the college drop-out decision. The model could be easily extended by introducing a "disutility shock" of attending college, i.e., a stochastic component of the psychic cost (κ^{CL}, see below), realized after the college attendance decision, whose distribution could depend on individual characteristics.

the child's initial asset level a_0. For tractability we do not model multiple transfers or endogenous timing of transfers. Our focus is on transfers made during the late teens and college years, which can be captured reasonably well as a one-off lump sum.

Individuals are also linked by the intergenerational transmission of ability. A parent with ability θ has a probability of having a child with ability less than or equal to $\hat{\theta}$ determined by the conditional c.d.f. $\Gamma_\theta\left(\hat{\theta}, \theta\right)$. Parents know the function Γ_θ, but only at age j^{TR} (just before the inter vivos transfer) the ability of the child is fully revealed to both.

A final intergenerational linkage arises from the dependence of a child's education financing opportunities (through loans and grants) on parental wealth. This is discussed in more detail below.

3.2.4 Individual Labor Productivity

Individual labor efficiency ε_j^e for an individual of education e at age j is the sum of three components in logs,

$$\log \varepsilon_j^e = \lambda^e \log \theta + \xi_j^e + z_j^e \tag{3.1}$$

where λ^e is an education-specific loading factor on (log-) ability, ξ_j^e is an education-specific age profile for productivity, and z_j^e is a stochastic component drawn from the education-specific c.d.f. $\Gamma_z^e\left(z_{j+1}, z_j\right)$ describing the conditional cumulative probability of a realization less than or equal to z_{j+1} at age $j+1$ when the idiosyncratic stochastic component at age j was z_j. Let Γ_0^e denote the initial distribution of productivity upon entry in the labor market with educational level e. Finally, we assume that the labor services of a college student are equivalent to those of the average high-school graduate of the same age and ability θ.[13]

3.2.5 Commodities, Technology, and Markets

Commodities: There are two kinds of commodities in the economy: (i) the final good, which can be used for private/public consumption, investment, education services, and intermediation services provided by the banking sector; and (ii) efficiency units of the three types of labor. They are all exchanged in competitive markets. We let the price of the final good act as the numeraire.

Production technology: The final good is produced by a representative firm which operates a constant returns to scale (CRS) technology

$$F\left(K, \mathcal{H}\left(H_{LH}, H_{HS}, H_{CL}\right)\right)$$

employing physical capital K, which depreciates at rate $\delta \in (0,1)$, and the three types of human capital bundled in the aggregator $\mathcal{H}$, also displaying CRS. Each human capital stock

[13] For simplicity we abstract from the fact that part time work may be paid less.

H_e is the sum of individual hours worked times efficiency units of labor, ε_j^e, over all working-age individuals within each education group. Recall that the stock H_{HS} is also augmented by the effective labor supply of the college students. We denote by w^e the equilibrium price of an effective hour of labor of type e.

Education sector: The education sector offers a range of college degrees. Each degree has the same pecuniary return, but different non-pecuniary attributes, and hence different operating costs ϕ per year of college, per student.[14] Since the sector is competitive, ϕ is also the price of attending a year of college faced by the student, i.e. the tuition fees (before grants and loans). We summarize this heterogeneity through the distribution $\Phi\left(\bar{\phi}, \sigma_\phi\right)$. High school education is financed by the government, and is included in government expenditures G.

Financial Assets and Markets: There are three financial assets, all risk-free, traded in competitive markets: (i) a claim on physical capital used as a vehicle for saving, with equilibrium interest rate r; (ii) one-period government bonds carrying the same interest rate r by no-arbitrage; and (iii) a one-period private loan contract exchanged among households through the banking system. Households with positive savings receive from banks an equilibrium interest rate which must equal r (again, by no-arbitrage). Banks lend the funds to other households with borrowing needs at the rate $r^p = r + \iota$, where the wedge between the two interest rates is the cost of overseeing the loans ($\iota > 0$ per unit of consumption intermediated).

Individuals face a private debt limit that varies with phases of the life-cycle. Retirees and high-school students cannot borrow. In the work-stage, agents can borrow in private markets up to a limit $\underline{a}$. A subset of college students –those whose parental net worth is above a given threshold a^{**}— can borrow privately up to $\underline{a}^p$, at the equilibrium interest rate rate r^p. We think of these students as having either an excellent credit score, or as being safe borrowers from the banks' viewpoint because of their parental wealth.[15]

Finally, there are perfect annuity markets insuring retired households' survival risk.

3.2.6 Education and Fiscal Policies

The government offers grants and loans to help students who are considering college education.[16] We model the financial aid program to reproduce the key features of the US system

[14] Relevant non-pecuniary attributes explaining differences in tuition fees are prestige, location, characteristics of student body, infrastructures, etc. Actual differences in tuition fees clearly also reflect differentials in the quality of the degree. The model could be extended to incorporate this additional dimension of heterogeneity, but to avoid further complexity we have abstracted from college quality.

[15] Actual interest rates on private educations loans depend on the credit score because of default risk. See Ionescu and Simpson (2012). As a result, poor families with low credit scores face high borrowing rates on private education loans. Implicitly, we are assuming that these rates are so high that these families choose not to use the private market to finance education of their kids. We choose a^{**} to replicate the fraction of households who borrow privately.

[16] Some states make large transfers to local colleges. Since we focus on federal policies, we exclude these transfers from the model's government budget constraint. These transfers also explain part of the variation in

which we describe in Appendix F. The government assesses parental wealth at the age of the inter vivos transfer j^{TR} to determine eligibility status q of the child for financial aid.

Education Loans: If parental wealth $a_{j^{TR}}$ is below the threshold a^*, then $q = 1$ and children qualify for subsidized government loans up to a limit $\underline{b}^s$. Interest on subsidized loans is forgiven during college, and cumulates at rate r^s during working life. Students of type $q = 1$ who have reached the borrowing limit for subsidized loans can access additional unsubsidized loans up to $\underline{b}^u$. Unsubsidized loans cumulate interest at rate r^u both during and after college.

If parental wealth $a_{j^{TR}}$ is between a^* and a^{**}, then $q = 2$ and children qualify only for unsubsidized loans up to the cumulative limit $\underline{b}^s+\underline{b}^u$. Recall that if parental wealth is above a^{**}, students can also borrow privately at the rate r^p. Because $r^u \geq r^p$, these students will use government loans only if they need to borrow beyond the private borrowing limit $\underline{a}^p$. For this third group of students, $q = 3$.

All government loans are subject to a fixed repayment scheme: for n periods after the start of employment, the individual repays an amount π every period until exhaustion of all the principal plus interest.[17] Therefore, the last period of repayment in the individual life cycle is $j^{CL} + n < j^{TR}$.[18] In summary, the key exogenous policy parameters of the education loan program are $\{n, a^*, r^s, r^u, \underline{b}^s, \underline{b}^u\}$.

If at the end of college the individual has an amount $b_{j^{CL}} < 0$ of education debt, π is determined by the actuarial formula

$$\pi = \begin{cases} -\frac{r^s}{1-(1+r^s)^{-n}} b_{j^{CL}} & \text{if } q = 1 \text{ and } -\underline{b}^s \leq b_{j^{CL}} < 0 \\ \frac{r^s}{1-(1+r^s)^{-n}} \underline{b}^s - \frac{r^u}{1-(1+r^u)^{-n}} \left(b_{j^{CL}} + \underline{b}^s\right) & \text{if } q = 1 \text{ and } b_{j^{CL}} < -\underline{b}^s \\ -\frac{r^u}{1-(1+r^u)^{-n}} b_{j^{CL}} & \text{if } q \in \{2, 3\} \text{ and } b_{j^{CL}} < 0 \end{cases} \tag{3.2}$$

which shows that, given the policy parameter triplet (r^s, r^u, n), there is a one-to-one mapping between the pair $\left(b_{j^{CL}}, q\right)$ and π.

Education Grants: Grants are awarded by the government through the formula $g\,(q, \theta)$ where the dependence on (q, θ) makes grants a function of both parental wealth and students' ability. Hence, we allow grants to be both need-based and merit-based.

Fiscal Policies: The government levies proportional taxes at rate τ_c on consumption, τ_w on labor earnings, and τ_k on capital income and pays a lump-sum transfer ψ which makes the

out-of-pocket fees for students, which we model explicitly.

[17] The fixed repayment schedule is another reason why an individual with type $q = 3$ who can borrow privately (and hence with a flexible repayment schedule) will prefer to do that before tapping into federal loans.

[18] Enforceability of government students loans is very high: student loans cannot be expunged by bankruptcy, and wage garnishments and tax offsets can be used as repayments. In light of these features we assume they are fully enforceable. The assumption that n is such that the individual must finish its repayment before the inter vivos transfer is made for tractability. This restriction is not binding when the model is calibrated to US data on typical repayment periods of this type of debt contracts, which is typically 20 years.

system progressive.[19] Tax revenues are used to finance non-valued government consumption G, transfers, education policies, a social security system that pays pension benefits p^e to all workers of type e, and interest to service debt rD, where D is the stock of outstanding government bonds.[20]

3.2.7 The Individual Problem in Recursive Form

It is convenient to describe the individual problem backward, from retirement to schooling.

Retirement Stage: From age $j^{WK}+1$ to age J, the individual solves:

$$\begin{aligned} \Omega_j\left(e, a_j\right) &= \max_{c_j, a_{j+1}} \; u\left(c_j, 1\right)+\beta \zeta_{j+1} \Omega_{j+1}\left(e, a_{j+1}\right) \\ & \quad s.t. \\ \left(1+\tau_c\right) c_j+a_{j+1} &= p^e+\psi+\left(\zeta_{j+1}\right)^{-1}\left[1+r\left(1-\tau_k\right)\right] a_j \\ a_{j+1} &\geq 0, \quad c_j \geq 0 \end{aligned} \tag{3.3}$$

where p^e is a social security benefit conditional on the education level (the reason why e remains a state variable of this problem besides wealth a_j). The term ζ_{j+1} in the budget constraint reflects the perfect annuity markets assumption. The retired agent does not work $(l_j = 1)$ and cannot borrow.

Work Stage After the Inter Vivos Transfer: From age $j^{TR}+1$ until retirement, the working individual solves:

$$\begin{aligned} W_j\left(e, a_j, \theta, z_j\right) &= \max_{c_j, l_j, a_{j+1}} \; u\left(c_j, l_j\right)+\beta \mathbb{E}_z W_{j+1}\left(e, a_{j+1}, \theta, z_{j+1}\right) \\ & \quad s.t. \\ \left(1+\tau_c\right) c_j+a_{j+1} &= \left(1-\tau_w\right) w^e \varepsilon_j^e\left(\theta, z_j\right)\left(1-l_j\right)+\psi+\left[1+r\left(1-\tau_k\right)\right] a_j \\ a_{j+1} &\geq -\underline{a}, \quad c_j \geq 0, \quad l_j \in[0,1] \\ z_{j+1} &\sim \Gamma_z^e\left(z_{j+1} \mid z_j\right) \end{aligned} \tag{3.4}$$

The individual state variables in this problem are education level e, asset holdings a_j, ability θ, and the productivity shock z_j. The variable w^e is the price of an effective hour ε_j^e of labor of type e. Workers can borrow up to an exogenously set debt limit $\underline{a}$ from private markets. In the last period of work before retirement $\left(j=j^{WK}\right)$ the continuation value is replaced by $\zeta_{j^{WK}+1} \Omega_{j^{WK}+1}\left(e, a_{j^{WK}+1}\right)$.

[19] The tax τ_k is levied only on positive capital income. We use τ_k throughout with the convention that if $a < 0$ (and therefore $r = r^p$) then $\tau_k = 0$.

[20] Since government debt needs no intermediation through the financial sector to be exchanged, the government can borrow at the cheaper rate r relative to the households' borrowing rate $r^p = r + \iota$. Alternative assumptions have no significant bearing on the results.

Work Stage in the Period of the Inter Vivos Transfer: At age j^{TR}, the individual problem reads:

$$
\begin{aligned}
W_j\left(e, a_j, \theta, z_j, \hat{\theta}\right) &= \max_{c_j, l_j, \hat{a}_0, a_{j+1}} u\left(c_j, l_j\right) + \beta\left[\mathbb{E}_z W_{j+1}\left(e, a_{j+1}, \theta, z_{j+1}\right)\right. \\
&\quad \left. + \omega \mathbb{E}_{\hat{z}_0} V^*\left(\hat{a}_0, \hat{\theta}, \hat{z}_0, q\right)\right] \\
&\quad s.t. \\
\left(1+\tau_c\right) c_j + a_{j+1} + \hat{a}_0 &= \left(1-\tau_w\right) w^e \varepsilon_j^e\left(\theta, z_j\right)\left(1-l_j\right) + \psi + \left[1 + r\left(1-\tau_k\right)\right] a_j \qquad (3.5) \\
a_{j+1} &\geq -\underline{a}, \quad \hat{a}_0 \geq 0, \quad c_j \geq 0, \quad l_j \in [0,1] \\
z_{j+1} &\sim \Gamma_z^e\left(z_{j+1} \mid z_j\right), \quad \hat{z}_0 \sim \Gamma_0^{LH} \\
q &= \begin{cases} 1 & \text{if } a_j \leq a^* \\ 2 & \text{if } a^* < a_j \leq a^{**} \\ 3 & \text{if } a_j > a^{**} \end{cases}
\end{aligned}
$$

The altruistic parent puts weight $\omega \in [0,1]$ on the discounted utility $V^*\left(\hat{a}_0, \hat{\theta}, \hat{z}_0, q\right)$ of her child. At this date, parents know their child's ability $\hat{\theta}$, but need to form expectations about the child's productivity next period in order to choose the transfer $\hat{a}_0$. The transfer determines the initial asset position of the child in the period when she becomes an independent decision maker. Parental wealth determines the child's eligibility status for financial aid (q). The constraint $\hat{a}_0 \geq 0$ means that parents cannot force kids to transfer resources to them.[21]

Work Stage Between Full Repayment of Government-Sponsored Loan & Inter Vivos Transfer: Over this period, the household's problem is exactly as in (3.4). The only difference being that, in the period just before the transfer (age $j = j^{TR} - 1$), the continuation value in (3.4) is replaced by $\mathbb{E}_{z,\hat{\theta}} W_{j^{TR}}\left(e, a_{j^{TR}}, \theta, z_{j^{TR}}, \hat{\theta}\right)$, defined above in equation (3.5) where the expectation over $\hat{\theta}$ is computed based on the conditional distribution $\Gamma_\theta\left(\hat{\theta}, \theta\right)$.

Work Stage Before Full Repayment of Government-Sponsored Loan: In this stage, the individual solves:

$$
\begin{aligned}
W_j\left(e, a_j, \theta, z_j, \pi\right) &= \max_{c_j, l_j, a_{j+1}} u\left(c_j, l_j\right) + \beta \mathbb{E}_z W_{j+1}\left(e, a_{j+1}, \theta, z_{j+1}, \pi\right) \qquad (3.6) \\
&\quad s.t. \\
\left(1+\tau_c\right) c_j + a_{j+1} &= \left(1-\tau_w\right) w^e \varepsilon_j^e\left(\theta, z_j\right)\left(1-l_j\right) + \psi + \left[1 + r\left(1-\tau_k\right)\right] a_j - \pi \\
a_{j+1} &\geq -\underline{a}, \quad c_j \geq 0, \quad l_j \in [0,1] \\
z_{j+1} &\sim \Gamma_z^e\left(z_{j+1} \mid z_j\right)
\end{aligned}
$$

[21] This constraint is here for clarity, but it is not necessary to restrict the solution to the optimization problem since, at age $j = 1$ (high-school), students cannot borrow.

where the main difference with problem (3.4) is the presence of the additional state variable π, the size of the fixed repayment of the government-sponsored education loan.

College Education: Let (a_j, b_j) be private net worth and government education debt, respectively. Furthermore, let ϕ be the idiosyncratic tuition cost faced by a student. College students between ages $j^{HS}+1$ and j^{CL} solve:

$$\begin{aligned} V_j\left(CL, a_j, b_j, \theta, q, \phi\right) &= \max_{c_j, l_j, a_{j+1}, b_{j+1}} u\left(c_j, l_j\right) - \kappa^{CL}\left(\theta, q\right) + \beta V_{j+1}\left(CL, a_{j+1}, b_{j+1}, \theta, q, \phi\right) \\ & \quad s.t. \\ c &\geq 0, \quad l_j \in \left[0, 1-\bar{t}\right] \end{aligned} \tag{3.7}$$

where $\kappa^{CL}(\theta, q)$ is the psychic cost of attending college. Their budget constraint depends on their eligibility status q. A student who qualifies for a subsidized government loan ($q=1$) faces the budget constraint:

$$\begin{aligned} &\left(1+\tau_c\right) c_j + a_{j+1} + b_{j+1} - \left(1-\tau_w\right) w^{HS} \varepsilon_j^{HS}\left(\theta, 0\right)\left(1-\bar{t}-l_j\right) - \psi + \phi - g\left(q, \theta\right) = \\ &= \begin{cases} \left[1 + r\left(1-\tau_k\right)\right] a_j & \text{if } a_j \geq 0, \quad b_j = 0 \\ b_j & \text{if } a_j = 0, \quad 0 > b_j \geq -\underline{b}^s \\ -\underline{b}^s + \left(1+r^u\right)\left(b_j + \underline{b}^s\right) & \text{if } a_j = 0, \quad b_j < -\underline{b}^s \end{cases} \\ &a_{j+1} \geq 0 \quad b_{j+1} \geq -\left(\underline{b}^s + \underline{b}^u\right) \end{aligned} \tag{3.8}$$

A student who qualifies only for unsubsidized government loans ($q=2$) faces the budget constraint:

$$\begin{aligned} &\left(1+\tau_c\right) c_j + a_{j+1} + b_{j+1} - \left(1-\tau_w\right) w^{HS} \varepsilon_j^{HS}\left(\theta, 0\right)\left(1-\bar{t}-l_j\right) - \psi + \phi - g\left(q, \theta\right) = \\ &= \begin{cases} \left[1 + r\left(1-\tau_k\right)\right] a_j & \text{if } a_j \geq 0, \quad b_j = 0 \\ \left(1+r^u\right) b_j & \text{if } a_j = 0, \quad b_j < 0 \end{cases} \\ &a_{j+1} \geq 0 \quad b_{j+1} \geq -\left(\underline{b}^s + \underline{b}^u\right) \end{aligned} \tag{3.9}$$

If the student's parental wealth is high enough that she can also borrow privately ($q=3$), she faces the following budget constraint:

$$
\begin{aligned}
& (1+\tau_c)\,c_j + a_{j+1} + b_{j+1} - (1-\tau_w)\,w^{HS}\varepsilon_j^{HS}(\theta,0)\,(1-\bar{t}-l_j) - \psi + \phi - g(q,\theta) = \\
= & \begin{cases} [1+r(1-\tau_k)]\,a_j & \text{if } a_j \geq 0, \quad b_j = 0 \\ (1+r^p)\,a_j & \text{if } 0 > a_j > -\underline{a}^p, \quad b_j = 0 \\ (1+r^p)\,\underline{a}^p + (1+r^u)\,b_j & \text{if } a_j = -\underline{a}^p, \quad b_j < 0 \end{cases} \qquad (3.10) \\
& a_{j+1} \geq -\underline{a}^p \quad b_{j+1} \geq -\left(\underline{b}^s + \underline{b}^u\right).
\end{aligned}
$$

Finally, the continuation value in the last period of college is replaced by $\mathbb{E}_z W_{j^{CL}+1}\left(CL, a_{j^{CL}+1}, \theta, z_{j^{CL}+1}, \pi\right)$ where $z_{j^{CL}+1} \sim \Gamma_0^{CL}$, and π is determined by equation (3.2) based on $b_{j^{CL}+1}$ and q.

College Decision: At age $j = j^{HS} + 1$, the student draws her tuition cost ϕ from the distribution $\Phi(\overline{\phi}, \sigma_\phi)$, and choose whether to attend college. Given this draw they solve

$$V^{**}(a_j, \theta, z_j, q, \phi) = \max \ \{V_j(CL, a_j, b_j, \theta, q, \phi), \mathbb{E}_z W_j(HS, a_j, \theta, z_j)\} \qquad (3.11)$$

The dummy variable $d^{CL} \in \{0,1\}$ reflects the college education decision.[22]

High-School Education: A high-school student solves:

$$
\begin{aligned}
V_j(HS, a_j, \theta, q) &= \max_{c_j, a_{j+1}} u(c_j, 1-\bar{t}) - \kappa^{HS}(\theta, q) + \beta V_{j+1}(HS, a_{j+1}, \theta, q) \qquad (3.12) \\
& \quad s.t. \\
(1+\tau_c)\,c_j + a_{j+1} &= [1+r(1-\tau_k)]\,a_j + \psi \\
a_{j+1} &\geq 0, \quad c_j \geq 0
\end{aligned}
$$

High-school students are not permitted to borrow. In the last period of high-school $\left(j = j^{HS}\right)$, the continuation value is $\mathbb{E}_{z,\phi} V^{**}(a_{j+1}, \theta, z_{j+1}, q, \phi)$ where V^{**} is defined above.

High-School Decision: At age $j = 0$, the student chooses whether to enter the labor market as a high-school dropout or stay in school. If she chooses to work she draws $z_0 \sim \Gamma_{z_0}^{LH}$, the initial productivity level. So, each individual solves:

$$V^*(\hat{a}_0, \theta, z_0, q) = \max \ \{V_0(HS, \hat{a}_0, \theta, q), \mathbb{E}_z W_0(LH, \hat{a}_0, \theta, z_0)\} \qquad (3.13)$$

where $\hat{a}_0$ is the transfer received from the parent, θ is innate ability, and q is eligibility status

[22] The presence of discrete education choices introduces non-convexities in the budget sets. This implies that standard results on uniqueness and continuity of optimal policy functions cannot be applied to this problem. For a discussion of related issues and the numerical solution of this problem see Gallipoli and Nesheim (2007).

for college support (grants and loans), which depends on parental wealth. The high-school enrollment decision is denoted $d^{HS} \in \{0, 1\}$.

3.2.8 Equilibrium

It is useful to introduce some additional notation to simplify the definition of an equilibrium. Let $\mathbf{s}_j \in S_j$ denote the age-specific state vector implicit in the recursive representation of the agents' problems above. We also define $\mathbf{s}_j^e$ to be the state vector minus the education level (the school cycle they are in for students), i.e., $\mathbf{s}_j^e \equiv \{\mathbf{s}_j \setminus e\} \in S_j^e$.

A *stationary recursive competitive equilibrium* for this economy is a collection of: (i) individual decision rules for consumption, leisure, wealth holdings, and college students' debt $\left\{c_j(\mathbf{s}_j), l_j(\mathbf{s}_j), a_{j+1}(\mathbf{s}_j), b_{j+1}\left(\mathbf{s}_j^{CL}\right)\right\}$, inter vivos transfers $\left\{\hat{a}_0\left(\mathbf{s}_{j^{TR}}\right)\right\}$, and education choices $\left\{d^{HS}(\mathbf{s}_0), d^{CL}\left(\mathbf{s}_{j^{HS}}\right)\right\}$; (ii) value functions $\{V_j(\mathbf{s}_j), W_j(\mathbf{s}_j), \Omega_j(\mathbf{s}_j)\}$; (iii) aggregate capital and labor inputs
$\{K, H_{LH}, H_{HS}, H_{CL}\}$; (iv) prices $\left\{r, w^{LH}, w^{HS}, w^{CL}\right\}$; (v) labor income tax $\{\tau_w\}$; (vi) age and education specific measures $\left\{\mu_j^e\right\}$ such that:

1. Decision rules $\left\{c_j(\mathbf{s}_j), l_j(\mathbf{s}_j), a_{j+1}(\mathbf{s}_j), b_{j+1}\left(\mathbf{s}_j^{CL}\right), \hat{a}_0\left(\mathbf{s}_{j^{TR}}\right), d^{HS}(\mathbf{s}_0), d^{CL}\left(\mathbf{s}_{j^{HS}+1}\right)\right\}$ solve their respective individual problems (3.3), (3.4), (3.5), (3.6), (3.7), (3.11), (3.12), and (3.13). And $\{V_j(\mathbf{s}_j), W_j(\mathbf{s}_j), \Omega_j(\mathbf{s}_j)\}$ are the associated value functions.

2. The representative firm optimally chooses factors of productions, and prices equate with marginal products

$$\begin{aligned} r + \delta &= F_K\left(K, \mathcal{H}\left(H^{LH}, H^{HS}, H^{CL}\right)\right) \\ w^e &= F_{H^e}\left(K, \mathcal{H}\left(H^{LH}, H^{HS}, H^{CL}\right)\right), \text{ for } e \in \{LH, HS, CL\}. \end{aligned} \tag{3.14}$$

3. The labor market for each educational level $e \in \{LH, HS, CL\}$ clears

$$H^e = \sum_{j=j^e+1}^{j^{WK}} \int_{S_j^e} \varepsilon^e\left[1 - l\left(e, \mathbf{s}_j^e\right)\right] d\mu_j^e + I_{\{e=HS\}} \cdot \sum_{j=j^{HS}+1}^{j^{CL}} \int_{S_j^{CL}} \varepsilon^{CL}(\theta, 0)\left[1 - \bar{t} - l\left(e, \mathbf{s}_j^{CL}\right)\right] d\mu_j^{CL}$$

   where the second term in the sum is the effective labor supply of college students.

4. The intermediation market clears at the price $r^p = r + \iota$.

5. The asset market clears

$$K + D = \sum_{\substack{e \in \{LH, HS\}, j \geq 0 \\ e = CL, j \geq j^{HS}+1}} \int_{S_j^e} a_j\left(e, \mathbf{s}_j^e\right) d\mu_j^e$$

and the aggregate net worth of all households (right-hand side) equals the capital stock plus government debt (left-hand side).

6. The goods market clears

$$\sum_{e,j} \int_{S_j^e} c_j \left(e, \mathbf{s}_j^e\right) d\mu_j^e + \delta K + G + \Phi + \Upsilon = F\left(K, \mathcal{H}\right)$$

where Φ is the aggregate amount of private expenditures in educational services by college students

$$\Phi = \sum_{j=j^{HS}+1}^{j^{CL}} \int_{S_j^{CL}} \phi d\mu_j^{CL}$$

and Υ is the revenue of the intermediating sector

$$\begin{aligned} \Upsilon &= \iota \cdot \sum_{\substack{e \in \{LH,HS\}, j \geq 0 \\ e=CL, j \geq j^{HS}+1}} \int_{S_j^e} I_{\{a_j<0\}} a_j \left(e, \mathbf{s}_j^e\right) d\mu_j^e \\ &+ \iota \cdot \sum_{j \geq j^{HS}+1}^{j^{CL}} \int_{S_j^{CL}} b_j d\mu_j^{CL} + \iota \cdot \sum_{j \geq j^{CL}+1}^{j^{CL}+n} \int_{S_j^{CL}} \left[b_{jCL} - \pi \cdot \left(j - j^{CL}\right)\right] d\mu_j^{CL} \end{aligned}$$

where the three terms are, respectively, the intermediation services for private loans, student loans during college, and student loans during the repayment phase of working life.

7. The government budget constraint holds

$$G + \left(\sum_e p^e \sum_{j=j^{WK}+1}^{J} \int_{S_j^e} d\mu_j^e\right) + \psi + rD + E = \tau_c \sum_{e,j} \int_{S_j^e} c_j \left(e, \mathbf{s}_j^e\right) d\mu_j^e + \tau_w \sum_e w^e H_e + \tau_k r K$$

where E are net government expenditures in college education:

$$\begin{aligned} E &= \sum_{j \geq j^{HS}+1}^{j^{CL}} \int_{S_j^{CL}} \left[g\left(q, \theta\right) - \Delta b_j\right] d\mu_j^{CL} + \iota \sum_{j \geq j^{HS}+1}^{j^{CL}} \int_{S_j^{CL}} b_j d\mu_j^{CL} \\ &- r^u \sum_{j \geq j^{HS}+1}^{j^{CL}} \int_{S_j^{CL}} \left[I_{\{q=1, b_j < -\underline{b}^s\}} \cdot \left(b_j + \underline{b}^s\right) + I_{\{q \geq 2\}} \cdot b_j\right] d\mu_j^{CL} \\ &- \sum_{j \geq j^{CL}+1}^{j^{CL}+n} \int_{S_j^{CL}} \pi d\mu_j^{CL}. \end{aligned}$$

Government outlays (first row) are determined by grants and the total amount of loans extended to college students which is equal to the sum of the Δb_j increments in each year, plus the intermediation cost (ι) incurred on all outstanding loans. Revenues (second and third rows) are determined by interest on unsubsidized loans during college and debt repayments after graduation.

8. Individual and aggregate behaviors are consistent: the vector of measures $\mu = \{\mu_0^e, ..., \mu_J^e\}_{e\in\{LH,HS,CL\}}$ is the fixed point of $\mu(S) = Q(S,\mu)$ where (i) $Q(S,\cdot)$ is a transition function generated by the individual decision rules, the exogenous laws of motion $\left\{\Gamma_\theta, \Gamma_z^e, \Gamma_{z_0}^e\right\}$, the distribution Φ, the institutional rules determining π, q, and p_e, and the survival rates $\{\zeta_j\}$; (ii) and S is the generic subset of the Borel-sigma algebra $\mathcal{B}_S$ defined over the state space $\mathbf{S}$, the Cartesian product of all S_j^e.

3.3 Parameterization of the Model

We describe below how we parameterize the model economy. Some of the parameters are calibrated "internally" from the equilibrium of the model, while others are estimated "externally" directly from data. All parameter values are reported in Table 3.5.

Demographics: A model period is one year. Individuals become adults at the real age of 16 (i.e. $j = 0$ in the model), and they can live up to age 99, after which death is certain. Retirement occurs at age 65. Inter vivos transfers are made at age 48. The conditional survival rates $\{\zeta_j\}$ are taken from the Actuarial Life Tables for the United States.

Preferences: We specify period utility over consumption and leisure as a CRRA function

$$u(c_j, l_j) = \frac{\left(c_j^\nu l_j^{1-\nu}\right)^{1-\gamma}}{1-\gamma}. \tag{3.15}$$

The parameters ν and γ jointly pin down (i) the level of labor supply over the life cycle, (ii) the inter-temporal elasticity of substitution of consumption (IES) $1/[1-\nu(1-\gamma)]$, and (iii) the Frisch labor supply elasticity $[1-\nu(1-\gamma)]/\gamma \cdot (l_j/(1-l_j))$ which, with this preference specification, depends on hours worked.

The weight of leisure in preferences, ν, is set to 0.385 to match average hours worked, estimated to be 35% of the time endowment. Hence, a value of $\gamma = 2$ is required to match an inter-temporal elasticity of substitution of 0.75 as estimated by Blundell, Browning, and Meghir (1994) and Attanasio and Weber (1993). The Frisch elasticity evaluated at the (non-stochastic) average hours worked $\bar{h} = \nu$ implied by this choice of ν and γ is 1.25.[23]

[23] This value is on the high end of micro-estimates for men, but on the low end for women. See Keane and Rogerson (2011) for a recent survey.

	Children					
Mothers	1	2	3	4	5	Total
1	45.5	23.8	19.7	6.5	4.7	100
2	25.8	24.2	24.2	15.7	11.0	100
3	16.0	22.3	27.1	19.0	15.7	100
4	11.4	17.1	25.7	20.9	24.9	100
5	7.2	7.6	19.5	24.2	41.5	100
Total	21.3	19.2	23.2	17.1	19.2	100

Notes: Ability transition, by quintile. Each cell reports the conditional probability in %. Quintile 1 is the lowest, quintile 5 is the highest. (NLSY79)

Table 3.1: Ability transition.

Patience and Altruism: The value for β is chosen to reproduce a ratio of median net worth to average income, which is estimated to be 1.64 using the 2001 Survey of Consumer Finances.[24]

To calibrate the degree of altruism towards offsprings ω, we target the size of intervivos transfers in the data. Because we model early inter vivos transfers as one-off gifts from parents to their child, we restrict attention to the cumulative transfer between age 16 and 22. The NLSY97 provides information on family transfers received by young individuals. In particular, it asks respondents about any gifts in the form of cash (not including loans) from parents. Appendix A describes the sample we construct and the methodology we use to measure early inter vivos transfers, and it reports basic facts about parental gifts to young individuals, as recorded in the NLSY97. In our calculations we also include imputed rents for students living in their parents' house.[25] The average inter vivos transfer over the seven-year period considered is \$30,566. The data show large heterogeneity of IVTs in the population: for example, college students receive twice as much as high-school dropouts.

Intergenerational Transmission of Ability: In our model "ability" θ represents a set of permanent characteristics which affect lifetime earnings as well as education attainment. For the purpose of measuring the distribution of ability over the population we use NLSY data. The NLSY79 provides IQ test scores for both mothers and children, which we link in order to estimate an ability transition matrix Γ_θ.

Using the "Children of the NLSY79" survey, we build pairs of mother and child test-score

[24]It is well known that the SCF oversamples the rich households relative to the CPS, PSID, and NLSY, the other surveys we use to parameterize the model. To make the SCF more comparable to the other surveys, we exclude the top 5% of households ranked by net worth. The sample selection is the same as in Kaplan and Violante (2012).

[25]As also emphasized by Johnson (2011), the co-residence component makes up a large fraction of the total IVT.

measurements. For mothers we use AFQT89 measurements, whereas for children we choose the PIAT Math test-scores.[26] Mothers and children are ranked based on their own test scores, and then split into quintiles. We then compute a "quintile-transition" matrix, which assigns a probability to the event that a child ends up in a given ability group, given the observed ability rank of the parent.[27] Note that in the model we allow for θ to be continuous by assuming a uniform distribution of abilities within each of the five bins.

The estimated ability transition matrix across quintiles is reported in Table 3.1. The matrix implies a great deal of upward and downward mobility in the middle of the distribution, but less so at the top and the bottom, where the diagonal element is larger.

Disutility of Schooling: The psychic costs of attending high-school and college consist of two additive components, respectively changing by ability and family background (wealth) : $\kappa^e(\theta, q) = \kappa^e_\theta(\cdot) + \kappa^e_q(\cdot)$. The first component (5 values, corresponding to the 5 ability quintiles) is set to reproduce high-school and college enrolment rates by ability. The AFQT89 scores (over the entire NLSY79 sample) can be matched with the education level of the subjects to measure education shares by ability level. The NLSY79 education shares by ability bin are reported in Table 3.2. [28]

Psychic costs based on family background, $\kappa^e_q(\cdot)$, are normalized to zero for $q = 1$ and set equal for $q = 2, 3$. These remaining two parameters (one for high-school students and one for college students) are set to match two ratios: (i) the average transfer received by a college graduate divided by the average transfer received by a high school graduate, which is 1.93 in the data, and (ii) the average transfer received by a high school graduate divided by the average transfer received by a high school drop-out, which is 1.05 in the data.[29] We estimate psychic costs of attending college of a magnitude comparable to that estimated by Heckman,

[26] No AFQT measure is available for children, and the PIAT Math score is generally considered the most accurate measure of future ability among the available test-scores. Often the test was administered at different ages to the same child. We use the latest available measurement, as we wish to approximate the distribution of ability at age 16. Details about the procedure used to compute the ability-transition matrix and the test scores used can be found in Appendix B.

[27] We also estimate transition matrices based on 10 ability bins, and results are similar. Moreover, the transition matrix is virtually identical when we use a smaller sample including only mother-child pairs in which the child was at least 13 years of age at the time of the test.

[28] Because the age and education structure of the NLSY cohort is not representative of the overall working population in any given year, we rescaled the size of each education/ability cell so that their aggregation yields the correct fraction of 16-65 years-old individuals in the three education groups for the year 2000, as estimated in the CPS (26.1 pct of college graduates, 59.8 pct of HS graduates and 14.1 of HS dropouts). The distribution by ability within each education group remains unchanged.

[29] It is important to match these statistics for the following reason. If we do not allow children of rich families to have low psychic cost from acquiring education, the wealth effect on their labor supply is so large that they choose to remain uneducated and their parents make a very large transfer. As a result, the distribution of initial wealth by education level becomes counterfactual (many low educated young individuals are very rich). The model, indeed, calls for a large utility gain of schooling for children of wealthy households. A more direct way to compute these psychic cost terms would be targeting enrollment by parental wealth. However, parental net worth (as reported by parents themselves in the 1997 wave of the NLSY97) is of poor quality, as in our sample, more than 40% of individuals have missing values or non-response codes.

	Quintile (AFQT89)					
Education	1	2	3	4	5	Total
Less than H.S.	32.0	9.2	3.9	1.0	0.3	9.5
H.S. Graduate	66.2	83.0	78.5	61.9	31.0	65.0
College Graduate	1.8	7.7	17.6	37.1	69.0	26.0

Table 3.2: Education shares (%) by AFQT89 quintile (NLSY79).

Lochner, and Todd (2006). Moreover, we find that they are decreasing in parental wealth. We return to this result below. Appendix C contains a detailed discussion of the estimated psychic costs.

Labor Supply of Students: A report by the National Center for Education Statistics (1998) documents that full-time college students work on average 15 hours in a typical week during college.[30] To reproduce this statistic in the model, we set $\bar{t} = 0.312$ in (3.7), i.e., attending college reduces the time endowment of students by 31%.

Individual Labor Productivity: The model implies the following reduced-form specification for the hourly wage W^e_{ijt} of an individual i of age j, and education level e at date t

$$\log W^e_{ijt} = \log w^e_t + \lambda^e \log \theta_i + \xi^e(j_{it}) + u^e_{ijt}, \tag{3.16}$$

where w^e_t is the marginal product of one efficiency unit of human capital of education-type e, λ^e is the gradient on permanent individual heterogeneity, $\xi^e(j_{it})$ is an education specific age-profile (approximated by a 4-th order polynomial), and u^e_{ijt} is a stochastic residual component.

We use the NLSY79 to estimate reduced-form education-specific wage equations like (3.16) because the NLSY test scores data (a proxy for θ) can be linked to wage data to quantify the effect of measured ability on lifetime earnings. To overcome the problem that the NLSY provides observations only for workers between age 14 and 45, we use wage data from the PSID 1968-2001 to estimate age polynomials for different education groups. After the age profiles have been used to filter out age effects from the log wage observations in the NLSY79 –assuming that the unobserved error term is uncorrelated with θ_i– we can identify the loading factors by running simple regressions. For each education group $e \in \{LH, HS, CL\}$ an OLS regression of log individual wages on time dummies and on log AFQT89 scores (as a proxy for θ_i) was fit in order to recover λ^e (see equation 3.16). [31] These reduced-form results, reported in Table 3.3, show a steep gradient by education; a 10% increase in ability implies, on average,

[30] This number is an average between 37.6% of full-time students not working while enrolled, 13.9% working 1-14 hours, 34.8% working 15-33 hours and 13.7% working 34 hours or more. See NCES (1998, Table 6).

[31] We estimate the above equation for the cross-sectional representative sample as well as the full sample of people in the NLSY79, which includes oversamples for minorities and disadvantaged groups. The two samples give essentially the same results and we report the ones for the larger sample.

Education group	Gradient (S.E.)	# of indiv.	# of obs.
Less than HS	0.36 (0.06)	1,341	8,982
HS Graduate	0.54 (0.03)	5,403	42,270
College Ggraduate	0.89 (0.09)	1,206	8,719
Pooled	0.71 (0.02)	7,954	60,009

Table 3.3: Estimated ability gradient λ^e (NLSY79)

a 8.9% increase in hourly wages for college graduates but only a 3.6% increase for dropouts.[32]

The residuals from this regression are a consistent estimate of u^e_{ijt}. We model the unobservable shock u^e_{ijt} as the sum of two independent components

$$u^e_{ijt} = z^e_{ijt} + m^e_{ijt} \tag{3.17}$$

where z^e_{ijt} is a (persistent) shock assumed to have an AR(1) structure

$$z^e_{ijt} = \rho_e z^e_{i,j-1,t-1} + \eta^e_{ijt}, \qquad \eta^e_{ijt} \overset{iid}{\sim} N\left(0, \sigma^e_{\eta t}\right),$$

and $m^e_{it} \overset{iid}{\sim} N\left(0, \sigma^e_{mt}\right)$ is measurement error (and hence noise from the point of view of the model). Finally, we let the initial draw $z^e_{i0t} \sim N\left(0, \sigma^e_{z_0}\right)$. To estimate the parameters of the error-component model $\left\{\rho_e, \sigma^e_{\eta t}, \sigma^e_{mt}, \sigma^e_{z_0}\right\}$, we use a Minimum Distance Estimator (see Rothenberg, 1971; Chamberlain, 1984; Heathcote, Storesletten, and Violante, 2010). Table 3.4 reports parameter estimates. Details about our sample selection, estimation of quartic age polynomials, and estimation of the error component model are reported in Appendix D.[33]

Overall, we confirm the finding of Meghir and Pistaferri (2006) that the persistent component of wage risk does not vary systematically across education groups.

Technology: The aggregate production function is Cobb-Douglas and constant returns to scale, i.e.,

$$Y = F(K, \mathcal{H}) = K^{\alpha}\mathcal{H}^{1-\alpha}. \tag{3.18}$$

We set $\alpha = 1/3$ and let the aggregate human capital stock $\mathcal{H}$ be given by the CES aggregator

$$\mathcal{H} = \left[s^{LH}\left(H^{LH}_t\right)^{\rho} + s^{HS}\left(H^{HS}_t\right)^{\rho} + s^{CL}\left(H^{CL}_t\right)^{\rho}\right]^{\frac{1}{\rho}} \tag{3.19}$$

[32] The unconditional log AFQT89 distribution is normalized to have mean zero in the regression.

[33] By using an observable variable as a proxy for permanent heterogeneity, we avoid selection bias in the estimation of the process for u^e_{ijt}. Moreover, if one estimates wage equations from individual panel data sets, as we do, selection bias attributable to persistent shocks becomes less severe. The issue of selection bias ensuing from persistent shocks is related to the so-called "incidental parameters problem" discussed in Heckman (1981). The severity of the incidental parameters problem becomes smaller as the number of panel observation for each given individual in a sample increases.

Less than HS		HS Graduates		College graduates	
ρ	0.936	ρ	0.951	ρ	0.945
σ_{z0}^2	0.105	σ_{z0}^2	0.101	σ_{z0}^2	0.128
σ_{η}^2	0.020	σ_{η}^2	0.017	σ_{η}^2	0.020
σ_{m}^2	0.070	σ_{m}^2	0.055	σ_{m}^2	0.052

Notes: Estimated parameters of the process for individual efficiency units u_{ijt} (NLSY79)

Table 3.4: Estimated idiosyncratic efficiency process

where H^e is the stock of human capital associated with education level e, and $s^{LH} + s^{HS} + s^{CL} = 1$. The elasticity of substitution between each pair of labor types is $1/(1-\rho) \in (0, \infty)$.

From the iso-elastic CES specification for the human capital aggregate in equation (3.19), and the equilibrium condition in the labor market (3.14), we can derive expressions for the wage bills ϖ_t^e. For education groups HS and CL, for example, we can write

$$\log\left(\frac{\varpi_t^{CL}}{\varpi_t^{HS}}\right) = \log\left(\frac{s_t^{CL}}{s_t^{HS}}\right) - (1-\rho)\log\left(\frac{H_t^{CL}}{H_t^{HS}}\right) \tag{3.20}$$

To estimate the above equation, we use wage and hours data from the Current Population Survey (CPS) for 1968-2001 (see Heckman, Lochner, and Taber, 1998a). After computing total aggregate wage bills year by year, we divide them by the (normalized) marginal products of the three types of human capital $\left\{w_t^{LH}, w_t^{HS}, w_t^{CL}\right\}$ estimated from from PSID data (see Appendix D) to obtain point estimates of aggregate efficiency-weighted labor supply (human capital aggregates) by education level and year. Because of the well documented relative demand shifts over the period considered, in the equation above we let share parameters vary over time, i.e., $s_t^e = \exp(s_0^e + g^e t)$, where t denotes calendar year and g^e captures the growth rate in each human capital share of type e.

To account for the endogeneity of schooling choices in the estimation of equation (3.20), we instrument the human capital aggregates in two ways: in the first, we use lagged variables; in the second, we use the total number of individuals with a given level of education regardless of their labor force status. The latter instruments do not depend on the serial correlation properties of the technology shocks. Results do not change much with the choice of instruments chosen nor with the specification. The estimated value for ρ ranges between 0.36 and 0.68, which corresponds to an elasticity of substitution between 1.6 and 3.1.[34]. See

[34]Many existing estimates in the literature are based on a two-type skilled/unskilled classification for labor. Katz and Murphy estimate the elasticity of substitution to be 1.41; Heckman, Lochner, and Taber (1998a) report a favorite estimate of 1.44; Card and Lemieux (2001) obtain an elasticity of substitution between college and high school workers of about 2.5. Finally, using a nested specification with three human capital types Goldin and Katz (2007) suggest a preferred elasticity between college and non-college workers of 1.64

Appendix E for additional details. As a baseline value for the elasticity, we choose its higher bound of 3.1 representing also a lower bound for the general equilibrium effects. The values of the shares used in the model's simulation are $s^{LH} = 0.16, s^{HS} = 0.39, s^{CL} = 0.45$, which are the estimated values for the year 2000.

This specification of aggregate technology together with the equilibrium selection mechanism of the model yields college and high school wage premia that are consistent with the data. Following the method of Goldin and Katz (2007) using model simulated data, the log college/high-school wage differential is 0.58, and the log high-school graduate/dropout wage differential is 0.37. These values are similar to the estimates presented in Goldin and Katz (2007, Table A8.1) for the year 2000 which place the college premium between 0.58 and 0.61, and the high-school premium between 0.26 and 0.37.

Finally, we set the annual depreciation rate of capital at 6.5% (see Heathcote, Storesletten, and Violante, 2010).

In what follows, we specify values for a number of parameters characterizing the government-sponsored loans and grants in our economy. In Appendix F we provide a detailed description of the federal system of financial aid to college students (as in the year 2000) that we aim to reproduce in this calibration.

Private Borrowing: We set the liquidity constraint $\underline{a}$ for working-age households to reproduce the fraction of the US population aged 16-65 with zero or negative net worth. From the SCF 2001, we estimate this fraction to be 11.2%.[35] We set the wedge ι on private borrowing (for both students and workers) and the limit to private students' loans a^p to match (i) the fraction of college students enrolled at 4-year institutions who borrow privately, which is 4.9%, and (ii) a ratio between the total volume of private loans and the total volume of federal loans of 0.12.

Cost of college: We define the cost of college as tuition fees plus the cost of books and other academic material net of institutional and private grants, and we compute an average across all full-time, full-year dependent students enrolled in private not-for-profit and public 4-year colleges in the year 2000. We obtain an average annual cost of $6,710 which we match to the ex-post average tuition cost in the model by appropriately setting the mean $\bar{\phi}$ of a log-Normal distribution for Φ.[36] We set the dispersion of tuition draws σ_ϕ to reproduce, in equilibrium,

[35] As reported in Table 3.5, this limit is close to $85,000 and hence quite generous. However, recall that this model has age and education-type "natural borrowing limits" that may be tighter, especially for the middle aged and the elderly workers. For the young workers, the presence of the guaranteed lump sum transfer ψ implies a sizable borrowing ability.

[36] Tuition charges vary considerably across institutions: average tuition and fees costs in private institutions were $15,000 versus $4,300 in public institutions. Since the model does not distinguish between them, we use a weighted average. We could allow tuition fees to vary by family income, but we find that variation (before grants and loans) is small. For example, families with annual income in the bracket $20,000-$40,000 face average tuition costs of $4,000 ($15,000) at public (private) colleges; families with income between $80,000 and $100,000 face average tuition costs of $4,400 ($17,000) at public (private) institutions. We chose not to include

the percentage of college student who pay lower than average tuition fees, which is 68% in the data.

Grants: Based on data summarized in Appendix F, we assign a grant of \$2,829 per year to the students of type $q = 1$, \$668 per year to students of type $q = 2$, while students of type $q = 3$ receive on average a grant of \$143 per year. This pattern reflects the need-based formula used in the vast majority of cases to award grants. In the baseline model we do not allow grants to vary by ability.

Education Loans: The seven policy parameters $\{n, a^*, a^{**}, r^s, r^u, \underline{b}^s, \underline{b}^u\}$ fully characterize the system of government education loans. We set $n = 20$, as repayment schedules of federal loans are easily extended to 20 years. The remaining six parameters are chosen to replicate as closely as possible the following cross-sectional moments in year 2000: (i) 37.3% of students have subsidized Stafford loans, (ii) 21.2% of students have unsubsidized Stafford loans, (iii) 44.9% of students have a Stafford loan, whether subsidized or not, (iv) the average cumulated amount of federal loans at graduation is \$17,016, (v) the ratio between the total volume of subsidized and unsubsidized loans is 1.36, (vi) the maximum cumulative amount of Stafford loans is \$35,125.[37]

Fiscal Policy: We use flat tax rates for labor and capital income, and consumption. Following Domeij and Heathcote (2003), we set $\tau_w = 0.27$, $\tau_k = 0.4$ and $\tau_c = 0.05$. The lump-sum subsidy ψ is set to replicate the degree of progressivity of the tax/transfer system. In particular, we target the ratio of the variance of log post-government income to the variance of log pre-government income, equal to 0.61 in the US (Heathcote, Perri, and Violante, 2010b). Domestic government debt D is set at 20 percent of GDP, its value in 2000, since only half of federal debt is held domestically. Pensions are a lump sum for all agents in a given education group. Table 1 in Mitchell and Phillips (2006) reports replacement rates for three types of workers (low, medium, and high earners) whose average labor income closely corresponds to that of our three educational groups. Expressed as a replacement ratio of average gross earnings for the economy, the pension for LH types is 0.26, for HS types is 0.42, and for CL types is 0.56.[38]

room and board in the calculation of the cost as these expenses are also borne by working individuals.

[37] Under the Stafford loan program, the student cannot borrow more than the full cost of schooling (tuition, books, room, and board). We did not impose this tighter constraint because our measure of college costs ϕ does not include room and board.

[38] These replacement rates show the progressivity of the social security system: even though average gross earnings of college graduates are roughly three times as large as earnings of high-school dropouts, their pension benefits are only twice as large.

Parameter	Value	Moment to Match	Data	Model
J	79	Length of life cycle (age 16-95)	79	79
j^{RET}	50	Length of working life (age 16-65)	50	50
$\{\zeta_j\}$	-	Survival rates (from US Life Tables)	-	-
ν	0.385	Average hours worked	0.35	0.31
γ	2.0	Elasticity of intertemporal substitution	0.75	0.75
β	0.979	Match median wealth-average income ratio	1.64	1.64
$\bar{t}$	0.312	Average student hours-average worker hours ratio	0.33	0.33
ω	0.131	Average inter vivos transfer	\$30,566	\$31,256
α	0.35	Capital share of GDP	0.35	0.35
δ	0.065	Depreciation rate	0.065	0.065
ρ	0.677	Elasticity of substitution	3.1	3.1
$\underline{a}$	1.39	Fraction of households with net worth ≤ 0	0.11	0.11
$\underline{a}^p$	0.29	Aggr. private/gov. student loans ratio	0.12	0.12
ι	0.0405	Fraction of students with private loan	0.05	0.05
$\overline{\phi}$	-1.7	Average yearly tuition fees	\$6,710	\$6,690
σ_ϕ	1.0	Fraction paying less than average tuition	0.68	0.65
n	20	Repayment period for student loans (years)	20	20
a^*	3.10	Fraction of students with subsidized loan	0.37	0.36
a^{**}	5.35	Fraction of students with unsubsidized loan	0.21	0.20
r^u	0.043	Average govt. student debt at graduation	\$17,016	\$18,905
r^s	0.0295	Fraction of students with govt. student loan	0.45	0.45
$\underline{b}^u$	0.295	Aggr. subs./unsubs. student loans ratio	1.36	1.30
$\underline{b}^s$	0.275	Total govt. loans borrowing limit	\$35,125	\$35,497
τ_w	0.27	Labor income tax rate	0.27	0.27
τ_c	0.05	Consumption tax rate	0.05	0.05
τ_k	0.40	Capital income tax rate	0.40	0.40
D	0.13	Domestic debt-GDP ratio	0.20	0.20
ψ	0.086	Var. of log post-tax inc./Var. of log pre-tax inc.	0.61	0.61
p^{LH}	0.26	Pension replacement rate - Less than HS	0.26	0.26
p^{HS}	0.42	Pension replacement rate - HS Graduates	0.42	0.42
p^{CL}	0.56	Pension replacement rate - College Graduates	0.56	0.56

Notes: Parameter values for the benchmark economy. In the model, the price (marginal product) of an efficiency unit of labor for a high school graduate w^{HS} is normalized to one, and hence 1=\$62,275 in the model. The model's period is one year.)

Table 3.5: Calibrated benchmark parameters

3.4 Assessing the Model's Behavior

We examine the behavior of the model along five dimensions. First, we analyze the implied cross-sectional age profiles for hours worked, earnings, consumption, and wealth. None of these moments is explicitly targeted in the parameterization (only those for wages are). Second, we study the determinants of parental transfers to children. Third, we measure the degree of intergenerational income persistence in the model (also, not targeted). Fourth, we examine the role of parental wealth in determining educational achievement. Fifth, we "validate" the model by reproducing, within our structural framework, a randomized experiment where a (treatment) group of high-school graduates receives a college tuition subsidy and a (control)

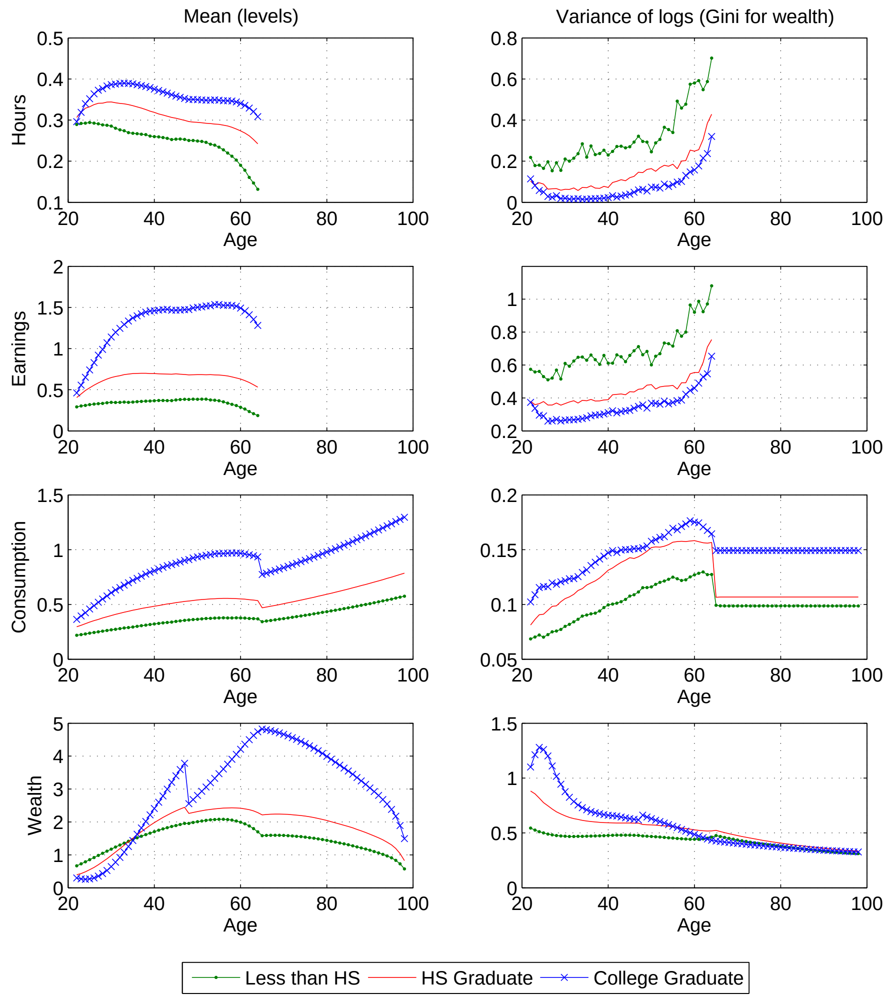

Figure 3.1: Means and dispersion of log hourly wages, earnings, consumption, and wealth.

group does not.

3.4.1 Life-Cycle Profiles

Figure 3.1 plots averages and dispersion of log hours worked, log earnings, log consumption, and wealth over the life cycle, for our three education groups.

Average hours worked increase in the level of education, which is a reflection of differences in the average return to work (the wage rate). For the same reason, hours drop much faster for the less educated groups over the life cycle. Hours dispersion is higher for the low-educated, who are the most marginally attached to the labor force, and rises over the life cycle for all three groups, following the dispersion in labor productivity. Quantitatively, the rise in the variance of hours is in line with the data (see Figure 15 in Heathcote, Perri, and Violante, 2010b).

The rise in average earnings over the life cycle is more pronounced for more educated workers and the late decline in earnings sharper for the less educated. The rise in the variance of log earnings between ages 25 and 60 (around 0.4 log points) is quantitatively consistent with its empirical counterpart (Guvenen, 2009, Figure 4).

A comparison between consumption and earnings paths (both their mean and dispersion) reveals that consumption smoothing through borrowing and saving is quite effective after the schooling phase. During working life the variance of log consumption grows by roughly 0.06 log points for all groups, compared to a rise four times larger in the variance of log earnings. The downward jump in average consumption at retirement reflects the nonseparability of consumption and leisure. The average drop in expenditures at retirement is around 14%, in line with the empirical evidence. For example, Aguiar and Hurst (2005) estimate a drop of 17%.[39]

Wealth accumulation features the typical hump-shaped pattern. In the model, the drop in household wealth at age 48 arises as a consequence of the inter vivos transfer to the children. The drop is much larger for the highly educated families, whose children are the most likely to attend college. Young college students and college graduates decumulate their wealth and borrow aggressively to enrol in college and to smooth consumption in their first years of working life. Finally, note that wealth inequality declines gradually over the life cycle. The magnitude of this decline is very close to its empirical counterpart, as documented in Kaplan (2011) from SCF data.

3.4.2 Determination of Inter Vivos Transfers

Two opposing forces shape the parent's decision of how much to transfer to their child. The first purpose is narrowing the gap between parent's and child's lifetime utilities, and the extent to which parents want to close this gap depends on the degree of altruism ω. This motive (intergenerational smoothing) is strongest for low ability (and low earnings potential) children, especially those with rich parents. The second purpose is that of alleviating the financial constraints of children in the event they choose to go to college. This second motive (college education financing) is strongest for high ability children whose return to attending college is the highest.

The left panel of Figure 3.2 shows that in the model inter vivos transfers (IVTs) increase monotonically with parental wealth at the age of the transfer (age 48). For many poor families the marginal cost of transferring to the children is too high in terms of their own foregone consumption, and they make no transfer. However, for the reasons discussed above, IVTs are not monotonic in child's ability (right panel). For low levels of ability the intergenerational smoothing motive dominates and IVTs decline in child's ability (most sharply for high parental wealth). At the high end of child's ability, IVTs rise again as the college education financing

[39]During retirement, the combination of annuity markets and discount rate above the interest rate implies a linear upward sloping pattern

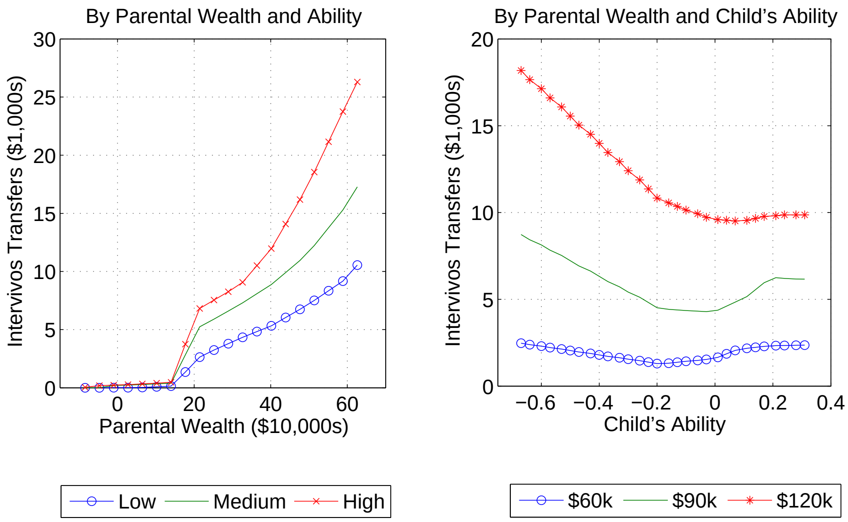

Figure 3.2: Parental transfers by wealth and child's ability

motive dominates. Note that this reversal of slope is most pronounced for intermediate wealth levels at which families can afford to make an extra sacrifice that allows their child to earn a college degree.

The right-panel of Figure 3.2 confirms that the optimal IVT rises fairly steeply with parental wealth. Note that IVTs are zero or very low for a wide range of parental wealth levels (left panel).[40] Finally this plot also shows that, for given wealth, high-ability parents save more for the IVT, as they expect their children to be on average of a high ability type as well, therefore with large gains from college education.

Parental IVTs determine the distribution of initial wealth in equilibrium. Meanwhile, the costs and returns to college education are jointly dictated by financial resources, and ability (directly, and through psychic costs). Therefore, it is important to ensure that the correlation among these two variables is consistent with the data. Zagorsky (2007) uses the 2004 module of the NLSY79 to estimate a correlation between income (net worth) and AFQT test scores of 0.30 (0.16, respectively) in a sample of individuals aged 40 and 47. In our benchmark simulation, the correlation between income (wealth) and ability θ for the same age range is 0.37 (0.10, respectively), hence empirically plausible.

In our model there is no insurance market for idiosyncratic productivity risk. As such, future income risk will have adverse welfare effects and will affect all aspects of behavior, including transfers and education choices. In this context transfers play a particularly important role. By offering an initial level of wealth to individuals, they improve their ability to self-insure against future shocks. To understand the interaction between risk, IVT and college enrollment in our model, we perturb the standard deviation of the idiosyncratic wage

[40] Indeed, in many cases parents would be better off with a negative transfer (i.e., receiving a transfer from their child) as they expect their child to earn more, eventually.

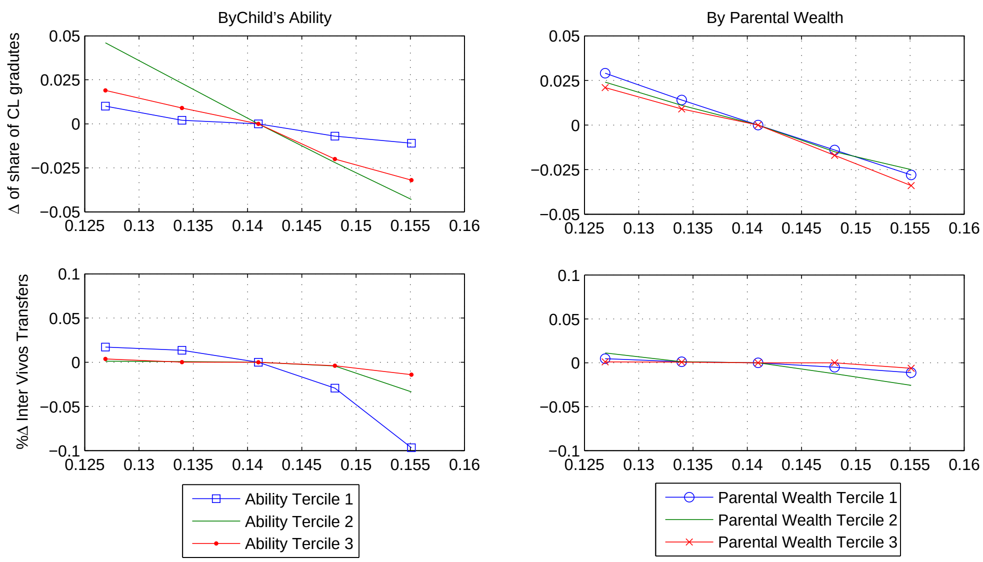

*Notes:*Inter vivos transfers and college attainment as functions of wage risk for college graduates (measured by the S.D. of idiosyncratic productivity innovation σ_η), by child's ability and parental wealth.

Figure 3.3: IVTs, Attainment and wage risk

shock for *new cohorts* of college graduates around its baseline value of 0.141 and plot the impact of this change on college attendance and transfers for different parental wealth levels and different levels of ability.[41] The results are presented in Figure 3.3.

College enrollment declines with an increase in risk: a rise of 10 percent in σ_η decreases enrollment by roughly 2.5 percentage points. The reason is that education choice is akin to a risky investment with a sunk cost: if the return becomes more uncertain, investment falls. The effect is strongest on middle-ability children (top-left panel), among which there are many who are marginal in the college decision. The impact of risk on enrollment is, instead, quite similar across wealth levels.[42]

Transfers are not very sensitive to changes in wage risk. On the one hand, a larger labor market risk for new college graduates induces parents whose child goes to college to a higher (precautionary) transfer. On the other hand, the marginal value of transferring resources declines because of lower college enrollment. These two forces tend to offset each other.[43].

[41]In this experiment, we keep prices and taxes fixed.

[42]These similar variations are, however, incremental to different benchmark attainment rates. For example, the share of college graduates from poorer families falls proportionally more with risk.

[43]The drop in inter-vivos transfers to low ability children is sizable in percentage terms, but the average magnitude of the transfer is small in absolute terms

	Data (NLSY97)	Model
Family Wealth Quartile 2	0.064	0.017
	(0.025)	(0.013)
Family Wealth Quartile 3	0.135	0.051
	(0.026)	(0.013)
Family Wealth Quartile 4	0.236	0.271
	(0.028)	(0.016)

Notes: Marginal effects of parental wealth on the probability of the child attending college. The column labeled 'Data' reproduces Table 5 in Belley and Lochner (2007). The column labeled 'Model' is the same regression from simulated data. As in Belley and Lochner, the model's regression also controls for parental education and ability quartiles (coefficients not shown).

Table 3.6: Marginal effects of parental wealth on attainment.

3.4.3 Intergenerational Income Correlation

Solon (1999) surveys the empirical evidence on intergenerational correlation in earnings for the U.S. and concludes that 0.4, or a bit higher, seems a reasonable estimate of the intergenerational elasticity in long-run earnings, once correcting for the attenuation bias due to the transitory component in earnings dynamics (either genuine or due to measurement error). When we compute the model's correlation between labor earnings of parent and child controlling for age (a moment not targeted in the calibration), we obtain 0.43, a value consistent with the empirical estimates.

The model also performs well in terms of intergenerational correlation of educational attainment. The model's correlation coefficient is 0.31. Mulligan (1999), Table 1, summarizes eight empirical studies and reports a range of 0.11-0.45 with an average correlation between parent's and child's education of 0.29.

3.4.4 Role of Family Wealth in Education Outcome

We have modeled government aid to college students through grants and loans in order to reproduce some salient features of the US federal aid system in the year 2000. We now verify that, under this set of institutional rules (and given the distribution of college costs in the same period), the role of parental resources in the determination of college enrollment in the model is consistent with that estimated in the data. Belley and Lochner (2007) measure the marginal effect of family wealth on the probability that their child attends college, conditional on child's ability and parental education from the NLSY97. We do the same in our model, from simulated data. Table 3.6 summarizes our findings. The model displays a steep parental wealth gradient on college enrollment, in line with the finding of Belley and Lochner (2007): controlling for child's ability, moving from the first to the fourth quartile of family wealth increases the probability of attending college by 27 percentage points.[44] One should keep in mind that in the data, as our model, parental wealth may be correlated with children's psychic costs of schooling, and therefore significance of parental wealth in this regression may not necessarily indicate a strong role for credit constraints. We return to this point in Section 3.5.

3.4.5 Validation Through a Simulated Randomized Trial

One of the objectives of our investigation is the comparison between partial and general equilibrium effects of various education policies. To further augment the credibility of our findings, we replicate within our model a randomized trial which is similar in nature to a quasi-experimental policy change: a (treated) group of high-school graduates receives an additional college tuition subsidy of $1,000 per year and a (control) group does not. Group membership is random. When we compute the additional college enrollment of the treated high-school graduates (which occurs in that same period) relative to that of the control group, keeping prices and taxes fixed, we obtain a rise of 3.3 percentage points (from 0.304 to 0.337).[45]

This large response is consistent with the existing empirical evidence. Kane (2003) and Deming and Dynarski (2009) provide a synopsis of the empirical estimates from similar quasi-natural experiments in which a discrete change in aid policy affects one group of individuals but not others, and conclude that enrollment in college of high-school graduates benefitting from an additional tuition grant of $1,000 rises between 3 and 5 percentage points.[46] Other studies use cross-state variation in tuition costs to estimate that enrollment would rise by 4 to 6 percentage points per $1,000 reduction in tuition costs (Cameron and Heckman, 1998; Kane, 1994).

[44] Adding income quartiles to the regression barely changes the marginal effects on wealth, and the coefficients on income are statistically insignificant. Also in Belley and Lochner the income quartile effects are insignificant, and the estimated wealth coefficients drop only by a couple of points. When we run the same regression on the probability of high-school completion, we find that the wealth gradient is much flatter, e.g., the effect of the fourth wealth quartile is only 0.093 (S.E. 0.005). Results in Belley and Lochner, Table 5, are similar. Their estimated coefficient on the fourth family wealth quartile in the HS completion regression is 0.095 (S.E. 0.019).

[45] In this experiment, the policy is announced to the parents and the children after the IVT. When we announce the subsidy before the IVT, and hence we allow a limited behavioral response from parents in that period, the rise in enrollment is just below 3 percentage points because the subsidy partially crowds out parental transfers.

[46] Among the policy changes surveyed in these two studies, the closest to our simulated experiment are the Georgia Hope Scholarship program, the Social Security Student Benefit program, the Washington DC Tuition Assistance Grant program, the Cal Grant program, and other similar examples of discontinuities in fellowship eligibility at individual institutions.

3.5 Policy Experiments

In this section we conduct two sets of policy experiments aimed at understanding the role of the federal financial aid system: the first focuses on government tuition grants, and the second on government-sponsored loan limits. We are especially interested in the difference between short-run partial equilibrium (PE) and long-run general equilibrium (GE) effects. The gap between the two is mediated by two major adjustments of the benchmark economy to the policy change: (i) the response of the endogenous distribution of inter vivos transfers (which determines the initial resources available to individuals); and (ii) the response of prices clearing input markets, and the fiscal adjustment following the additional educational expenditures. In all our policy experiments, the fiscal instrument that adjusts in equilibrium is the labor income tax τ^w.

We therefore present our results sequentially in three steps. The first step (PE Short run) computes changes in outcomes of interest one year ahead for the cohort treated by the policy. The policy announcement is made just before parents make their IVT, and hence this experiment incorporates only the short-run response of IVTs to the policy. In the second step (PE Long-run), we compute the new steady-state after the policy change where the distribution of parental IVTs (and initial assets for the newborn) has converged, but prices and taxes are fixed at their initial steady-state values. This step fully incorporates the adjustment of parental IVTs to the policy, but abstracts from GE feedbacks. In the third and final step (GE Long run), we incorporate these GE feedbacks and compute the new long run steady-state with the new market clearing prices and the new government budget-balancing taxes. In GE we also compute welfare gains, expressed as changes in expected lifetime consumption of a newborn individual after accounting for all costs associated with the policy.

3.5.1 Tuition Grants

Results of the grants experiments are summarized in Table 3.7. More detail is provided in Tables G1-G4 in Appendix G.

Removal of grants: We begin by assessing the value of the system of tuition grants currently in place in the US, as described in Section 3.3. When we remove grants altogether, in the short run college enrollment has a sharp drop of over 5 percentage points (PE Short run).[47] As shown in Table G1, the students who would suffer most are the children of low and middle wealth parents whose college attendance rate falls by 7 points. Children of high wealth parents are not much affected. This reduction in tuition grants is partially compensated by two mechanisms: a significant rise in IVT (which offsets 32% of the decline in grants) among parents whose kids go to college, and a small increase in hours worked by college students. In

[47] The college enrollment rate is expressed as a fraction of the the treated cohort in the PE short run experiment, and as a fraction of any cohort in the long-run steady-state experiments.

the long-run PE economy output falls by almost 2% because of the lower supply of educated labor, and this negative wealth effect reduces IVTs and contributes to an additional rise in hours worked by college students.

As the number of college graduates falls the relative price of their labor services goes up, which induces higher enrollment –especially among the wealthy (see Table G1). This GE feedback partially offsets the short-run PE forces, and the final share of college graduates in the long run GE economy is only 1.2 points below the initial steady state share. We conclude that long-run GE effects are weaker than short-run PE ones (roughly 4 times smaller). The GE impact on output associated to the long-run loss in productivity (due to the less skilled educated labor input) is negative and sizable (-1.4%). The welfare loss from scrapping the grant program, computed as the consumption-equivalent change in the expected lifetime utility of a newborn individual, is 0.38%, or roughly \$150 per year.

General Expansion of Tuition Grant: In this experiment, the government provides an extra \$1,000 of yearly tuition subsidy to all college students, independently of their needs or ability. In the short-run PE economy, college enrollment rises by almost 3 percentage points.[48] Table G2 shows that the composition of the college graduates shifts in favor of middle and high ability students, but is roughly unchanged by wealth level. The experiment reveals a substantial short-run crowding out of IVT in response to the grant: lower IVT of parents whose children go to college offset 30% of the grant.[49]. In the long-run PE experiment output increases substantially (+1.5%) because of the additional college-educated labor input in production. This adds wealth to the economy, which causes IVTs to increase and labor supply of college students to fall.

The long-run GE response to the general expansion of tuition subsidies is much more muted, roughly 0.6 points, i.e. five times smaller than its short-run PE counterpart. The main margin of adjustment is the change in relative labor price of college labor, which is larger than the the equilibrium fiscal adjustment or the interest rate adjustment. Note that in the long-run GE economy IVTs fall somewhat because the subsidy is now paid for by higher taxes, which is in contrast to the long-run PE economy. Finally, aggregate output rises by 0.4% and welfare of the newborn by 0.15% in lifetime consumption equivalent units, when the government expands its tuition subsidy program this way.

Means-Tested Expansion of Tuition Grant: In this experiment, grants are expanded proportionally by 56% for each student in an attempt to amplify the gap (in absolute value) between college tuition subsidies for poor families relative to those for rich families.[50] The

[48] The first column of Table G2 in the Appendix labeled 'Treatment' corresponds to the simulation in Section 3.4.5 where we have validated the model based on quasi-experimental evidence.

[49] Table G2 shows that the crowding out is especially sharp for wealthy parents, while for poor families in the lowest wealth tercile there is a small crowding-in effect as, for some parents, raising the IVT means now being able to afford the tuition fees and send their kid to college

[50] The additional cumulative 4-year grant for a student with $q = 1$ ($q = 3$) is \$4,740 (\$243, respectively).

Removal of grants				
	Benchmark	P.E. Short-run	P.E. Long-run	G.E. Long-run
College enrollment	0.261	0.209	0.223	0.249
Price of CL/HS labor	1.335	–	–	1.364
Crowding out of IVTs	–	+32%	-34%	+17%
Student labor supply	–	+3.1%	+5.4%	+4.5%
Aggregate output	–	–	-1.9%	-1.4%
Welfare gain	–	–	–	-0.38%

General tuition grant expansion ($1,000)				
	Benchmark	P.E. Short-run	P.E. Long-run	G.E. Long-run
College enrollment	0.261	0.290	0.295	0.267
Price of CL/HS labor	1.335	–	–	1.319
Crowding out of IVTs	–	-30%	+30%	-17%
Student labor supply	–	-0.4%	-4.8%	-4.3%
Aggregate output	–	–	1.5%	0.4%
Welfare gain	–	–	–	0.15%

Means-tested grant expansion				
	Benchmark	P.E. Short-run	P.E. Long-run	G.E. Long-run
College enrollment	0.261	0.294	0.297	0.273
Price of CL/HS labor	1.335	–	–	1.313
Crowding out of IVTs	–	-28%	+19%	-25%
Student labor supply	–	-2.8%	-2.8%	-2.1%
Aggregate output	–	–	1.5%	0.7%
Welfare gain	–	–	–	0.20%

Ability-tested grant expansion				
	Benchmark	P.E. Short-run	P.E. Long-run	G.E. Long-run
College enrollment	0.261	0.287	0.286	0.268
Price of CL/HS labor	1.335	–	–	1.317
Crowding out of IVTs	–	-33%	+23%	-24%
Student labor supply	–	-0.6%	-4.9%	-3.6%
Aggregate output	–	–	1.1%	0.4%
Welfare gain	–	–	–	0.14%

Notes: 'PE Short-run' incorporates only the short-run response of parental IVTs. 'PE long-run' is the new steady state, with the new ergodic distribution of parental IVTs, but where prices and taxes are fixed at their initial values. 'GE Long-run' is the new steady-state where prices clear input markets and tax rates balance the government budget constraint. "%" denote percentage changes from the benchmark value. The welfare gain is expressed in terms of changes in expected lifetime consumption for a newborn agent. 'Crowding Out' of IVT is defined as the average change in IVTs as a percentage of the average change in the grant (in absolute value) for all the individuals who attend college in the experiment. See Appendix G for more detailed results.

Table 3.7: Results of tuition grants experiments.

factor of proportionality is chosen so that the total cost of this policy to the government equals that of the general grant expansion of $1,000 in the short run PE economy.

Relative to the general grant expansion, the impact of this policy on college enrollment is somewhat larger. The short-run PE economy shows a rise in college enrollment of 3.3 percentage points. Table G3 reveals that it is mostly kids from low and middle wealth families, and middle ability levels, who increase their rate of college attendance. It remains true, also in this experiment, that relative price effects tend to strongly counteract the forces at work in PE, and in the long run GE economy the share of college graduates increases only by 1.2 percentage points. Overall, the impact of this need-based grant expansion on enrollment, output, and welfare is larger compared to the general grant expansion.

Ability-Tested Expansion of Tuition Grant: In this experiment grants are expanded by amounts proportional to students' abilities. A student in the bottom decile of the ability distribution receives $700 extra per year in federal grants, while someone in the top decile receives almost twice as much.This experiment is also designed so that its cost equals the cost of the general grant expansion in the short run PE economy.

This policy has smaller effects on enrollment, output, and welfare than the means tested expansion. The main reason is that, even in the absence of this additional transfer from the government, a large share of the high ability students (those targeted by the subsidy) would choose to attend college and reap the high return for college graduates, by borrowing through the federal loan system.

Taking stock of these results yields several important lessons. First, the existing federal grant program is welfare improving in our model and adds about 1.5% to aggregate output. Second, every additional dollar of government tuition grants crowds out 20-30 cents of parental IVTs. Third, expansions of this program may have sizeable effects on enrollment in the short run, but long-run GE effects are 3-4 times smaller because of relative price adjustments in the labor market. Interestingly, strong GE effects are triggered by small changes in relative skill prices.

3.5.2 Education Loans

Removal of Government-Sponsored Loans: We begin, as we did for grants, by assessing the value of the federal loan program. We remove from the economy all government-sponsored loans (i.e., we set both $\underline{b}^u$ and $\underline{b}^s$ to zero). From Table 3.8 it appears that the short run PE effects of eliminating federal loans are huge: college enrollment drops from 26% to 11%. Table G5 shows that college attendance rates collapse especially among high-ability students (from 56.8% to 27.4% in the highest ability tercile): it is the high ability students facing high tuition fees who borrow the most to attend college and take advantage of the high market returns for their type.

Removal of government loans				
	Benchmark	P.E. Short-run	P.E. Long-run	G.E. Long-run
College enrollment	0.261	0.113	0.237	0.258
Price of CL/HS labor	1.335	–	–	1.353
Crowding out of IVTs	–	+13%	-39%	-29%
Student labor supply	–	+3.1%	+33.1%	+32.3%
Aggregate output	–	–	-2.4%	-1.5%
Welfare gain	–	–	–	-0.38%

Expansion of government loans limits ($10,000)				
	Benchmark	P.E. Short-run	P.E. Long-run	G.E. Long-run
College enrollment	0.261	0.281	0.263	0.263
Price of CL/HS labor	1.335	–	–	1.334
Crowding out of IVTs	–	-29%	-2%	+1%
Student labor supply	–	-21.1%	-2.0%	-2.4%
Aggregate output	–	–	+0.2%	+0.1%
Welfare gain	–	–	–	+0.02%

Notes: 'PE Short-run' incorporates only the short-run response of parental IVTs. 'PE long-run' is the new steady state, with the new ergodic distribution of parental IVTs, but where prices and taxes are fixed at their initial values. 'GE Long-run' is the new steady-state where prices clear input markets and tax rates balance the government budget constraint. "%" denote percentage changes from the benchmark value. The welfare gain is expressed in terms of changes in expected lifetime consumption for a newborn agent. 'Crowding Out' of IVT is defined as the average change in IVTs as a percentage of the average change in the borrowing limit (in absolute value) for all the individuals who attend college in the experiment. See Appendix G for more detailed results.

Table 3.8: Results of loan limits experiments.

Enrollment in the long-run PE economy is already back to 0.237. This adjustment takes place largely through a much higher labor supply in college (which increases by 1/3). Although IVTs fall on average, this is entirely driven by reduced net worth at the high end of the parental wealth distribution. In particular, college-educated parents transfer nearly 7 times the average amount transferred by their HS-educated counterparts, and the long-run drop in their number results in a strong negative impact on average transfers in the economy. Moreover, aggregate output falls by 2.4% due to the large loss in labor productivity from the lower enrollment, especially of high-ability types. The relative price adjustment in GE contributes to setting the new long-run steady-state share of college graduates just 0.3 percentage points below its initial value, but output remains 1.5% lower, and the welfare of a newborn is reduced by 0.4%. Average inter vivos transfers also recover in GE, although they are substantially lower than in the benchmark: this derives from the changes in the composition of the college educated group, which now includes relatively more low ability individuals. In fact, despite recovery of aggregate enrollment in GE, output, welfare and IVTs are all below their benchmark level because of the reduced sorting in ability by education.

Because we assume that the private lenders do not respond to this policy, this experiment yields an upper bound. In a separate experiment, we let the private market pick up the demand for loans unmet by the government, i.e., we allow everyone to borrow from the private market at the rate $r + \iota$ and increase the private education loan limit $\underline{a}^p$ up to $\underline{b}^u + \underline{b}^s$. In this case, effects are smaller but still sizable. For example, in the long run GE economy, output falls by 1 percentage point.[51]

Finally, when we remove the entire financial aid system of loans and grants at the same time, in the long-run GE economy, enrollment falls by over 2 percentage points, output falls by 2.5%, and welfare by 0.5%. Hence the effects of the two separate experiments cumulate to some extent.

Expansion of Government-Sponsored Loans: In this experiment, the government expands unsubsidized loan limits by \$10,000 (but keeps interest rates r^u and r^s fixed at their initial values).[52] In the short-run, enrollment increases by 2 percentage points.[53] The higher ability to borrow induces students to work significantly less during college, and their parents to reduce their IVTs by 30 cents for each dollar of additional borrowing capacity. However, the

[51]This experiment is likely to be a lower bound for the value of government-sponsored loans. In an extended model where private lenders are profit-maximizing agents who take into account students' default risk, the removal of government loans would lead to an expansion of private lending, but possibly at a higher cost, which would reflect the entry of more risky borrowers into the private education loan market. Simpson and Ionescu (2010) model the private education loan market with some of these features.

[52]For the reasons discussed above (no reaction from private markets), this experiment is an upper bound on the effects of expanding public credit. In simulations of a model with endogenous private credit limits, Lochner and Monje-Naranjo (2011) find small crowding out, which offers a justification for our assumption.

[53]In line with this result, Johnson (2011) finds that the virtual elimination of limits on student loans has small effects on attainment.

long-run GE price adjustments offset entirely these forces, and enrollment and output remain virtually unchanged. How does one reconcile this finding with the result highlighted in Section 3.4.4 that parental wealth is a significant determinant of college attendance, conditional on child's ability? In our model parental wealth is negatively correlated with the non-pecuniary psychic costs of schooling. The effect of family resources on college enrollment outcomes is, indeed, largely determined by the fact that they proxy for psychic costs, an omitted variable in the regression of Table 3.6. When we add them to the regression, the marginal effect of parental wealth drop by half.

In summary, the current government-sponsored loan program is welfare improving in our model, especially for high-ability children, and adds 1.5 percentage points to aggregate output. Moreover, borrowing limits for college students appear to be generous: expanding them would have negligible consequences on enrollment and output in the long run. The key reason is that able students from low-wealth households –who could potentially benefit from relaxed federal loan limits– choose not to attend college because they have a large disutility cost of schooling. While we do not model explicitly the relationship between family assets and taste for college, a natural candidate is the fact that parental wealth determines the neighborhood and the quality of local public schools in which students mature and acquire the taste and capacity for further learning (see, e.g., Fernandez and Rogerson (1998)).

3.6 Discussion

The design of education policy is an issue at the top of the policy and research agenda. Over the years we have learnt a lot about the sensitivity of individual education outcomes to small policy changes as well as about the incidence and impact of liquidity constraints that justify policy interventions. However, most of what we know can only contribute partially to the design of policy because longer term adjustments, that take some time to play out, have been largely neglected. Moreover, the magnitude and direction of these adjustments differ substantially across the distribution of wealth and ability.

We have stressed two key channels of adjustment that are not allowed for in the usual policy evaluation literature. The first is the adjustment of relative wages and returns to different levels of education, an issue that has been studied before by Heckman, Lochner, and Taber (1998b,c), Lee (2005), Lee and Wolpin (2006). We confirm that allowing for equilibrium prices to change is crucial for quantifying the long run effects of policy. The second channel is the adjustment of funding by parents, which is a major source of support for college and reacts to policy. Its level, and the way it responds to public funding interventions, varies both with wealth and with child ability and as a result induces substantial heterogeneity in the effects of policy. For example, we find that while for wealthy parents with high ability children public subsidies crowd out private transfers, poorer parents tend to reinforce government subsidies since the expected return to their transfers increases when college becomes more attainable, particularly for those with high ability children. These two channels of adjustment interact strongly with each other, resulting in differences in policy outcomes both cross-sectionally and over time. To buttress the credibility of our estimates and to understand the extent to which liquidity constraints can affect education outcomes we have modelled the public grant and loan system in some detail. This, together with funding from parents, are central elements of a student's budget constraint.

Our policy experiments illustrate the role of an array of existing policy instruments and the complexity of their long run effects. The final outcomes after transfers have adjusted and further, after prices have been allowed to move to the new equilibrium, are substantially different from both the direct treatment effect and the short run PE effect before the entire output and distributional consequences have played out. Indeed any conclusions about the magnitude of the impact of policy based on immediate impacts are likely to overestimate the impacts in a serious way.[54] Based on our estimates, the largest long-term adjustments come from changes in returns rather than from transfers' behavior; in part this is because the latter are affected in counteracting fashion by wealth effects due to aggregate output changes. However, it is important to stress that the way IVTs change differs substantially depending on the level of parental wealth and on the ability of the child, determining large distributional impacts of policy.

The other conclusion of our policy experiments is that educational subsidies are welfare improving in our model even after accounting for changes in taxation to fund them and after allowing the economy to adjust to its new equilibrium. This is a clear indication that there are still individuals who are prevented from attending because of liquidity constraints and that this does have welfare implications, potentially justifying interventions. Importantly, the result is demonstrated in a context where we have allowed for the existing set of public programs as well as parental financing and self-financing through work while in college. As we might expect, the most effective policies are means tested grants, a finding consistent with the presence of some liquidity constraints. The gains are particularly large when returns to education are kept fixed (as in an open economy with free trade and factor price equalization) and are mitigated when prices adjust, but do not disappear. Our welfare results might understate the true value of government support during college because we restrict transfers to take place at a given age over the lifecycle. In reality families may have children at different points in their lifecycle and if poorer families were to have kids earlier they might need to transfer resources when they are relatively younger. With imperfect credit markets this additional heterogeneity would make it harder for them to help their kids, implying even larger welfare

[54] As emphasized, given our conservative choice for the elasticity of substitution in production, our GE effects are a lower bound.

gains from government subsidies.[55]

Finally, we also find that a non-trivial margin for education financing at the individual level is work while in college, which responds strongly when all other types of education finance are exhausted. This adjustment is different from the others, as it is relatively less dependent on ability and family background, and therefore available to more people; its importance however should be investigated further, allowing also for potential disruptions to schooling effort associated with working while in college.[56]

3.7 Conclusions

The capacity of people to optimally invest in education is crucial for economic prosperity and social mobility, and is an important determinant of the cross-sectional and intergenerational distribution of income (see Becker and Tomes, 1979; Loury, 1981). Education investments respond to the prevailing economic environment and contribute to shape it. The long-run nature of these relationships motivates the study of education policies within a general equilibrium framework.

To analyse these lifecycle decisions we set up a dynamic lifecycle model of individual behaviour in a context of altruistically linked overlapping generations. Our parameter estimates are based on estimation from various data sets, and when traditional microeconometric estimation is not feasible, we use calibration to make sure the model can replicate key features of the US economy.

We contribute to the debate on the increasing importance of family resources by modelling inter vivos transfers as equilibrium outcomes that respond to changes in market returns and to prevailing policy and credit conditions. In particular, we show that increases in college grants can displace IVTs in different proportions depending on family wealth, with the transfers made by wealthy families being generally crowded out the most. This finding suggests that means-tested policy expansions may be preferable also because they displace parental transfers to a lesser extent than more general expansions of aid.

Our analysis provides a sequence of snapshots of the economy after a policy change, revealing why the short-run partial equilibrium effects of education policies can be very different from their long-run general equilibrium effects. By accounting for the intricate patchwork of policies that compose the federal financial aid system, we are able to assess its value. Our results indicate that existing grant and loan programs improve welfare, and the changes in the distribution of human capital induced by these policies increase GDP in the long run.

[55] To correctly identify these additional transfer margins one would arguably need a much richer model allowing for quantity, quality and timing of fertility choices and transfers over the entire life cycle, including at death.

[56] See Garriga and Keightley (2009) for a model where time devoted to work competes with time needed to cumulate credits in college.

Our results also indicate that increased progressivity in grant programs would imply greater benefits than a generalized transfer expansion, although the incremental benefits from any aid expansion would be smaller than the gains from the current system. A means-tested grant expansion would generate greater selection into education among high-ability individuals than a general expansion, implying greater welfare gains.

Some of the methods and findings of this paper are promising for future research. Recent work (see Ionescu and Simpson, 2012; Lochner and Monje-Naranjo, 2011) has emphasized the expansion of private provision of student credit. Nesting endogenous borrowing constraints within an equilibrium framework, similar to the one developed in this paper, would allow for explicit co-determination of all credit and skill prices. Such a model, while more complex, could answer interesting questions about how private markets should be designed and regulated. Another promising direction for future research would involve looking at the role of early skill investment. We study college-age policies which take the ability distribution at age 16 as given. Different research (e.g., Caucutt and Lochner, 2012) stresses the importance of complementarities between college-age policies and interventions that release parental constraints in the critical phase of early skill accumulation, arguing that early investments may improve the effectiveness of tertiary education policies. Explicitly modelling sequential human capital investments within a rich environment, with endogenous skill and credit prices, might offer a natural way to relate the existence of early credit constraints with the observation that parental financial resources at the time of college choice matter much more now than they did in the past. Finally, one interesting generalization would account for heterogeneity in college types (e.g. Fu, 2012) allowing for the endogenous determination of returns based on demand and supply of different college types. This would require the careful modelling of the supply side with many heterogeneous types of education providers, as well as the possibility of differential credit access, to account for variation in the riskiness and returns of alternative education choices.

Bibliography

Adda, J., and R. Cooper (2003): *Dynamic Economics*. MIT Press, Cambridge, MA.

Aguiar, M., and E. Hurst (2005): "Cosumption versus Expenditure," *Journal of Political Economy*, 113(5), 919–948.

Attanasio, O., R. Bottazzi, H. Low, L. Neisham, and M. Wakefield (2012): "Modelling the Demand for Housing over the life-cycle," *Review of Economic Dynamics*, 15(1), 10–53.

Attanasio, O., and G. Weber (1993): "Consumption Growth, the Interest Rate and Aggregation," *Review of Economic Studies*, 60(3), 631–649.

Bajari, P., P. Chan, D. Krueger, and D. Miller (2013): "A Dynamic Model of Housing Demand: Estimation and Policy Implications," *International Economic Review*.

Barlevy, G., and J. D. Fisher (2011): "Mortgage Choices and Housing Speculation," *Federal Reserve Bank of Chicago*, WP 2010-12.

Becker, G. S. (1964): *Human Capital*. NBER.

Becker, G. S., and N. Tomes (1979): "An Equilibrium Theory of the Distribution of Income and Intergenerational Mobility," *Journal of Political Economy*, 87, 1153–89.

Belley, P., and L. Lochner (2007): "The Changing Role of Family Income and Ability in Determining Educational Achievement," *Journal of Human Capital*, 1, 37–89.

Ben-Porath, Y. (1967): "The Production of Human Capital and the Life Cycle of Earnings," *Journal of Political Economy*, 75(4), 352–365.

Blundell, R., M. Browning, and C. Meghir (1994): "Consumer Demand and the Life-Cycle Allocation of Household Expenditure," *Review of Economic Studies*, 61(1), 57–80.

Blundell, R., L. Pistaferri, and I. Preston (2008a): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

——— (2008b): "Consumption Inequality and Partial Insurance," *American Economic Review*, 98(5), 1887–1921.

Bohacek, R., and M. Kapicka (2012): "A Quantitative Analysis of Educational Reforms in a Dynastic Framework," Manuscript, University of California Santa Barbara.

Brown, M., J. K. Scholz, and A. Seshadri (2012): "A New Test of Borrowing Constraints for Education," *Review of Economic Studies*, 79, 511–538.

Burnside, C., M. Eichenbaum, and S. Rebelo (2011): "Understanding Booms and Busts in Housing Markets," *NBER Working Paper No. 61734*.

Calem, P., S. Firestone, and S. Wachter (2010): "Credit impairment and housing tenure status," *Journal of Housing Economics*, 19, 219–232.

Cameron, S. V., and J. J. Heckman (1998): "Life Cycle Schooling and Dynamic Selection Bias: Models and Evidence for Five Cohorts of American Males," *Journal of Political Economy*, 106(2), 262–311.

Cameron, S. V., and C. Taber (2004): "Estimation of Educational Borrowing Constraints Using Returns to Schooling," *Journal of Political Economy*, 112(1), 132–182.

Card, D., and T. Lemieux (2001): "Can Falling Supply Explain The Rising Return To College For Younger Men? A Cohort-Based Analysis," *The Quarterly Journal of Economics*, 116(2), 705–746.

Carneiro, P., and J. J. Heckman (2002): "The Evidence on Credit Constraints in Post-Secondary Schooling," *The Economics Journal*, 112, 705–734.

Caucutt, E., and L. Lochner (2012): "Early and Late Human Capital Investments, Borrowing Constraints, and the Family," Unpublished manuscript, University of Western Ontario.

Chamberlain, G. (1984): "Panel data," in *Handbook of Econometrics*, ed. by Z. Griliches, and M. D. Intriligator, vol. 2 of *Handbook of Econometrics*, chap. 22, pp. 1247–1318. Elsevier.

Chambers, M., C. Garriga, and D. Schlagenhauf (2009a): "Accounting for Changes in the Home Ownership Rate," *International Economics Review*, 50(3), 677–726.

——— (2009b): "The loan structure and housing tenure decisions in an equilibrium model of mortgage choice," *Review of Economic Dynamics*, 12, 444–468.

Coleman IV, M., M. LaCour-Little, and K. D. Vandell (2008): "Subprime Lending and the Housing Bubble: Tail wags dog?," *Journal of Housing Economcis*, 17, 272–290.

Cunha, F., J. Heckman, and S. Navarro (2005): "Separating uncertainty from heterogeneity in life cycle earnings," *Oxford Economic Papers*, 57(2), 191–261.

Deming, D., and S. Dynarski (2009): "Into College, Out of Poverty? Policies to Increase the Postsecondary Attainment of the Poor," NBER Working Paper 15387.

Díaz, A., and M. J. Luengo-Prado (2008): "On the user cost and homeownership," *Review of Economic Dynamics*, 11, 548–613.

Diaz-Serrano, L. (2005): "On the Negative Relationship between Labor Income Uncertainty and Home Ownership: Risk Aversion vs. Credit Constraints," *Journal of Housing Economics*, 14, 109126.

Domeij, D., and J. Heathcote (2003): "On The Distributional Effects Of Reducing Capital Taxes," Mimeo.

Favilukis, J., D. Kohn, S. C. Ludvigson, and S. Van Nieuwerburgh (Forthcoming): "International Capital Flows and House Prices: Theory and Evidence," in *Housing and The Financial Crisis*. Univesity of Chicago Press.

Favilukis, J., S. C. Ludvigson, and S. Van Nieuwerburgh (2012): "The Macroeconomic Effects of Housing Wealth, Housing FInance and Limited Risk Sharing in General Equilibrium.," *Working Paper*.

Fernandez, R., and R. Rogerson (1998): "Public Education and Income Distribution: A Dynamic Quantitative Evaluation of Education-Finance Reform," *American Economic Review*, 88(4), 813–833.

Fisher, J. D. M., and M. Gervais (2011): "Why has home ownership fallen among the young?," *International Economic Review*, 52(3), 883–912.

Flavin, M., and S. Nakagawa (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.

Fu, C. (2012): "Equilibrium Tuition, Applications, Admissions and Enrollment in the College Market," Mimeo, University of Wisconsin.

Gale, W., and J. Scholz (1994): "Intergenerational Transfers and the Accumulation of Wealth," *Journal of Economic Perspectives*, 8(4), 145–160.

Gallipoli, G., and L. Nesheim (2007): "The problem of non convexities in life-cycle models," Mimeo.

Gallipoli, G., and L. Turner (2011): "Household Responses to Individual Shocks: Disability and Labour Supply," Mimeo.

Garriga, C., and M. Keightley (2009): "A General Equilibrium Theory of College with Education Subsidies, In-School Labor Supply, and Borrowing Constraints," FRB St. Louis Working Paper 2007-015A.

Glaeser, E. L., J. D. Gottlieb, and J. Gyourko (2010): "Can Cheap Credit Explain the Housing Boom?," *Working Paper*.

Goldin, C., and L. F. Katz (2007): "The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005.," NBER Working Paper 12984.

Gouriéroux, C., A. Monfort, and E. Renault (1993): "Indirect Inference," *Journal of Applied Econometrics*, 8, s85–s118.

Guvenen, F. (2009): "An Empirical Investigation of Labor Income Processes," *Review of Economic Dynamics*, 12(1), 58–79.

Hanushek, E., and J. Quigley (1980): "What is the price elasticity of housing demand?," *The Review of Economics and Statistics*, 62(3), 449–454.

Haughwout, A., D. Lee, J. Tracey, and W. van der Klaauw (2011): "Real Estate Investors, the Leverage Cycle, and the Housing Market Crisis," *FRB of New York Staff Report 514*.

Heathcote, J., F. Perri, and G. L. Violante (2010a): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States 1967-2006," *Review of Economic Dynamics*, 13(1), 15–51.

——— (2010b): "Unequal We Stand: An Empirical Analysis of Economic Inequality in the United States, 19672006," *Review of Economic Dynamics*, 13(1), 15–51.

Heathcote, J., K. Storesletten, and G. L. Violante (2010): "The Macroeconomic Implications of Rising Wage Inequality in the US," *Journal of Political Economy*, 118(4), 681–72.

——— (2012): "Consumption and Labor Supply with Partial Insurance: An Analytical Framework," Mimeo.

Heckman, J. (1981): "The Incidental Parameters Problem and the Problem of Initial Condition in Estimating a Discrete Time-Discrete Data Stochastic Process," in *Structural Analysis of Discrete Data and Econometric Applications*, ed. by C. F. Manski, and D. L. McFadden. MIT Press.

HECKMAN, J., L. LOCHNER, AND C. TABER (1998a): "Explaining Rising Wage Inequality: Explorations with a Dynamic General Equilibrium Model of Labor Earnings with Heterogeneous Agents," *Review of Economic Dynamics*, 1(1), 1–58.

——— (1998b): "General Equilibrium Treatment Effects: A Study of Tuition Policy," *American Economic Review, Papers and Proceedings*, 88(2), 381–386.

——— (1998c): "Tax Policy and Human Capital Formation," *American Economic Review, Papers and Proceedings*, 88(2), 293–297.

HECKMAN, J., L. LOCHNER, AND P. TODD (2006): "Earnings functions, rates of return and treatment effects: The Mincer equation and beyond," *Handbook of the Economics of Education*, 1, 307–458.

HUGGETT, M. (1993): "The Risk Free Rate in Heterogeneous-Agents, Incomplete Insurance Economies," *Journal of Economic Dynamics and Control*, 17(5/6), 953–970.

IACOVIELLO, M., AND M. PAVAN (2013): "Housing and Debt over the life-cycle and over the business cycle," *Journal of Monetary Economics*, Forthcoming.

IONESCU, F., AND N. SIMPSON (2012): "Credit Scores and College Investment," mimeo.

JOHNSON, M. (2010): "Borrowing Constraints, College Enrollment, and delayed Entry," *Working Paper*.

——— (2011): "Borrowing Constraints, College Enrollment, and Delayed Entry," Human Capital and Economic Opportunity Working Paper 2011-006.

JOHNSON, M., AND M. KEANE (2013): "A Dynamic Equilibrium Model of the US Wage Structure, 1968-1996," *Journal of Labor Economics*, 31(1), 1–49.

KANE, T. (1994): "College entry by blacks since 1970: The role of college costs, family background, and the returns to education," *Journal of Political Economy*, 102(5), 878–911.

KANE, T. (2003): "A Quasi-Experimental Estimate of the Impact of Financial Aid on College-Going," NBER Working Paper 9703.

KAPLAN, G. (2011): "Inequality and the Lifecycle," *Quantitative Economics*, Forthcoming.

KAPLAN, G., AND G. L. VIOLANTE (2012): "A Model of the Consumption Response to Fiscal Stimulus Payments," Mimeo.

KEANE, M., AND R. ROGERSON (2011): "Reconciling Micro and Macro Labor Supply Elasticities: A Structural Perspective," NBER Working Paper 17430.

KEANE, M., AND K. WOLPIN (2001): “The Effect of Parentnal Transfers and Borrowing Constraints on Educational Attainment,” *International Economic Review*, 42(4), 1051–1103.

KORTEWEG, A., AND M. SORENSEN (2011): “Estimating Loan to Value and Foreclosure Behavior,” Working Paper.

KRUEGER, D., AND A. LUDWIG (2013): “Optimal Progressive Taxation and Education Subsidies in a Model of Endogenous Human Capital Formation,” Mimeo.

KRUEGER, D., AND F. PERRI (2011): “How do Households Respond to Income Shocks?,” Mimeo.

LAITNER, J. P. (1992): “Random Earnings Differences, Lifetime Liquidity Constraints, and Altruistic Intergenerational Transfers,” *Journal of Economic Theory*, 58, 135–170.

LANDVOIGHT, T., M. PIAZZESI, AND M. SCHNEIDER (2012): “The Housing Market(s) of San Diego,” NBER Working Paper No. 17723.

LANDVOIGT, T., M. PIAZZESI, AND M. SCHNEIDER (2012): “The Housing Market(s) of San Diego,” Working Paper.

LEE, D. (2005): “An Estimable Dynamic General Equilibrium Model of Work, Schooling and Occupational Choice,” *International Economic Review*, 46, 1–34.

LEE, D., AND K. WOLPIN (2006): “Intersectoral Labor Mobility and the Growth of the Service Sector,” *Econometrica*, 47, 1–46.

LEVHARI, D., AND Y. WEISS (1974): “The effect of risk on the investment in human capital,” *American Economic Review*, 64, 950–63.

LI, W., H. LIU, AND R. YAO (2009): “Housing over time and over the life cycle: A structural estimation,” *FRB of Philadelphia Working Paper*.

LI, W., AND R. YAO (2007): “The Life-Cycle Effects of House Price Changes,” *Journal of Money, Credit and Banking*, 39(6), 1375–1409.

LINNEMAN, P., AND S. WACHTER (1989): “The impact of borrowing constraints on homeownership,” *American Real Estate and Urban Economics Association Journal*, 17, 389–402.

LOCHNER, L., AND A. MONJE-NARANJO (2011): “The Nature of Credit Constraints and Human Capital,” *American Economic Review*, 101(6), 2487–2529.

LOURY, G. C. (1981): “Inter-generational transfers and the distribution of earnings,” *Economica*, 49(4), 843–67.

Low, H., C. Meghir, and L. Pistaferri (2010): "Wage Risk and Employment Risk over the Life Cycle," *American Economic Review*, 100(5), 1432–67.

McFadden, D. (1989): "A method of simulated moments for estimation of discrete response models without integration.," *Econometrica*, 57, 531–42.

Megboluge, I., and P. Linneman (1993): "Home Ownership," *Urban Studies*, 30(4/5), 659–682.

Meghir, C., and L. Pistaferri (2006): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.

Mian, A., and A. Sufi (2009): "The Consequences of Mortgage Credit Expansion: Evidence From the U.S. Mortgage Default Crisis," *The Quarterly Journal of Economics*, 124, 1449–1496.

Mincer, J. (1958): "Investment in Human Capital and Personal Income Distribution," *Journal of Political Economy*, 66(4), 281–302.

Mitchell, O., and J. Phillips (2006): "Social Security Replacement Rates for Own Earnings Benchmarks," Pension Research Council Working Paper 2006-6.

Mulligan, C. (1999): "Galton versus the Human Capital Approach to Inheritance," *Journal of Political Economy*, 107(S6), S184–S224.

National Center for Education Statistics (1998): "Postsecondary Financing Strategies:How Undergraduates Combine Work, Borrowing, and Attendance," Statistical Analysis Report.

——— (2000): "Trends in College Pricing, 2000," Report.

Pakes, A., and D. Pollard (1989): "Simulation and the asymptotics of optimization estimators," *Econometrica*, 57(5), 1027–1057.

Pistaferri, L., and C. Meghir (2004): "Income Variance Dynamics and Heterogeneity," *Econometrica*, 72(1), 1–32.

Polivka, A. E. (2000): "Using Earnings Data from the Monthly Current Population Survey," mimeo, Bureau of Labor Statistics.

Restuccia, D., and C. Urrutia (2004): "Intergenerational Persistence of Earnings: The Role of Early and College Education," *American Economic Review*, 94, 1354–1378.

Ríos-Rull, J.-V. (1995): "Models with Heterogenous Agents," in *Frontiers of Business Cycle Research*, ed. by T. F. Cooley, chap. 4. Princeton University Press, Princeton.

Rosen, S. (1977): "Human Capital: Relations between Education and Earnings," in *Frontiers of Quantitative Economics*, ed. by M. D. Intriligator, vol. 3b. Amsterdam: North Holland.

Rothenberg, T. (1971): "Identification in parametric models," *Econometrica*, pp. 577–591.

Seigal, S. (2008): "Consumption-Based Asset Pricing: Durable Goods, Adjustment Costs, and Aggregation," Working Paper.

Simpson, N., and F. Ionescu (2010): "Credit Scores and College Investment," Mimeo.

Solon, G. (1999): "Intergenerational Mobility in the Labor Market," in *Handbook of Labor Economics*, ed. by O. C. Ashenfelter, and D. Card, vol. 3A, pp. 1761–1800. New-Holland, Amsterdam.

Storesletten, K., C. I. Telmer, and A. Yaron (2000): "Consumption and Risk Sharing Over the Life Cycle," *NBER Working Paper*.

——— (2004): "Consumption and Risk Sharing Over the Life Cycle," *Journal of Monetary Economics*, 51, 609–33.

Su, C.-L., and K. L. Judd (2012): "Constrained Optimization Approaches to Estimation of Structural Models," *Econometrica*, 80, 2213–2230.

United States Census Bureau (2012): "Table 982," in *The 2012 Statistical Abstract*. U.S. Census Bureau.

Willis, R., and S. Rosen (1979): "Education and Self-Selection," *Journal of Political Economy*, 87(5), s7–s36, pt. 2.

Winter, C. (2012): "Accounting for the Changing Role of Family Income in Determining College Entry," mimeo.

Yang, F. (2009): "Consumption over the Life-Cycle: How Different is Housing?," *Review of Economic Dynamics*, 12(3), 423–443.

Zagorsky, J. L. (2007): "Do you have to be smart to be rich? The impact of IQ on wealth, income and financial distress," *Intelligence*, 35, 489–501.

Appendix A

Inter-Vivos Transfers

Our source of information on inter-vivos transfers (i.e., gifts from parents to their children) is the NLSY97. We mostly use measures from the 'Income' subsection of the survey, complemented with information from the College Experience section.[57]

Transfers measured in the Income section refer to *all* income transferred from parents or guardians to youth that are neither loans nor regular allowance. This data is elicited through a series of questions, which also assess whether the individual lives with both, one or none of their parents. Our measure of inter-vivos transfers uses the inter vivos transfer variable from youth who live with both parents, when it is available. When the youth reports not living with both parents we sum the inter-vivos transfers from both living mother/mother figure and father/father figure.[58] If any of these values are missing (e.g. mother's transfer) then we include only the non-missing value (in this example, father's transfer). Observations which have missing values for all three possibilities to report inter-vivos transfers are dropped from the sample.[59] For youth living at home we also compute the implicit transfer corresponding to the value of rent, which is based on the estimated average rent paid by independent youth of the same age.

We use waves from 1997 to 2003.[60] This gives us an initial sample of 8,984 youths who were between age 12 and 16 in 1997. Only respondents that are part of the cross-sectional (representative) sample are kept, which leaves 6,748 individuals. We compute the cumulative transfers received between ages 16 and 22. When we drop observations for youth below age 16 in 1997, and 13 cases of obvious mis-reporting, we obtain a final sample of 6,346 youths and a total number of observations equal to 21,136. In this final sample, approximately 75% of youth report living in households with at least one (biological or adoptive) parent as guardian.[61]

[57]The College Experience section has information about parental transfers earmarked for financial aid while attending a post-secondary academic institution. These transfers are not fully consistent with the information in the 'Income' section, contain many skips and, most importantly, they do not cover all transfers. For this reason we only use limited information from this section and make sure to include it so as to minimize reports' error.

[58]Those individuals who do not live with a mother/mother figure or a father/father figure, and whose biological mother and father are not alive, are not asked questions on transfers.

[59]Additional details are available from the authors.

[60]Data for 2004 are dropped as there are no comparable inter-vivos amounts available after that year.

[61] In principle, observations should be weighted when tabulating population characteristics. However, as suggested by the BLS, the use of weights is inappropriate in samples generated after dropping observations

	Not in College		In College	Share
Less than HS	\$3,385		N/A	14.1%
HS Graduate	\$3,534		N/A	59.8%
College Graduate	\$5,469		\$7,807	26.1%
Average		\$4,366		

*Notes:*Yearly inter-vivos transfers by educational attainment and by current college enrollment status (for college graduates) of the child. Amounts are expressed in year 2000 dollars and include allowances.

Table A.1: Yearly inter-vivos transfers.

In the final sample from the Income section, one third of observations (32.4%) report positive cash transfers elicited from the relevant survey questions, meaning 67.6% reported not receiving any such transfers. However, when imputed rent is included, 75.1% of observations have positive transfers. The value of imputed rent varies from age to age with a minimum of \$4,966 per year for kids aged 16 and a maximum of \$6,615 for 22 year old youth.

In the College Experience section questions about financial help from parents are asked for each term in College and refer to transfers specifically provided for school.[62] The sampling restrictions are the same as the ones used for the Income section. Parental aid variables are categorized by year for each respondent, and then summed up to generate an average variable for each year between 1997 and 2003. Given the way questions were designed and asked, the transfers recorded in the College section should be a subset of the transfers recorded in the Income section. However in a large number of cases, especially for students enrolled in 4-year Colleges, the transfer measures in the College section are larger than those in the Income section. Following some correspondence with the BLS, we concluded that transfer measures from the College section are generally less reliable than those in the Income section. However, it is also possible that respondents included parental payments of tuition fees in the College section transfer (for instance, if the parents paid tuition fees directly and respondents chose not to report such amounts in the Income section).

To calculate inter-vivos transfers, we chose to use both sources of data. More specifically, we use completed schooling by survey year 2009 to classify individuals within three groups: (1) those who have completed a 4-year college degree, and those who are currently enrolled

reporting item non-responses. Nonetheless we also experiment using the BLS custom weighting engine to construct specific weights for our sample, with results changing only marginally. In what follows we use only results from the un-weighted sample.

[62] After one term has been reported, the respondent is asked if the information for the next term has changed from the previous term, and if it has not, the information is not recollected.

in, or have completed, a graduate degree; (2) those who have completed a high school degree, but are not in group (1); (3) those who have not completed a high school degree.

Table A.1 summarizes the average yearly transfer received by people with different education achievement (in survey year 2009); in the case of College graduates we distinguish between transfers received while in College and transfers received in other years. For the years of College attendance we approximate the total inter-vivos transfer as the maximum between transfers recorded in the Income section and transfers recorded in the College Experience section.[63]

Using the average transfers in Table A.1 we compute the total amount received by youth with less than a College degree over a 7 year period by simply multiplying estimated yearly transfers and allowances by seven (note that all amounts are expressed in year 2000 dollars). In the case of College graduates we compute the total transfer received over 7 years by summing up the average amount received while in College multiplied by four (which is the College duration in the model) and the amount received while out of College multiplied by three.

This procedure results in a total transfer of $23,697 for HS drop-outs, $24,735 for HS graduates and $47,637 for College graduates. These figures are used to target transfers-by-education-achievement in the benchmark economy. Weighting each transfer level by the share of workers of a given education type in the benchmark economy results in an economy-wide average transfer of $30,566. Transfers reported in Table A.1 include allowances. Reported allowances are small, adding only $135 dollars to the average yearly transfer.

[63] An alternative way to approximate transfers during College years is to sum the measures from the two sections, rather than taking the higher one. This results in very similar average yearly transfers.

Appendix B

Ability Transmission and Distribution

Intergenerational Transmission: To estimate the Markov transition matrix that dictates the intergenerational transmission of ability between parents and children, we use the expanded mother-child data collection from the NLSY79. The NLSY79 is a representative sample of 12,686 American young men and women who were 14-22 years old when they were first surveyed in 1979. Data was collected yearly from 1979 to 1994, and biennially from 1996 to the present. The "Children of the NLSY79" survey began in 1986 and has occurred biennially since then. This survey consists of detailed information on the development of children born to NLSY79 women. A battery of cognitive, socio-emotional, and physiological assessments are administered to these children at various ages and scores recorded.

There are 11,340 children born to the total 4,890 female respondents of the NLSY79 who are mothers of at least one child. We link the children's file to the main data file using the individual identifier for mothers. Each child has test scores taken in different years. However, many child/year combinations do not have any test score observations. The child test scores reported are the PIAT Math, the PIAT reading comprehension, the PIAT Reading Recognition, and the PPVT score. We use the latest PIAT Math test scores to rank children's ability: in particular, we use standardized scores of the PIAT Math test, which are derived on an age-specific basis from the child's raw score and are comparable across ages. The Peabody Individual Achievement Test (PIAT) is a wide-ranging measure of academic achievement which is well known and used in applied research. For details of the way the PIAT is computed and "normed" by age, see Chapter 2 (page 89 and up) of the "NLSY79 Child and Young Adult Users Guide". In general, the PIAT Math is a highly reliable and valid assessment. As described in the "NLSY Child Handbook:1986-1990" and "The NLSY Children 1992", it correlates closely with other cognitive measures, and it is both predicted by and predicts scores on a variety of the other assessments.

This leaves us with 3,389 mothers and 7,589 mother-child pairs. We restrict our attention only to mothers who are part of the cross-sectional (nationally representative) sample of the NLSY79, which further reduces our mother-child pairs to 4,455 and the total number of mothers to 2,087. Table B1 reports the distribution of children's age at the time of test in

Age	Number	Percent	Age	Number	Percent
5	98	2.2	12	331	7.4
6	202	4.5	13	1,208	27.1
7	194	4.4	14	1,081	24.3
8	231	5.2	15	87	2.0
9	251	5.6	16	49	1.1
10	301	6.8	17	45	1.0
11	368	8.3	18	9	0.2

*Notes:*Child's age at time of test (relative frequency).
Total number of mother-child pairs: 4,455 (NLSY).

Table B1: Child's age at time of test.

our final sample.

The fact that children took the PIAT test at different ages should have no relevance because we use standardized scores which control for the age of the test-subject. In a robustness check we also computed ability transition matrices using a smaller sample including only mother-child pairs in which the child was at least 13 years of age at the time of the test and results were virtually the same.

To measure ability of mothers, we use AFQT scores. During the summer and fall of 1980, NLSY79 respondents participated in an effort of the U.S. Departments of Defense and Military Services to update the norms of the Armed Services Vocational Aptitude Battery (ASVAB). A total of 11,914 civilian and military NLSY79 respondents completed this battery of tests.[64] A composite score derived from selected sections of the battery can be used to construct an approximate and unofficial Armed Forces Qualifications Test score (AFQT) for each youth. The AFQT is a general measure of trainability and a primary criterion of enlistment eligibility for the Armed Forces. Two methods of calculating AFQT scores, developed by the U.S. Department of Defense, have been used by CHRR to create two percentile scores, an AFQT80 and an AFQT89, for each respondent. We use the latter score in our analysis, because it is also the ability measure used in the estimation of the wage equations (see below).

Test-scores (AFQT89 for mothers, PIAT Math for children) are used to assign mothers and children to quintiles, according to their relative ranking in the sample. After splitting mothers and children into these quintiles, we compute the conditional probabilities of transiting from a given mother's quintile to her child's quintile. Results are reported in Table 3.1 in the main text.

[64] The ASVAB consists of 10 tests that measure knowledge and skill in the following areas: (1) general science; (2) arithmetic reasoning; (3) word knowledge; (4) paragraph comprehension; (5) numerical operations; (6) coding speed; (7) auto and shop information; (8) mathematics knowledge; (9) mechanical comprehension; and (10) electronics information.

For each maternal quintile, the first row reports the number of sample children in each quintile, the second row reports the conditional probability of ending up in that quintile.

Empirical Distribution: The distribution of AFQT scores among mothers is extremely similar to the distribution of AFQT scores in the entire cross-sectional sample, which we use in the estimation of the wage-ability gradient.

Appendix C

Psychic Costs

The psychic costs entering the decision problems of potential high-school and college students consist of two additive components: a preference for education by ability, and a preference for education by family background. Thus we can write $\kappa^e(\theta, q) = \kappa^e_\theta(\cdot) + \kappa^e_q(\cdot)$.

Table C1 reports the the consumption equivalent values (CEV's) of the psychic costs associated to graduating from High School and to graduating from College for each pair of parental wealth (summarized by the variable q) and ability θ. The CEV's are expressed, respectively, as a share of average lifetime consumption of high school graduates and college graduates after education is completed. The CEV incurred by individuals to graduate from college includes both high school and college costs, and should be interpreted as the total psychic cost faced by college graduates over the life cycle.

By comparing the top and the bottom table, it appears that college graduates of wealthy families ($q = 2, 3$) face psychic costs of schooling mostly during high-school, whereas the additional cost of college is small, or even negative for some types. In contrast, children of low-wealth families incur a large additional psychic cost of college, on top of that incurred during high-school. One could interpret this finding as saying that the quality of high-school education for children of low-wealth households is poor and they are less well prepared for college, or as saying that children of low-wealth parents have been inculcated with less of a taste for education.

The CEV of the average psychic costs of attending college implied by the model gives us an average value of $335,000. This magnitude is comparable to the values suggested by Cunha, Heckman, and Navarro (2005) and Heckman, Lochner, and Todd (2006), who estimate average psychic costs of graduating from college at around $375,000 (see, e.g., Table 19 in Heckman, Lochner, and Todd, 2006).

	High School graduates				
	$\theta(1)$	$\theta(2)$	$\theta(3)$	$\theta(4)$	$\theta(5)$
$q(1)$	15.5	17.4	18.7	18.7	24.7
$q(2)$	10.2	12.5	14.0	13.9	20.4
$q(3)$	11.0	13.1	14.6	14.5	21.8
	College graduates				
	$\theta(1)$	$\theta(2)$	$\theta(3)$	$\theta(4)$	$\theta(5)$
$q(1)$	30.2	33.7	35.7	35.0	35.0
$q(2)$	9.2	14.3	16.3	14.8	14.6
$q(3)$	9.5	14.6	16.6	15.2	15.3

*Notes:*Estimates of psychic costs $\kappa(\theta, q)$ by ability and parental wealth (measured by the indicator q). We report them here as a percentage of average lifetime consumption of, respectively, high-school and college graduates, after graduation.

Table C1: Psychic costs

Appendix D

Individual Productivity Dynamics

Wage-Age Profiles from the PSID: The Panel Study of Income Dynamics (PSID) is a longitudinal survey of the US population. We use data for the waves from 1968 to 2002 (referring to calendar years 1967 to 2001). Since 1997 the PSID has become biannual. We follow closely the sampling criteria of Meghir and Pistaferri (2006) and restrict attention only to heads of household in the the SRC sample, which was originally nationally representative, so we use no sample weights in the calculations. By selecting only heads of household (mostly men or single women) we restrict our attention to individuals with relatively strong attachment to the labor force. After selecting the observations on household heads we are left with 19,583 individuals. Dropping people younger than 25 or older than 65 leaves us with 18,186 individuals. Dropping the self-employed leaves 14,866 persons in the sample. We then select only individuals with at least 8 (possibly non continuous) observations, which further reduces the individuals in the sample to 6,228. Dropping individuals with unclear education records leaves 6,213 people in sample. Disposing of individuals with missing, top-coded, or zero earnings reduces the sample to 5,671 individuals and dropping those with zero, missing or more than 5,840 annual work hours brings the sample size to 5,660 individuals. We then eliminate individuals with outliers in earnings growth, defined as changes in log-earnings larger than 4 or less than -2, which leaves 5,477 individuals in the sample. Finally, dropping people connected with the original SEO sample (which oversamples low income households) reduces the number of individuals to 3,085.

The wage variable we use for our calculations is the hourly earnings (total labor income divided by total hours worked) for the head of the household expressed in 1992 dollars by deflating nominal wages through the CPI-U for all urban consumers.[65] Information on the highest grade completed is used to allocate individuals to three education groups: high school drop-outs, high school graduates, and college graduates.

The estimated age polynomials are presented in Table D1 for different education groups and for the pooled sample. As commonly found, these estimates imply a steeper efficiency profile for more educated workers: between ages 22-52, productivity doubles for HS dropouts,

[65] In the PSID the head of the household is a male whenever there is a cohabiting male/female couple. The earnings variable includes the labor part of both farm and business income, wages, bonuses, overtime, commissions, professional practice and others.

Dependent variable: Log hourly wages				
	Less than HS	**HS Graduate**	**College Graduate**	**Pooled**
	Coefficient	Coefficient	Coefficient	Coefficient
	(S.E.)	(S.E.)	(S.E.)	(S.E.)
Age	0.2	0.41	0.67	0.46
	(0.015)	(0.06)	(0.10)	(0.05)
Age^2	-0.01	-0.013	-0.02	-0.014
	(0.001)	(0.002)	(0.004)	(0.002)
Age^3	1.e-4	2.e-4	3.e-4	2.e-4
	(1.e-5)	(4.e-5)	(6.e-5)	(3.e-5)
Age^4	-8.e-7	-1.e-6	-1.6e-6	-1.2e-6
	(2.e-7)	(2.e-7)	(3.7e-7)	(1.8e-7)

Table D1: Estimated age polynomials' coefficients

but it triples for college graduates.

Price of Labor Inputs from PSID: Once we filter out age effects from hourly wages, we can construct first-differences in logs and also filter out ability, since it enters linearly in the log-wage equation. Performing this estimation in first differences is essential because the average ability by education group is not constant over time due to composition changes within the group. Therefore, we can easily estimate, through time dummies, the time series of price growth in each education group, i.e., the term $\Delta \log w_t^e$. Given a normalization one can recover spot prices year by year.[66]

Wage-Ability Gradient from NLSY: Ability is approximated by the AFQT89. For hourly wages, we use the wage variable corresponding to the hourly rate of pay on the current or most recent job, available only from 1979 to 1994. We start with the 11,878 individuals for which we have AFQT89 scores. We drop those individuals who are unemployed or out of the labor force, or employed but reporting zero wage or with annual work hours missing, below 400 or larger than 5,840: this reduces the sample to 10,592 individuals. Dropping individuals who report (at least once) hourly wages above \$400 or below \$1 further reduces the sample to 10,202. We also eliminate individuals who report log wage increases larger than 4 or smaller than -2, which leaves 10,056 workers in the sample. Finally, we drop individuals who change their education level during their working life, which gives us a final sample of 7,954 individuals. When we split this sample in 3 education groups, we get a HS drop-outs' sample of 1,341

[66] We use a normalization based on the relative hourly wages observed in our PSID sample in 1989. First we compute average wages by education group for 1989, and next we correct for ability composition using information from the NLSY79 (AFQT test scores distribution together with their education-specific gradient on wages). We choose 1989 because people from the NLSY79 are between age 23 and 31, which means most of them are already working. Additional details on the normalization and the ability adjustment are available upon request.

individuals, a HS graduates' sample of 5,403 individuals and a college graduates' sample of 1,206 individuals. [67] Table 3.3 reports estimates of the ability gradient by education group, and for the pooled sample. All standard errors are corrected for individual clustering.

We use specifications with time dummies to control for time variation in market wages, but estimates are almost identical to those obtained without time dummies. We also run specifications based on wages which are not purged of the estimated PSID age-effects: again, results based on these measures are similar to those obtained for the age-free wages reported below.

The NLSY contains two additional measures of wages: (i) a variable corresponding to the hourly rate of pay in the first reported job, available only from 1979 to 2002; (ii) a hourly wage rate obtained dividing total earnings by total hours worked in the previous calendar year. The latter variable can be constructed for each wave between 1979 and 2002. The earnings' measure includes wages, salary, commissions or tips from all jobs, before deductions for taxes. The ability gradient estimated from our preferred wage measure is very close to, and falls between, the estimated ability gradients estimated using these two alternative definitions of hourly wages. Differences are statistically insignificant and confirm the robustness of the estimated reduced-form ability gradients.

Estimation of Error Component Model for Wage Residuals: The final step is estimating the parameters of the persistent-transitory shocks model for wage residuals. Wage residuals are obtained from NLSY data purging from individual log wages time dummies, the age component and the ability component, calculated as explained above. For estimation, we use the Minimum Distance Estimator originally proposed by Chamberlain (1984), as implemented in Heathcote, Storesletten, and Violante (2010). In a nutshell, as moments we use the covariances of wage residuals at various lags for different age groups. Table 3.4 reports the estimates of these parameters obtained for the 15-year period between 1979 and 1993.[68]

[67] We use all workers including NLSY79 over-samples in our estimation to maximize the number of observations: a dummy is introduced to control for possible hourly wage differences of workers from the over-samples. Over-sample dummies are mostly not significant. Even when significant they are very small.

[68] More details are available from the authors upon request.

Appendix E

Aggregate Technology Parameters

The estimation of the aggregate technology parameters is based on data from the Current Population Survey (CPS) March supplement, a survey conducted by the Bureau of the Census for the Bureau of Labor Statistics. The sample includes the adult universe (i.e., the population of marriageable age, with all individuals aged 15 and over unless they have missing or zero earnings, or missing educational attainment information). Since earnings data are top-coded in the CPS, we extrapolate the average of the top-coded values by using a tail approximations based on a Pareto distribution.[69] We compute total wage bills in billions of dollars for the three education groups. Dividing the wage bills by the (normalized) marginal products of human capital estimated from PSID data (see discussion in Appendix D), we obtain point estimates of total, efficiency-weighted, labor supply (human capital aggregates) by education and year.

With wage bills and human capital aggregates in hand, we can estimate the elasticity of substitution among labor inputs, using equations like (3.20), for the three relative wage bills. We use two different specifications: the first one is based on first-differences of equation (3.20), while the second is in levels. In both cases we control for possible endogeneity of human capital inputs in the production function through an IV approach. We experiment with different sets of instruments. First, we use lagged regressors (lags up to 5 periods back are included in the first step, depending on the specification). Alternatively, and as a robustness check, we also instrument using the total number of people in each education group in a given year, including those people not working. This latter instrument, being a stock, is independent of the serial correlation properties of the technology shock.

Table E1 reports results for both specifications (first-differences and levels) and both types of instruments. The estimation procedure is based on a stacking method which allows one to test for differences in the elasticity of substitution across different types of labor (like in a Chow test). Panel (A) reports the results using as instruments, respectively, lags (columns 1 to 4) or education 'stocks' (levels in column 5 and relative growth rates in column 6). Panel (B) reports tests of the null hypothesis of iso-elasticity for a set of specifications (more specifications and

[69] Polivka (2000) provides evidence that this method closely approximates the average of the top-coded tails by validating the fitted data through undisclosed and confidential non top-coded data available only at the BLS.

Panel (A): Estimation						
Specification	Growth rates		Levels			
	(1)	(2)	(3)	(4)	(5)	(6)
First stage IV	Up to 4 lags	Up to 3 lags	Up to 5 lags	Up to 4 lags	Edu. stock (L)	Edu. stock (G)
Number of obs.	75	78	75	78	90	87
	Coefficient	Coefficient	Coefficient	Coefficient	Coefficient	Coefficient
	(S.E.)	(S.E.)	(S.E.)	(S.E.)	(S.E.)	(S.E.)
$\rho^{HS,LH}$	.54	.14	.58	.47	-9.84	1.23
	(.18)	(.32)	(.20)	(.22)	(10.51)	(.37)
$\rho^{CL,HS}$	.58	.54	.50	.44	.26	.42
	(.35)	(.35)	(.11)	(.11)	(.64)	(.20)
$\rho^{CL,LH}$	.45	.39	.89	.90	.93	.49
	(.19)	(.26)	(.11)	(.11)	(.65)	(.16)
$\rho^{CL,HS,LH}$	.51	.35	.67	.64	.57	.48
	(.12)	(.17)	(.07)	(.07)	(.07)	(.03)

Panel (B): Hypothesis Testing				
Specification	Growth rates	Levels		
	(1)	(3)	(5)	(6)
First stage IV	Up to 4 lags	Up to 5 lags	Edu. stock (L)	Edu. stock (G)
Null Hypothesis	F-stat.	F-stat.	F-stat.	F-stat.
$\rho^{HS,LH} = \rho^{CL,HS}$	$F_{(1,69)} = .01$	$F_{(1,66)} = .12$	$F_{(1,81)} = .92$	$F_{(1,78)} = .32$
	$Pr. > F = .92$	$Pr. > F = .73$	$Pr. > F = .34$	$Pr. > F = .57$
$\rho^{CL,HS} = \rho^{CL,LH}$	$F_{(1,69)} = .10$	$F_{(1,66)} = 5.54$	$F_{(1,81)} = .54$	$F_{(1,78)} = .02$
	$Pr. > F = .75$	$Pr. > F = .02$	$Pr. > F = .47$	$Pr. > F = .88$
$\rho^{HS,LH} = \rho^{CL,LH}$	$F_{(1,69)} = .10$	$F_{(1,66)} = 1.63$	$F_{(1,81)} = 1.05$	$F_{(1,78)} = .27$
	$Pr. > F = .75$	$Pr. > F = .21$	$Pr. > F = .31$	$Pr. > F = .60$
$\rho^{CL,LH} = \rho^{CL,HS} = \rho^{HS,LH}$	$F_{(2,69)} = .08$	$F_{(2,66)} = 2.87$	$F_{(2,81)} = .76$	$F_{(2,78)} = .16$
	$Pr. > F = .93$	$Pr. > F = .07$	$Pr. > F = .47$	$Pr. > F = .84$

*Notes:*Panel (A): Estimates of ρ for various specifications. ρ^{e_1,e_2} denotes the parameter determining the elasticity of substitution between groups e_1 and e_2 estimated with the corresponding wage-bill ratio equation. $\rho^{CL,HS,LH}$ denotes the estimate from the restricted (iso-elastic) model. (L) and (G) in columns (5) and (6) indicate whether the education stock enters in Level or Growth rate in the estimated equation, respectively. (B): Tests for equality of elasticities of substitution among labor inputs. P-values are reported below the F-statistic.

Table E1: Production function estimates

tests are available from the authors). Overall, all specifications give remarkably similar results and we are unable to reject the null hypothesis that the aggregate technology is iso-elastic at 5% level of significance (see Table E1). The restricted model with a unique ρ improves the efficiency of the estimator, which is particularly valuable since we are using a relatively short time series (approximately 30 observations). Estimated shares of different human capital types in production (unreported) are also remarkably robust across specifications.

Appendix F

Cost of College Attendance, Grants, and Loans

To calculate the price of college attendance and the extent of government aid to higher education financing through grants we focus on the sample of full-time full year (FTFY) students enrolled in public and private not-for-profit 4-year post-secondary institutions. This group of students is the closest counterpart to students in the model. All our statistics refer to the year 2000 and nominal amounts are in 2000 dollars. According to the "Student Financing of Undergraduate Education: 1999-2000" (SFUE, thereafter), a report published by the National Center for Education Statistics (NCES), 65% of these students were enrolled in public colleges and 35% were enrolled in private not-for-profit colleges (Table 1.10).

Cost of College: The cost of college attendance has three components: (i) tuition and fees, (ii) non-tuition expenses that would not only be incurred by a college-student, and (iii) institutional and private grants which reduce the cost to families. The publication "Trends in College Pricing, 2000" published by the College Board, reports that average tuition and fees in public institutions in 2000-2001 were \$3,510 in public institutions and \$16,332 in private ones. We add non-tuition expenses, which includes books and other supplies, amounting to \$704 and \$730, respectively, in the two types of colleges. We also add an additional \$500 to account for any commuting or room and board expenses that would not be incurred by a worker. Average tuition and non-tuition expenses (before grants) amount to \$9,210. According to the SFUE, average tuition and fees did not differ by income level of the family in public institutions. In private institutions (where only 1/4 of students are enrolled), average fees were only roughly 20% lower for families whose income was between \$20,000-40,000 compared to fees faced by families whose income exceeded \$100,000 (Table 2.2-B).

Institutional and private grants are effectively a way to reduce the cost of attendance. Roughly half of these grants are based on pure merit and half are based on need. This fact, together with the negative empirical correlation between family need and students' merit, explains why both the fraction of students receiving grants and their amount is not strongly correlated with family income, as reported in Table F1 which is based on the SFUE, Table 1.2-G.

	Institutional Grants		Private Grants		Average Amount
Income	% receiving	Amount	% receiving	Amount	
<$30k	0.36	$4,077	0.12	$2,061	$1,715
$30k-80k	0.34	$5,474	0.16	$2,281	$2,234
>$80k	0.28	$5,383	0.14	$2,338	$1,855

*Notes:*Summary of institutional and private grants data used for the computation of the net tuition fees (NCES)

Table F1: Summary of institutional and private grants

To arrive at our estimate of average net tuition ($6,710) we subtract average private and institutional grants from average tuition expenses. Our measure of dispersion comes from Figure 12 of National Center for Education Statistics (2000), which provides the entire distribution of college costs. From this figure we determine that only 68% of college students pay less than the average amount of net tuition. To calibrate the standard deviation of tuition draws in the model we match this proportion. It turns out that a standard deviation of 1.0 captures this. Importantly, it is the ex-post distribution of college costs for which 68% of the mass is below $6,710, not the actual distribution of costs. Selection on realized costs causes the ex-post distribution of costs to have a smaller average than the actual distribution of draws. The mean of the ex-post distribution of costs is $6,710, but the mean ex-ante tuition draw is higher.

Federal and State Grants: Based on the "Guide to U.S. Department of Education Programs" (GDEP thereafter) published by the US Department of Education, we identify three main federal grant programs. The *Federal Pell Grant Program* is the largest single source of grants to undergraduates. It provides need-based grants to individuals to access post-secondary education. It is especially targeted to the lowest-income students. In 2000 it provided $7.3 billion to 3.8 million students, with a maximum grant of $3,125. The *Federal Supplemental Education Opportunity Grant* has a more modest endowment (approximately 15 times smaller). These are grants which supplement the amount received through Pell up to a maximum of $2,100. *Smart Grants* are awarded to needy student who are enrolled in certain technical fields and maintain a cumulative GPA of at least 3.0 in the first year – and so they're partly merit based. The program is approximately as big as the Supplemental Opportunity grant program. State funding is very diverse, but most of the funds available are concentrated in 10 "high-aid" states. Only a very small fraction of state grant awards are merit-based (less that 18%). The fraction of students receiving federal and state grants and their average amount by family income levels (from Table 1.2-G of the SFUE) is summarized in Table F2.

	Federal Grants		State Grants		Average Amount
Income	% receiving	Amount	% receiving	Amount	
<\$30k	0.72	\$2,753	0.38	\$826	\$2,820
\$30k-80k	0.14	\$1,579	0.21	\$455	\$668
<\$80k	0.01	\$1,605	0.07	\$133	\$143

*Notes:*Summary of federal and state grants by family income level (US Department of Education)

Table F2: Federal and state grants

We use the average amount for these three income levels, and the joint distribution of income and wealth in the model, to calibrate the dependence of the transfers' function $g(q, \theta)$ on assets (through the state variable q). In the baseline experiment, we do not allow g to depend on θ. However in one of our policy experiments we consider an expansion of merit-based grants.

Federal Loans: While grants are administered by both federal government and states, loans are almost entirely administered by the federal government (less than 1% of the total loan volume is state-based). The largest federal loan program in the US is the *Federal Family Education Loan Program.* The total volume of loans available in 2000 through this program was around \$40 billion, extended to around 10 million students. The program includes two main types of loans to students, *Subsidized* and *Unsubsidized Stafford Loans.* A third form of loan offered by the *Federal Family Education Loan Program* are *Parent PLUS loans.* These are loans made to the parents on behalf of a child to help pay for tuition by covering up to the cost of attendance less other aid. Eligibility for the PLUS Loan depends on a credit check and interest rates are similar to those in the private sector. Since this type of loan is equivalent to parents borrowing and then making a transfer to their child, we do not model them explicitly (Johnson, 2010, makes the same modelling choice). The other major source of financial aid for undergraduates, beyond the *Federal Family Education Loan Program,* is the *William D. Ford Federal Direct Loan Program.* This is, in essence, an alternative source of funding for Stafford loans whose total size is roughly half of that available through the *Federal Family Education Loan Program.* Finally, the *Federal Perkins Loan Program* provides low-interest loans to help needy students to finance undergraduate education whose conditions are similar to those of the subsidized Stafford loans. Its total funding is small though, roughly 3% of Stafford loans. Because of their nature, we aggregate these loans with subsidized Stafford loans in our calculations. In light of this discussion, in calibrating the features of the Federal loan program, we focus on (subsidized and unsubsidized) Stafford loans only.

Subsidized Stafford Loans are loans to students who meet a financial needs test (based on family income and assets), with the interest paid by the government on behalf of borrowers while the student is in school. Interest payments after school are subsidized. In 2000, the total cumulative borrowing limit for subsidized loans over the four years of college was $17,125.

Unsubsidized Stafford Loans are loans available to students who either do not meet a financial needs test or do qualify, but need to supplement their subsidized loans. The interest on the unsubsidized Stafford loan cumulates when in school, it is added to the principal, and the student starts repaying her debt after graduation. In 2000, the cumulative unsubsidized Stafford loans limit over the four years of college was $18,000.

Since it is largely up to the institution and the federal government to determine the maximum subsidized amount that each student can borrow, in the calibration we use $\underline{b}^s$ to match some moments of the cross-sectional distribution of loans (see below), and we make sure it does not exceed $17,125. We do fix the total cumulative (subsidized and unsubsidized) Stafford debt limit $\underline{b}^s+\underline{b}^u$ to $35,125. Repayment plans for Stafford loans typically impose fixed monthly amount for a loan term of up to 10 years. But extended repayment periods can be obtained. In the model, we set a fixed repayment plan with duration $n = 20$ years.

According to the SFUE, in the year 2000, 44.9% of students in 4-year institutions had (subsidized or unsubsidized) Stafford loans, 37.3% had subsidized Stafford or Perkins loans, and 21.2% had unsubsidized Stafford loans (Tables 1.5A, 1.6-A). Moreover, among borrowers, the average cumulated amount of student loans at graduation was $17,016 (Table 1.1-A). Finally, the College Board (1998) reports that the ratio of total volume of subsidized to unsubsidized federal loans is 1.36.

We have six parameters related to the federal loan program $(a^*, a^{**}, \underline{b}^s, r^s, r^u, \underline{b}^s, \underline{b}^u)$ that we use to target six moments: (i) 37.3% of students have subsidized Stafford loans, (ii) 21.2% of students have unsubsidized Stafford loans, (iii) 44.9% of students have any Stafford loans, (iv) the average cumulated amount of federal loans at graduation is $17,016, (v) the ratio between the total volume of subsidized and unsubsidized loans is 1.36, (vi) the maximum cumulative amount of Stafford loans is $35,125.

Private Loans: The report "Private Loans and Choice in Financing Higher Education" published by the Institute for Higher Education Policy (2003) contains useful information on private borrowing with the purpose of funding post-secondary education. Available estimates suggest that private loans at that time composed only 12 percent of the total volume of Federal loans (page 9). For many student borrowers, a poor credit rating often is the largest barrier to obtaining a private loan. Less than 1% of private loan products were credit-blind, or available without a credit check (page 15). However, for those who qualify, interest rates on private loans are often more advantageous of those on Stafford Loans (Figure 2.2). In 2000, 4.9% of students enrolled in 4-year institutions received private loans (Figure 4.2), and the average amount received was $4,767 (Table A.4).

We calibrate two parameters related to private borrowing: the wedge ι, and the wealth threshold $\underline{a}^p$. To capture some general features of access to private loans we target two moments: (i) 4.9% of students have private loans, and (ii) the ratio of the volume of private to federal loans is 12%.

Appendix G

Policy Experiments

This Appendix reports additional outcomes of all the policy experiments on government financial aid to college students described in Section 3.5 in the main text.

Removal of Grants

		Benchmark	Treatment	P.E. Short-run	P.E. Long-run	G.E. Long-run
Policy Cost		n/a	-100%	-100%	-100%	-100%
College	Ability Tercile 1	0.038	0.035	0.034	0.037	0.047
Attainment	Ability Tercile 2	0.178	0.112	0.113	0.132	0.159
Rates	Ability Tercile 3	0.568	0.476	0.48	0.500	0.541
	Par. Wealth Tercile 1	0.197	0.128	0.127	0.133	0.166
	Par. Wealth Tercile 2	0.191	0.120	0.124	0.138	0.164
	Par. Wealth Tercile 3	0.396	0.375	0.375	0.397	0.417
	Aggregate	0.261	0.208	0.209	0.223	0.249
Labor Tax Rate		0.270	0.270	0.270	0.270	0.274
Price of CL/HS Labor		1.335	–	–	–	1.364
Log of College Premium		0.62	n/a	n/a	+3.1%	+3.4%
Student Labor Supply		n/a	+3.5%	+3.1%	+5.4%	+4.5%
Aggregate Output		n/a	n/a	n/a	-1.9%	-1.2%
Interest Rate		0.0438	0.0438	0.0438	0.0438	0.0439
Intergenerational Correlation of Education		0.308	–	–	0.327	0.326
Welfare (CEV)		n/a	n/a	n/a	n/a	-0.384%
Average Inter Vivos ($)		31,256	31,256	30,967	29,585	31,202
Crowding-Out	Aggregate	n/a	n/a	+32%	-34%	+17%
	Par. Wealth Tercile 1	n/a	n/a	11%	-11%	1%
	Par. Wealth Tercile 2	n/a	n/a	37%	-33%	-8%
	Par. Wealth Tercile 3	n/a	n/a	65%	-76%	125%
	Ability Tercile 1	n/a	n/a	127%	-136%	118%
	Ability Tercile 2	n/a	n/a	62%	-39%	39%
	Ability Tercile 3	n/a	n/a	28%	-31%	17%

*Notes:*Response to the elimination of all government grants. 'Treatment' identifies the direct immediate effect of the policy on the child's education decision. 'PE Short-run' incorporates the short-run response of parental IVTs. 'PE long-run' is the new steady state, with the new ergodic distribution of parental IVTs, but where prices and taxes are fixed at their initial values. 'GE Long-run' is the new steady-state where prices clear input markets and tax rates balance the government budget constraint. "%" denote percentage changes from the benchmark value. The welfare gain is expressed in terms of changes in expected lifetime consumption for a newborn agent. 'Crowding Out' of IVT is defined as the average change in IVTs as a percentage of the average change in the grant (in absolute value) for all the individuals who enroll in college in the experiment.

Table G1: Removal of grants experiment

General Tuition Subsidy

		Benchmark	Treatment	P.E. Short-run	P.E. Long-run	G.E. Long-run
Policy Cost		n/a	90%	87%	88%	67%
College	Ability Tercile 1	0.038	0.052	0.051	0.053	0.040
Attainment	Ability Tercile 2	0.178	0.229	0.230	0.236	0.184
Rates	Ability Tercile 3	0.568	0.601	0.589	0.596	0.577
	Par. Wealth Tercile 1	0.197	0.230	0.230	0.230	0.201
	Par. Wealth Tercile 2	0.191	0.223	0.217	0.221	0.194
	Par. Wealth Tercile 3	0.396	0.430	0.423	0.434	0.405
	Aggregate	0.261	0.294	0.290	0.295	0.267
Labor Tax Rate		0.270	0.270	0.270	0.270	0.271
Price of CL/HS Labor		1.335	–	–	–	1.319
Log of College Premium		0.62	-2.1%	-2.9%	-2.1%	-2.3%
Student Labor Supply		n/a	-4.2%	-0.4%	-4.8%	-4.3%
Aggregate Output		n/a	–	–	1.5%	0.4%
Interest Rate		0.0438	0.0438	0.0438	0.0438	0.0439
Intergenerational Correlation of Education		0.308	–	–	0.304	.305
Welfare (CEV)		n/a	n/a	n/a	n/a	0.25%
Average Inter Vivos ($)		31,256	31,256	31,127	32,332	30,930
Crowding-Out	Aggregate	n/a	n/a	-30%	+30%	-17%
	Par. Wealth Tercile 1	n/a	n/a	+1%	+8%	+2%
	Par. Wealth Tercile 2	n/a	n/a	-41%	+14%	-2%
	Par. Wealth Tercile 3	n/a	n/a	-43%	+57%	-31%
	Ability Tercile 1	n/a	n/a	-17%	+61%	-8%
	Ability Tercile 2	n/a	n/a	-21%	+22%	-19%
	Ability Tercile 3	n/a	n/a	-36%	+32%	-14%

*Notes:*Responses to an additional $1,000 per-year tuition subsidy. 'Treatment' identifies the direct immediate effect of the policy on the child's education decision. 'PE Short-run' incorporates the short-run response of parental IVTs. 'PE long-run' is the new steady state, with the new ergodic distribution of parental IVTs, but where prices and taxes are fixed at their initial values. 'GE Long-run' is the new steady-state where prices clear input markets and tax rates balance the government budget constraint. "%" denote percentage changes from the benchmark value. The welfare gain is expressed in terms of changes in expected lifetime consumption for a newborn agent. 'Crowding Out' of IVT is defined as the average change in IVTs as a percentage of the average change in the grant (in absolute value) for all the individuals who enroll in college in the experiment.

Table G2: $1000 Tuition subsidy experiment

Means-Tested Grants

		Benchmark	Treatment	P.E. Short-run	P.E. Long-run	G.E. Long-run
Policy Cost		n/a	91%	87%	87%	66%
College Attainment Rates	Ability Tercile 1	0.038	0.05	0.049	0.050	0.035
	Ability Tercile 2	0.178	0.240	0.241	0.242	0.200
	Ability Tercile 3	0.568	0.604	0.592	0.599	0.585
	Par. Wealth Tercile 1	0.197	0.244	0.243	0.242	0.216
	Par. Wealth Tercile 2	0.191	0.236	0.230	0.233	0.208
	Par. Wealth Tercile 3	0.396	0.414	0.409	0.415	0.396
	Aggregate	0.261	0.298	0.294	0.297	0.273
Labor Tax Rate		0.270	0.270	0.270	0.270	0.269
Price of CL/HS Labor		1.335	–	–	–	1.313
Log of College Premium		0.62	-2.7%	-3.1%	-3.1%	-2.7%
Student Labor Supply		n/a	+0.3%	-2.8%	-2.8%	-2.1%
Aggregate Output		n/a	1.6%	1.3%	1.5%	0.7%
Interest Rate		0.0438	0.0438	0.0438	0.0438	0.0439
Intergenerational Correlation of Education		0.308	–	–	0.304	0.304
Welfare (CEV)		n/a	n/a	n/a	–	0.23%
Average Inter Vivos ($)		31,256	31,256	31,167	31,893	30,790
Crowding-Out	Aggregate	n/a	n/a	-28%	+19%	-25%
	Par. Wealth Tercile 1	n/a	n/a	-0%	+4%	+4%
	Par. Wealth Tercile 2	n/a	n/a	-29%	+11%	-1%
	Par. Wealth Tercile 3	n/a	n/a	-97%	+89%	-151%
	Ability Tercile 1	n/a	n/a	-27%	+27%	-211%
	Ability Tercile 2	n/a	n/a	-23%	+17%	-31%
	Ability Tercile 3	n/a	n/a	-30%	+20%	-13%

*Notes:*Responses to a 56% increase in tuition subsidy for every student (i.e., extra $1,580, $374 and $81 per year for students of type $q = 1, 2, 3$, respectively). The cost of this policy matches that of the $1,000 general grant expansion in the 'PE Short-run'. 'Treatment' identifies the direct immediate effect of the policy on the child's education decision. 'PE Short-run' incorporates the short-run response of parental IVTs. 'PE long-run' is the new steady state, with the new ergodic distribution of parental IVTs, but where prices and taxes are fixed at their initial values. 'GE Long-run' is the new steady-state where prices clear input markets and tax rates balance the government budget constraint. "%" denote percentage changes from the benchmark value. The welfare gain is expressed in terms of changes in expected lifetime consumption for a newborn agent. 'Crowding Out' of IVT is defined as the average change in IVTs as a percentage of the average change in the grant (in absolute value) for all the individuals who enroll in college in the experiment.

Table G3: Means-tested grant expansion

Ability-Tested Grants

		Benchmark	Treatment	P.E. Short-run	P.E. Long-run	G.E. Long-run
Policy Cost		n/a	92%	87%	82%	68%
College	Ability Tercile 1	0.038	0.044	0.046	0.044	0.041
Attainment	Ability Tercile 2	0.178	0.233	0.227	0.216	0.187
Rates	Ability Tercile 3	0.568	0.604	0.590	0.598	0.577
	Par. Wealth Tercile 1	0.197	0.239	0.230	0.222	0.203
	Par. Wealth Tercile 2	0.191	0.231	0.216	0.215	0.194
	Par. Wealth Tercile 3	0.396	0.412	0.417	0.422	0.407
	Aggregate	0.261	0.294	0.287	0.286	0.268
Labor Tax Rate		0.270	0.270	0.270	0.270	0.271
Price of CL/HS Labor		1.335	–	–	–	1.317
Log of College Premium		0.62	-2.1%	-2.4%	-1.1%	-1.8%
Student Labor Supply		n/a	-4.7%	-0.6%	-4.9%	-3.6%
Aggregate Output		n/a	1.5%	1.0%	1.1%	0.3%
Interest Rate		0.0438	0.0438	0.0438	0.0438	0.0440
Intergenerational Correlation of Education		0.308	–	–	0.309	0.307
Welfare (CEV)		n/a	n/a	n/a	–	0.14%
Average Inter Vivos ($)		31,256	31,256	31,086	31,925	30,636
Crowding-Out	Aggregate	n/a	n/a	-33%	23%	-24%
	Par. Wealth Tercile 1	n/a	n/a	+1%	5%	-1%
	Par. Wealth Tercile 2	n/a	n/a	-42%	12%	-8%
	Par. Wealth Tercile 3	n/a	n/a	-51%	42%	-46%
	Ability Tercile 1	n/a	n/a	-35%	59%	-61%
	Ability Tercile 2	n/a	n/a	-25%	22%	-24%
	Ability Tercile 3	n/a	n/a	-36%	21%	-22%

*Notes:*Responses to an increase in tuition subsidies equal to 0.054*exp(θ), e.g., extra $723, $979 and $1,098 per year for students in the bottom, middle, and top ability tercile, respectively). The cost of this policy matches that of the $1,000 general grant expansion in the 'PE Short-run'. 'Treatment' identifies the direct immediate effect of the policy on the child's education decision. 'PE Short-run' incorporates the short-run response of parental IVTs. 'PE long-run' is the new steady state, with the new ergodic distribution of parental IVTs, but where prices and taxes are fixed at their initial values. 'GE Long-run' is the new steady-state where prices clear input markets and tax rates balance the government budget constraint. "%" denote percentage changes from the benchmark value. The welfare gain is expressed in terms of changes in expected lifetime consumption for a newborn agent. 'Crowding Out' of IVT is defined as the average change in IVTs as a percentage of the average change in the grant (in absolute value) for all the individuals who enroll in college in the experiment.

Table G4: Ability tested grants expansion

Removal of Government-Sponsored Loans

		Benchmark	Treatment	P.E. Short-run	P.E. Long-run	G.E. Long-run
College	Ability Tercile 1	0.038	0.017	0.016	0.035	0.047
Attainment	Ability Tercile 2	0.178	0.077	0.050	0.184	0.218
Rates	Ability Tercile 3	0.568	0.217	0.274	0.490	0.515
	Par. Wealth Tercile 1	0.197	0.003	0.032	0.160	0.187
	Par. Wealth Tercile 2	0.191	0.014	0.042	0.172	0.198
	Par. Wealth Tercile 3	0.396	0.292	0.266	0.377	0.395
	Aggregate	0.261	0.103	0.113	0.237	0.258
Labor Tax Rate		0.270	0.270	0.270	0.270	0.277
Price of CL/HS Labor		1.335	–	–	–	1.353
Log of College Premium		0.62	n/a	n/a	-0.9%	-1.6%
Student Labor Supply		n/a	-19%	+3.1%	+33.1%	+32.3%
Aggregate Output		n/a	n/a	n/a	-2.4%	-1.5%
Interest Rate		0.4038	0.0438	0.0438	0.0438	0.0444
Intergenerational Correlation of Education		0.308	–	–	0.326	0.323
Welfare (CEV)		n/a	n/a	n/a	n/a	-0.377%
Average Inter Vivos ($)		31,256	31,256	29,618	25,647	27,265
Crowding-Out	Aggregate	n/a	n/a	13%	-39%	-29%
	Par. Wealth Tercile 1	n/a	n/a	5%	-0.1%	1.5%
	Par. Wealth Tercile 2	n/a	n/a	15%	0.2%	6.3%
	Par. Wealth Tercile 3	n/a	n/a	14%	-79%	-63%
	Ability Tercile 1	n/a	n/a	17%	-18%	-7.8%
	Ability Tercile 2	n/a	n/a	30%	-24%	-17%
	Ability Tercile 3	n/a	n/a	10%	-51%	-37%

*Notes:*Elimination of all government-sponsored loans. 'Treatment' identifies the direct immediate effect of the policy on the child's education decision. 'PE Short-run' incorporates the short-run response of parental IVTs. 'PE long-run' is the new steady state, with the new ergodic distribution of parental IVTs, but where prices and taxes are fixed at their initial values. 'GE Long-run' is the new steady-state where prices clear input markets and tax rates balance the government budget constraint. "%" denote percentage changes from the benchmark value. The welfare gain is expressed in terms of changes in expected lifetime consumption for a newborn agent. 'Crowding Out' of IVT is defined as the average change in IVTs as a percentage of the change in the borrowing limit (in absolute value), i.e., $35,497, for all the individuals who enroll in college in the experiment.

Table G5: Removal of government loans experiment

Expansion of Government-Sponsored Loan Limit

		Benchmark	Treatment	P.E. Short-run	P.E. Long-run	G.E. Long-run
College Attainment Rates	Ability Tercile 1	0.038	0.040	0.038	0.037	0.037
	Ability Tercile 2	0.178	0.214	0.210	0.179	0.178
	Ability Tercile 3	0.568	0.613	0.602	0.574	0.573
	Par. Wealth Tercile 1	0.197	0.228	0.228	0.200	0.199
	Par. Wealth Tercile 2	0.191	0.220	0.213	0.187	0.191
	Par. Wealth Tercile 3	0.396	0.418	0.408	0.403	0.399
	Aggregate	0.261	0.289	0.281	0.263	0.263
Labor Tax Rate		0.270	0.270	0.270	0.270	0.270
Price of CL/HS Labor		1.335	–	–	–	1.334
Log of College Premium		0.62	n/a	n/a	-0.1%	-0.1%
Student Labor Supply		n/a	-24.2%	-21.1%	-2.0%	-2.4%
Aggregate Output		n/a	n/a	n/a	+0.2%	+0.1%
Interest Rate		0.0438%	0.0438%	0.0438%	0.0438%	0.0439%
Intergenerational Correlation of Education		0.308	–	–	0.302	0.304
Welfare (CEV)		n/a	n/a	n/a	n/a	0.002%
Average Inter Vivos ($)		31,256	31,256	30,469	31,335	31,220
Crowding-Out	Aggregate	n/a	n/a	-29%	-1.7%	0.5%
	Par. Wealth Tercile 1	n/a	n/a	-3.3%	0.6%	0.1%
	Par. Wealth Tercile 2	n/a	n/a	-37%	-3.8%	0.6%
	Par. Wealth Tercile 3	n/a	n/a	-39%	-1.9%	0.7%
	Ability Tercile 1	n/a	n/a	-19%	-30%	2.4%
	Ability Tercile 2	n/a	n/a	-17%	2.2%	-0.2%
	Ability Tercile 3	n/a	n/a	-34%	-1.1%	0.6%

*Notes:*Responses to a $10,000 expansion of the unsubsidized government sponsored loan cumulative limit. 'Treatment' identifies the direct immediate effect of the policy on the child's education decision. 'PE Short-run' incorporates the short-run response of parental IVTs. 'PE long-run' is the new steady state, with the new ergodic distribution of parental IVTs, but where prices and taxes are fixed at their initial values. 'GE Long-run' is the new steady-state where prices clear input markets and tax rates balance the government budget constraint. "%" denote percentage changes from the benchmark value. The welfare gain is expressed in terms of changes in expected lifetime consumption for a newborn agent. 'Crowding Out' of IVT is defined as the average change in IVTs as a percentage of the change in the borrowing limit (in absolute value), i.e., $10,000, for all the individuals who enroll in college in the experiment.

Table G6: Loan limits expansion experiment

Appendix H

Computational Algorithm

This appendix describes the solution method for our long-run GE economy. The usual nested fixed point approach is extended in order to accommodate the novel features of our model. That is, the essence of our approach is to guess a set of prices and taxes, compute decision rules (given prices and taxes) to simulate the economy, and finally verify whether those are the equilibrium prices and taxes. To accommodate endogenous inter vivos transfers we must also begin with guesses of the decision rules of age zero agents and the initial distribution of wealth.

Specifically, we execute the following steps:

1. Make an initial guess for the wage vector, $\tilde{\mathbf{w}}$, and the real interest rate, $\tilde{r}$. Also make an initial guess for the age zero consumption decision rule, $\tilde{c}_0$, and the initial distribution of wealth $\tilde{\mathbf{a}}_0$. In the policy experiments, an initial guess for the labor tax rate is also required.

2. Solve the household dynamic programming problem described in Section 2.6 at the prices $\tilde{\mathbf{w}}$ and $\tilde{r}$. This is a finite horizon problem easily solved by backward induction using Euler equation methods. At the age inter vivos transfers are given, the intergenerational Euler equation requires the optimal consumption decision of the age 0 child. The guess $\tilde{c}_0$ is used here. The solution yields optimal decision rules for education, take-up of student loans, consumption, leisure, private saving/borrowing, and inter vivos transfers.

3. Simulate the life-cycles of 10,000 agents who start with initial wealths given by $\tilde{\mathbf{a}}_0$. Each of the 10,000 simulated agents is exogenously matched with another agent who represents her child. The abilities of the parents and children in these matches are consistent with the transition matrix for ability. Importantly, these matches are fixed across iterations so that the inter vivos transfer given by the parent in the match converges to the initial wealth of the child in the match.

4. This step consists of four sub-steps:

   (i) Aggregate the decisions of the 10,000 simulated agents to check market clearing conditions and update prices appropriately.

(ii) Compare simulated inter vivos transfers to $\tilde{\mathbf{a}}_0$ and update appropriately.

(iii) Compare the age zero consumption rule to $\tilde{c}_0$ and update appropriately.

(iv) If computing the benchmark economy, adjust residual government expenditure G to solve the government budget constraint. If computing a policy experiment, update the labor tax rate appropriately if the government budget constraint is not satisfied.

5. If updates were required in any of sub-steps (i)-(iii) of step (4) (i-iv for an experiment) return to step (2) and proceed with the updated guesses. Otherwise, exit because a fixed point of the algorithm has been achieved.

Once the fixed point has been attained simulated data from the economy can be used to compute the various moments, tables, and figures of interest.