

Essays in Macroeconomics

by

Adam Bergeron

Department of Economics
Duke University

Date: _____

Approved:

A. Craig Burnside, Supervisor

Andrea Lanteri

Kyle Jurado

Michelle Connolly

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2018

ABSTRACT

Essays in Macroeconomics

by

Adam Bergeron

Department of Economics
Duke University

Date: _____

Approved:

A. Craig Burnside, Supervisor

Andrea Lanteri

Kyle Jurado

Michelle Connolly

An abstract of a dissertation submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2018

Copyright © 2018 by Adam Bergeron
All rights reserved except the rights granted by the
Creative Commons Attribution-Noncommercial Licence

Abstract

This paper consists of three essays, with a focus on household behavior and decision making. In the first essay, I consider a two asset lifecycle model with transaction costs and show that allowing individuals to learn about their idiosyncratic skill (and therefore their future income distribution) over time yields starkly different asset allocation and overall savings decisions by households as compared to a more restricted income process typically considered in the literature. Using parameters estimated from the same underlying income data as the standard process, I show that a model with learning generates more liquid saving over the entire lifetime, suppresses illiquid saving early on, and increases illiquid saving later in the lifecycle.

In the second essay, I examine household financial data from the Survey of Consumer Finances. I show that a model with learning is more easily able to generate sufficient liquid assets to match the ratio of average holdings of illiquid to liquid assets over the lifetime found in the data. I further show that a model without learning requires a risk aversion parameter nearly six times that of the learning model to generate the same ratio. Finally, I show that for a select group of individuals, the marginally propensity to consume out of extra income in the learning model is only 56% of that of the model without learning. This suggests that introducing learning into the household portfolio problem potentially has large policy implications.

In the third essay, I consider a city-level environment in which individuals search simultaneously for both housing and employment. I construct a model in which

the unemployment rate, per capita income, population growth, home prices, home construction rates, and rent are all endogenously determined. I then compare the response of the variables from a labor demand shock to their empirical counterparts. The model has some success qualitatively matching impulse responses from the data.

To Jimmy, a good boy

Contents

Abstract	iv
List of Tables	x
List of Figures	xi
List of Abbreviations and Symbols	xii
Acknowledgements	xiii
1 Introduction	1
2 A Model of Learning and Portfolio Choice	9
2.1 Model	9
2.1.1 Income Process and Learning	9
2.1.2 Household Problem	13
2.1.3 Solving the Model	15
2.1.4 Calibration	18
2.2 Intuition	20
2.2.1 Uncertainty	20
2.2.2 Forecasts	23
2.3 Results	29
2.3.1 Set-Up	29
3 The Ramifications of Incorporating Learning Into A Portfolio Choice Model	36

3.1	Household Portfolio Data	36
3.2	Comparison to the Data	43
3.3	Implications	47
4	City-Level Housing and Employment Dynamics	51
4.1	The Model	51
4.1.1	Households	51
4.1.2	Firms	56
4.1.3	Housing	57
4.2	Equilibrium	58
4.2.1	Bargaining	59
4.2.2	Summarizing the Equilibrium	61
4.3	Calibration	63
4.3.1	Functional Forms	63
4.3.2	Parameter Values	64
4.4	Empirical Model	66
4.5	Results	68
4.5.1	Empirical	68
4.5.2	Theoretic	70
5	Conclusion	74
A	Appendix for Chapters 2 & 3	76
A.1	Benchmark Model and Computation	76
A.2	Data	77
A.2.1	Survey of Consumer Finances (SCF)	77
A.3	Derivations	79
A.3.1	Lifetime Income Distributions	79

A.3.2	Kalman Filter	80
A.3.3	Mean Square Error	80
B	Appendix for Chapter 4	82
B.1	Data	82
B.2	Population Laws of Motion	83
B.3	Normalizing Variables	85
B.4	Optimal Maintenance Problem	86
	Bibliography	88
	Biography	91

List of Tables

2.1	Income Process Calibration	19
3.1	Average Share of Liquid Assets	38
3.2	Average Share of Illiquid Assets	38
3.3	% of Households With Minimum Share of Each Liquid Asset	40
3.4	% of Households With Minimum Share of Each Illiquid Asset	41
4.1	Parameter Values	65

List of Figures

2.1	Present Discounted Value of Forecast Errors	22
2.2	Income Expectations	25
2.3	Comparison of Expectations at Different Income Levels	28
2.4	Average Savings Over the Lifecycle	30
3.1	Assets Over the Lifecycle	42
3.2	Ratio of Illiquid to Liquid Savings	44
3.3	Median Ratio Over the Lifecycle	46
4.1	Movement Between Nodes	52
4.2	Data: Impulse Response Functions	69
4.3	Theoretic Impulse Responses	71
A.1	Median Illiquid-Liquid Ratio For Select Years	78
A.2	Distribution of Income Over the Lifecycle	78

List of Abbreviations and Symbols

AR	Autoregressive
BEA	Bureau of Economic Analysis
CD	Certificate of Deposit
CPI	Consumer Price Index
CRRA	Constant Relative Risk Aversion
FHFA	Federal Housing Finance Agency
FMR	Fair Market Rents
GMM	Generalized Method of Moments
HPI	Home Price Index
HUD	(Department of) Housing and Urban Development
IRA	Individual Retirement Account
IRF	Impulse Response Function
MSA	Metropolitan Statistical Area
MPC	Marginal Propensity to Consume
PSID	Panel Survey of Income Dynamics
SCF	Survey of Consumer Finances
VAR	Vector Auto-Regression

Acknowledgements

I am grateful to Craig Burnside, Andrea Lanteri, Kyle Jurado, and Michelle Connolly for their invaluable advice and continual support. I also thank Pietro Peretto, Cosmin Ilut, Matthias Kehrig, Francesco Bianchi, Nir Jaimovich, Curtis Taylor, Vasco Botelho, Soroush Ghazi, Linxi Chen, Marat Kussainov, Hasan Arik, and the participants of the Duke Macroeconomics Workshop and the Trade Dynamics Macroeconomics Workshop Seminar series for their helpful comments and discussions.

1

Introduction

This paper is comprised of three essays, with a focus on household behavior and decision making. The first two essays are concerned with household portfolio decisions over the lifetime. The household portfolio allocation problem is not trivial. Household portfolios consist of a diverse array of assets, differing in many dimensions such as risk, maturity, and liquidity. I focus on the decision the household faces between liquid and illiquid assets. The third essay shifts gears and considers the dynamics of housing and labor markets at the city level. Households are once again at the forefront of the analysis, being crucial participants in both markets. This essay examines the relationship between labor demand, income, population growth, and housing in the context of a city where individuals can move to and from the rest of the world.

Recent work by Kaplan and Violante (2014) has shown that the liquidity margin is essential for understanding the efficacy of fiscal policy. Individuals who are liquidity constrained desire, but are unable to finance, more consumption. These individuals will spend more out of a one off transfer payment than other agents. However, solely considering a household's overall assets belies just how many households may be constrained. Many households, though rich in wealth, lack liquidity as

their net worth is tied up in difficult to adjust assets such as housing or retirement accounts. By decomposing household portfolios and considering an alternative measure of constrained households using only liquid assets, the authors are able to provide an explanation for the substantial consumption response to tax rebates documented in Johnson et al. (2006). Thus, to understand the impact of fiscal policy, it is important to understand how and why individuals make decisions on the liquid-illiquid margin.

I consider a two asset model and show that the choice of exogenous income process is fundamental in household portfolio decisions. The functional form typically assumed¹ in the literature is one in which income draws are a result of a very persistent shock, a transitory shock, and a deterministic trend component. Some papers, such as Kaplan and Violante (2014), allow for an idiosyncratic level component that is drawn at birth, fully known by the individual, and constant throughout the lifetime. This specification has the benefit of being highly tractable, as well as having some support in the data.²

I consider an alternative specification, one in which individuals have both an idiosyncratic level and growth component, in addition to facing persistent and transitory shocks. These idiosyncratic components are drawn at birth and are constant throughout the lifetime, and can be thought of as an individual's innate "skill." This interpretation is both intuitively appealing and supported empirically; estimations using income panel data in Baker (1997) and Guvenen (2009) lend credence to this formulation. As in Guvenen (2007), I assume that skill is not directly observable by individuals and that they must learn about their earning potential over time. Individuals update their beliefs according to received income draws.

I incorporate this learning mechanism into an otherwise standard portfolio choice

¹ See Cocco (2004), Scholz et al. (2006), Polkovnichenko (2007), Lynch and Tan (2011), among many others

² See MaCurdy (1982), Abowd and Card (1989), Hryshko (2012)

model. Individuals receive exogenous income draws throughout their working life and must save for consumption smoothing and/or precautionary reasons. Individuals have two accounts in which to save: a liquid account which can be adjusted freely and an illiquid account which bears a cost upon adjustment. To compensate for its lack of liquidity, the illiquid asset yields a higher return. I compare the portfolio decision implications of this model to one using the standard income setup, with no individual growth and the composition of all draws fully known by the individual.

In simulating the asset allocation decisions of individuals, I use parameters for each process that are estimated off of the same income data. Using a stylized calibration in which illiquid assets are fully irreversible until retirement, I show that despite the underlying income data being identical, the choice of income specification implies sharp differences in overall saving and portfolio allocation over the lifecycle. Overall saving is higher in the learning model throughout the lifetime. Further, individuals in the learning model shift their illiquid savings to later in life and begin decumulating liquid savings earlier in life as compared to their benchmark counterparts. Average liquid asset holdings are higher throughout the entire working lifecycle in the learning model, as well. I argue that the benchmark model generates too little liquid holdings as compared to illiquid holdings. Briefly comparing these findings to data from 2013, I show that the learning model's overall ratio of average illiquid assets over the lifecycle to average liquid assets over the lifecycle is more in line with the ratio in the data. Moreover, if one considers the median individual's ratio of illiquid to liquid assets throughout their lifetime, the benchmark ratio is always higher than in the data. The learning model does better at matching this ratio at the beginning of the lifecycle, though the model does underpredict this ratio in the middle of the median individual's life.

These results are driven by the way the learning model affects lifetime income uncertainty and the mean forecasts of agents. Lifetime uncertainty is higher when

individuals must learn their skill level, particularly early in the life. Further, due to the negative population covariance between the level and growth components of income, prosperous individuals in the learning model are highly pessimistic early in their life. Later in life, individuals who receive high income eventually become more optimistic than their counterparts in the standard model. These two factors contribute to the different savings patterns in the two models.

Thus, the model with heterogeneous skill and learning has potentially different policy implications than standard models. The increase in liquid holdings in the learning model implies that individuals will be less liquidity constrained, potentially diminishing any stimulative fiscal policy. If one is calibrating the model to match the number of constrained households, as is done in Kaplan and Violante (2014), then other parameters may have to adjust, such as the coefficient of risk aversion, which may then in turn affect the response to policy. The message of this paper extends to other applications as well. Housing is an obvious case, since it dominates household balance sheets. Iacoviello and Pavan (2013) create a general equilibrium lifecycle model incorporating housing. They consider the effects of changes in income risk and down payments on various macroeconomic variables and model the effect of countercyclical financial conditions on the housing market and household balance sheets. Fisher and Gervais (2011) use a life cycle model to understand why home ownership has fallen among young individuals, and conclude that an increase in earnings risk and decrease in marriage accounts for the trend. Yang (2009) uses borrowing constraints and transaction costs to model household consumption over the lifecycle. These papers incorporate the standard model of income risk. If one considers the illiquid asset to consist of just housing, then this paper implies that there may be additional insights to be gleaned from the exercises considered in the literature.

This paper also fits into the literature seeking to understand how incorporating an

income process with learning into a lifecycle model affects household decision making. Guvenen (2007) shows that in a one asset environment, allowing households to learn their skill level over time has success in matching the inequality of consumption over the lifecycle. Guvenen and Smith (2014) use indirect inference to estimate a one asset model of consumption and savings where individuals learn their idiosyncratic growth rate. They find that the persistent component mean reverts much faster than commonly believed and that individuals know around half of their skill level immediately upon entering the work force. As in this paper, Chang et al. (2016) consider a two asset model. However, their margin of interest is the one between “safe” assets and riskier assets with a stochastic return. They combine learning with unemployment risk and occupational switching to produce a model that can reproduce the documented finding that the risky share of an individual’s portfolio increases with age. Though similar in spirit to this paper, their margin has different policy implications compared to the liquid-illiquid margin. Finally, to the extent that this paper provides some evidence for the learning process being an adequate model for another aspect of household behavior, this paper adds to the debate on the appropriate way to think about income.

In Chapter 2 I introduce the model of household lifetime portfolio choice in a context where individual’s can learn about their earning potential over time. I discuss relevant mechanisms and examine the results of the model. In Chapter 3 I relate this model to available data on household finances and discuss the broader theoretical and policy implications of the results of this model.

In Chapter 4 I introduce a model of joint housing and employment search, in the context of a single city where individuals move to or from the rest of the world. Individuals within the city can either rent or own a home; however prospective homeowners must search to find a good match. The probability of a match is driven by the number of other home searchers and vacant homes available for sale; in other

words, the “tightness” of the market. This tightness in turn affects the value of putting a house on the market. This affects the value of building a new housing unit, and construction companies respond to shifting valuations by increasing or cutting down on their construction activity. At the same time, individuals also search for employment. There is a “tightness” to this market as well, driven by the number of job seekers and vacancies. Individuals can match with a firm and earn an endogenous wage that is bargained for by the individual and the firm. In this way individuals can be at various states within the city, as they can be any combination of homeowner or renter and employed or unemployed. The value of moving to the city is endogenous, as well. Population growth reflects changes in the desirability of the city.

I use this model to examine the housing and employment dynamics at the city-level, focusing on the unemployment rate, per capita income, population growth, home prices, housing construction, and rents. I compare the impulse responses from a shock to labor demand in the model to those from a panel VAR using city level data. Though this shock is a bit different in the model and in the data— in the data this is taken to be a shock to the unemployment rate, while in the model this is a shock to matching efficiency in the labor market— care was taken to make the scenarios as similar as possible in both cases. The model has some success in replicating the empirical impulse responses, at least qualitatively.

Though perhaps not as extensive as the literature on labor markets alone, housing search and matching has been the subject of a fair amount of research. The first to employ this method was Wheaton (1990), who was able to establish a theoretical link between vacancies and prices in a static environment with a fixed number of houses. More recently, Caplin and Leahy (2011) use a search model that can replicate certain features of the joint process of home prices, sales, and inventory. Krainer (2001) builds a theoretical foundation linking time to sell, housing volume, and prices. His search model is able to replicate the occurrence that the volume of transactions tends

to be high when prices are rising and liquidity is high, and low when prices are falling and liquidity is low.

Other papers incorporate both home and labor search. Head and Lloyd-Ellis (2012) also use search to explore the link between home liquidity, geographical mobility, and unemployment. They conduct an equilibrium analysis of a two city model in which agents consider the liquidity-driven value of their homes when deciding whether to accept a job and move to another city. The authors compare different home ownership rates and examine if there are any conditions in which the home-ownership rate affects the overall unemployment rate. Rupert and Wasmer (2012) explore the link between mobility and unemployment. They envision a single city model in which agents receive job offers with varying commuting times. Agents must decide whether to accept or decline, and if they accept they face a tradeoff between moving costs and transportation costs. These papers focus on comparisons between equilibria rather than an analysis of dynamics which is the emphasis in my paper. Karahan and Rhee (2013) utilize a model with different (but finite) locations to attempt to model mobility and unemployment in the Great Recession. They create a model with directed search in both labor and housing markets and run a counterfactual simulation to account for the amount of the increase in unemployment in the Great Recession due to the decline in house prices.

While not constructing a model, Blanchard and Katz (1992) offers an examination of state-level data. Among other relationships, they examine the link between employment and house prices and income. Though this is a statewide analysis, the results are reminiscent of those found in this paper. A negative shock to employment leads to a decrease in per capita income, though they find that the percentage deviation of income at peak is about half that of unemployment. Further, a negative shock to employment leads to a fall in the median house price, which has a peak deviation greater in absolute value than employment while converging back to the

mean much more quickly.

My paper owes a fair amount to Head et al. (2014). Their model is the basis for the model presented in Chapter 4, with a few notable differences. They do not look at employment, and therefore have no need for a labor search component. There are no goods producing firms, though there are outside-owned construction companies. Finally, all residents of the city receive an exogenous amount of income, though they supplement this with wages earned from constructing houses. This income process is the driver of the dynamics of their model.

In Chapter 5, I conclude the paper.

A Model of Learning and Portfolio Choice

2.1 Model

2.1.1 Income Process and Learning

The most general version of the income process can be thought of as a combination of three factors. The first is deterministic demographic considerations, such as number of children or experience, and is denoted $D_{i,e}$, where e is the individual's level of experience. This component is fully predictable, completely observable, and affects everyone in the same way. The second factor is an idiosyncratic component, such as inherent "skill", denoted $X_{i,e}$. This factor is unobservable to the econometrician, and perhaps even to the individual themselves. Note that this individual component can potentially change as the individual gains experience. The last component is random chance, i.e. "shocks", and is denoted by $\Xi_{i,e}$. To summarize, one can write that individual i in time t has income $Y_{i,e}$:

$$Y_{i,e} = D_{i,e}X_{i,e}\Xi_{i,e} \quad (2.1)$$

Taking logs, one gets the following:

$$y_{i,e} = d_{i,e} + x_{i,e} + \xi_{i,e} \quad (2.2)$$

I assume that demographic effects are linear in their arguments, that is that $d_{i,e} = \Phi_t demo_{i,e}$. Further, I subtract this from log income to get $\tilde{y}_{i,e}$, which is the portion of an individual's income draw that is due to factors unobservable by the econometrician. As mentioned, these factors can either be due to individual-specific factors or mere randomness.

I follow Guvenen (2007) and the subsequent literature in assuming that the individual component is a linear function of experience:

$$x_{i,e} = \alpha^i + \beta^i e \quad (2.3)$$

Notice that neither the intercept term, α^i , nor the growth term, β^i , have experience subscripts. These objects are parameters for each individual that are drawn at birth and are constant throughout the individual's lifetime. These parameters are assumed to be jointly normal and may be correlated:

$$(\alpha^i, \beta^i) \sim N\left([0 \ 0], \begin{bmatrix} \sigma_\alpha^2 & \sigma_{\alpha\beta} \\ \sigma_{\alpha\beta} & \sigma_\beta^2 \end{bmatrix}\right)$$

The shock process, ξ , can be thought of as an ARMA(p,q). For clarity, I separate it into two components, a moving average process ε and an autoregressive process z . Though ε could be modelled as any length of MA(q), for tractability I assume that ε is white noise. For similar reasons I take z to be an AR(1). To summarize:

$$z_{i,e} = \rho z_{i,e-1} + \eta_{i,e} \quad (2.4)$$

$$\eta \sim N(0, \sigma_\eta^2)$$

$$\varepsilon \sim N(0, \sigma_\varepsilon^2)$$

Thus, one can rewrite the portion of income that is due to factors unobservable to an outside observer as the following:

$$\tilde{y}_{i,e} = \alpha^i + \beta^i e + z_{i,e} + \varepsilon_{i,e} \quad (2.5)$$

Notice that this process nests the more straightforward “near”-random walk. This typical case restricts $\beta^i = 0$, and often holds $\alpha^i = 0$ as well. In the case that α^i is drawn from some distribution, individuals are fully aware of this value. In either case, any individual uncertainty in \tilde{y} is driven completely by the unpredictable persistent and transitory shocks. For the remainder of this paper, I consider the “benchmark” model to be the one that allows for heterogeneity in individual intercepts. Allowing for this generality affects the distribution of initial income, but not the uncertainty the individual faces per se.¹

The individual has knowledge about his earning potential to the extent that he knows α^i , β^i , and the current value of the persistent component $z_{i,e}$ (and therefore implicitly the size of the transitory shock, $\varepsilon_{i,e}$). However his understanding of these components may be imperfect. Therefore, the individual uses his income draws to gradually learn about his earning potential over time. The individual understands all of the parameters of the income process, such as the underlying variances, and of course observes his income. He never observes $\{\alpha^i, \beta^i, z_{i,e}, \varepsilon_{i,e}\}$, but rather updates his expectations and uncertainty surrounding these objects through observable income draws.

The individual may have some prior information about his idiosyncratic components. Priors may be set up in a similar way to Guvenen (2007). One can partition the α^i and β^i into components the individual has full information on, α_k^i and β_k^i , and components for which the individual has no information, α_u^i and β_u^i , such that:

$$\alpha^i = \alpha_k^i + \alpha_u^i \tag{2.6}$$

$$\beta^i = \beta_k^i + \beta_u^i \tag{2.7}$$

Let the within-parameter components be mean zero and independently drawn, that

¹ Note however that any estimation that allows for $\alpha^i \neq 0$ can potentially return different estimates of the variance of the persistent and transitory components than a more restrictive estimation. The estimates of these parameters may affect implied uncertainty.

is, $\alpha_k^i \perp\!\!\!\perp \alpha_u^i$ and $\beta_k^i \perp\!\!\!\perp \beta_u^i$, and let the variances of $\alpha_k^i, \alpha_u^i, \beta_k^i$, and β_u^i be $\sigma_{\alpha,k}^2, \sigma_{\alpha,u}^2, \sigma_{\beta,k}^2$, and $\sigma_{\beta,u}^2$, respectively. Then the variances of α^i and β^i are

$$\sigma_\alpha^2 = \sigma_{\alpha,k}^2 + \sigma_{\alpha,u}^2 \quad (2.8)$$

$$\sigma_\beta^2 = \sigma_{\beta,k}^2 + \sigma_{\beta,u}^2 \quad (2.9)$$

Denote λ_α and λ_β as the fraction of variance of the intercept and growth component, respectively, that is immediately resolved for the individual. In other words, before the individual has even received his first income draw, his prior uncertainty about $q \in \{\alpha^i, \beta^i\}$ is $\sigma_{q,u}^2 = (1 - \lambda)\sigma_q^2$. It follows that his prior for the covariance between the components is $\sqrt{(1 - \lambda_\alpha)}\sqrt{(1 - \lambda_\beta)}\sigma_{\alpha\beta}$.

The individual updates his beliefs via Kalman Filtering. The agent is concerned with the components $[\alpha^i \ \beta^i \ z_{i,e}]'$ that govern his income. Based on the draws up to and including his current level of experience, the individual forms expectations of this vector, $\hat{S}_{e|e}$. Further, his uncertainty about these parameters is denoted by $P_{e|e}$, the variance-covariance matrix of the components given his experience level. The individual's prior variance-covariance matrix is:

$$P_{0|0} = \begin{bmatrix} (1 - \lambda_\alpha)\sigma_\alpha^2 & \sqrt{(1 - \lambda_\alpha)}\sqrt{(1 - \lambda_\beta)}\sigma_{\alpha\beta} & 0 \\ \sqrt{(1 - \lambda_\alpha)}\sqrt{(1 - \lambda_\beta)}\sigma_{\alpha\beta} & (1 - \lambda_\beta)\sigma_\beta^2 & 0 \\ 0 & 0 & 0 \end{bmatrix}$$

Prior expectations are based on the initially resolved, known portion of the component partition: $\hat{S}_{0|0} = [\alpha_k^i \ \beta_k^i \ 0]'$. Note that the individual's prior belief about the persistent component will always be zero as this is a random shock that is completely independent of his idiosyncratic and demographic characteristics. The evolution of his expectations and the variance-covariance matrix thereafter is standard and therefore the details are left to the Appendix.

2.1.2 Household Problem

I consider a relatively standard two asset lifecycle model. Individuals work and receive exogenously generated income, Y_t , for T years, and then retire for an additional T_{Ret} years, during which they receive no income. Individuals supply labor inelastically during their working years and receive utility only from the consumption good, c_t . I assume that utility is time-separable. That is, in each period s , agents make decisions to maximize the following:

$$U(c_t, c_{t+1}, \dots) = \sum_{t=s}^{T+R_t} \delta^{t-s} u(c_t) \quad (2.10)$$

where δ is the household's discount factor.

Households can save either through a liquid or illiquid asset, denoted by m_{t+1} and b_{t+1} , respectively. Households are free to accumulate or decumulate the liquid asset without incurring any cost. Adjustment of the illiquid asset is costly, however, and carries a proportional transaction cost. A proportional transaction cost was chosen for this environment to mimic costs faced in adjusting a household's housing stock and retirement accounts (such as 401ks or IRAs), the two most important illiquid assets on most household's balance sheets. Housing transaction costs can include attorney fees, real estate agent commissions, and transfer taxes, all of which scale with the value of the home transacted. Retirement accounts face a 10% withdrawal penalty². Of course, generalizing all of the transaction costs associated with these assets in this way is not perfect: one can easily argue for the existence of fixed costs for housing purchases, and individuals can freely deposit into retirement accounts. However, for the illustrative purposes of this stylized model, proportional costs will suffice.

² There are a few exceptions to the withdrawal penalty, such as when the withdrawal is used to purchase a first home or pay for certain medical expenses, among other things. I abstract away from these exceptions for the rest of the paper.

Transaction costs must be reinvested in the illiquid asset or those returns will incur a cost, as is the case for any investments made within a retirement account. This also matches how returns are accessed in the case of housing. Households must sell their home, and pay the corresponding transaction costs, to permanently withdraw any equity from capital gains. Total transactions costs, the fraction γ of the change in the illiquid account, can therefore be written the following way:

$$tc_t = \gamma |(1 + r^b)b_t - b_{t+1}|$$

As briefly alluded to previously, the transaction costs associated with expansion and reduction of these assets may differ. To allow for this generality, $\gamma \in \{\underline{\gamma}, \bar{\gamma}\}$, where $\underline{\gamma}$ is applied to withdrawals ($b_{t+1} < (1 + r^b)b_t$) and $\bar{\gamma}$ is charged for deposits ($b_{t+1} > (1 + r^b)b_t$). All individuals begin their life with zero assets; that is $b_0 = 0 = m_0$. Further, individuals cannot borrow against any asset at any point in time: $b_{t+1} \geq 0$; $m_{t+1} \geq 0 \forall t$.

Liquid assets yield a net return of r^m . To compensate for their relative lack of liquidity, illiquid assets carry a higher return than that of the liquid asset: $r^b > r^m$. These returns are time-invariant and do not carry any risk.

When making consumption and savings choices, households care not only about current conditions but also the distribution of future income realizations. Since the current income realization Y_t is a combination of different factors, which all have very different implications for future income, this object alone is not very informative about future prosperity. Thus, the household's value function must take into account the full vector of individual beliefs about the underlying structure of his current income draw, $\hat{S}_{t|t}$, as this value influences his expectations about the future, E^S . Further, since beliefs evolve with additional information, the individual must internalize how his beliefs with change from an income draw in the next period when

understanding his expected continuation value. So, the household's value function can be written as follows:

$$V_t(Y_t, \hat{S}_{t|t}, m_t, b_t) = \max_{c_t, b_t, m_t} u(c_t) + \delta E_t^S [V_{t+1}(Y_{t+1}, \hat{S}_{t+1|t+1}, m_{t+1}, b_{t+1})] \quad (2.11)$$

In retirement, the problem is somewhat simplified. Since there are no transaction costs in retirement, the individual no longer has use for the liquid asset and therefore dumps all of his savings into the illiquid asset in his final working year. Further, in retirement the individual no longer receives any income and therefore faces no uncertainty. The dynamic problem becomes:

$$V_t(b_t) = \max_{c_t, b_t} u(c_t) + V_{t+1}(b_{t+1}) \quad (2.12)$$

$$(1 + r^b)b_t = c_t + b_{t+1} \quad (2.13)$$

In other words, this is a cake eating problem in which the individual must determine the optimal amount of his initial retirement assets to consume in each retirement period. This can be solved for analytically and necessarily satisfies the following set of equations:

$$c_{T+t} = c_{T+1} [\delta(1 + r^b)]^{\frac{t-1}{\sigma}}, \forall t \quad (2.14)$$

$$\sum_{t=1}^{Rt} \frac{1}{1 + r^b} c_{T+t} = b_{T+1} \quad (2.15)$$

where the first set of equations are the Euler equations and the second equation is the budget constraint for the entire retirement period.

2.1.3 Solving the Model

The model is solved via backwards induction. For a given level of assets, accumulated during the working periods, the total utility derived from the retirement periods can be solved for analytically and summarized as one value, $V^{Ret}(b_{T+1})$. Note that

$m_{T+1} = 0$ as there are no transaction costs in retirement, but the illiquid asset still yields a higher return. Further, due to the lack of uncertainty the only relevant state variable is the level of illiquid assets, b_{T+1} .

Numerical methods are necessary to retrieve the decision rules throughout the individual's working life. Individuals make their portfolio selections, $\{m_{t+1}, b_{t+1}\}$, off of a $M \times B$ grid, where $M = 100$ and $B = 100$. The grid is not evenly spaced. To make sure that individuals do not have to adjust their illiquid holdings unless they desire to, each dimension of the grid is spaced such that $(1 + r^b)q_i = q_{i+1}$, where $q \in \{m, b\}$ and q_i and q_{i+1} are sequential grid points. This allows individuals to choose their previous holdings (plus accrued interest) as their current asset holdings, without having to incur a cost or commit more of their assets to a less liquid position. In every period, the choice set for each asset is the same, so differences between grid spacing does not bias individual decision making towards any one asset.

The value function of the agent depends not only on their current income, y_t , but also their beliefs about how much of that income can be attributed to α^i , β^i , and z_i . Thus, there are additional state variables that one must keep track of while computing the decision rules. Notice that for a long enough time horizon y_t is unbounded, even if income shocks are truncated. Thus, it is of some use to replace y_t in the individual's problem with the forecast error, ϵ_t .

Notice that, conditional on his expectations about the state today, from the individual's perspective his next-period income follows the following distribution:

$$y_{t+1}|s_{t|t} \sim N(H'_t s_{t+1|t}, H'_{t+1} P_{t|t} H_{t+1} + R)$$

where H, P, and R are defined in the Appendix. Defining $\epsilon_{t+1} = y_{t+1} - s_{t+1|t}^\alpha - s_{t+1|t}^\beta (t + 1) - s_{t+1|t}^z$, it is clear that ϵ_{t+1} is normally distributed with mean zero and a variance of $H'_{t+1} P_{t|t} H_{t+1} + R$. Notice that the variance here does not depend on individual

income draws, as $P_{t|t}$ evolves solely according to the variances of the components of the income process. Thus, the variance for ϵ_t can be solved separately from backward induction. For a given period's variance, I create a 5-spaced grid for ϵ_t , adapting the method introduced in Tauchen (1986) for the case where $\rho = 1$. Extreme values are given by three standard deviations above or below the mean, zero.

The value function is a function of the individual's beliefs in that period. Thus, when computing the expected value of the following period's value, one must account for how beliefs will change between periods. Next period beliefs are a function of current beliefs and the next period income draw, restated above as ϵ_{t+1} . Thus, for each of the gridpoints of ϵ , I calculate the values of $\hat{S}_{t+1|t+1}$ given $\hat{S}_{t|t}$:

$$s_{t+1|t+1}^\alpha = s_{t+1|t}^\alpha + K_{t,\alpha}\epsilon_{t+1} \quad (2.16)$$

$$s_{t+1|t+1}^\beta = s_{t+1|t}^\beta + K_{t,\beta}\epsilon_{t+1} \quad (2.17)$$

$$s_{t+1|t+1}^z = s_{t+1|t}^z + K_{t,z}\epsilon_{t+1} \quad (2.18)$$

where $K_{t,q}$ is the cell of the time-t Kalman Gain K_t corresponding to component q and $\hat{S}_{t|t} = \exp(\hat{s}_{t|t})$.

I use a $5x5x5xT$ grid to store the values at various beliefs. Since beliefs evolve on a continuum, typically the individual's next period beliefs do not fall on the grid. I use trilinear interpolation in these cases. The endpoint values of the grid change with the time period, t, and reflect the most extreme values the grid can take given the possible income draws of ϵ in the current period. Thus, there is no need for extrapolation.

The computational strategy of the benchmark process is fairly straightforward and discussed in the Appendix.

2.1.4 Calibration

The household has constant relative risk-aversion (CRRA) preferences over consumption: $u(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma}$, where σ is chosen to be 2. I consider each period to be a year. I set total working years, $T = 43$. This aligns with the longevity of an individual who enters the labor force upon graduation from college at age 22 and retires at the age of 65. I set $T_{Ret} = 15$, which implies that individuals die at age 80, which is only slightly higher than the unconditional life expectancy in the US of 79.³ Since I consider time to be annual, I set the discount factor to $\delta = .96$. I set the interest rate for the liquid asset such that it completely offsets the discount factor; that is $r^m = \frac{1}{\delta} - 1$. To compensate for the increased risk in the illiquid asset, I set a positive wedge between the interest rates such that $r^b = r^m + .02$. For the remainder of the paper I consider the “no-prior” case; that is, $\lambda_\alpha = 0 = \lambda_\beta$, and assume that median income in each period, D_t is normalized to 1.

My parameterization of the adjustment cost is highly stylized, which is done primarily for expositional purposes. I set $\bar{\gamma} = 0$; in other words I completely waive costs for households adjusting their illiquid assets upward. This is not necessarily an unrealistic assumption for many types of illiquid assets. For example, retirement accounts levy a penalty only if funds are withdrawn before a certain age, and do not penalize deposits.⁴ The main deviation from reality is my assumption of $\underline{\gamma} = 1$, or complete irreversibility, which is perhaps only applicable to a highly specific subset of assets, such as irrevocable trusts. While this choice has the effect of increasing the risk of holding the illiquid asset, it also makes the choice between liquid and illiquid assets quite stark. Because of their complete irreversibility under this parameterization, illiquid assets are accumulated solely for use in retirement periods as they

³ In 2015, according to the World Bank

⁴ Though the tax-exemption deposits receive ends at a certain threshold. Since I abstract from this feature of retirement accounts entirely, I ignore this distinction.

are unavailable to smooth any severity of shocks during working periods. Thus, in this setting liquid assets are the only way to save for use in later working periods. Thus this parameterization makes it very clear what amount of savings the individual plans to use for retirement (illiquid savings) and what amount of savings is held either for precautionary purposes or to use later in the lifecycle, perhaps if the individual expects income to fall (liquid savings).

Table 2.1: Income Process Calibration

Model	ρ	σ_α^2	σ_β^2	$\sigma_{\alpha\beta}$	σ_η^2	σ_ε^2
Learning	.805	.023	.00049	-.0024	.025	.032
Benchmark	.979	.031	—	—	.0099	.047

I calibrate the income process to parameter point estimates found in Guvenen (2009). First isolating and removing the deterministic trend, Guvenen then uses GMM to match moments from the income residuals. In this way he is able to estimate each model in turn using income data alone. Both processes are estimated from the same male earnings data, taken from the Panel Survey of Income Dynamics (PSID). See that paper for more details on sample selection and method.

There are significant differences between the parameterizations of the two processes. Dispersion of the intercept is much higher in the benchmark case. However, note that this high level of dispersion does not contribute in any way to uncertainty in the benchmark model, as individuals receive this draw at birth and are fully aware of its value. This value merely affects the initial variance of income draws and each individual’s idiosyncratic mean income. The transitory component does impact uncertainty in both models, and is higher in the benchmark. Further, the benchmark has a higher persistence parameter. The difference is substantial; only $\approx 14\%$ of a persistent shock remains after 10 periods in the learning process as opposed to a little over 82% in the benchmark process. However, the variance of the persistent shock is higher in the learning model, so how the individual feels about this trade-off

depends on how heavily the individual discounts future periods.

Crucially, individuals in the learning model have an idiosyncratic growth component whereas those in the benchmark model do not. The small variance of this component belies its importance in determining income, especially later in life. An individual who has a growth parameter equal to one standard deviation above the mean has $\beta^i = .0221$, or a level of income $\approx 2\%$ higher in the first period of working and $\approx 25\%$ higher after the 10th period of working, when compared to an individual who has a mean growth of 0. Finally, notice that the population covariance between the intercept and growth component is actually negative. The correlation between these two components is considerable: $-.7059$.

2.2 Intuition

The household's asset decisions depend crucially on the type of exogenous income process they are facing. This is because a model with learning changes both the structure of uncertainty over the lifetime as well as the individual's forecast of their mean income. I consider both of these mechanisms in turn.

2.2.1 Uncertainty

To get a grasp on how uncertainty resolves itself over the lifetime⁵ I consider the present discounted value of the square root of the individual's mean squared error of log income in both the learning model and under the benchmark. This is because the mean squared error of log income can be written as follows:

$$\begin{aligned} MSE_{t+s|t} &= E_t[(y_{t+s} - y_{t+s|t})^2] \\ &= E_t[(\tilde{y}_{t+s} - \tilde{y}_{t+s|t})^2] \end{aligned}$$

⁵ This analysis is reminiscent to that of Guvenen (2007), Guvenen and Smith (2014), and Chang et al. (2016). They focus on forecast error at different horizons, while I attempt to quantify the total amount of uncertain left in the lifetime at different ages.

$$\begin{aligned}
&= Var_t(\tilde{y}_{t+s} - \tilde{y}_{t+s|t}) \\
&\approx Var_t\left(\frac{\tilde{Y}_{t+s} - \tilde{Y}_{t+s|t}}{\tilde{Y}_{t+s}}\right)
\end{aligned}$$

where $\tilde{Y} = \frac{Y}{D}$ and the final approximation follows from the convenient properties of logs, as long as forecast error is small. Thus, one can think of the mean squared error of the forecast in period (t+s) from the vantage point of period t as the variance of the percentage forecast error. The square root of this object gives the percentage deviation from mean income of a one standard deviation increase in percentage forecast error.

In order to get a sense of lifetime uncertainty, I then multiply the percentage deviation by median income, which is normalized to $D_t = 1$ in each period. This gives the dollar amount of a one standard deviation increase in forecast error. For each period t, I then discount each of these deviations by the household's discount factor, and sum them for each future period (t+s):

$$\sum_{s=1}^T \delta^s D_{t+s} \sqrt{MSE_{t+s|t}}$$

This gives the total (discounted) change in income if the household's forecast error was one standard deviation above the mean in every future period.

In Figure 2.1a I plot this amount for each period for both the learning case as well as the benchmark, using the values from the calibration section. Figure 2.1b is the same information, just normalized so that initial uncertainty is equal to one in both cases. It is clear that over the entire lifecycle uncertainty is greater in the learning model, with the greatest disparity being towards the beginning of the lifecycle. Further, from the normalized graph it is clear that initially uncertainty is resolved much quicker in the benchmark case, while the pace of uncertainty resolution

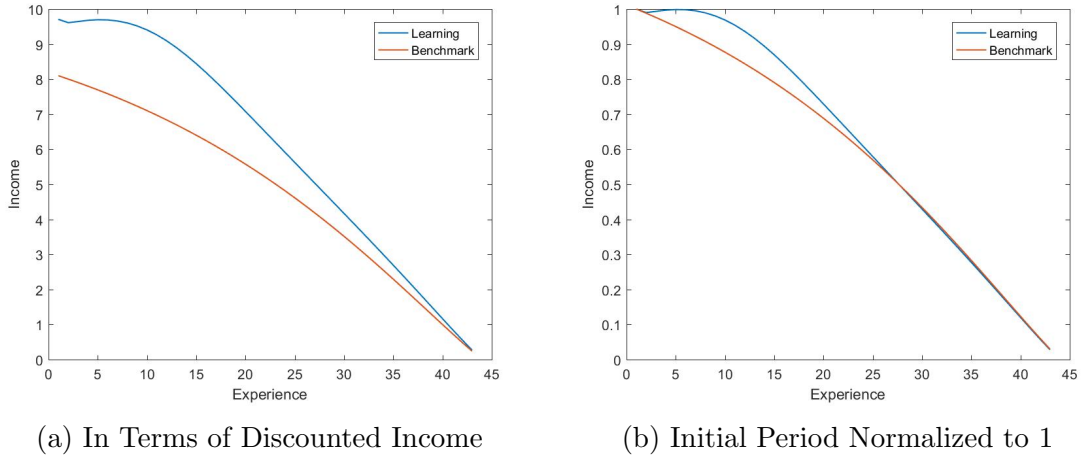


FIGURE 2.1: Present Discounted Value of Forecast Errors

is higher towards mid-life in the learning case. After period 25 or so both models imply nearly equivalent uncertainty resolution for the rest of the lifecycle.

For the benchmark case, uncertainty resolves itself almost linearly. This results entirely from the individual's lifespan shortening over time. As they gain experience, there are simply fewer periods left for the persistent shock to accumulate; thus their uncertainty falls as the variance of the transitory component is constant over the lifecycle. Since the periods that drop off are highly discounted at the beginning of life and are only slightly discounted latter on, there is some concavity, albeit slight.

The graph depicting the learning model similarly features declining uncertainty through the attenuation of the accumulation of persistent shocks. The shape of the uncertainty profile in this model is of course also driven by the individual's lack of information about his own skill parameters. Early in life, the individual knows little about his true level of skill. Though his lack of knowledge about his intercept contributes to initial uncertainty, this high level of uncertainty is mostly due to his lack of information on his growth potential. Income in mid-life and later years is dominated by the growth potential of the individual, and therefore even small changes in this draw can yield large changes in total lifetime income.

Further, since this component does not contribute much to income draws early in life, it takes a while before the individual can tighten the distribution of this component. Thus, uncertainty resolves itself very slowly early in life under the learning model, and even increases slightly as individuals age closer to (and therefore discount less) the periods where the growth rate matters faster than they learn about it. Eventually income draws become more informative, and the individual is able to get a better sense of what his growth component truly is. This is reflected in each of the above graphs by a quick decline in lifetime uncertainty in mid-life.

2.2.2 Forecasts

An individual's conditional mean forecasts of future income are also different over the lifetime between the two models. To illustrate, I consider demeaned log income, \tilde{y} . In the benchmark case, forecasts are fairly straightforward. Individual know how much of their income is due to each component. Thus, they know the value of their intercept and consider it fixed; they allow for mean reversion of the persistent component; and they assume that the transitory component will completely dissipate. Thus, given $\tilde{y}_t = \alpha + z_t + \varepsilon_t$, conditional expectations for period (t+s) are:

$$E_t[\tilde{y}_{t+s}] = \alpha + \rho^s z_t$$

In the learning model individuals of course do not know the value of their intercept, growth rate, persistent component, or transitory component, but merely have beliefs about these components based on their priors and information in the form of income draws. Thus, the actual composition of the current income draw does not matter for forecasts; only the composition of the individual's beliefs. The individual assumes that their intercept and growth terms will remain constant, and use the later component to extrapolate growth in future income. Further, the individual assumes the transitory component will immediately disappear and the persistent component

will deteriorate geometrically with time. Thus, one can write the conditional forecast as so:

$$E_t[\tilde{y}_{t+s}] = s_{t|t}^\alpha + (t+s)s_{t|t}^\beta + \rho^s s_{t|t}^z$$

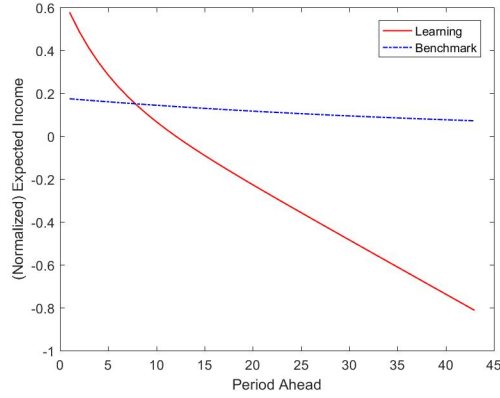
Notice that, while the forecasts of the benchmark model evolve the same way irrespective of starting period t , this is not the case in the learning model. This is because in the learning model forecasts are based on current beliefs, and current beliefs can vary depending on the timing within the lifecycle. To illustrate this point I compare the forecasts for individuals whose current income draw is one standard deviation above the mean at various points in the lifecycle.

For ease of exposition I assume that in the benchmark case $\alpha = 0$. If α were allowed to vary the forecast would still revert at the same rate; it would merely revert to α rather than 0. This would also affect the size of ε_t for any given income y_t . However, for my purposes here these distinctions are largely unimportant. Thus, in the benchmark case I consider the forecast of an individual who has income one standard deviation above the mean, conditional on $\alpha = 0$. An individual who just experienced a positive shock will first forecast an immediate, discrete drop in income due to the purely transitory component reverting back to its mean. The individual will then forecast a constant rate of mean reversion, as forecasted income tracks the gradual decline of the accumulated persistent component. For any initial time t , the share of the total variance $Var_{tot} = \sigma_\varepsilon^2 + \frac{1-\rho^{2e}}{1-\rho^2}\sigma_\eta^2$, and therefore the share of any deviation from mean income, due to the transitory component is $\frac{\sigma_\varepsilon^2}{Var_{tot}}$. What this means is that as time increases, more of the deviation of income from the mean is due to the accumulation of permanent shocks. Thus, for later periods, the immediate discrete drop is smaller relative to current income.


Because contemporaneous beliefs depend on the entire sequence of income draws,

FIGURE 2.2: Income Expectations

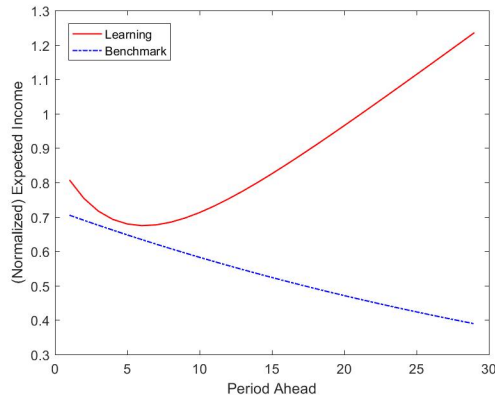
(a) Mean Forecast at Period 1




(b) Mean Forecast at Period 5



(c) Mean Forecast at Period 15



(d) Mean Forecast at Period 30



there is no one “true” forecast for an individual in the learning model. Therefore I simulate income draws for 100,000 individuals, and average the expectation path of the 100 individuals immediately above and 100 individuals immediate below the one standard deviation mark. The graphs are normalized so that the initial period’s income is equal to 1.

The most striking result here is that although benchmark forecasts are largely invariant to the period, forecasts in the learning model depend heavily on the individual’s experience level. As mentioned before, the benchmark forecasts revert to

the mean at rate ρ in all periods, after an initial discrete decline due to the dissipation of the transitory shock. For the learning model, the entire expectation path changes shape as the expectation level increases. In period 1, mean forecasts are highly pessimistic. Individuals predict their income will not only fall immediately, but will continue to decline over their entire lifetime. In fact, individuals believe that their income will eventually fall to less than mean income.⁶ Forecasts in period 5 decline rapidly as well, though there is some hint of increases in income after 20 periods. By period 15, forecasts attain more of a U-shape, declining for the first 5 or so periods and then actually increasing for the rest of the lifecycle. The trough is at around 70% of initial income. Forecasts in period 30 have a much shallower U-shape; these forecasts decline to only about 95% of initial income.

The forecasts are a result of which components the individual attributes his positive income draw to. The individual knows the population distributions of each of the parameters, and understands the potential of each component to affect his income throughout his life. Due to the low variance of the growth component, the individual understands that this component will most likely not account for much of his early income draws, but will dominate his income path later in life. The opposite is true of the intercept component, which will be important early in life but have very little effect later on, relative to the growth component. Further, he understands that the intercept and growth components negatively covary, and that a positive intercept will tend to be accompanied by a negative growth rate.

Thus, an individual who receives a very high income draw in the first period knows that it is likely not due to the growth component. Therefore, he attributes some of this draw to having a positive intercept, some of the draw to receiving a positive persistent shock, and the rest to receiving a positive transitory shock. Further,

⁶ Note that this exercise deals with de-measured income, and that the mean income profile is increasing over most of the lifecycle. Thus, these individuals are not necessarily predicting that their actual income will fall over their lifetime; only that their income relative to the mean will decline.

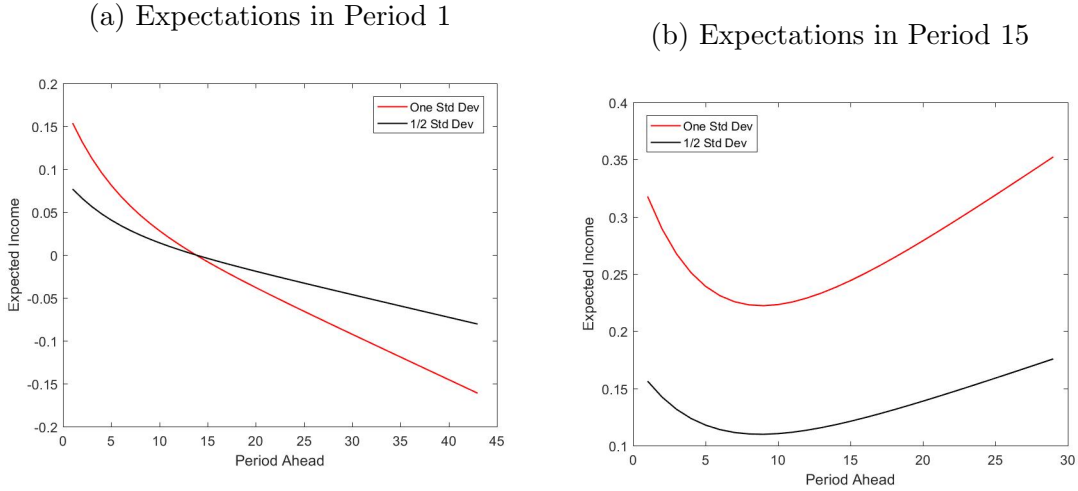
he believes that the growth component is actually negative, due to the negative covariance in the population between the intercept and growth components. Thus, the individual's forecast is a combination of mean reversion and an expectation of negative growth, relative to the mean. This leads to low expected income throughout the lifecycle, despite receiving a high income draw.

By period 15, individual forecasts change dramatically as compared to those earlier in life. At this point the growth component does account for a non-negligible portion of income draws, and individuals receiving high income can assume that their growth component is positive. They still believe that some portion of their income is due to a positive persistent shock; this accounts for the decline in expected income over the first few subsequent periods as the shock diminishes faster than their income increases due to the growth component. Eventually, around 5 periods later the growth component wins out and individuals begin forecasting increasing income. After around 20 periods the individual expects income to be higher than it is in the initial period, and the individual anticipates even higher income thereafter.

Forecasts in period 30 tell a similar story. However, consider that, due to the widening of the income distribution over the lifecycle, an income draw one standard deviation above the mean in period 30 is much higher than it is in period 15. Further, on average individuals who receive a high draw in period 30 have received more high draws throughout their lifetime than an individual who receives a high draw in period 15. Both of these factors lead individuals who receive this draw to attribute proportionally more of their income to a high growth component and less to the persistent component. Thus, the U-shape is much more shallow, as the mean reversion of the persistent shock is mostly offset by the higher growth rate. Individuals expect to recover their relative income level after around 8 periods and to have higher income thereafter.

Individuals at different income levels within each period have different forecasts,

FIGURE 2.3: Comparison of Expectations at Different Income Levels



as well. In Figure 2.3 I plot out the forecasts of individuals at one standard deviation and one half standard deviation above the mean in period 1 and again for period 15. One can tell with even a cursory look at these graphs that individuals with income draws closer to the mean tend to have less “extreme” forecasts; that is, their forecasts deviate less from their current income than individuals who have current income values further away from the mean. This is because, with smaller income draws, individual’s have lower beliefs about the growth and persistent components, which drive the forecast dynamics.

For example, in period 1 a lower draw leads individuals to believe both the persistent and intercept components are lower. Though they still believe the growth component is negative, because the intercept component is closer to the mean value of zero their expectation for growth is closer to zero as well. Thus, individuals who receive less income expect less mean reversion (since the base is smaller) and less negative growth in the future. So their expectations, while still pessimistic, are less drastic than individuals who receive higher income in this period. In period 15, individuals who receive an income draw only one half standard deviation above the

mean have lower beliefs of the persistent and growth component. This leads them to expect less mean reversion in the near term and lower growth over the rest of the lifecycle. This implies a shallower U-shape than that of the individuals who receive a higher draw.

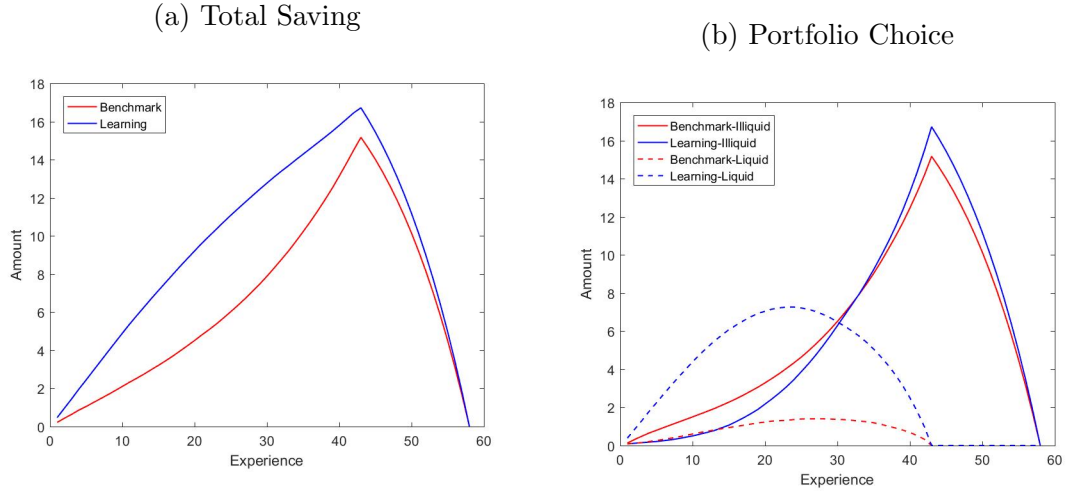
Though I have omitted them from the graphs, expectations are symmetric around the mean. That is, individuals who receive an income draw one standard deviation below the mean have the opposite forecast from those who receive a draw one standard deviation above the mean. This implies that individuals with a poor draw in period 1 are actually fairly optimistic and believe that eventually they will receive income above the mean. Individuals who continue to receive negative draws throughout their lifetime become more and more pessimistic, however.

2.3 Results

2.3.1 Set-Up

I simulate the lifecycle paths of income, consumption, and assets for 1000 individuals in both models using decision rules found via backward induction. There may be some concern that, because these two processes yield different income distributions over the lifecycle (see Appendix), the results discussed here may be more of a function of the distribution rather than from the mechanisms emphasized above. Thus, I use the same income draws, generated from the benchmark process, for each simulation. The individuals in the learning simulation still act as if they are receiving draws from the learning process and update their beliefs accordingly. Though not reported here, the results discussed in this section are robust to using different income draws for each model generated by their respective parameter estimates. Each individual has an unconditional expected income of 1 in the first period.

FIGURE 2.4: Average Savings Over the Lifecycle



Comparison of the Two Models

Figure 2.4 shows the population average of various measures of household savings over the lifecycle. All individuals begin their life with zero assets, and choose to save for two reasons. The first is precautionary, to buffer consumption from unpredictable shocks. Due to the curvature of the utility function, expectations over the marginal utility of uncertain draws is greater than the marginal utility of the expectation of the draws. Thus, the more uncertainty they face, the more individuals suppress consumption today in favor of greater saving. The second savings motive is to make up for future expected shortfalls in income. This includes the lengthy retirement period, in which individuals receive zero income (and therefore infinitely negative utility if they have no savings; i.e. death) as well any anticipation of declining earnings during the individual’s working periods.

In Figure 2.4a , I graph the total assets- liquid and illiquid- that the average household has accumulated at each experience level. In both models average savings increase monotonically until retirement, and then decline thereafter as individuals draw down their savings to fund retirement consumption. Due to the absence of any

bequest motive (or possibility of an early, unanticipated death), all individuals end their lifetime with zero net assets. Average savings in the learning model are higher throughout the lifecycle as compared to the benchmark case. Much of this can be attributed to the higher level of uncertainty individuals face throughout their lifetime, increasing the precautionary motive and therefore the average level of savings. There are two things to note here. First, though income uncertainty is gradually resolved throughout the lifetime in both models, theoretically decreasing the amount of desired savings, average savings never falls during working periods. This is because any precautionary savings that are freed up due to less uncertainty must then be spread over the rest of the individual's working life and retirement, so little of it translates into consumption and most remains held as savings. Further, while the disparity in uncertainty between the two models falls as individuals age, at every point there is more uncertainty in the model with learning. Thus, at every experience level the precautionary motive is smaller with the benchmark process.

The change in expectations from the benchmark to the learning model also affects savings. High income individuals who learn are more pessimistic about their future income path early in life than their benchmark counterparts, and therefore save more of their income to smooth consumption. Similarly, low income individuals are more optimistic than if they had the benchmark process, and therefore would like to dissave and consume more. As individuals gain experience, these expectations reverse and high income individuals become more optimistic while low income individuals become more pessimistic. Therefore high income individuals lower their desired savings while low income individuals raise theirs, relative to their benchmark counterparts.

Note that at all times, the expectations channel generates opposing behavior for individuals with income above and below the median. Of course, whichever group dominates will determine the aggregate effect of the expectations channel. If one

assumes that the behavior of the high income individuals dominates⁷, then this implies that this channel raises the average savings rate early in life while decreasing average savings rates later on.⁸

Notice that the peculiarities of the transaction cost restrict the aforementioned savings motives to specific assets. Because illiquid savings cannot be withdrawn before retirement, any savings accumulated for use during the individual's work life must be deposited into the liquid account. This means that any savings done for precautionary reasons will be placed in the liquid account, as income is deterministic in retirement. Further, any saving undertaken in anticipation of a decline in earnings during the individual's working periods must be stored in the liquid account as well as that is the only way they can be accessed when needed. Similarly, since the illiquid asset can be freely accessed in retirement and earns a higher return, any funds saved specifically for retirement will be placed into the illiquid account.

In 2.4b I decompose total savings up into its constituent parts for each model. The general, qualitative pattern of saving is similar in each model. Average illiquid savings are strictly increasing until retirement, after which time they are slowly drawn down. Average liquid savings begin to rise immediately and are hump shaped over the working years, peaking at around twenty or thirty years of experience, depending

⁷ This is a fairly safe assumption for two reasons. The first is that incomes are distributed lognormally. This means that deviations from the median are more extreme at the top of the income distribution than at the bottom. Thus, changes in expectations from the benchmark to the learning case are larger for incomes above the median than below. In a certainty equivalent world, where individuals perfectly consumption smooth, larger changes in future income realizations translate into larger changes in desired savings. While individuals in both models do experience uncertainty, they face the same amount regardless of their current income. The second reason is that borrowing constraints introduce an asymmetry, at least in the early periods of life. While high income individuals do not face constraints on adjusting their savings upwards, low income individuals face a situation in which they would like to dissave but do not have much savings built up already; thus the borrowing constraint may be binding. This is less of an issue in later periods, where individuals will have a higher level of savings on average.

⁸ Note that this does not imply that this channel causes overall savings to decrease towards the end of the lifetime. Because the amount of average resources will be higher later in life (due to higher savings at lower levels of experience), there may be higher savings overall even with a lower savings rate.

on the model. In the final period of work individuals dump all of their liquid assets into the illiquid account and maintain a zero liquid balance thereafter. Quantitatively, the two models exhibit stark differences in average portfolio allocation. Individuals who learn about their income over time hold on average less savings in their illiquid account until they have around 30 years of experience, despite saving more overall. Afterwards, this relationship flips for the rest of the lifecycle with learning individuals holding more illiquid assets on average. At the time of retirement, average illiquid holdings in the learning model are around 2 units higher than in the benchmark case. This is fairly substantial, given that median income is normalized to 1 unit. Throughout all working periods, the learning model produces a higher level of liquid savings. This difference between the models can be large: at its peak, average liquid savings in the learning model are more than four times the peak in the benchmark model. Further, average liquid holdings reach their maximum point at different experience levels. Individuals who learn begin decumulating liquid assets on average in their early twenties. For the benchmark agents, this does not occur until their late twenties.

This discrepancy in portfolio composition is driven by the way the two income processes influence risk and expectations. The higher precautionary motive not only causes individuals in the learning model to save more, but also to direct more of their savings into the liquid fund, which they can access in case of poor earnings in the future. Similarly, the high income individuals who are pessimistic early in their life must save more of their funds in the liquid account. They expect their income to immediately decline and remain low throughout the remainder of their working years, and therefore for current savings to be used to smooth consumption during the rest of their working years in addition to retirement. Assuming once more that the behavior of the high earners determines aggregate changes, the expectations channel then causes liquid savings to be higher in the learning case than the benchmark case;

the less pessimistic benchmark individuals can put more of their savings towards retirement, in the illiquid account, because they have to save less for their working years.

As individuals age in both models, they eventually begin to decumulate liquid assets in favor of illiquid, due primarily to the decrease in lifetime uncertainty reducing the precautionary motive. As the necessity of precautionary savings decreases, the need for liquid assets falls as well. Individuals then optimally distribute these funds over their remaining lifetime. Thus, a portion of the funds are put towards the illiquid asset for retirement, some are immediately consumed today, and then the rest are earmarked for future working periods and re-deposited in the liquid account. This decumulation begins slightly earlier in the learning model. This may be partly due to how much of the lifetime uncertainty is resolved in the first half of the agent's working life. From that point on, the amount of residual uncertainty in the learning model is not that much higher than in the benchmark model. However, the amount of liquid savings the average individual has accumulated up to that point in the learning model vastly exceeds their benchmark counterpart. Thus the point at which the learning individual no longer needs to expand their liquid holdings occurs a bit sooner than in the benchmark model. Further, because they saved so much early in their lifetime, learning individuals have a higher stock of savings to reappportion as the precautionary motive falls. They therefore end up with a higher stock of illiquid savings at the end of their working life (as well as higher levels of consumption).

Finally, the expectations channel at the end of the income earning portion of life works in the reverse as compared to at the beginning of life. High income individuals in the learning model are now very optimistic about their future earnings. Thus there is less need for liquid savings, as future working years will be more prosperous than the present and retirement. Thus, these individuals take some of the funds they were saving to bolster consumption during future working periods and divide

it between consumption today and retirement. Low income individuals are now very pessimistic, and would like to increase their liquid savings to cover expected future shortfalls in income. However, their current income is low, meaning that they cannot save much in absolute terms. Further, while individuals are able to transfer from the liquid to the illiquid account, they are unable to transfer the opposite way. Therefore, it is plausible that the asymmetry of savings opportunity contributes to liquid dissaving and increased illiquid saving in the aggregate in later periods. Note that this reversal is not present at all in the benchmark model, as all individuals still expect mean reversion.

The Ramifications of Incorporating Learning Into A Portfolio Choice Model

3.1 Household Portfolio Data

To get a sense of household portfolio composition, I consider data from the Survey of Consumer Finances (SCF). The SCF is a triennial survey conducted by the Federal Reserve since 1983. The survey contains very limited data on income, demographics, and expenditures, though as the name might suggest it is noted for its rich and exhaustive consumer financial data. I partition this data into groups, liquid and illiquid, only slightly deviating from the methodology of Kaplan et al. (2014). Liquid assets consist of checking and savings accounts, money market funds, call accounts, directly held mutual funds or stocks, liquid bonds (such as Treasury bills or corporate bonds), and miscellaneous financial assets such as futures, oil gas, and mineral investments, and non-public stock¹. Cash is lumped in with this latter group. I follow Kaplan et al. (2014) by imputing a household's cash holdings us-

¹ Some types of assets with this designation are arguably more illiquid than liquid, such as outstanding loans to someone else, future proceeds from lawsuits, etc. However, this asset class is fairly insubstantial to the typical household, and including this group as an illiquid asset does not affect the analysis of this paper

ing the ratio of average cash to checking, savings, money market, and call accounts found in Foster et al. (2013), .055. For the sake of brevity, for the remainder of this analysis I group checking, savings and cash together, money market and call accounts together, and combine directly held mutual funds and stocks as “stocks”, all liquid bonds into “bonds”, and miscellaneous financial assets into “other liquid assets”².

Illiquid assets consist of residential and nonresidential properties, retirement accounts (such as a 401K or IRA) and annuities, certificates of deposit (CDs), the total cash value of any life insurance policies, trusts, savings bonds (which cannot be transferred), any stake in a non-actively managed business, and miscellaneous non-financial assets such as jewelry, expensive paintings, baseball cards, etc. I exclude vehicles, as I consider these short-term consumption items, rather than long-term investments. Further, I exclude the value of actively managed businesses of entrepreneurs. I group together secondary residential and nonresidential properties as “non-primary property”, and consider primary housing as a separate entity. I group retirement accounts and annuities together as “retirement assets”, and combine CDs, life insurance, trusts, savings bonds and stake in non-actively managed businesses into “illiquid financial” assets. Miscellaneous illiquid assets I denote “illiquid nonfinancial” assets.

Because the focus of this paper is on learning about ability while working, I put everything in terms of experience, rather than age. “Experience” is not provided by the SCF. Therefore, I create an experience variable by subtracting the highest grade completed (plus six, the standard age to begin schooling) from the age of the head of household. For individuals who have less than ten years of schooling, I subtract sixteen from the head, as this is the minimum age an individual can begin working.

² I do not remove imputed cash holdings from this asset class; thus its importance may be somewhat overstated in the analysis of this paper. However, I show that this asset class is fairly unimportant to the typical household, even with cash included.

These surveys are fairly vague when it comes to postgraduate work. There is no way of telling from the data how long an individual has spent in graduate school, or when they graduated; one can only tell whether the individual attended graduate school at some point in their life. I consider any graduate school attendance as one more year of schooling, and assume that the individual attended immediately after completing their undergraduate work.

Table 3.1: Average Share of Liquid Assets

Exp.	Tot. Value	Cash/Chk/Sav	MM/Call	Stocks	Bonds	Other
All	90,852.96	70.11	8.18	14.77	.55	3.46
≤ 4	6,527.043	79.23	4.49	8.4	0	4.76
5 - 9	10,902.03	75.68	5.57	10.33	0	5.5
10 - 14	20,219.41	76.06	5.58	10.82	0	5.02
15 - 19	41,210.54	72.18	7.61	12.94	0	4.38
20 - 24	60,990.56	70.61	7.06	15.41	.19	3.83
25 - 29	93,701.5	69.15	7.98	15.67	.52	3.85
30 - 34	128,618.1	67.84	9.04	16.72	.41	3.29
35 - 39	157,850.5	66.97	8.74	18.34	.74	2.66
40 - 44	184,044.4	64.16	9.97	18.28	1.13	4.18
45 +	143,585	67.34	10.5	15.85	1.65	1.79

Table 3.2: Average Share of Illiquid Assets

Exp.	Tot. Value	Residence	Oth. Prop.	Ret.	Oth. Fin	Oth. Nonfin
All	260,677.9	51.88	7.6	24.42	13.06	.75
≤ 4	14,797.75	23.71	3.49	28.77	33.79	9
5 - 9	42,920.9	37.40	5.18	33.19	18.82	2.43
10 - 14	76,493.99	45.95	5.03	30.68	14.47	1.12
15 - 19	140,625.5	49.00	6.91	28.26	12.64	.97
20 - 24	206,609.5	50.71	6.32	29.31	9.73	.75
25 - 29	292,519.3	51.58	8.34	27.72	9.24	.64
30 - 34	379,477.1	51.99	9.07	27.61	8.17	.62
35 - 39	449,188.4	52.19	9.44	26.27	8.51	.66
40 - 44	493,504.5	53.52	9.92	23.07	9.99	.60
45 +	352,829.2	62.66	7.94	11.16	15.41	.46

I construct Tables 3.1 - 3.4 using data from 1989 to 2013 to get at the relative importance of each asset in the typical household's portfolio. In Tables 3.1 and

3.2 I calculate the average share of each asset class as a percentage of total liquid and illiquid assets for different levels of experience, as well as for the entire sample. Splitting the sample up in this way gives some insight into whether some assets gain importance at certain points in the lifetime, which could potentially be lost by just looking at the sample overall. In order to attenuate the impact of outliers, after computing shares I Winsorize each asset at the 5th and 95th percentile for each group. In the tables I also include the average holdings of each type of asset, to give some sense of how large each asset class is in the average household's portfolio.

The implications of these tables are clear. The average household balance sheet is dominated by a few assets. Households hold an average of 65 - 70 % of their liquid assets in easily accessible checking or savings accounts, or directly as currency, though the share of this asset declines over the lifecycle. They hold an additional 5 - 10 % in money market accounts or call accounts, which bear little risk and are nearly as easy to access as checking accounts. Put together, these near-riskless, highly liquid assets make up nearly 80% of the average household's liquid assets. Bonds play nearly no role in the average household's balance sheet at any point throughout the lifetime; the average household holds just .5% of their liquid assets in non-savings bonds. Miscellaneous financial assets, such as precious metals, are nearly as unimportant to the average household, accounting for only 3.5 % of their liquid assets. Directly held stocks and mutual funds increase as a percentage of liquid assets throughout most of the lifecycle, with a slight dip for individuals who have been out of school for more than forty-five years. This asset class is the second most important type of liquid assets, accounting for around 15% of the average household's liquid portfolio. This amount is still dwarfed, however, by near-cash assets.

Illiquid assets are similarly dominated by only a few types. Primary residences account for over half of all illiquid holdings; secondary residences and non-residential structures account for another $\approx 8\%$. Further, retirement accounts and annuities

comprise on average nearly 25% of a household’s illiquid assets. Miscellaneous non-financial assets, such as valuable personal collections, are nearly a non-factor for the average household, at less than 1%. Non-retirement accounts, such as CDs or trusts, are the largest asset on average for individuals at the very beginning of their careers. However, for the vast majority of their working life these assets are small compared to retirement accounts and properties. Overall, the average household holds 13% of their illiquid wealth in these assets.

It may seem curious to consider the mean, rather than the median here. This is because, at least for the entire sample, the median holding of all assets other than housing, retirement accounts and checking accounts is zero, or nearly zero. Nevertheless, it may be the case that certain assets, while not important for the median household, may comprise a large percentage of the assets of a minor but substantial subset of households. Therefore, in Tables 3.3 and 3.4 I compute the percentage of households for whom each asset comprises a certain percentage of each asset type.

Table 3.3: % of Households With Minimum Share of Each Liquid Asset

Minimum Share	Cash/Chk/Sav	MM/Call	Stocks	Bonds	Other
1	96.82	21.27	25.9	2.9	10.77
5	90.99	18.71	24.88	2.62	9.62
10	84.24	16.47	23.75	2.26	8.74
25	76.01	12.44	20.57	1.59	6.88
50	68.30	8.10	15.39	.86	4.83
75	62.15	4.78	9.66	.39	3.37
90	58.68	2.55	4.88	.15	2.22
99	56.45	0	.56	.02	1.22


Tables 3.3 and 3.4 do not contradict the implications of Tables 3.1 and 3.2 regarding the relative importance of each asset to households. On the liquid side, it is clear that near-cash assets dominate the portfolios of most households. 97% of households have at least some of their assets as near-cash, and a majority of households hold

Table 3.4: % of Households With Minimum Share of Each Illiquid Asset

Minimum Share	Residence	Oth. Prop.	Ret.	Oth. Fin	Oth. Nonfin
1	79.37	21.7	57.73	46.39	8.72
5	78.56	19.65	52.00	33.81	6.31
10	76.72	17.56	46.86	26.36	4.61
25	68.78	12.3	35.24	16.34	2.65
50	52.13	6.40	20.84	9.98	1.63
75	36.24	3.03	11.43	7.10	1.15
90	26.29	1.78	8.11	6.023	.94
99	17.83	1.12	6.72	5.46	.79

nearly all of their liquid assets (>99%) as near-cash. Money market and call accounts are less popular; only 21% of households have any funds in these accounts and less than a tenth of households keep the majority of their liquid assets in these accounts. Few households have any exposure to bonds; only 2.9% of households have even 1% of their liquid portfolio as bonds. Miscellaneous liquid assets are fairly insubstantial; a tenth of household own some of these assets, but they only comprise a majority of liquid assets for less than 5% of households. For a subset of the population, stocks play a large role in their liquid portfolio. For 10% of households, stocks are over 75% of their liquid holdings. However, stocks are unimportant for the majority of households, as nearly 75% of household keep less than 1% of their liquid portfolio in equities.

Illiquid portfolios are dominated by primary residences. Nearly 80% of households have some share of their illiquid holdings in their primary residence. The majority of households hold the majority of their illiquid holdings in their home, and 18% of households hold nearly their entire illiquid portfolio in their home. A fifth of households have some equity in other properties, and these properties account for over a quarter of illiquid wealth for a tenth of households. Retirement accounts, too, are important. A majority of households have assets earmarked for retirement, and a fifth of households hold the majority of their illiquid holdings in these accounts.



(a) Average Net Assets

(b) Ratio of Illiquid to Liquid Assets

FIGURE 3.1: Assets Over the Lifecycle

A tenth of households have more than 75% of their illiquid portfolio in retirement accounts. Non-retirement financial instruments are important to households as well. 47% of households have some exposure to these assets and 17% of households have at least a quarter of their illiquid assets in these instruments. Miscellaneous nonfinancial assets are not crucial to understanding household balance sheets. Less than a tenth of households have any of these assets, and less than 3% of households even hold a quarter of their illiquid assets in these objects.

In Figure 3.1, I graph aggregate measures of portfolio composition. I consider data from the 2013 SCF. Because the illiquid asset data are for the most part reported as *net* asset data, I subtract from liquid assets any outstanding debts, such as credit card debt, educational loans, or installment loans. I exclude vehicle debt since I have excluded vehicle worth from each category. Thus, in Figure 3.1, it is best to think about these measures as assets net of any corresponding debts.

In Figure 3.1a, I plot the average holdings of liquid and illiquid assets over the lifecycle. The average holding of illiquid assets is always higher than that of liquid. Further, both series are single peaked at around 45 years of experience, which is slightly after retirement for individuals who begin working after college. Aver-

age liquid assets are actually negative at the beginning of working life, though this measure becomes positive fairly quickly. Illiquid assets are decumulated rapidly in retirement, while average liquid assets decline at a much slower pace.

In Figure 3.1b, I consider the the ratio of illiquid to liquid assets, to get a sense of the distribution of portfolio composition. I plot the ratio for 25th percentile, the 75th percentile, and the median individual at each level of experience. Households tend to being life with little illiquid assets; thus the ratio for each percentile is near zero at the beginning of their career. The 75th percentile steadily rises over the lifecycle, almost linearly. By retirement the 75th percentile has 25 times illiquid net worth as compared to liquid net worth. These ratios are not nearly as large for lower percentiles. The median individual increases his relative illiquid holdings over the lifecycle as well, though by sixty years of experience this ratio only rises to 6. The 25th percentile's ratio actually becomes negative during mid-life, implying negative holdings of one type of asset. Nevertheless, after 60 years of experience the ratio is positive. However, it is less than one, implying that liquid holdings are higher than illiquid holdings.

3.2 Comparison to the Data

Given the very stylized nature of the model, it is precarious to compare it too seriously with the data. However, it is useful to examine the data to understand the relative successes of the model. I again use data from the 2013 SCF. First, consider average holdings of each asset over the lifecycle. It is clear that, on average, the benchmark model generates too little liquid holdings relative to illiquid holdings. Consider the peaks of the average holdings for each asset in each model. The benchmark model implies the average individual's maximum holdings of the illiquid asset is 10.85 times greater than their maximum holdings of the liquid asset. The ratio for the learning model is 2.30, which is much closer to the empirical ratio of 1.82. Further, if one

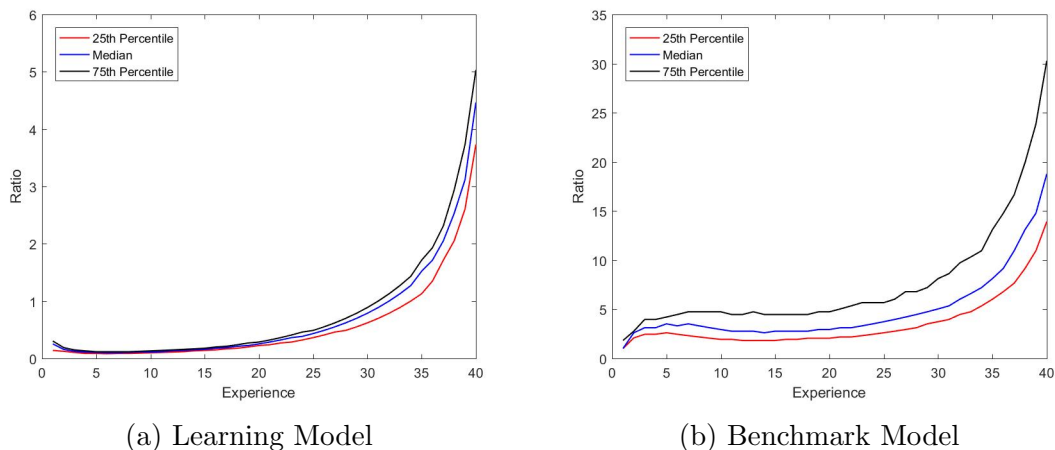


FIGURE 3.2: Ratio of Illiquid to Liquid Savings

considers the ratio of holdings of illiquid to liquid assets, averaging over the entire lifecycle, it again becomes clear that the benchmark understates holdings of the liquid asset. In order to prevent demographics from biasing the ratio, in the data I first take average holdings of each asset by experience level. I then average all of the experience levels together (in this way each level of experience is weighted the same, regardless of the distribution of the population over experience levels), and then take the ratio. In the data, the average experience level holds 2.63 more times illiquid assets than liquid assets. The ratio for the benchmark model is 9.08, while the learning model comes much closer at 1.66.³⁴ Neither model does terribly well matching the shapes of average holdings (though they do capture the single-peakedness of both types of

³ Attempting similar comparisons with illiquid holdings and income yield ratios that are way too high in both models. This is partially due to the models having no income in retirement, as well as rates of return that are potentially too high (see Kaplan and Violante (2014)).

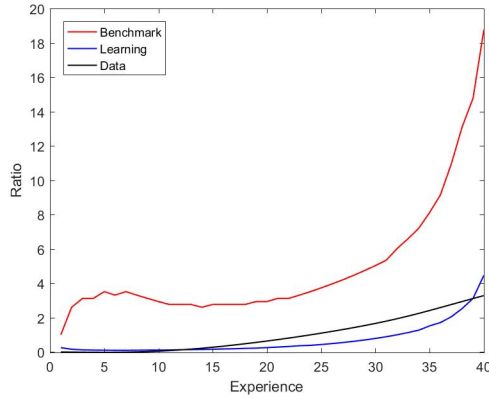
⁴ Since liquid assets mechanically fall to zero in both models, it may be of some use to compare the ratios over the working life. The ratios only considering the first 43 periods in the model—those corresponding to working periods—fall to .98 and 5.75 for the learning and benchmark model, respectively. Considering the same levels of experience in the data yields a ratio of 2.81. While this measure now seems to favor neither model, it is worth remembering that the full irreversibility of illiquid assets lowers the benefit of these assets. A more realistic adjustment cost would make illiquid assets more desirable, presumably increasing both model ratios. In this case the benchmark ratio would drift further from the data.

assets).


One object of particular interest is the proportion of assets that individuals choose to hold. Suppose we were calibrating the model to match the total amount of savings. If the composition of those savings depended on the underlying income process, then just by changing the income process we could potentially have different conclusions for the number of liquidity constrained individuals, the relative amount of housing services, or some other metric of interest, even keeping overall savings for each household fixed. I show in Figure 3.2 that the entire distribution of portfolio composition depends on the income specification. Note that I only plot up until period 40, since the ratio explodes as the amount of liquid assets mechanically approaches zero.

In Figure 3.2a I plot the 25th percentile, the median, and the 75th percentile of the ratio of illiquid to liquid assets for the learning model. In Figure 3.2b I plot the same statistics for the benchmark. In both figures this ratio is fairly stable over the lifetime until individuals start decumulating liquid savings as they approach retirement. The main difference is the size of the ratios. For the learning model, the ratio stays at nearly zero until mid-way through the working lifetime. The median individual only holds as much illiquid as liquid assets at 30 years into his working life. This is in contrast to the median benchmark individual, who holds around 3 times as much illiquid as liquid assets throughout most of his lifetime. It is worth noting that these distributions are fairly tight. The 75th percentile of the learning model is below the 25th percentile for the benchmark model throughout the entire lifecycle.

In Figure 3.3, I compare the median ratio throughout the lifecycle found in the data to its model counterparts. Figure 3.3a considers both the benchmark and learning model, while 3.3b ignores the benchmark model to zoom in on the other two series. The benchmark model implies that the median individual holds too many illiquid assets relative to liquid assets throughout the lifecycle. In fact, early in the



(a) Both Models Compared to Data



(b) Learning and Data Only

FIGURE 3.3: Median Ratio Over the Lifecycle

lifecycle the benchmark model suggests that this ratio is nearly 4, whereas in the data the median individual holds little to no illiquid assets at all, relative to the liquid asset. The learning model predicts positive illiquid assets for the median individual early in the lifecycle, but this amount is very low relative to liquid assets with a ratio below .5. The data and learning model match well at the end of the lifecycle as well. Where the model does more poorly is mid-age, from around period 10 to 35. Here the curvature is much higher in the learning model: the ratio remains very low until around experience level 20 before slowly increasing in a convex manner. In the data, this ratio increases almost linearly beginning at age 10.

To guard against these results being driven purely by time effects, in the Appendix I plot this series for multiple years. Though this series does have some between year variation, the major conclusions remain the same. Over the entire lifecycle, the benchmark model predicts that the median individual's illiquid holdings are too high relative to his liquid holdings and the learning model does a good job matching the ratio at the beginning and towards the end of working life. The main differences between years in the data are the point in the lifecycle when the ratio begins to increase and size of the ratio over the middle of the lifecycle. For years other

than 2013, this increase begins closer to 5 years of experience than 10. Moreover, the ratio is higher in the middle of the median individual’s life, which makes the difference between the learning model and the data more pronounced and diminishes the difference between the benchmark model and the data. Depending on the year, the benchmark ratio may be closer than the learning ratio for some periods in the middle of the lifecycle. For instance, the benchmark model is only slightly higher than the 2007 ratio for periods surrounding 20 years of experience.

Both models have difficulty matching the entire distribution. As mentioned, both models feature a very tight distribution, which is not the case in the data where the asset choices are more diffuse across percentiles. Thus, we could say that the benchmark model does better “matching” the asset choices of the higher percentiles in the data, while the learning model does better with the lower percentiles, although these matches get worse as the percentiles get further away from the median.⁵ In the case of the learning model, this deterioration of versimilitude at lower percentiles may be partially mechanical: lower percentiles actually have negative ratios throughout the lifecycle, which is not possible in either model.

3.3 Implications

Clearly, the choice of income process has stark implications for household portfolio decisions over the lifecycle. This could potentially impact the efficacy of fiscal policy. An example of this would be a temporary fiscal stimulus, such as a one-off payment to households. Kaplan and Violante (2014) show that such a one-off payment can generate large consumption changes in a two asset model with liquid and illiquid assets. This is due to the presence of what is termed the “wealthy hand-to-mouth”; that is, individuals who hold positive net worth, yet hold most of it in accounts that

⁵ It is worth noting that the 75th percentile of the benchmark model does an acceptable job matching the corresponding series in the data

are difficult to adjust. These individuals cannot consume out of savings without incurring a fixed cost, and therefore may choose not to dissave even if they would like to increase consumption. These individuals will tend to consume a large portion of a transfer from the government.

This result depends crucially on the number of individuals who are constrained; in other words, the number of individuals who choose to hold low levels of liquid assets. What this paper shows is that switching from the standard approach of modelling income, as was done in Kaplan and Violante (2014), to a plausible alternative can generate higher proportion of liquid savings over the entire lifetime for individuals throughout the distribution. This implies that individuals should be less constrained, which would diminish any consumption response to a fiscal transfer.

To show the affect this alternative specification has, I calculate the marginal propensity to consume out of extra income for individuals in both the benchmark and learning models. The marginal propensity to consume is defined as the change in consumption resulting from a concomitant change in income; in other words $\frac{\Delta C_t}{\Delta Y_t}$ for some period t . This information can be gathered by solving the model under the original deterministic sequence of income and then once more with a new deterministic income sequence, where income for period t has been raised. This generates decision rules for all periods under both scenarios.

To ensure that this extra income is truly unanticipated, one can then take the 1 to $(t - 1)$ decision rules from the original sequence and append the t to T decision rules from the new sequence. Thus, even though the individual receives a different amount of income in period t , they act as if they are going to receive the typical amount of income in that period until it is actually received—so they act as if truly is a “surprise.” If one then uses the same generated stochastic income draws then any simulations will only differ by the extra income gained in period t . This makes it easy to back out the change in consumption that results from this change in income,

as all other factors, such as changes to expectations, have been held fixed.

Because of the discrete nature of the savings grids, individuals may differ in their MPC for reasons unrelated to economic fundamentals. That is, some individuals may want to adjust a little already, but cannot due to the discrete nature of the environment making their only possible choice either a large adjustment or none at all. It will take these individuals less extra income to change their consumption than an individual who would not want to change their consumption even with a finer grid. Thus, it makes sense to consider the average response of individuals to an income increase.

I consider the average response of 1000 individuals at 20 years of experience to a 5% increase in their income, relative to what their income is in the baseline scenario. Since I generated the same income draws for both models, the $\Delta Y_{20} = 56$ in both scenarios. In the benchmark case, the $\Delta C_{20} = 4.64$, while the learning case has a much lower change in consumption of $\Delta C_{20} = 2.61$. In other words, the benchmark MPC for these individuals is 8.3%, while the learning MPC is 4.7%, or a little more than half. To put this another way, the MPC in the learning case is 56.6% of that of the benchmark case, which is a fairly significant difference. Note that some of this change is almost surely due to differences in individual expectations and uncertainty between the two types of individuals, and not their asset holdings per se. Nevertheless, allowing for the alternative learning process has a large impact on consumption responses to exogenous income, at least in this setting.

Even if one calibrates parameters to match the total number of hand-to-mouth individuals, this still implies that other parameters of the calibration, such as the coefficient of relative risk aversion, must adjust in response. To understand the magnitude of adjustment necessary, I examined the effect of risk aversion in the benchmark model on a related measure: the ratio of average illiquid to liquid assets over the lifetime. I used a combination of bisection and secant methods to find the

measure of risk aversion required in the benchmark model to match this ratio in the learning model with the baseline level of risk aversion, $\sigma = 2$. The value of risk aversion that generated this much lower ratio was $\sigma = 11.9$, which is clearly a significant difference.

Finally, there are potentially other policy implications if one thinks about individual assets, rather than the broad measures used in this paper. For example, suppose one was interested in housing policy, specifically on regulations that tighten or relax down payment constraints. In this environment we can consider the illiquid asset to be housing, which is not a terrible simplification given how much housing dominates household balance sheets, and down payments as a fixed cost, γ_F , that only applies when $b_{t+1} > b_{min}$ and $b_t = 0$, for some minimum housing size, b_{min} .

Consider the benchmark model, where the typical individual (both average and median) holds little liquid assets. Introducing a non-zero downpayment should cause many individuals to have to defer illiquid purchases, as they will not have enough saved up to pay for the fixed cost. This may be less of an issue in the learning model, as the typical individual holds more liquid assets and chooses to invest relatively less in the illiquid asset already, due to the mechanisms outlined above. This same reasoning applies to changes in the downpayment. In an environment with a downpayment, benchmark agents will increase their liquid savings such that they can afford the downpayment, plus have a little left over. Learning individuals will have a large excess of accessible savings. Thus, for a small increase in γ_F , many more benchmark agents should be priced out of making an immediate illiquid purchase, likely affecting overall savings decisions more in the benchmark model than in the learning model. Similar policy experiments, such as adjusting government subsidies to retirement savings, may yield different results depending on the assumed income process as well.

City-Level Housing and Employment Dynamics

4.1 The Model

4.1.1 Households

Figure 4.1 gives a visual depiction of the environment¹. Individuals can either be a resident of the world or of the city. While in the city, individuals are either employed or unemployed and actively looking for a job. Further, individuals can either rent or own a home. Upon arrival in the city, all individuals initially rent. A fraction ψ actively search for housing while the rest $(1 - \psi)$ decide to rent permanently. Job-seekers match with an employer at a rate of μ_t , while individuals searching for housing succeed at rate λ_t . Individuals exogenously separate from both their job and their home at rates δ and φ , respectively.

Individuals can search for, match with, and separate from, housing and employment simultaneously. Matching and separation occur independently, and therefore the probability of finding both a job and home simultaneously is given by $\lambda_t\mu_t$, while

¹ This chapter uses the Stata program “pvar2” to implement a panel VAR. This program was written by Ryan Decker, who adapted the program “pvar”, which was introduced with Love and Zicchino (2006) and further developed and documented by Abrigo and Love (2016).




FIGURE 4.1: Movement Between Nodes

the rate of simultaneous separation is given by $\varphi\delta$. Individuals who are searching for a home will remain in the city until they find housing. Current homeowners may exogenously decide to relocate to the rest of the world with probability π_n , while perpetual renters may move away with probability π_f . These probabilities are invariant to the employment status of the individual.

Thus, once an individual moves to a city they can be at one of six states of existence. Individuals who permanently rent may either be unemployed, where they receive a valuation of W_t^f , or they may be employed, in which case they receive a valuation of D_t^f . Similarly, individuals who are searching for a home may be unemployed and draw a value of W_t^U or be employed with value D_t . Employed homeowners receive a value of J_t while unemployed homeowners have a value of J_t^U . Moving to the city from the rest of the world carries a value of \bar{W}_t .

Individuals in the rest of the world have an idiosyncratic value of living outside of the city, ε . These values are distributed with a cumulative density function of G , and a support of $[0, \bar{\varepsilon}]$. Let e_t^C denote the value at which individuals are indifferent

between living in the city or the rest of the world. It immediately follows that individuals with valuations below e_t^C will choose to move into the city, and that the proportion of movers is given by $G(e_t^C)$. The world population is given by Q_t , and grows exogenously at a rate of Δ_P . I assume that only new agents have the option of moving to the city; in other words, only $\Delta_P Q_t$ of individuals are eligible to move. This is done to obviate the necessity of keeping track of the proportion of the population that is currently outside of the city as an additional state variable. The time-invariant average value of living outside of the city is Z .

It is important to keep track of the mass of individuals in each state within the city. Let F_t^U and F_t be the number of unemployed and employed perpetual renters, respectively; B_t^U and B_t be the number of unemployed and employed renters looking for housing, respectively; and N_t^U and N_t be the number of unemployed and employed homeowners, respectively. In any period t , $G(e_t^C)\Delta_P Q_t$ individuals decide to relocate within the city. These individuals arrive in period $t + 1$, so B_{t+1}^U incorporates ψ of them while $(1 - \psi)$ are added to F_{t+1}^U . In any period, only $(1 - \pi^n)$ of the N_t^U and N_t individuals remain inside the city by the following period; similarly only $(1 - \pi^f)$ of the F_t^U and F_t individuals remain by time $t+1$. Movement between “nodes” within the city is dictated by the aforementioned matching and separation probabilities. I leave the explicit expressions of the laws of motion for the population measures to Appendix B.

The value of being at each state within the city can be written explicitly. The value of moving to the city is the expected value individuals will receive.

$$\bar{W}_t = \psi W_t + (1 - \psi) W_t^f \quad (4.1)$$

Each period renters pay rent, r_t , which adjusts due to market conditions. Homeowners have a constant level of maintenance they must upkeep each period, m , and receive value from homeownership z^h . Employed individuals receive a wage, which

they bargain for each period. Because perpetual renters, seeking renters, and homeowners face very different living situations, their bargained wages potentially differ and are denoted w_t^f , w_t^d , and w_t^j respectively. Unemployed individuals receive constant benefits b . Individual valuations are a combination of current factors as well as the expected valuation tomorrow, which is discounted by β .

Perpetual renter valuations are fairly simple. Today renters either receive a wage w_t^f or unemployment benefits b , pay rent r_t , and then consider their continuing valuation. Renters know there is a π^f chance they relocate elsewhere next period and receive valuation Z and a $(1 - \pi^f)$ chance of remaining in the city where they will either be employed or unemployed. Thus, these valuations can be written as follows:

$$W_t^f = -r_t + b + \beta[\pi^f Z + (1 - \pi^f)(\mu_t D_{t+1}^f + (1 - \mu_t)W_{t+1}^f)] \quad (4.2)$$

$$D_t^f = -r_t + w_t^f + \beta[\pi^f Z + (1 - \pi^f)((1 - \delta)D_{t+1}^f + \delta W_{t+1}^f)] \quad (4.3)$$

The continuing valuation for home seekers is a little more complicated. Again, these individuals rent and therefore must pay the market rent of r_t and either receive a wage w_t^d or unemployment benefits, depending on their current labor market status. Unemployed home seekers have a λ_t probability of matching with a home seller and a μ_t chance of matching with an employer. Employed individuals have a δ probability of separating from their employer. An individual who is successful in finding a home will not move in until the following period. However, they will negotiate and pay the price of the home in the current period. As I will discuss later on, currently employed and unemployed individuals do not necessarily negotiate for the same price.

$$W_t = -r_t + b + \beta[(1 - \mu_t)(1 - \lambda_t)W_{t+1} + \mu_t(1 - \lambda_t)D_{t+1} + (1 - \mu_t)\lambda_t(J_{t+1}^U - P_t^U) + \mu_t\lambda_t(J_{t+1} - P_t^U)] \quad (4.4)$$

$$D_t = -r_t + w_t^d + \beta[\delta(1 - \lambda_t)W_{t+1} + (1 - \delta)(1 - \lambda_t)D_{t+1} + \delta\lambda_t(J_{t+1}^U - P_t^E) + (1 - \delta)\lambda_t(J_{t+1} - P_t^E)] \quad (4.5)$$

Homeowners receive some value, z^h , from occupying their home. This can be thought of as anything from being able to make personal modifications not allowed in a rental unit to the satisfaction of being a homeowner and taking part in the “American Dream.” The individuals are also responsible for routine maintenance on their home, m , and either receive a wage, w_t^j or unemployment benefits. There is a π_n chance that they will need to relocate out of the city. In this case in the next period they will receive the average world value, Z , as well as value from their unoccupied home, \tilde{V}_{t+1} . The rest of the time they will need to relocate within the city with probability φ . This means that they will need to rent in the following period, though they will still have an unoccupied structure that they can derive value from, as well. Job finding and separation rates are as described above.

$$J_t^U = z^h - m + b + \beta[\pi^n(Z + \tilde{V}_{t+1}) + (1 - \pi^n)((1 - \mu_t)\varphi(W_{t+1} + \tilde{V}_{t+1}) + \mu_t\varphi(D_{t+1} + \tilde{V}_{t+1}) + (1 - \mu_t)(1 - \varphi)J_{t+1}^U + \mu_t(1 - \varphi)J_{t+1})] \quad (4.6)$$

$$J_t = z^h - m + w_t^j + \beta[\pi^n(Z + \tilde{V}_{t+1}) + (1 - \pi^n)(\delta\varphi(W_{t+1} + \tilde{V}_{t+1}) + (1 - \delta)\varphi(D_{t+1} + \tilde{V}_{t+1}) + \delta(1 - \varphi)J_{t+1}^U + (1 - \delta)(1 - \varphi)J_{t+1})] \quad (4.7)$$

The number of matches made in the labor market in a given period is $M^e(u_t, v_t)$, where u_t is the number of unemployed individuals and v_t is the number of vacancies. Given an equal probability of each potential match happening, $\mu_t = \frac{M^e(u_t, v_t)}{u_t}$ denotes the probability of an individual matching with a firm while $q_t = \frac{M^e(u_t, v_t)}{v_t}$ is the probability of a firm matching with an individual. Similarly, the number of housing matches is given by $M^h(B_t^U, B_t, S_t)$. The probability of a home buyer finding a

home is $\lambda_t = \frac{M^h(B_t^U, B_t, S_t)}{B_t}$ while the probability of a home seller finding a buyer is $\gamma_t = \frac{M^h(B_t^U, B_t, S_t)}{S_t}$

4.1.2 Firms

Firms employ workers and produce output. A matched firm is able to produce 1 unit of output per period, while an unmatched firm produces nothing. Unmatched firms must pay a search cost of κ and are randomly matched with a job seeker. Once matched, the firm and the worker attempt to negotiate for a wage. If no agreement can be reached, then the firm and individual go their separate ways and both the firm and individual cease action for the current period.

Recall that the negotiating power of the individual depends on their current situation; therefore there are potentially three negotiating scenarios the firm may find itself in. Let T_t^x be the value of matching with an individual in situation “ x ”, where $x \in X = \{f, d, j\}$, corresponding to perpetual renters, house searching renters, and homeowners, respectively. The firm has no way of knowing which individual it will be matched with before the match actually happens. Thus, they must factor in the population composition when calculating the value of posting a vacancy.

Let R_t be the value of posting a vacancy in time t . The cost of posting is the aforementioned search cost. The benefit of posting is the discounted weighted expected value of making a match. With probability $(1 - q_t)$ no match will be made. In this case, the firm can either post another vacancy with value R_{t+1} , or do nothing and receive a value of 0. With probability q_t a match is made and in expectation the firm receives a weighted value of the possible matches, weighted by the relative number of each type of job seeker.² Thus we can write the value of posting a vacancy

² To be precise, the expected value will be a weighted sum of the maximum of each continuing value, zero, or the value of posting a vacancy the next period. This is because the firm can always walk away from any match and therefore either do nothing the following period or post another vacancy. However, in equilibrium it will always be beneficial for the firm to come to a successful negotiation given a match, so I ignore this subtlety in this section in order to ease notation and

as

$$R_t = -\kappa + \beta[q_t \sum_{x \in X} \Psi_{t+1}^x T_{t+1}^x + (1 - q_t) \max\{R_{t+1}, 0\}] \quad (4.8)$$

where Ψ_t^x is the population weight of each type of individual in period t .

Regardless of the type of individual the firm is matched with, the firm produces one unit of a general output good. This good is sold on the world market and, as the city itself is small relative to the world market, the total number of goods produced in the city has no effect on world prices. Thus, the price can be assumed fixed. For simplicity, I normalize the price to 1. The generalized goods can be costlessly converted to “construction goods”, and can be used to construct houses by construction companies.

Matched firms must also pay the negotiated wage. With a probability of δ they then separate from the worker and face the decision between posting a vacancy or exiting the market in the following period. With a probability of $(1 - \delta)$ they remain matched and renegotiate the wage for the following period. Thus a firm matched with an individual of type x has the following value:

$$T_t^x = 1 - w_t^x + \beta[(1 - \delta)T_{t+1}^x + \delta \max\{R_{t+1}, 0\}] \quad (4.9)$$

4.1.3 Housing

Each period ϱ of the housing stock is lost to depreciation. New housing is added to the city by construction companies. Housing requires one parcel of land and ϕ units of the general good to produce. Current land usage is denoted by L_t , and total available land is given by \bar{L}_t . Total developable land grows exogenously at rate Δ_L . Available land, $\bar{L}_t - L_t$, has outside agricultural and recreational value given by a distribution with a probability density function of Λ . Land is sold to a developer if

make the valuation clearer. I follow this same convention when writing out the valuation of a matched firm.

the price of land, p_t , is higher than the value of its alternative use. These parcels of land are converted to available housing units for the next period. Thus the law of motion of total housing is as follows:

$$H_{t+1} = (1 - \varrho)H_t + (\bar{L}_t - L_t) \int_0^{p_t} \Lambda(x) dx \quad (4.10)$$

Once a home is built, it can be either occupied by a homeowner, or unoccupied. Unoccupied housing can either be rented or sold, and carries a value of \tilde{V}_{t+1} . Houses that are put on the market have a γ_t probability of being sold for the prevailing market price P_t , and a $1 - \gamma_t$ probability of remaining unoccupied in the following period. Therefore, letting V_t be the value of putting a home up for sale, we have:

$$V_t = \gamma_t P_t + (1 - \gamma_t) \beta \tilde{V}_{t+1} \quad (4.11)$$

There are no search frictions with renting, so rentiers earn rental income r_t immediately. However, they must also cover the maintenance costs of the building, m . The rentier also gets value from knowing that the building will remain in their possession in the following period, where they can either rent the building out again or attempt to sell it. Thus, the value of possessing an unoccupied home is the following:

$$\tilde{V}_t = \max[r_t - m + \beta \tilde{V}_{t+1}, V_t] \quad (4.12)$$

4.2 Equilibrium

Because there are no barriers to posting a vacancy, firms will enter (or exit) the market until the value of posting a vacancy equals 0. Thus, in equilibrium the firm's problem simplifies a fair amount. One can rewrite Equation 4.8 as

$$\kappa = q_t \beta \sum_{x \in X} T_{t+1}^x \Psi_{t+1}^x \quad (4.13)$$

Similarly, Equation 4.9 can be simplified as

$$T_t^x = 1 - w_t^x + (1 - \delta)\beta T_{t+1}^x \quad (4.14)$$

In the housing sector, it must be the case that the return on renting or attempting to sell a home will be the same. Using Equations 4.11 and 4.12, the immediate implication is that the value of an unoccupied structure is equal to the value of trying to sell it; that is, $\tilde{V}_t = V_t$. Combining the two equations, we have:

$$V_t = r_t - m + \beta V_{t+1} \quad (4.15)$$

Notice that this also implies the following relationship:

$$r_t - m = \gamma_t(P_t - \beta V_{t+1}) \quad (4.16)$$

Finally, the value to the builder of building a home is βV_{t+1} , while the cost is the price of 1 parcel of land and ϕ units of the general good. Therefore, in equilibrium we must have

$$\beta V_{t+1} = \frac{1}{\phi} + p_t \quad (4.17)$$

4.2.1 Bargaining

Both wages and home prices are set through Nash bargaining. Firms and households bargain over the total surplus in the labor market, while home buyers and home sellers negotiate over total surplus in the housing market. Bargaining power is determined by the outside option each agent has, and an individual's outside option depends on the state they are in. For agents bargaining over wages, the outside option depends on their homeowner status. For agents bargaining over home prices, their outside option depends upon whether or not they are employed. Therefore, as there are three possible home positions— permanent renter, searching renter, and homeowner— there must be three possible wages. Similarly, as agents can be in one of two employment states there must be two equilibrium market prices.

Because there are two negotiations occurring in this model, it is crucial to clarify the timing. Agents enter each period knowing their housing status for that period, having already paid the price of the house in the previous period. At the beginning of the current period agents negotiate their wage for that period. Thus, letting τ^w be the bargaining power of the worker and noting that the outside option of the firm is $R_t = 0$, we have:

$$w_t^x = \operatorname{argmax}(T_t^x)^{1-\tau^w} (X_t - X_t^U)^{\tau^w} \quad (4.18)$$

where x is in the set defined above, $X \in \{D^f, D, J\}$, and X^U is the unemployment equivalent of X . Solving, wages can be stated as:

$$w_t^f = \tau^w + (1 - \tau^w)b + (1 - \tau^w)\beta[(1 - \delta) - (1 - \pi^f)(1 - \delta - \mu_t)](D_{t+1}^f - W_{t+1}^f) \quad (4.19)$$

$$w_t^d = \tau^w + (1 - \tau^w)b + (1 - \tau^w)\beta[(\lambda_t - \delta\lambda_t + \mu_t - \lambda_t\mu_t)(D_{t+1} - W_{t+1}) + (\mu_t\lambda_t + \delta\lambda_t - \lambda_t)(J_{t+1} - J_{t+1}^U) + \lambda_t(P_t^E - P_t^U)] \quad (4.20)$$

$$w_t^j = \tau^w + (1 - \tau^w)b + (1 - \tau^w)\beta[(1 - \pi^n)(\mu_t\varphi + \delta\varphi - \varphi)(D_{t+1} - W_{t+1}) + ((1 - \delta) - (1 - \pi^n)(1 - \delta - \varphi + \delta\varphi - \mu_t + \mu_t\varphi))(J_{t+1} - J_{t+1}^U)] \quad (4.21)$$

For accounting purposes, aggregate wage w is just a weighted summation of the bargained wages. The weights are the percentage of employed agents drawing each wage. For computational simplicity, the wage is re-bargained every period for every agent-firm match.

Note that because firms can only match with one agent per period, it is always beneficial for them to bargain with whomever they are matched. Further, exogenous separation remains a viable imposition on the model, as separation yields a valuation of zero in the following period, while the valuation of a continued match will be positive.

Following any wage negotiations, agents are matched with a home. Though they cannot move in until the following period, agents negotiate in the current period and immediately pay the negotiated price. However, at this point agents are unsure of their employment status in the following period. Thus, their outside options depend partially on the probabilities of being employed in the following period. For home sellers, the value of selling the house is the current price and the outside option is the value of the unoccupied home in the next period. Thus, letting τ^h be the bargaining power of the prospective home buyer, equilibrium housing prices are given by the following two equations:

$$P_t^U = \operatorname{argmax}([\mu_t J_{t+1} + (1 - \mu_t) J_{t+1}^U - P_t^U] - [\mu_t D_{t+1} + (1 - \mu_t) V_{t+1}])^{\tau^h} (P_t^U - V_{t+1})^{1-\tau^h} \quad (4.22)$$

$$P_t^E = \operatorname{argmax}([(1 - \delta) J_{t+1} + \delta J_{t+1}^U - P_t^E] - [(1 - \delta) D_{t+1} + \delta W_{t+1}])^{\tau^h} (P_t^E - V_{t+1})^{1-\tau^h} \quad (4.23)$$

This problem is solved in an identical manner to that of the equilibrium wages. As with wages, the aggregate home price level is a weighted average of home prices across home buyers. The prices that solve the above equations are as follows:

$$P_t^U = (1 - \tau^h) \beta [(\mu_t J_{t+1} + (1 - \mu_t) J_{t+1}^U) - (\mu_t D_{t+1} + (1 - \mu_t) W_{t+1})] + \beta \tau^h V_{t+1} \quad (4.24)$$

$$P_t^E = (1 - \tau^h) \beta [((1 - \delta) J_{t+1} + \delta J_{t+1}^U) - ((1 - \delta) D_{t+1} + \delta W_{t+1})] + \beta \tau^h V_{t+1} \quad (4.25)$$

Finally, the period ends with agents matching with and separating from firms, determining their employment status for the following period.

4.2.2 Summarizing the Equilibrium

The dual search equilibrium of this model is a collection of values of the overall state, s_t . To achieve a stationary equilibrium, it is necessary to first normalize

the variables by total population. The population measures $F_t^U, F_t, B_t^U, B_t, N_t^U, N_t$ satisfy the normalized versions of the laws of motion of Equations B.7 - B.12. The values of each of the homeowner and job status nodes $W_t^f, D_t^f, W_t, D_t, J_t^U, J_t$ evolve according to 4.2 - 4.7. The entry value of the city, \bar{W}_t is determined by 4.1, and population growth is such that $\bar{W}_t = e_t^C$.

Wages w_t^f, w_t^d, w_t^j solve the Nash bargaining problems 4.19 - 4.21. Home prices P_t^U, P_t^E solve the Nash bargaining problems in 4.24 and 4.25. The value of a vacant home is determined by 4.15, while rent is determined by equation 4.16.

The value to a firm of being matched depends on the negotiated wage and is given by 4.14. In turn, the number of vacancies adjusts such that 4.13 clears. The price of land adjusts according to 4.17 and the level of housing evolves according to the normalized version of the law of motion in Equation B.13. The measure of unoccupied housing adjusts so that

$$S_t = H_{t-1} - F_t^U - F_t - B_t^U - B_t - N_t^U - N_t \quad (4.26)$$

Finally, the probability of an individual finding a match in the labor and housing market are given by Equations 4.27 and 4.29, respectively.

Note that much of the basic structure of this model is owed to Head et al. (2014). That paper has a similar city-world set-up, where moving away from the city is exogenous but entry is endogenous. In that model, as in mine, individuals are exogenously separated into permanent and temporary renters, and the housing market operates in the same way it does in this paper. Further, the home construction sector operates in essentially the same fashion in both papers. The main difference comes from the existence of a labor market and matching with general goods producing firms in this paper. This allows me to look at employment rates. Further, both papers handle income differently. In my model, individuals receive income from endogenous wages determined through matching with a general goods producing firm; in their

model individuals receive an exogenous income stream as well as endogenous wages from working in the construction sector. The labor market is modelled after a fairly standard version of the Diamond-Mortensen-Pissarides model.³

4.3 Calibration

4.3.1 Functional Forms

I follow the search-and-matching literature in assuming the following functional form for the number of matches in the labor market: $M^e(u_t, v_t) = \chi^e v_t^\nu u_t^{1-\nu}$. If one denotes $\theta_t = \frac{v_t}{u_t}$ as the tightness of the labor market, then the matching probabilities can conveniently be written as

$$\mu_t = \chi^e \theta_t^\nu \tag{4.27}$$

$$q_t = \chi^e \theta_t^{\nu-1} \tag{4.28}$$

One can perform a similar exercise in the housing market. Here I take the functional form from Head et al. (2014), where $M^h(B_t, B_t^U, S_t) = \chi^h (B_t + B_t^U)^\nu S_t^{1-\nu}$. Denoting tightness in the housing market as $\vartheta_t = \frac{B_t^U + B_t}{S_t}$, then one can rewrite the matching probabilities for the housing market, as well:

$$\lambda_t = \chi^h \vartheta_t^{\nu-1} \tag{4.29}$$

$$\gamma_t = \chi^h \vartheta_t^\nu \tag{4.30}$$

In subsequent sections, I compare the impulse responses of the model to that of the data following a labor demand shocks. In the data, this is taken as an exogenous shock to the unemployment rate. This cannot be done in the theoretic model, as there are numerous variables that factor into the unemployment rate. The most practical source of exogenous variation in the unemployment rate comes from the labor matching function constant, χ^e , as this parameter only directly affects the

³ See Mortensen (2011)

number of employed individuals. I model the exogenous process as a simple AR(1) with a drift:

$$\chi_t^e = \bar{\chi}^e + \rho\chi_t^e + \sigma^e e \quad (4.31)$$

where $e \sim N(0, 1)$ and $\bar{\chi}^e$ is the steady state value of χ^e .

I follow the Head et al. (2014) convention of setting $G(e) = (\frac{e}{\bar{e}})^{\alpha_e}$. Further, I make the extra assumption that the steady state value of the city, \bar{W} , represents the median value of the distribution. This is a simplifying assumption that allows me to rewrite $G(e) = \frac{1}{2}(\frac{e}{\bar{W}})^{\alpha_e}$, rendering the value of \bar{e} irrelevant. I also set the value of the rest of the world such that $Z = \bar{W}$. Therefore, in the steady state the city is just as appealing as the rest of the world. I make a similar choice in the functional form of Λ , such that $\int_0^{p_t} \Lambda(x)dx = \frac{1}{2}(\frac{p_t}{\bar{p}})^{\alpha_e}$, where \bar{p} is the steady state value of p_t .

4.3.2 Parameter Values

The model is tuned to a monthly frequency, and the parameters are calibrated accordingly. β is set to .997 to match an annual real interest rate of 3.6%. The rest of the parameters are set to match statistics found in either Shimer (2005) or Head et al. (2014).

δ is set to .0333 to match the Shimer (2005) analysis of a quarterly separation rate of .1. I also adopt his values for the elasticity of the matching function, $v = .72$, vacancy cost, $\kappa = .324$, and household bargaining power, $\tau^w = .72$. The steady state value of the labor matching coefficient, $\bar{\chi}^e$, is set such that $\bar{\mu} = .45$, to match the value presented that paper.

$\pi^f = .01$ and $\pi^n = .002666$ are set to match the annual mobility of renters (12%) and homeowners (3.2%), respectively. As 60% of moving homeowners remain withing their original county, I set φ such that $\frac{\varphi}{\pi^n + \varphi} = .6$. ϕ is set to .1 to match the ratio of housing permits to the number of construction employees, as in the model it only takes one employee to create one general good. I set $\Delta_P = .001$ to

Table 4.1: Parameter Values

Parameter	Value
β	.997
δ	.0333
κ	.324
τ^w	.72
$\bar{\chi}^e$.6102
π^f	.01
π^n	.002666
φ	.00405
ϕ	.1
Δ_P	.001
Δ_L	.001
α_e	7.3
ν	.86
ϱ	0
α_c	876.75
ψ	.4487
m	.0235
b	.3902
z^h	.0517
τ^h	.1238

set an annual population growth rate of 1.2%; Δ_L is set to the same rate to ensure model stationarity. These statistics are found in Head et al. (2014) and I use their parameter values for $\alpha_e = 7.3$ and $\nu = .86$. α_c is chosen using a similar methodology as Head et al. (2014) used to determine α_e . For simplicity, ϱ is set to 0.

The remainder of the parameters, ψ, b, m, z^h, τ^h are jointly determined. To achieve identification, I include additional moment conditions besides those implied by the equations outlined in the body of the paper. Unemployment compensation, b , is set so in the steady state it is equal to $.397w$, similar to the figure in Shimer (2005) of a quarterly unemployment compensation of 40% of total income. I also target a steady state of renters equal to the national average fraction of renting households of 32%, as documented in Head et al. (2014). That paper also documents an average rental rate equal to 13.7% of income. Thus I target a steady state real interest rate,

$r = .137w$. Further, I target their finding that home prices are 38.4 times average income.

Following Head et al. (2014), one can pin down a value for z^h by considering the household's maintenance problem more carefully. Assuming that housing utility, z_t^h , is a function of housing quality, a_t , then one can show that if the household chooses its level of maintenance m_t optimally then $z_t^h \approx (1 + \frac{1-\beta}{\beta d})m_t$. This is shown in more detail in Appendix B. Note that d here corresponds to depreciation of the individual structure and is chosen to be .00025, for an annual depreciation rate of 3% matching the Harding et al. (2007) estimated value for the median aged home. Since m_t is assumed to be constant in all periods, it follows that z_t^h is constant as well.

4.4 Empirical Model

In order to gauge the success of the theoretic model, I compare the impulse responses generated from a labor demand shock in the model to the empirical responses found in the data. I consider the unemployment rate to be a useful proxy for labor demand, and consider the effects of a shock to the unemployment rate on per capita income, population growth, home prices, the construction rate, and rent. Data is taken from a data set constructed by Head et al. (2014), which I briefly describe in Appendix B.

Because I am interested in city-level responses, it is necessary to conduct a panel analysis. To that end, I use the “pvar2” Stata package, written by Ryan Decker⁴ and adapted from the “pvar” package introduced in Love and Zicchino (2006) and further developed and documented in Abrigo and Love (2016), to compute the desired impulse responses. This package implements the system GMM estimator developed by Arellano and Bover (1995) and Blundell and Bond (1998). The panel

⁴ Program can be found at <http://econweb.umd.edu/decker/code.html>

vector autoregression (VAR) can be written as follows:

$$BX_{c,t} = \sum_{i=1}^T A_i X_{c,t-1} + F_c + e_{c,t} \quad (4.32)$$

$X_{c,t}$ is a vector of the endogenous variables for city c at time t , F_c is a vector of city fixed effects, and $e_{c,t}$ is a vector of error terms for each city at time t . Here, $X_{c,t} = [ur_{c,t} \quad y_{c,t} \quad g_{c,t}^{pop} \quad p_{c,t} \quad g_{c,t}^h \quad r_{c,t}]'$, where the rows are the unemployment rate, per capita income, the population growth rate, house prices, the construction rate, and rent. The vector of structural shocks, $e_{c,t}$, is set accordingly.

Identification requires orthogonalized shocks, wherein the variables earlier in the ordering affect latter variables contemporaneously, but not vice versa. It is plausible that this assumption is justified based on the ordering in the model. Rent is a market clearing price that is driven by the current supply of buildings and demand for livable space within a city. Therefore it is a function of the overall population and construction of residences. One may argue that low rent would induce individuals to move into a city, thus affecting population growth. However, this population growth would not be contemporaneous, as individuals would make their moving decisions today and relocate by the following period. This would affect the next period's population growth, which is not ruled out by the identification strategy.

Similarly, construction rates are influenced today by housing prices, as the market clearing housing price determines the benefit of constructing an additional house. New housing units may affect the market clearing price of housing in the following period, of course, but since the units will not go on the market today the price today will be unaffected. The market clearing house price is partially a result of overall population, which is affected by the current growth rate of the population. Population growth is driven by opportunities within the city. This may reflect expectations of housing opportunities (driven by previous housing and rental prices),

but more importantly it reflects current earning opportunities, which is measured by per capita income. Again, one may make the argument that population growth can affect income, but this effect would come in subsequent periods, not contemporaneously. Earning opportunities today are driven by labor demand, which is proxied here by the unemployment rate.

This model outlined above is nearly identical to the one used by Head et al. (2014), with the crucial exception that I also include the unemployment rate in my analysis. As with that paper, I let the number of lags equal two.

4.5 Results

4.5.1 Empirical

The results of the empirical exercise are displayed in Figure 4.2. This figure depicts the impulse responses of the unemployment rate, per capita income, house prices, the housing construction rate, population growth, and rent to a negative labor demand shock. This labor demand shock is considered to be a positive one standard deviation shock unemployment rate. Responses are given in terms of percent deviations from the mean. Confidence intervals are depicted by dotted lines. The standard errors are fairly large, so much so that it is fairly difficult to interpret too much at longer horizons. However, there is enough precision to understand qualitatively the general path that each series takes following the shock in question.

The first series, the unemployment rate, decreases following the initial shock. After six years, one is unable to reject that the shock has fully dissipated. Interestingly enough, the unemployment rate continues to decline after reaching the steady state and for a brief period around 12 years after the labor demand shock one can reject that the unemployment rate is equal to its steady state value. After this time the unemployment rate increases back to its steady state value.

Per capita income immediately falls following a labor demand shock, and stays

persistently below its steady state value. This is consistent with a worsening labor market decreasing the bargaining power of workers. The population growth rate also immediately falls, as the city has become much less attractive to move into and more attractive to move out of. Home prices immediately fall as well, and remain persistently below their steady state value for quite a while.

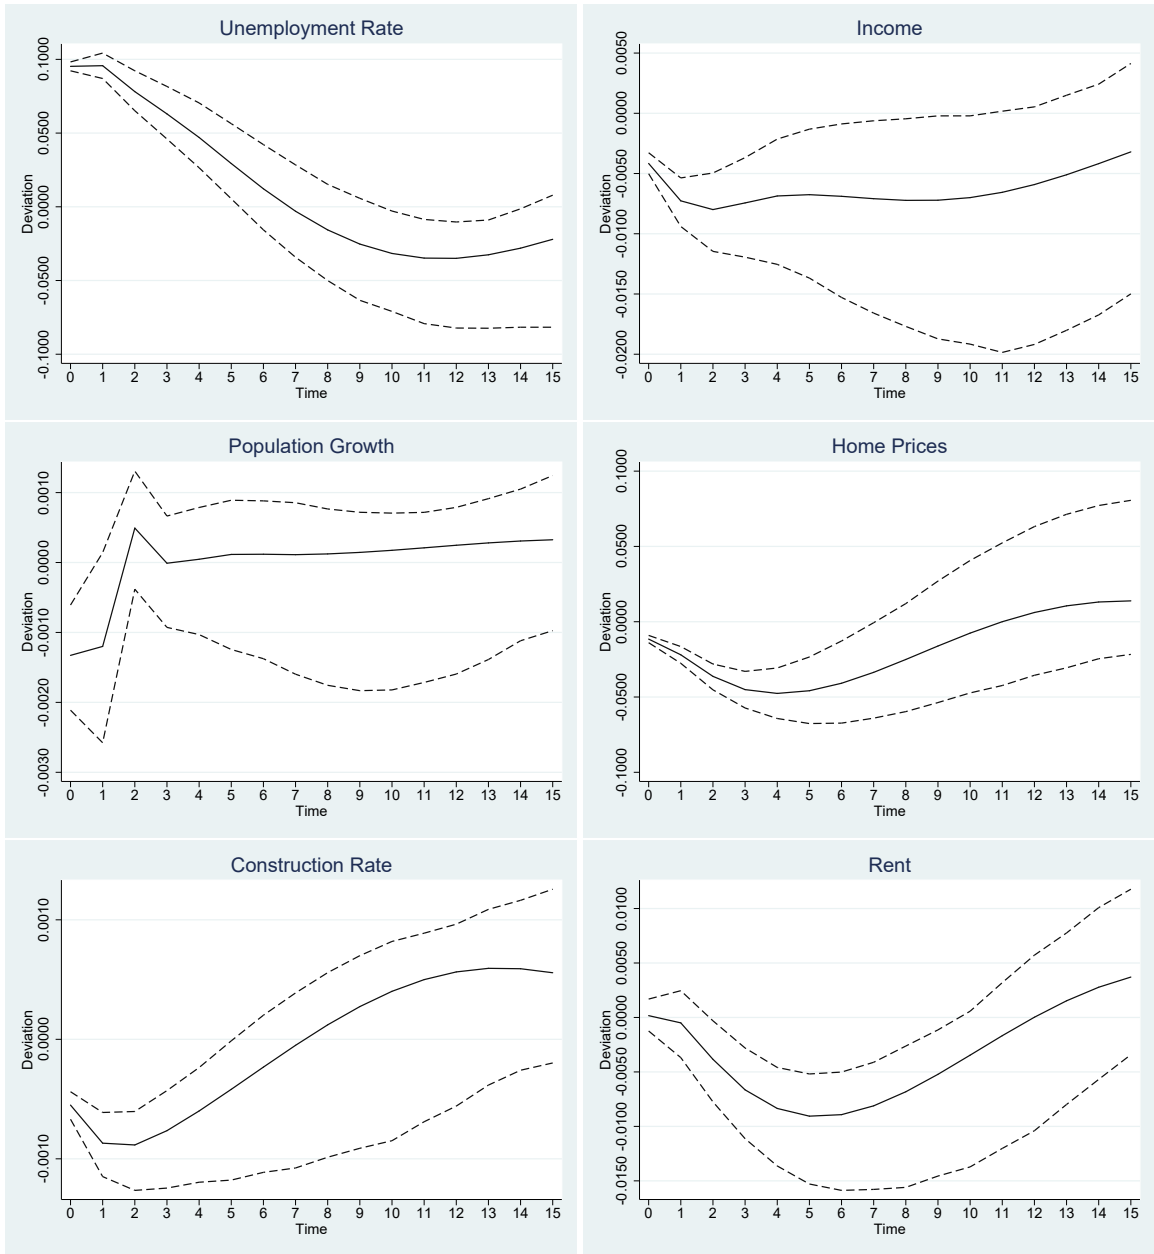


FIGURE 4.2: Data: Impulse Response Functions

This fall in home prices may be due to two primary factors. First, home demand for individuals within the cities fall, as individuals are less willing to pay high prices to purchase a home amidst declining labor market prospects. Further, as population growth falls, the demand for housing falls with it. Note that there may be feedback effects at play here, as well. Though population growth declines initially, it quickly recovers and cannot be distinguished from its steady state value after the first year. This may partially be due to declining home prices making it less costly to move into the city and inducing individuals to move despite the declining job market to take advantage of the temporarily depressed home prices.

Finally, both construction rates and rents decline. Construction rates immediately fall, and take some time to recover. This is most likely due to the fall in home prices, making it much less beneficially to construct additional homes in the near term. It is interesting to note that, while rents decline over the medium term, it takes some time before they begin to fall. This may potentially be due to the demand for rental units increasing relative to housing resulting from a decrease in labor demand. If renting is seen as a cheaper, or easier (due to the lack of a down payment requirement), alternative to owning a home, then declining labor market prospects may induce individuals to switch from home owning to renting. This would keep average rents up within the city, at least initially.

4.5.2 Theoretic

The impulse responses from the theoretic model were generated from an indirect shock to the unemployment rate via a direct shock to the matching coefficient. Dynare was used to solve the model and generate the impulse responses— documentation of this program is provided by Adjemian et al. (2011). As with the impulse responses in the previous section, all of the responses here measure percentage deviation from the steady state. As the model was calibrated to monthly

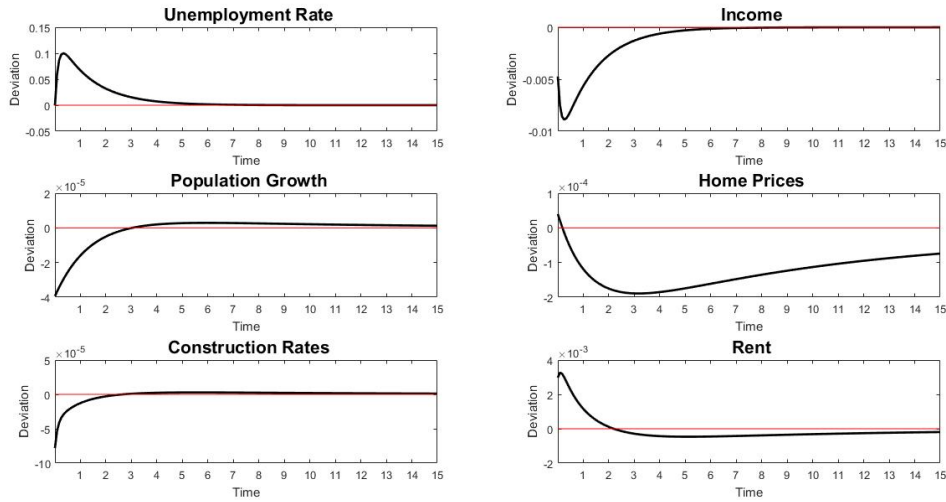


FIGURE 4.3: Theoretic Impulse Responses

statistics, the model responses are plotted out along one hundred and eighty periods to match the fifteen year length of the VAR responses. The tick marks of the x-axis of the responses of Figure 4.3 display the 12-period “yearly” intervals. Finally, note that the per capita income measurement in the model captures both wages and unemployment benefits.

Parameters were chosen so that the shock to the unemployment rate mimicked that of the empirical impulse response. The magnitude of the initial shock was chosen so that the maximum deviation of the unemployment rate from the steady state was similar in both sets of impulse responses ($\approx .1$). Note that while this does not occur at the same time in both models, this maximum value is still reached almost immediately in both models. Further, the magnitude of the persistence coefficient ρ was chosen so that the unemployment rate returned to the steady state at roughly the same time in both the theoretic and empirical model (around 7 years). Thus, by construction, the impulse response of the unemployment rate to a labor demand shock in the theoretic model approximates the shape of that of the VAR. One significant difference, however, is that the theoretic model lacks the overcorrection found

in the VAR: while the unemployment rate briefly dips below the steady state value in the VAR, in the theoretic model this variable monotonically approaches the steady state from above.

On the other measures the theoretic model does a fair job of matching the VAR, at least qualitatively. The model does well to capture the immediate fall and subsequent decline in income before the eventual recovery found in the data. There is a caveat here, as the shape of this response, at least initially, may reflect the fact that the unemployment rate increases before decreasing in the model. Nevertheless, income reaches steady state at roughly the same time (within 1 -2 years) in the model as when the confidence interval crosses zero in the data. Further, the magnitude of the immediate fall in income in both the model and the data are quite close as both are near .005.

Population growth initially declines in both the model and in the data. However, the magnitude of this decline in the data is much higher than what is implied in the model. The recovery is much faster as well, as it takes the model around three to four years to reach steady state as opposed to one in the data. Population growth actually rises for some time after initially reaching the steady state in the model. Impulse responses from the data suggest that population growth may eventually rise above steady state levels in the data as well, though the size of the confidence intervals preclude any more precise analysis.

In both the model and the data home prices are persistently below steady state after the labor demand shock. The shape of these responses are quite different, however. The magnitude of the effect in the data is much higher than in the model. Further, the model features more persistence as home prices do not reach their steady state value for at least fifteen years, while in the data this occurs closer to ten years. Finally, home prices in the model are actually initially positive, whereas there is an immediate drop in the data. This positive response is both very small and brief.

In the data, the construction rate response follows a shallow “u”-shaped pattern, and takes around seven years to return to the steady state value. In the model, the construction rate has an initial fall, followed by a monotonic increase that converges back to steady state very quickly. However, there is some success with this series as the construction rate in both scenarios falls as a result of a labor demand shock, and the magnitude of the maximum deviation from the steady state in both the model and the data are very close, albeit reached at different times.

Finally, the rent series perhaps displays the most disparities when compared to the data. While there is initially no change in rent in the data, in the model a negative labor demand shock results in an increase in rents. Rents do eventually become negative around the same time that one can reject rents being equal to their steady state value in the data. However, the magnitude of the deviation of rents from their steady state value is much smaller than in the data. Further, rents remain negative in the model for longer than they do in the data.

Conclusion

When not explicitly examining labor markets, the literature largely considers labor income processes consisting solely of a highly persistent component, a transitory component, and occasionally a known idiosyncratic intercept when incorporating exogenous income into a model. Recent work has shown that income data also supports an alternative specification, where individuals have some idiosyncratic factor that drives income realizations over their lifetime. Using parameters estimated from income data, I show that this choice of income process is not trivial when considering household portfolio decisions, where households must decide between investing in a liquid or illiquid asset.

I consider a two asset model where individuals can gradually learn about their future income potential over time. I show that in a stylized model of portfolio choice, this approach to modeling income yields drastically different outcomes for household asset decisions compared to a “near”-random walk benchmark specification without learning. In addition to generating more average total savings over the entire lifetime, learning generates more liquid savings over the lifetime as well and delays illiquid saving.

I then show that these results have ramifications for both modelling choices and policy initiatives. A model with learning has a much easier time generating enough liquid assets to match the ratio of average illiquid to liquid assets found in the data. If this is a moment that one is attempting to match with the benchmark model, it would take a coefficient of risk aversion nearly six times that of the learning model to achieve the same result. Risk aversion affects many facets of the household's decisionmaking, including their rate of intertemporal substitution, depending on the functional form of the utility function. Thus, this result potentially has implications for a wide variety of models.

Further, I show that, for individuals with 20 years of experience, the consumption response to an unanticipated change in income differs drastically depending on the underlying income process. Individuals in a model with learning have around half of the consumption response to an increase in income as individuals in a more standard model. Though it is important to emphasize that these numbers should be treated with some caution, due to the simplified nature of the model, these results nevertheless indicate that allowing individuals to learn about their income may have implications for the efficacy of certain fiscal policies, such as stimulative tax rebates. Thus, this paper advances the notion that an assumption typically unemphasized by the literature, that of the specification of underlying income process, has crucial implications for portfolio choice models and any conclusions drawn from them.

Finally, I consider the dynamics of various aggregate variables at the city-level. I derive the impulse responses of the unemployment rate, income, population growth, home prices, construction rates, and rent to a shock to local labor demand. To do this I construct a model in which individuals face search and matching frictions in both the labor market and the housing market. I show that a model with these features does a respectable job replicating qualitative, and some quantitative, features of empirical responses to a similar type of disturbance.

Appendix A

Appendix for Chapters 2 & 3

A.1 Benchmark Model and Computation

Retirement periods are handled identically in the benchmark model as in the model with learning. The differences arise in the working periods, where the model with the benchmark process is more straightforward than the model with learning. Since the individual observes not just his income draw but also the composition of his income, the individual's beliefs coincide fully with reality, obviating the need for the Kalman Filtering apparatus. However, since each component of income has different implications for the individual, each must be kept track of separately. Thus, there are still five state variables for the following dynamic problem:

$$V_t(\alpha, z_t, \varepsilon_t, m_t, b_t) = \max_{c_t, b_t, m_t} u(c_t) + \delta E_t[V_{t+1}(\alpha, z_{t+1}, \varepsilon_{t+1}, m_t, b_t)] \quad (\text{A.1})$$

In order to simplify the problem computationally, I use a normalization popularized by Carroll (2011) for dynamic programming problems featuring CRRA preferences. When iterating backwards to solve for decision rules, I divide the problem by α . This affects only the levels of assets, consumption, and income (and therefore also the value function), but does not affect the ratios of any of these variables. Thus,

this normalization removes a state variable while retaining all important information. When simulating individual decisions, I draw α for each household separately and then multiply the choice variables (and income) by these draws. This recovers the decisions of households who receive the full income draw (including the individual intercept).

A.2 Data

A.2.1 *Survey of Consumer Finances (SCF)*

The SCF is a triennial survey running from 1983 to 2016. Post-1986 the SCF uses multiple imputations to fill in missing data. Thus, in the raw data there are five records for every household. While averages and frequencies should be unaffected by the imputations, medians may be slightly off. Therefore, for each statistic of interest I follow the advice of the Federal Reserve Board and first calculate that statistic by the imputation, before averaging the five values together. For the 1983 and 1986 surveys only one set of data is provided, and therefore the aforementioned process is unnecessary. Since the SCF oversamples wealthier households, to better capture their more diverse portfolios, I use the provided weights while computing all of the statistics in this paper.

Median Portfolio Comparison Between Years

In Figure A.1, I plot the smoothed experience profile of the ratio of illiquid to liquid savings for the median individual for the years 2004, 2007, and 2013. Series from other years are roughly spanned by the three depicted here. Notice that there is a fair bit of variation between the three series, both in size and shape. The 2013 series stays roughly constant at near zero until around ten years experience, where it first becomes concave and then convex. The 2004 and 2007 series begin to increase around an experience level of 5. The 2004 series is convex while the 2007 series begins




FIGURE A.1: Median Illiquid-Liquid Ratio For Select Years

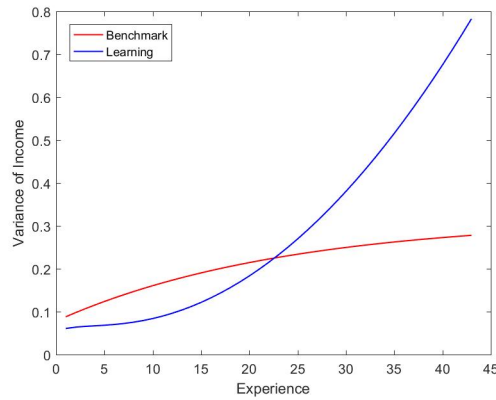


FIGURE A.2: Distribution of Income Over the Lifecycle

concave, ends convex, and changes inflection at around 30 years of experience. There are nontrivial differences between the value of the ratio over the lifetime between the three curves, as well. For example, at age 20, the ratio in 2013 is around .5 while the ratio in 2007 is above 2. Further, the median individual first has more illiquid net worth than liquid net worth after 10 years of experience in 2007, compared with nearly 25 years of experience in 2013.

A.3 Derivations

A.3.1 Lifetime Income Distributions

Taking the variance of the general form of the de-meanded income process, we get the following:

$$Var(\tilde{y}_{i,e}) = \sigma_{\alpha}^2 + 2cov_{\alpha\beta}e + \sigma_{\beta}^2e^2 + \sigma_{\varepsilon}^2 + \sigma_{\eta}^2 \sum_{j=0}^{e-1} \rho^{2j} \quad (\text{A.2})$$

Now, of course the sum in the final term depends on the value of ρ , which is typically taken to be between 0 and 1, inclusive. Thus:

$$\begin{aligned} \rho = 0 &\implies \sigma_{\alpha}^2 + 2cov_{\alpha\beta}e + \sigma_{\beta}^2e^2 + \sigma_{\varepsilon}^2 + \sigma_{\eta}^2 \\ \rho \in (0, 1) &\implies \sigma_{\alpha}^2 + 2cov_{\alpha\beta}e + \sigma_{\beta}^2e^2 + \sigma_{\varepsilon}^2 + \sigma_{\eta}^2 \left[\frac{1 - \rho^{2e}}{1 - \rho^2} \right] \\ \rho = 1 &\implies \sigma_{\alpha}^2 + 2cov_{\alpha\beta}e + \sigma_{\beta}^2e^2 + \sigma_{\varepsilon}^2 + e\sigma_{\eta}^2 \end{aligned}$$

In the benchmark model, with $\rho \in (0, 1)$ and β restricted to zero, the variance simplifies:

$$\sigma_{\alpha}^2 + \sigma_{\varepsilon}^2 + \sigma_{\eta}^2 \left[\frac{1 - \rho^{2e}}{1 - \rho^2} \right]$$

Notice in Figure A.2 that the two profiles have a distinct shape, even with parameters estimated from the same income data. The experience-variance profile is concave in the benchmark parameterization while the learning parameterization yields a profile that is convex after the first few working years. The population-wide variance of income is slightly higher with the benchmark parameters early in the lifecycle, but the variance is much higher later in the lifecycle using the learning estimates.

A.3.2 Kalman Filter

Let $S_t = [\alpha \quad \beta \quad z_t]'$. If we let $F = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & \rho \end{bmatrix}$ and $\nu_t = [0 \quad 0 \quad \eta_t]'$, then the law of

motion, understood by all agents in the model, is:

$$S_t = FS_{t-1} + \nu_t \quad (\text{A.3})$$

Conditional expectations of the parameters at time t are denoted by $\hat{S}_{t|t}$, while one-period ahead forecasts of the parameters are given by $\hat{S}_{t+1|t} = F\hat{S}_{t|t}$. Further, letting $H = [1 \quad t \quad 1]'$ and $w_t = \varepsilon_t$, one can write the income process as:

$$y_t = H_t' S_t + w_t \quad (\text{A.4})$$

Denote $Q = E(\nu_t \nu_t')$ and $R = E(w_t w_t')$. It can then be shown that beliefs update according to the following relation:

$$\hat{S}_{t|t} = \hat{S}_{t|t-1} + K_t(y_t - H_t' \hat{S}_{t|t-1}) \quad (\text{A.5})$$

where $K_t = P_{t|t-1} H_t (H_t' P_{t|t-1} H_t + R)^{-1}$ is the Kalman Gain.

Finally, the perceived variance-covariance matrix of the parameters of interest evolves according to the following two equations:

$$P_{t+1|t} = FP_{t|t}F' + Q \quad (\text{A.6})$$

$$P_{t|t} = P_{t|t-1} - P_{t|t-1} H_t (H_t' P_{t|t-1} H_t + R)^{-1} H_t' P_{t|t-1} \quad (\text{A.7})$$

A.3.3 Mean Square Error

For the benchmark case, deriving the mean squared error is fairly straightforward:

$$\begin{aligned} MSE_{t+s|t} &= E_t[(y_{t+s} - y_{t+s|t})^2] \\ &= E_t[(d_{t+s} + x_{t+s} - \varepsilon_{t+s} - d_{t+s} - \rho^s x_t)^2] \end{aligned}$$

$$= E_t[(\varepsilon_{t+s} + \sum_{j=0}^{s-1} \rho^j \eta_{t+s-j})^2]$$

Since all of these shocks are assumed to be i.i.d., this can be rewritten as:

$$\begin{aligned} &= \sigma_\eta^2 \sum_{j=0}^{s-1} \rho^{2j} + \sigma_\varepsilon^2 \\ &= \sigma_\eta^2 \frac{1 - (\rho^2)^s}{1 - \rho^2} + \sigma_\varepsilon^2 \end{aligned}$$

where the fraction is replaced by 1 if $\rho = 0$ and s if $\rho = 1$.

The learning case is a little more involved. However, it can be shown that:

$$MSE_{t+s|t} = H'_{t+s} P_{t+s|t} H_{t+s} + R \quad (\text{A.8})$$

where

$$P_{t+s|t} = F^s P_{t|t} F'^{s} + \sum_{j=0}^{s-1} F^j Q F'^{j} \quad (\text{A.9})$$

Note that this error will depend greatly on individual priors, as those will determine $P_{t|t}$. Of course, the more tight the priors, the less uncertain the individual will be over his entire lifecycle.

Appendix B

Appendix for Chapter 4

B.1 Data

The data used in Chapter 4 is from a dataset constructed by Head et al. (2014). The data set contains information on per capita income, population, rent, construction rates, home prices, labor force, and employment information for 106 Metropolitan Statistical Areas (MSAs) for the years 1980 - 2008.

There are a few differences between the final data set I utilize and the set used in that paper. Due to various discrepancies between the MSA codes in the unemployment data and the rest of the data set, I had to remove six MSAs for my analysis. Further, my analysis was limited to 1990 - 2008 as the employment data was unavailable for earlier years. Finally, the original data set also contains data on housing sales; however, due to the proprietary nature of this data it was necessarily excluded from my analysis. The rest of the data were not altered for my paper.

Income and population figures come from the Bureau of Economic Analysis (BEA) Regional Economic Accounts. Income is defined as total income from all sources; in addition to wages, it includes employer contributions to pensions and in-

insurance as well as government transfers. Nominal income is deflated by the consumer price index (CPI), obtained from the Bureau of Labor Statistics (BLS). Employment and labor force data were also obtained from the BLS.

Rent data were originally constructed by Nieuwerburgh and Weill (2010). They used rental data from the Fair Market Rents (FMR) database, which lists annual information for a certain percentile of the rent distribution at the sub-MSA level. This reported percentile changes over time, complicating matters somewhat. The authors are able to correct for these shifting percentile values and aggregate up to the MSA level.¹

Construction rates were created using permit data, released by the Department of Housing and Urban Development (HUD), and estimated 2000 Census data on housing stocks.² From these two data sets it is possible to extrapolate the overall housing stock for each MSA per year. Construction rates are then the number of permits issued each year divided by the overall housing stock.

Finally, home prices were created in a similar fashion to construction rates. Head et al. (2014) are able to use home price information obtained from the 2000 Census and the Home Price Index (HPI), from the Federal Housing Finance Agency (FHFA), to generate annual nominal home prices for each MSA in each year. The CPI is then used to convert these prices into real terms.

B.2 Population Laws of Motion

Each period $G(e_{t-1})\Delta_P Q_{t-1}$ individuals enter the city. They enter the city unemployed and $(1 - \psi)$ of them become perpetual renters. Of the perpetual renters, π^f of them will decide to leave the city each period. δ of the employed perpetual renters

¹ For more information, see Appendix D.3 in the Online Supplement of Nieuwerburgh and Weill (2010).

² The Census data must be estimated for the MSA level, as the Census only gives housing stock data from counties.

will separate with their firm and become unemployed, while $(1 - \delta)$ will remain with their firm the following period. μ_t of the unemployed perpetual renters will match with a firm and work in the following period while $(1 - \mu_t)$ will remain unemployed. Thus, one can write the laws of motion of the number of unemployed and employed perpetual renters, respectively, as:

$$F_t^U = (1 - \pi^f)[(1 - \mu_t)F_{t-1}^U + \delta F_{t-1}] + (1 - \psi)G(e_{t-1}^C)Q_{t-1} \quad (\text{B.1})$$

$$F_t = (1 - \pi^f)[(1 - \delta)F_{t-1} + \mu_t F_{t-1}^U] \quad (\text{B.2})$$

A proportion ψ of the new entrants to the city become unemployed renters who are also searching for a home. Of the number of unemployed perpetual renters, $(1 - \lambda_t)(1 - \delta)$ do not find a housing match but do find employment by the following period; $(1 - \lambda_t)(1 - \mu_t)$ neither find a house nor a job; $\lambda_t(1 - \mu_t)$ find a house but not a job; and $\lambda_t\mu_t$ find both a house and a job. Of the individuals who are unemployed and looking for a home in time t , by time $t+1$ $(1 - \lambda_t)\delta$ have lost their job and yet not found a home; $(1 - \lambda_t)(1 - \delta)$ have neither found a home nor lost their job; $(\lambda_t)\delta$ found a home but lost their job; and $\lambda(1 - \delta)$ found a home and kept their job.

Each period, π^n homeowners leave the city. Of the fraction of unemployed homeowners that remain, $\varphi(1 - \mu_t)$ do not find a job and decide to relocate within the city; $\varphi\mu_t$ find a job and relocate within the city; $(1 - \varphi)(1 - \mu_t)$ remain in their home without a job; and $(1 - \varphi)\mu_t$ find a job and remain in their home. Of the employed homeowners in period t who remain in the city, $\varphi\delta$ leave their house and lose their job; $\varphi(1 - \delta)$ relocate within the city while keeping their job; $(1 - \varphi)\delta$ keep their house but lose their job; and $(1 - \varphi)(1 - \delta)$ keep both their home and their job. One can then write the laws of motion of the number of unemployed and employed home

seekers as:

$$B_t^U = (1 - \lambda_t)(1 - \mu_t)B_{t-1}^U + (1 - \lambda_t)\delta B_{t-1} + (1 - \pi^n)[\varphi\delta N_{t-1} + \varphi(1 - \mu_t)N_{t-1}^U] + \psi G(e_{t-1}^C)Q_{t-1} \quad (\text{B.3})$$

$$B_t = (1 - \lambda_t)(1 - \delta)B_{t-1} + (1 - \lambda_t)\mu_t B_{t-1}^U + (1 - \pi^n)[\varphi(1 - \delta)N_{t-1} + \varphi\mu_t N_{t-1}^U] \quad (\text{B.4})$$

Further, the law of motion for the number of unemployed and employed home owners can be written the following way:

$$N_t^U = \lambda_t(1 - \mu_t)B_{t-1}^U + \lambda_t\delta B_{t-1} + (1 - \pi^n)[(1 - \varphi)(1 - \mu_t)N_{t-1}^U + (1 - \varphi)\delta N_{t-1}] \quad (\text{B.5})$$

$$N_t = \lambda_t\mu_t B_{t-1}^U + \lambda_t(1 - \delta)B_{t-1} + (1 - \pi^n)[(1 - \varphi)\mu_t N_{t-1}^U + (1 - \varphi)(1 - \delta)N_{t-1}] \quad (\text{B.6})$$

B.3 Normalizing Variables

Let $\Delta_P = \Delta_L = \Delta$. Further, denote $x_t = \frac{X_t}{Q_t}$. Then, using the functional forms detailed in Section 4.3.1 and noting that in equilibrium $e_t^C = \bar{W}_t$, one can rewrite Equations B.1 - B.6:

$$(1 + \Delta)f_t^U = (1 - \pi^f)[(1 - \mu_{t-1})f_{t-1}^U + \delta f_{t-1}] + \Delta(1 - \psi)\frac{1}{2}\left(\frac{\bar{W}_t}{\bar{W}}\right)^{\alpha_e} \quad (\text{B.7})$$

$$(1 + \Delta)f_t = (1 - \pi^f)[\mu_{t-1}f_{t-1}^U + (1 - \delta)f_{t-1}] \quad (\text{B.8})$$

$$(1 + \Delta)b_t^U = (1 - \lambda_{t-1})(1 - \mu_{t-1})b_{t-1}^U + (1 - \lambda_{t-1})\delta b_{t-1} + (1 - \pi^n)\varphi[(1 - \mu_{t-1})n_{t-1}^U + \delta n_{t-1}] + \Delta\psi\frac{1}{2}\left(\frac{\bar{W}_t}{\bar{W}}\right)^{\alpha_e} \quad (\text{B.9})$$

$$(1 + \Delta)b_t = (1 - \lambda_{t-1})\mu_{t-1}b_{t-1}^U + (1 - \lambda_{t-1})(1 - \delta)b_{t-1} + (1 - \pi^n)\varphi[\mu_{t-1}n_{t-1}^U + (1 - \delta)n_{t-1}] \quad (\text{B.10})$$

$$(1 + \Delta)n_t^U = \lambda_{t-1}(1 - \mu_{t-1})b_{t-1}^U + \lambda_{t-1}\delta b_{t-1} + (1 - \pi^n)(1 - \varphi)[(1 - \mu_{t-1})n_{t-1}^U + \delta n_{t-1}] \quad (\text{B.11})$$

$$(1 + \Delta)n_t = \lambda_{t-1}\mu_{t-1}b_{t-1}^U + \lambda_{t-1}(1 - \delta)b_{t-1} + (1 - \pi^n)(1 - \varphi)[\mu_{t-1}n_{t-1}^U + (1 - \delta)n_{t-1}] \quad (\text{B.12})$$

Note that, because $\Delta_P = \Delta_L$, the population and available land grow at the same rate. So $\frac{\bar{L}_t}{\bar{Q}_t} = \bar{l}$ is time-invariant. Thus, I can similarly rewrite Equation 4.10 as:

$$(1 + \Delta)h_t = (1 - \varrho)h_{t-1} + \frac{1}{2}\left(\frac{p_t}{\bar{p}}\right)^{\alpha_c}(\bar{l} - l_{t-1}) \quad (\text{B.13})$$

B.4 Optimal Maintenance Problem

Let a_t be the quality of a housing structure at time t , m_t be the amount of repairs and maintenance a household undertakes at time t , and d be the constant monthly rate of depreciation of a structure. Finally, assume that the utility of a structure, z_t^h is a function of the structure's quality, such that $z^h(a_t)$. Then the household faces the following optimal maintenance problem:

$$V(a_t) = \max_{m_t} z^h(a_t) - m_t + \beta V(a_{t+1}) \quad (\text{B.14})$$

$$a_{t+1} = (1 - d)a_t + m_t \quad (\text{B.15})$$

Taking the first order condition with respect to m_t and the envelope condition, and then solving, one gets

$$z^{h,\prime}(a_t) = \frac{1 - \beta(1 - d)}{\beta} \quad (\text{B.16})$$

Equation B.16 suggests that $z^h(a_t)$ is linear with respect to a_t . It is then the case that $z^{h,\prime}(a_t)a_t \approx z^h(a_t)$. Further, notice that in the steady state $a = \frac{m}{d}$. It then follows that

$$z^h = z^h(a) \approx z^{h,\prime}(a)a = \frac{1 - \beta(1 - d)}{\beta d}m = \left(1 + \frac{1 - \beta}{\beta d}\right)m \quad (\text{B.17})$$

Bibliography

- Abowd, J. M. and Card, D. (1989), “On the Covariance Structure of Earnings and Hours Changes,” *Econometrica*, 57, 411 – 445.
- Abrigo, M. R. and Love, I. (2016), “Estimation of Panel Vector Autoregression in Stata: a Package of Programs,” *Working Paper*, pp. 1 – 30.
- Adjemian, S., Bastani, H., Juillard, M., Karamé, F., Mihoubi, F., Perendia, G., Pfeifer, J., Ratto, M., and Villemot, S. (2011), “Dynare: Reference Manual, Version 4,” *Dynare Working Papers*, 1, 1 – 208.
- Arellano, M. and Bover, O. (1995), “Another Look at the Instrumental Variable Estimation of Error-Components Models,” *Journal of Econometrics*, 68, 29 – 51.
- Baker, M. (1997), “Growth-Rate Heterogeneity and the Covariance Structure of Life-Cycle Earnings,” *Journal of Labor Economics*, 15, 338 – 375.
- Blanchard, O. J. and Katz, L. F. (1992), “Regional Evolutions,” *Brookings Papers on Economic Activity*, 1, 1 – 75.
- Blundell, R. and Bond, S. (1998), “Initial Conditions and Moment Restrictions in Dynamic Panel Data Models,” *Journal of Econometrics*, 87, 115 – 143.
- Caplin, A. and Leahy, J. (2011), “Trading Frictions and House Price Dynamics,” *Journal of Money, Credit, and Banking*, 43, 283 – 303.
- Carroll (2011), “Theoretical Foundations of Buffer Stock Saving,” *Working Paper*.
- Chang, Y., Hong, J., and Karabarbounis, M. (2016), “Labor-Market Uncertainty and Portfolio Choice Puzzles,” *Working Paper*.
- Cocco, J. F. (2004), “Portfolio Choice in the Presence of Housing,” *The Review of Financial Studies*, 18, 535 – 567.
- Fisher, J. D. and Gervais, M. (2011), “Why Has Home Ownership Fallen Among the Young,” *International Economic Review*, 52, 883 – 912.
- Foster, K., Schuh, S., and Zhang, H. (2013), “The 2010 Survey of Consumer Payment Choice,” *Research Data Report no. 13-2, Federal Reserve Bank of Boston*.

- Guvenen, F. (2007), “Learning Your Earning: Are Labor Income Shocks Really Very Persistent,” *American Economic Review*, 97, 687 – 712.
- Guvenen, F. (2009), “An empirical investigation of labor income processes,” *Review of Economic Dynamics*, 12, 58 – 79.
- Guvenen, F. and Smith, A. (2014), “Inferring Labor Income Risk and Partial Insurance from Economic Models,” *Econometrica*, 82, 2085 – 2129.
- Harding, J. P., Rosenthal, S. S., and Sirmans, C. F. (2007), “Depreciation of Housing Capital, Maintenance, and House Price Inflation,” *Journal of Urban Economics*, 2, 193 – 217.
- Head, A. and Lloyd-Ellis, H. (2012), “Housing Liquidity, Mobility, and the Labour Market,” *The Review of Economic Studies*, 79, 1559 – 1589.
- Head, A., Lloyd-Ellis, H., and Sun, H. (2014), “Search, Liquidity, and Dynamics of House Prices and Construction,” *American Economic Review*, 104, 1172 – 1210.
- Hryshko, D. (2012), “Labor income profiles are not heterogeneous: Evidence from income growth rates,” *Quantitative Economics*, 3, 177 – 209.
- Iacoviello, M. and Pavan, M. (2013), “Housing and debt over the life cycle and over the business cycle,” *Journal of Monetary Economics*, 60, 221 – 238.
- Johnson, D. S., Parker, J. A., and Souleles, N. S. (2006), “Household Expenditure and the Income Tax Rebates of 2001,” *The American Economic Review*, 96, 1589 – 1610.
- Kaplan, G. and Violante, G. L. (2014), “A Model of the Consumption Response to Fiscal Stimulus Payments,” *Econometrica*, 82, 1199 – 1239.
- Kaplan, G., Violante, G. L., and Weidner, J. (2014), “The Wealthy Hand-to-Mouth,” *Brookings Papers on Economic Activity*, pp. 77 – 138.
- Karahan, F. and Rhee, S. (2013), “Geographical Reallocation and Unemployment during the Great Recession: The Role of the Housing Bust,” *Federal Reserve Bank of New York Staff Report*, 605, 1 – 47.
- Krainer, J. (2001), “A Theory of Liquidity in Residential Real Estate Markets,” *Journal of Urban Economics*, 49, 32 – 53.
- Love, I. and Zicchino, L. (2006), “Financial development and dynamic investment behavior: Evidence from panel VAR,” *The Quarterly Review of Economics and Finance*, 46, 190 – 210.

- Lynch, A. W. and Tan, S. (2011), “Labor income dynamics at business-cycle frequencies: Implications for portfolio choice,” *Journal of Financial Economics*, 101, 333 – 359.
- MaCurdy, T. E. (1982), “The Use of Time Series Processes to Model the Error Structure of Earnings in a Longitudinal Data Analysis,” *Journal of Econometrics*, 18, 83 – 114.
- Mortensen, D. T. (2011), “Markets with Search Friction and the DMP Model,” *American Economic Review*, 101, 1073 – 1091.
- Nieuwerburgh, S. V. and Weill, P.-O. (2010), “Why Has House Price Dispersion Gone Up?” *The Review of Economic Studies*, 77, 1567 – 1606.
- Polkovnichenko, V. (2007), “Life-Cycle Portfolio Choice with Additive Habit Formation Preferences and Uninsurable Labor Income Risk,” *The Review of Financial Studies*, 20, 83 – 124.
- Rupert, P. and Wasmer, E. (2012), “Housing and the labor market: Time to move and aggregate unemployment,” *Journal of Monetary Economics*, 59, 24 – 36.
- Scholz, J. K., Seshadri, A., and Khitatrakun, S. (2006), “Are Americans Saving ”Optimally” for Retirement?” *Journal of Political Economy*, 114, 607 – 643.
- Shimer, R. (2005), “The Cyclical Behavior of Equilibrium Unemployment and Vacancies,” *American Economic Review*, 95, 25 – 49.
- Tauchen, G. (1986), “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 20, 177 – 181.
- Wheaton, W. C. (1990), “Vacancy, Search, and Prices in a Housing Market Matching Model,” *Journal of Political Economy*, 98, 1270 – 1292.
- Yang, F. (2009), “Consumption over the life cycle: How different is housing?” *Review of Economic Dynamics*, 12, 423 – 443.

Biography

Adam Marshall Bergeron was born on December 9, 1989 in Leesville, Louisiana. He earned a Ph.D in Economics from Duke University in 2018, an MA in Economics from Duke University in 2015, and a BA in Economics and Mathematics from the University of Florida in 2012.