Essays on Social Insurance and Labor Markets

A DISSERTATION
SUBMITTED TO THE FACULTY OF THE GRADUATE SCHOOL
OF THE UNIVERSITY OF MINNESOTA
BY

Kurt See

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
Doctor of Philosophy

Ellen McGrattan, Advisor

April, 2019

Acknowledgements

I am deeply indebted to Naoki Aizawa, Anmol Bhandari, Kyle Herkenhoff, Loukas Karabarbounis, and Ellen McGrattan for their guidance and support throughout this thesis. I also want to express my deep appreciation to my great friend and amazing co-author Serdar Birinci for all the countless discussions, invaluable feedback, and excellent collaboration.

Dedication

I dedicate my work to my loving parents, and siblings Karl and Kyra. They have been my source of inspiration and strength.

Abstract

This dissertation consists of three chapters. In the first chapter, I study the macroeconomic and welfare consequences of introducing a universal healthcare system to replace the existing employer-based set-up in the U.S., paying particular attention to the reform's labor market effects. To study the policy, I develop an incomplete asset markets model with labor market frictions and medical expenditure risk over the life cycle. First, I compare the model-implied partial equilibrium employment responses to public health insurance generosity to existing empirical evidence. The model partially reconciles the puzzlingly wide range of estimates found in three microeconomic experiments conducted in Tennessee, Oregon, and Wisconsin. Next, I use the model to understand the general equilibrium effects of switching to universal healthcare. I find that it results in higher reservation wages and a corresponding reduction in firm vacancy creation, both of which lead to a quantitatively large decline in the job finding rate. The negative impact of the lower job finding rate outweighs the insurance benefits of generous public health coverage, resulting in substantial welfare losses among low-wealth households for whom employment is most valuable. In the second chapter (joint with Serdar Birinci), we investigate how unemployment insurance generosity should vary with the business cycle. We find that the optimal policy is countercyclical. Not only does the policy smooth the consumption of job losers, but also provides insurance against aggregate risk by reducing the need for excess savings during recessions. Meanwhile, the moral hazard effects of generous benefits are attenuated in recessions because jobs are scarce and thus, the forgone value of job search is low. In the third chapter (joint with Anmol Bhandari, Serdar Birinci, and Ellen McGrattan), we study survey data used for measuring business income and valuations. We document large inconsistencies between survey data and aggregated administrative data for statistics such as the level and distribution of business income, and the number of returns. These inconsistencies are attributable to both non-representative samples and measurement errors. Non-representativeness results from undersampling businesses with low income owners. Measurement errors emerge because respondents do not use relevant financial documents as basis for their responses while some survey questions suffer from framing problems.

Contents

List of Tables

List of Figures

Chapter 1

"Medicare for All": The Macroeconomic and Welfare Consequences of Expanding Public Health Insurance

1.1 Introduction

The United States health insurance system relies heavily on coverage provided through employer-sponsored plans, which is in stark contrast to other industrialized countries that have long since introduced a publicly financed universal healthcare system. Given the high costs associated with purchasing individual non-group private plans and the limited coverage of public health insurance such as Medicaid, being employed is essentially a prerequisite to having access to affordable health insurance for most Americans. Furthermore, the current system compounds the effects of employment risk by exposing job losers not only to earnings losses but also a potential loss of health insurance coverage. Recently, proposals to transition toward a universal healthcare system ("Medicare-for-all") have gained much attention in the policy agenda. However, the tight link between health insurance and employment implies that any expansion of public health insurance will inevitably have effects on the labor market behavior of individuals and firms. The purpose of this paper is to quantify the macroeconomic and welfare effects of shifting to a universal healthcare system, paying particular attention to the reform's effects on the labor market. I use a structural model whose features can jointly account for the puzzlingly wide range of estimates on labor supply responses to public health insurance generosity found in three microeconomic experiments conducted in Tennessee, Oregon, and Wisconsin. I find that transitioning to a universal healthcare system results in a large decline in job-finding rates due to lower job creation by firms. The negative impact of a lower job-finding rate results in substantial welfare losses for low-wealth households for whom employment is most valuable.

The main contribution of this paper is to study the interaction of household labor supply and firm labor demand decisions in response to government provision of health insurance. The expansion of public health insurance alters the value of non-employment relative to employment, which in turn governs non-employed individuals' job search and acceptance decisions, and employed individuals' quit decisions. At the same time, any changes in reservation wages, job acceptance rates, and turnover rates affect the value of a vacancy and a filled job. Firms react to this by endogenously adjusting the level of job creation. As a result, while universal healthcare can insure against employment risk and health expenditure risk, these insurance benefits may be offset by potential incentive costs in the labor market.

To study the reform, I develop a general equilibrium lifecycle model with incomplete asset markets with the following key features. First, individuals are subject to medical expenditure risk which can be partially insured primarily through employer-provided health insurance (EPHI), supplemented by a limited government-funded means-tested program (Medicaid). This provides a suitable framework to quantify the extent of medical expenditure risk over the lifecycle as well as mechanisms in place to insure against them. Second, jobs are found in a frictional labor market characterized by random search where households make endogenous job acceptance and quit decisions while firms offer a compensation package comprising of a wage and possibly a health plan and decide on the level of vacancies to post. This general equilibrium labor market set-up accounts for the possibility that generous public health insurance can raise reservation wages, lower job-filling rates, and raise quit rates, thereby inducing firms to respond by lowering vacancies. Furthermore, the framework can be used to study the effects of health insurance reform on aggregate job finding rates, equilibrium unemployment, and capital stock. Last, the model features complementarity in production and two-sided heterogeneity where workers differ in terms of their skill and match with jobs that differ in productivity. This feature allows me to study how healthcare reform affects not only the amount of job vacancies in equilibrium but also the allocation of workers to firms among vacancies that are filled. These features are important to the welfare evaluation of a universal care system that aims to delink health insurance coverage from employment.

The model is calibrated to a benchmark United States health insurance system pre Affordable Care Act (ACA). Importantly, this involves targeting the distribution of individuals who are eligible for Medicaid and the joint distribution of wages and access to employer-provided health insurance since these directly affect the insurance benefits and incentive costs of expanding publicly-provided healthcare. The model is then validated against quasi-experimental evidence on labor supply elasticities to changes in the generosity of public health insurance programs. Garthwaite, Gross, and Notowidigdo (2014) exploit the unexpected discontinuation of Tennessee's Medicaid (TennCare) expansion program in 2005 to identify the causal effect of public health insurance and they find large effects on labor supply. Dague, DeLeire, and Leninger (2017) use the Wisconsin BadgerCare enrollment cap in 2009 and find large but more modest effects, while Baicker, Finkelstein, Song, and Taubman (2014) use the Oregon Medicaid expansion in 2008 and find no effect. Conducting the same experiments using my model by introducing an expansion/discontinuation of coverage to a subpopulation that closely matches the enrollee/disenrollee demographics in each experiment, my model predicts effects in between the range of estimates in the literature. Furthermore, the model generates a reasonably large amount of dispersion across the experiments because of differences in the income- and asset- distribution of the targeted population. Finally, I also find substantial heterogeneity in labor supply responses to public health insurance generosity that are qualitatively consistent with the empirical findings, where older, less healthy, and unemployed workers are more responsive to policy. It is important for the model to generate a reasonable labor supply elasticity with respect to health insurance coverage in order to properly measure the incentive costs of the policy reform.

Next, I use the calibrated model to understand the general equilibrium effects of switching to a single-payer system. I conduct the policy experiment by introducing a universal healthcare system that is funded through a proportional increase in taxes. I find that the reform results in a substantial decline in the aggregate job finding rate from 45 percent to 31 percent and concomitantly, a higher non-employment rate and longer unemployment spell durations. The decline in the aggregate job finding rate is caused, in part, by higher reservation wages of individuals that become more selective in the range of jobs they accept, and a slightly higher quit rate among workers who either look for better jobs or leave the labor force. This results in a general equilibrium effect coming from lower vacancy creation by firms for whom the value of posting a vacancy is now much lower due to lower acceptance and higher turnover rates. Along the transition, however, there is an initial increase in vacancies resulting from increased firm profits coming from the withdrawal of employer-provided health insurance plans from the compensation package. Over time, reservation wages and quit rates rise to overturn this and eventually, job-finding rates decline substantially. Finally, due to increased public insurance against health risk, there is a sizable decrease in the capital stock and an increase in interest rates resulting from lower precautionary savings. Lower capital stock results in a net decrease in output per worker, despite productivity gains brought by higher average match quality due to increased job selectivity and switching.

The universal healthcare system results in a small ex-ante welfare gain of less than one percent additional lifetime consumption. The reason why aggregate welfare gains are small despite the large scale of the policy change is the substantial heterogeneity in ex-post welfare gains and losses in the population alive during the policy change. I find welfare losses accrue mostly to wealth-poor individuals for whom lower aggregate job-finding rates and longer unemployment durations are most costly since these households have severely limited ability to self-insure against consumption fluctuations caused by unemployment risk. In addition to wealth-poor individuals, welfare losses also accrue mostly to younger, healthier, and Medicaid-eligible households for whom the insurance benefits of universal care is small compared to its distortionary effects on taxation and the labor market.

Given that welfare losses mostly accrue to wealth-poor and low-income individuals, a natural question to ask then is if there are welfare improving ways to fund the universal healthcare reform. I explore alternative ways to fund the welfare reform by varying the tax burden of funding the policy reform across individuals with different income levels. Unsurprisingly, I find that the optimal revenue-generation scheme involves a highly progressive system where high income households bear the brunt of the additional taxes.

Related literature This paper contributes to a growing literature that studies the effects of health insurance reform on labor markets and the macroeconomy. These studies have focused on the effects of various ACA provisions (Pashchenko and Porapakkarm 2013, Aizawa and Fang 2015, Tsujiyama 2015, Jung and Tran 2016a, and Nakajima and Tuzemen 2016, Aizawa 2017), an evaluation of Medicaid for the poor and Medicare for the elderly (Attanasio, Kitao, and Violante 2010; Hansen, French and Jones 2011; Hansen, Hsu and Lee 2014; De Nardi, French and Jones 2016; Conesa et al. 2018), as well as the introduction of socialized health insurance (Jung and Tran 2016b). My paper contributes to this literature along two dimensions. First, I emphasize how the policy reform changes household labor market decisions and elicits a corresponding change in firm vacancy creation. I show that the reform's general equilibrium effect in the labor market is critical in determining the distribution of welfare gains and losses. Second, incorporating capital accumulation and sorting between worker skill and firm productivity allows me to study how health insurance reform endogenously affects productivity. I show that the change in aggregate productivity is interestingly small due to offsetting effects coming from better sorting but lower capital accumulation and job finding rates.

This paper is also related to a strand of literature that quantifies the effects of an employer-based health insurance system on and individual job choices. Dey and Flinn (2005) estimate an equilibrium model of employer-provided health insurance and wage determination and find little evidence inefficiencies in job-mobility. However, the role of health in their model is limited to a utility benefit and increased job productivity, whereas I explicitly incorporate health risk in a model of risk-averse households so that employer-provided benefits have an important consumption-smoothing role against these shocks. This mechanism drastically amplifies the value of employer-provided benefits for the households, and as a result, they accept jobs even if the wage offer is relatively lower. Chivers, Feng, and Villamil (2016) build a model of occupational choice to study the effects of employer benefits on entrepreneurship. They find that agents with high managerial ability but poor health status choose to be workers rather than entrepreneurs due to the high insurance value of employer-provided health insurance. Instead of focusing entrepreneurial choice, my framework analyzes how public health insurance coverage affects workers' allocation into different jobs and how this affects firm labor demand.

Another active area of research is the use of quasi-experimental evidence to measure the effects of public health insurance on the labor supply and earnings of households. Garthwaite, Gross, and Notowidigdo (2014) exploit the sudden loss of Medicaid coverage of over 170,000 Tennessee residents to measure the effects of employer-based health benefits on employment and earnings. They find that the disenrollment resulted in a 4.6 percentage points increase in the employment rate of all childless adults, equivalent to around a 60 percentage points change in employment among disenrollees. Other papers have found smaller effects. Dague, DeLaire, and Leininger (2017) study a sudden imposition of an enrollment cap in Wisconsin's BadgerCare Plus Core Plan. Specifically, they compare the employment outcomes of Medicaid applicants before and after the enrollment cut-off. They find that applicants who received Medicaid were 5 percentage points less likely to be employed relative to the control group of waitlisted applicants who eventually did not receive coverage. Baicker, Finkelstein, Song, and Taubman (2014) conduct a randomized study using data from the 2008 Oregon Medicaid expansion implemented via a lottery. They find that Medicaid enrollment leads to a (statistically insignificant) decline in employment of 1.6 percentage points relative to the control group.[1] In this paper, I run similar experiments that underlie the results of existing empirical evidence by holding taxes and prices fixed and targeting a similar subpopulation in the model. The model predicts large labor supply elasticities that are lower than those found by Garthwaite, Gross, and Notowidigdo (2014) but much higher than those found by Dague, DeLaire, and Leininger (2017) and Baicker, Finkelstein, Song, and Taubman (2014). Consistent with their findings on heterogeneous elasticities, I also find that older, less healthy, and unemployed households are more responsive to health insurance generosity. Furthermore, my paper builds on this empirical

[1] Another strand of the empirical literature measures the effect of health insurance coverage on worker turnover and job choice. Madrian (1994) uses the 1987 National Medical Expenditure Survey to estimate the relationship between employee turnover and health insurance coverage and finds that workers with employer-provided health benefits are 25 percent less likely to switch jobs. This result is also supported by Gruber and Madrian (1994) who show that policies that mandate extended coverage after job separation alleviate this "job lock" phenomenon.

literature by being able to study the effects of large-scale health insurance reform using a unified framework that endogenizes job-to-job and non-employment-to-employment transitions as well as general equilibrium responses involving firm vacancy posting decisions. Thus, the framework is designed to both generate key empirical predictions on the heterogeneous effects of employer-based health coverage on labor market transitions and at the same time understand how they endogenously add up to affect macroeconomic aggregates such as aggregate productivity and employment.

This paper is organized as follows. Section 1.2 presents the model. Section 1.3 describes the data and the calibration of the model. Section 1.4 discusses the the model's validation against quasi-experimental evidence and other non-targeted moments. Section 1.5 explains the welfare implications of the policy reform, and describes welfare-improving revenue generation through progressive taxation. In Section 1.6, I describe extensions to the model and various robustness checks. Finally, Section 1.7 concludes.

1.2 Model

1.2.1 Households

Time is discrete and runs forever. The economy is populated by $T+T^R$ overlapping generations of households. In each period, a new generation of households is born. Households can participate in the labor market for T periods, after which they spend T^R periods permanently in retirement. Agents survive up until the next period with conditional probability φ_t, which depends on age t. Death occurs deterministically at age $T+T^R+1$.

Agents are risk-averse and have access to a risk-free asset used to insure against idiosyncratic employment shocks and medical expenditure shocks, and to fund consumption during retirement. Households own and rent out capital to firms for use in production. Capital depreciates at rate δ and has a return r. The state vector of working-age households contains their age $t \in [1, T]$, employment status $l \in \{W, U, N\}$ where W denotes being employed, U denotes being unemployed, and N denotes non-participation, wealth holdings $a \in \mathcal{A} \equiv [\underline{a}, \bar{a}] \subseteq \mathbb{R}$, health status $h \in \mathcal{H} \equiv \left\{\underline{h}, \ldots, \bar{h}\right\}$, Medicaid non-financial eligibility $n \in \{0,1\}$ where $n=1$ indicates fulfillment of these requirements, and skill

$x \in \mathcal{X} = \{\underline{x}, \ldots, \bar{x}\}$. In addition, employed households are further characterized by having employer-provided group insurance $g \in \{0, 1\}$, where $g = 1$ indicates EPHI coverage, and the productivity of the firm they are matched with $y \in \mathcal{Y} \equiv \{\underline{y}, \ldots, \bar{y}\}$. The state vector of retired households consists only of their age $t \in [T + 1, T + T^R]$, wealth a, and health status h.

Preferences

Agents discount the future at rate $\beta \in (0, 1)$. Preferences of households at age t are given by

$$U(c_t, l_t, h_t) = u(c_t) - \mathbf{1}_{l_t \in \{W,U\}} \nu_t(h_t)$$

where c_t is consumption, l_t is employment status, and h_t is health in age t. The term ν is an age- and health-dependent utility cost of either working or actively looking for employment. Thus, non-participants do not incur this utility cost for all h and t.

Health and health insurance

Health h is stochastic in the model and evolves according to the Markov chain $\Gamma_t^h(h')$. The probability of transitioning into good or bath health depends on both age and current health status. Health affects the utility cost of participating in the labor market $\nu_t(h_t)$, as described in the previous section, and required medical expenditures $m_t(h)$ which depend on both health status and age.

Households have access to various insurance mechanisms to partially insure against medical expenditure risk. Households that are offered a job also receive an offer of EPHI with probability $\rho(x)$ which depends on their skill type. Employed households matched with jobs with no EPHI and non-employed households may be eligible for public health insurance through Medicaid if they fall below an asset threshold a^{mc} and income threshold z^{mc} , and meet non-financial eligibility requirements. Non-financial eligibility requirements capture the fact that Medicaid eligibility requires applicants to fall under designated eligibility categories such as being disabled, pregnant, or having children. Non-financial eligibility $n \in \{0, 1\}$ is met with probability γ^n in each period. Working-age households who do not have access to EPHI or Medicaid are considered uninsured.[2] Retirees are automatically

[2] In the baseline model, I abstract from private health insurance markets as the main focus of this study

enrolled in publicly-funded Medicare program. The insurance status i of agents can thus be summarized by

$$i\left(n,g,a,z\right)=\begin{cases}\text{Medicaid} & \text{if} \quad a<a^{mc},\, z<z^{mc},\, n=1,\, g=0,\, t\leq T\\ \text{EPHI} & \text{if} \quad g=1,\, t\leq T\\ \text{Medicare} & \text{if} \quad t>T\\ \text{uninsured} & \text{otherwise}\end{cases}$$

where a is the asset holdings and z is the total household gross income, that will be defined later, while a_{mc} and z_{mc} are the respective asset and income thresholds of the Medicaid program.

Let Ψ^{j} be the fraction of medical expenses covered by insurance type $j\in\{gr,\, mc,\, mr\}$ where gr denotes group EPHI, mc denotes Medicaid, and mr denotes Medicare. Out-of-pocket medical expenses can thus be summarized by:

$$oop_{t}\left(h,n,g,a,z\right)=\begin{cases}m_{t}\left(h\right)\left(1-\Psi^{mc}\right) & \text{if} \quad i\left(n,g,a,z\right)=\text{Medicaid}\\ m_{t}\left(h\right)\left(1-\Psi^{gr}\right) & \text{if} \quad i\left(n,g,a,z\right)=\text{EPHI}\\ m_{t}\left(h\right)\left(1-\Psi^{mr}\right) & \text{if} \quad i\left(n,g,a,z\right)=\text{Medicare}\\ m_{t}\left(h\right) & \text{if} \quad i\left(n,g,a,z\right)=\text{uninsured}\end{cases}$$

Taxes, transfers, and public programs

The government runs five insurance programs: Medicaid, Medicare, social security, unemployment benefits, and a safety net program to guarantee a consumption floor. The Medicaid program is available to working-aged households that do not have access to private health insurance coverage and meet eligibility requirements. The Medicare program, on the other hand, automatically enrolls retirees. The Social Security program pays retirees a benefit b^{R} until the household's death. Unemployment benefits b are paid to active job-seekers. Finally, the government guarantees households a consumption floor $\underline{c}$ using transfer T^{flr} to represent the option of households to resort to other safety net

is to understand effects of public and employer-provided health insurance policies on the labor market. Moreover, the fraction is people who purchase non-group private health insurance is small in the United States due to expensive prices of it.

government programs when they are subject to extraordinarily high medical expenditures. Government programs are funded with a consumption tax τ_c and a progressive income tax schedule $T^p(z)$ which depends on total household gross income z.

Evolution of skill

A household is born with skill x drawn from distribution Γ^x. Employed and non-employed households experience stochastic accumulation or depreciation of skills as in Ljungqvist and Sargent (1998). In particular, an employed household's skill increases by Δx with probability π_t^W while a non-employed household's skills depreciate by Δ_U with probability π^U in each period. Formally,

$$x' = \begin{cases} x + \Delta x & \text{with probability } \pi_t^W \\ x & \text{with probability } 1 - \pi_t^W \end{cases}$$

when employed.

$$x' = \begin{cases} x - \Delta x & \text{with probability } \pi^U \\ x & \text{with probability } 1 - \pi^U \end{cases}$$

when unemployed. Relative to Ljungqvist and Sargent (1998), I add an extra feature by allowing the probability of human capital accumulation when employed π_t^W to be age-dependent. It follows the recursion $\pi_t^W = \left(1 - \chi^W\right) \pi_{t-1}^W$ to capture a potential slowdown in skill accumulation over the lifecycle. Introducing household skill heterogeneity and on-the-job accumulation generates a motive for households to decide which jobs to accept and which jobs to leave. These decisions ultimately determine the distribution of match quality in equilibrium.

Matching in the labor market

The labor market is characterized by random search. There is a continuum of profit-maximizing firms owned by risk-neutral entrepreneurs that post vacancies.[3] Each period,

[3]Alternatively, one could assume that firms are owned by the risk averse households. In that case, we would have to solve for i) a portfolio choice problem of the households, and ii) fixed-point of distributing

vacancies and unemployed workers are randomly matched according to an aggregate matching function. The aggregate matching function $M(u,v)$ represents that number of matches when there are u unemployed workers searching for jobs and v vacancies posted. Let $\theta = \frac{v}{u}$ be the market tightness. Assuming that $M(\cdot)$ exhibits constant-returns-to-scale, we can define the probability of filling a vacancy as $q(\theta) = \frac{M(u,v)}{v} = M\left(\frac{1}{\theta}, 1\right)$ and the probability of receiving a job offer as $p(\theta) = \frac{M(u,v)}{u} = M(1,\theta)$. When a firm-worker pair meet in the labor market, match quality y is drawn from distribution Γ^y and is known to both parties. A worker may then decide to accept or reject the job offer. Each period a match exogenously dissolves with probability γ.

Household optimization

Working-age households can either be employed, unemployed, or out of the labor force (non-participants).

Define $s_N = s_U = (a, h, n, x)$ be the relevant state vector for a non-employed household of age t. The value of non-employment for such household is given by

$$\overline{U}_t(s_U) = \max_{d_l \in \{0,1\}} \{d_l U_t(s_U) + (1 - d_l) N_t(s_U)\}$$

where U_t is the value of unemployment, N_t is the value of non-participation, and d_l represents the decision to participate in or leave the labor force.

Define $s_W = (a, h, n, x, g, y)$ be the relevant state vector of an employed household of age t. The value of having a job offer with characteristics (y, g) in hand is given by

$$\overline{W}_t(s_W) = \max_{d_a \in \{0,1\}} \left\{d_a W_t(s_W) + (1 - d_a) \overline{U}_t(s_U)\right\}$$

where W_t is the value of employment and d_a represents the decision to accept or reject the job-offer.

An unemployed household receives unemployment benefits b. It chooses consumption vs savings and incurs disutility from actively looking for work. In the next period, it enters the labor market to search for jobs. With probability $p(\theta)$, an job offer characterized by match quality y and EPHI offer g is received and a choice of whether to accept or reject

dividends to them. Assuming that firms are owned by risk-neutral entrepreneurs simplifies the computation of the model along these dimensions.

the offer is made. If the household decides to reject the job offer, it can either remain unemployed or leave the labor force in the next period. With $1 - p(\theta)$ probability, no job offer arrives, and again the household can choose to continue looking for jobs as an unemployed next period or leave the labor force. We can now write the problem of the unemployed as:

$$\begin{aligned} U_t(s_U) = \max_{a' \geq \underline{a}} \quad & u(c) - \nu_t^U(h) + \varphi_t \beta \mathbb{E}_{h',n',x'} \Big[p(\theta) \mathbb{E}_{y',g'} \overline{W}_{t+1}(s'_W) \\ & + (1 - p(\theta)) \overline{U}_{t+1}(s'_U) \,|h, n, x \Big] \\ \text{s.t.} \quad & (1 + \tau_c) c + a' \leq b - oop_t(h, n, g, a, z) + (1 + r - \delta) a \\ & \qquad - T^p(z) + T^{flr}(h, n, g, a, z) \\ & z = b + (r - \delta) a \\ & T^{flr}(h, n, g, a, z) = \max \Big\{ 0, (1 + \tau_c) \underline{c} + oop_t(h, n, g, a, z) - b + T^p(z) \\ & \qquad - (1 + r - \delta) a + \underline{a} \Big\} \end{aligned} \tag{1.1}$$

A household that does not participate in the labor force is not eligible for unemployment benefits b. After choosing consumption vs savings, it can decide to enter unemployment or remain out of the labor force. The problem of the non-participant is given by

$$\begin{aligned} N_t(s_N) = \max_{a' \geq \underline{a}} \quad & u(c) + \varphi_t \beta \mathbb{E}_{h',n',x'} \left[\overline{U}_{t+1}(s'_U) \,|h, n, x \right] \\ \text{s.t.} \quad & (1 + \tau_c) c + a' \leq oop_t(h, n, g, a, z) + (1 + r - \delta) a - T^p(z) + T^{flr}(h, n, g, a, z) \\ & z = (r - \delta) a \\ & T^{flr}(h, n, g, a, z) = \max \Big\{ 0, (1 + \tau_c) \underline{c} + oop_t(h, n, g, a, z) + T^p(z) \\ & \qquad - (1 + r - \delta) a + \underline{a} \Big\} \end{aligned} \tag{1.2}$$

Finally, an employed household earns a wage income of $w(x, y)$ which depends on his skill as well as the match quality of the job he is employed in. The household chooses consumption vs savings this period and incurs a utility cost associated with working. At the beginning of the next period, with probability γ, the match is dissolved and the household must decide to either search for a job or move out of the labor force. If the match does not exogenously dissolve, the household can decide to either remain in the same job or quit into either unemployment or non-participation. The value of not separating from a match is given by

$$\widetilde{W}_t\left(s_W\right)=\max_{d_s\in\{0,1\}}\left\{d_sW_t\left(s_W\right)+\left(1-d_s\right)\bar{U}_t\left(s_U\right)\right\}$$

where d_s represents the choice of whether to stay in the current job or quit into non-employment. Then, the problem of the employed household can be written as:

$$W_t\left(s_W\right)=\max_{a'\geq\underline{a}}\quad u\left(c\right)-\nu_t^W\left(h\right)+\varphi_t\beta\mathbb{E}_{h',n',x'}\left[\left(1-\gamma\right)\widetilde{W}_{t+1}\left(s_W'\right)+\gamma\overline{U}_{t+1}\left(s_U'\right)|h,n,x\right] \tag{1.3}$$

$$\begin{aligned}\text{s.t.}\quad & \left(1+\tau_c\right)c+a'\leq w\left(x,y\right)-oop_t\left(h,n,g,a,z\right)+\left(1+r-\delta\right)a\\ &\qquad\qquad -T^p\left(z\right)+T^{flr}\left(h,n,g,a,z\right)\\ & z=w\left(x,y\right)\left(1-\tau\right)+\left(r-\delta\right)a\\ & T^{flr}\left(h,n,g,a,z\right)=\max\Big\{0,\left(1+\tau_c\right)\underline{c}+oop_t\left(h,n,g,a,z\right)-w\left(x,y\right)+T^p\left(z\right)\\ &\qquad\qquad\qquad\qquad -\left(1+r-\delta\right)a+\underline{a}\Big\}\end{aligned} \tag{1.4}$$

Households automatically enter retirement at age $T+1$. Let $s_R=(a,h)$ be the relevant state for retired households aged $t>T$. Retirees receive social security payments b^R every period until death. They also decide consumption vs savings but are no longer able to re-enter the labor market. Retirees are automatically enrolled into Medicare and are assumed to be ineligible for Medicaid.[4] The problem of the retired household for $T<t\leq T+T^R$ can be written as follows:

$$R_t\left(s_R\right)=\max_{a'\geq\underline{a}}\quad u\left(c\right)+\varphi_t\beta\mathbb{E}_{h'}\left[R_{t+1}\left(s_R'\right)|h\right] \tag{1.5}$$

$$\begin{aligned}\text{s.t.}\quad & c+a'\leq b^R+oop_t\left(h,0,0,a,z\right)+\left(1+r-\delta\right)a\\ & z=b^R+\left(r-\delta\right)a\end{aligned}$$

while for $t=T+T^R+1$, $R_t=0$.

[4]There are cases when Medicaid can supplement Medicare but we abstract from this case.

1.2.2 Firms

Value of a matched firm

Firms post vacancies in order to match with the unemployed seeking for jobs. A firm offers a compensation package consisting of predetermined wage rule $w(x, y)$ and health insurance plan $g \in \{0, 1\}$. Firms that offer health insurance pay a fraction $(1 - \omega_h)$ of the group plan premium w_h and pass on the rest of the cost to the worker. Firms face idiosyncratic costs to providing health insurance $\epsilon_i \in \{1,\ \epsilon_H\}$ where $\epsilon_H > 1$. This cost is modeled to be multiplicative to group health insurance premium w_h and is drawn from a distribution $\Gamma^{\epsilon}(x)$ which depends on the worker skill x that the firm is matched with.[5] The firm internalizes the job acceptance decision of the worker when it decides on whether or not it will offer health insurance. A match is formed when the worker accepts the job. Then, the firm rents capital and uses one unit of labor to produce $F(x, y, k)$ units of output. In the next period, if the worker decides to stay in the job, the match persists with $1 - \gamma$ probability.

The value of a firm with productivity y and health insurance administrative cost ϵ_i making an offer to an unemployed household aged t with state s_U can then be written as follows:

$$\bar{J}_t\big(s_U, y, \epsilon_i\big) = \max_{g \in \{0,1\}} \quad \Big\{ d_a\left(t,\ s_U,\ g=0\right) J_t\left(s_W,\ g=0\right),\ d_a\left(t,\ s_W,\ g=1\right) J_t\left(s_W,\ g=1\right) \Big\} \tag{1.6}$$

$$J_{T+1} = 0$$

where g represents the firm's decisions to offer health insurance or not. The value of a

[5] This is a rudimentary way to capture the fact that (1) some firms do not offer health insurance and (2) low-paying jobs are less likely to have EPHI. Dey and Flinn (2005) incorporate heterogeneity in households' strength of preference for health insurance coverage while Nakajima and Tuzemen (2016) allow firms to have idiosyncratic preference for offering health insurance. In the calibration of the model, ϵ_H is set to be large enough such that firms with a high administrative cost choose not to offer health insurance.

matched firm that offered health insurance is given by

$$\begin{aligned}
J_t\left(s_W,\ g=1\right) = \max_k \quad & F\left(x,y,k\right) - \tilde{w}\left(x,y\right) - w_h\epsilon_i - rk \\
& + \varphi_t \mathbb{E}_{h',n',x'}\left[\frac{1}{1+r}\left(1-\gamma\right) d_s\left(t+1,\ s_W'\right) J_{t+1}\left(s_W', g=1\right)\Big| h,n,x\right] \\
\tilde{w}\left(x,y\right) = & \ w\left(x,y\right) - \omega_h w_h \epsilon_i \\
a' = & \ a_t'\left(s_W\right) \\
J_{T+1} = & \ 0
\end{aligned} \tag{1.7}$$

while a matched firm that did not offer health insurance is given by

$$\begin{aligned}
J_t\left(s_W,\ g=0\right) = \max_k \quad & F\left(x,y,k\right) - w\left(x,y\right) - rk \\
& + \varphi_t \mathbb{E}_{h',n',x'}\left[\frac{1}{1+r}\left(1-\gamma\right) d_s\left(t+1,\ s_W'\right) J_{t+1}\left(s_W', g=0\right)\Big| h,n,x\right] \\
a' = & \ a_t'\left(s_W\right) \\
J_{T+1} = & \ 0
\end{aligned} \tag{1.8}$$

where $F\left(x,y,k\right) = A\left(x,y\right) f\left(k\right) = \left[\alpha x^\rho + \left(1-\alpha\right) y^\rho\right]^{\frac{1}{\rho}} k^\psi$. Here, I assume that health insurance contracts are with commitment and that output is shared between the firm and the worker so that $w\left(x,y\right) = \omega F\left(x,y,k\right)$ where $\omega \in \left(0,1\right)$ is the worker's *piece-rate*.[6]

Value of posting a vacancy

Firms pay a fixed cost κ in order to post a vacancy. With $q\left(\theta\right)$ probability (the job-filling rate), the firm meets a worker from the pool of job-seekers. The match quality y is then revealed to the firm-work pair and the firm draws a health insurance administrative cost shock ϵ_i. After the firm makes a compensation offer, the worker can decide to either accept the job or not as indicated by decision rule $d_{a,t}\left(s_W\right)$. The value of posting a vacancy can

[6] I make this assumption to simplify wage determination as in Nakajima (2012), Menzio, Telyukova, and Visschers (2016) and Herkenhoff (2017). In a set-up where workers and firms bargain over wages, these wages become a function of all individual characteristics of the worker and firm. This is computationally burdensome as it requires another fixed-point problem to be solved as in Krusell, Mukoyama, and Sahin (2012).

be written as follows:

$$V_f = -\kappa + \int\limits_{(\epsilon_i, y)} \left[q(\theta) \int\limits_{s_U} \frac{1}{1+r} \bar{J}_t(s_U, y, \epsilon_i)\, d\mu^U(s_U) \right] d\mu^J(\epsilon_i, y) \tag{1.9}$$

where μ^U is the distribution of unemployed households. Unmatched firms can enter the labor market and post vacancies by paying the cost, and so the free entry condition implies that in equilibrium, the value of a posting vacancies is driven down to zero $V_f = 0$.

Group health insurance premium

Firms purchase health insurance plans from insurance firms that act competitively. I assume that all active firms offering health insurance pool the medical risk of all covered employees across all firms. Thus, the zero-profit condition for insurance firms implies that the group insurance premium can be written as:

$$w_h = \frac{\sum_t \sum_j \mu^W(t, h_j, g=1)\, m_t(h)\, \Psi^{gr}}{\mu^W(g=1)}$$

where $\mu^W(t, h_j, g=1)$ is the measure of employed households aged t, with health status h_j, and possess employer-provided health insurance. This implies that total premia collected is equal to the average medical expenditure of all covered employees.

1.2.3 Government

Let $s \equiv (t, a, h, n.g, x, y)$ be the household's state vector. The government budget constraint is given by

$$\begin{aligned} \tau_c \int c(s)\, d\mu(s) + \int z(s)\, T^p(z(s))\, d\mu(s) = & \int_{t \le T} \Big[b \times \mathbf{1}_{\{l=U\}} \\ & + \Psi^{mc} m_t(h) \times \mathbf{1}_{\{a<a_{mc},\ z<z_{mc}\}} d\mu(s) \Big] \\ & + \int_{t>T} \left[b_R(e) + \Psi^{mr} m_t(h) \right] d\mu(s) \\ & + \int T^{flr}(s)\, d\mu(s) \end{aligned} \tag{1.10}$$

where the left-hand-side represents total consumption tax and income tax collected and the right-hand-side represents government spending on unemployment insurance and Medicaid for working-age households, social security and Medicare fore retired households, and the consumption floor guarantee for all households.

1.2.4 Stationary equilibrium

A recursive stationary equilibrium consists of a set of value functions
$\{W_t(s_W),\ U_t(s_U),\ N_t(s_N)\}_{t\in[0,T]}$, $\{R_t(s_R)\}_{t\in[T+1,T+T^R]}$,
V_f, $\left\{\bar{J}_t(s_U, y, \epsilon_i),\ J_t(s_W)\right\}_{t\in[0,T]}$,
a set of policy functions $\left\{a'_t(s_W),\ a'_t(s_U),\ a'_t(s_N),\ d_{s,t}(s_W),\ d_{a,t}(s_W),\ d_{l,t}(s_N)\right\}_{t\in[0,T]}$,
$\{a'_t(s_R)\}_{t\in[T+1,T+T^R]}$, interest rate r, health insurance premium w_h, vacancies v, unemployment rate u, market tightness θ, and the distribution of agents across individual states μ such that:

1. Household optimization: Given aggregate market tightness θ, interest rate r, wage rule $w(x,y)$, and government policy $\left\{T^p,\ \tau_c,\ b,\ b^R,\ \Psi^{mc},\ a^{mc},\ z^{mc},\ \gamma^n,\ \Psi^{mr},\ T^{flr}\right\}$, the value functions
$\{W_t(s_W),\ U_t(s_U),\ N_t(s_N)\}_{t\in[0,T]}$, $\{R_t(s_R)\}_{t\in[T+1,T+T^R]}$, solve (1.1), (1.2), (1.3), and (1.5), and
$\{a'_t(s_W),\ a'_t(s_U),\ a'_t(s_N),\ d_{s,t}(s_W),\ d_{a,t}(s_W),\ d_{l,t}(s_N)\}_{t\in[0,T]}$,
$\{a'_t(s_R)\}_{t\in[T+1,T+T^R]}$ are the associated optimal decision rules.

2. Firm optimization: Given interest rate r, wage rule $w(x,y)$, insurance premium w_h, the unemployed worker's job acceptance decision rule $\{d_{a,t}(s_W)\}_{t\in[0,T]}$, and the employed worker's quit decision rule $\{d_{s,t}(s_W)\}_{t\in[0,T]}$, the value functions
$\left\{\bar{J}_t(s_U, y, \epsilon_i),\ J_t(s_W)\right\}_{t\in[0,T]}$ solves (1.6), (1.7), and (1.8) with associated health insurance offer decision rule $d_g(s_U, y, \epsilon_i)$. Given job-finding rate q, the distribution of unemployed μ^U, V satisfies (1.9).

3. Free entry of firms: The number of vacancies posted v satisfies the free entry condition $V = 0$.

4. Asset market clearing: The asset market equilibrium condition

$$K = \int a_i di$$

holds where $K = \int_{x,y} k(x,y)\, d\mu^J(x,y)$ and $k(x,y)$ satisfies firm optimization $r = F_k'(x,y,k)$ given r.

5. Matching: The job-finding rate p and job-filling rate q is determined by market tightness θ as described in the model section.

6. Health insurance market: The employer health insurance premium is given by

$$w_h = \frac{\sum_t \sum_j \mu^W(t, h_j, g=1)\, m_t(h)\, \Psi^{gr}}{\mu^W(g=1)}$$

7. Government budget constraint: Consumption tax τ_c and income tax $T^p(z)$ balances the government budget constraint given in (1.10).

8. Consistency: $\mu(s)$ is the invariant distribution implied by job-finding rate $p(\theta)$, exogenous job separation probability γ, and household optimal decision rules.

1.3 Data and Calibration

1.3.1 Data

MEPS

To correctly measure the level of health risk individuals face, I need data on the distribution and transition of medical expenditures as well as health status over the lifecycle. To obtain this, I use waves 1999 to 2014 of the Household Component of the Medical Expenditure Panel Survey (MEPS) where in each panel, biannual interviews are conducted for up to two and a half years. This provides detailed information not only on health insurance coverage (both public or private) but also self-reported health status and medical expenditures. Medical expenditures are further broken down into various sources such as private insurance, Medicaid, etc. I use this data to calibrate the level of medical risk as well as health insurance coverage ratios over the life cycle in the model. Appendix A.1 provides detailed explanation on sample selection and construction of each of these data moments.

SIPP

The United States Census Bureau's Survey of Income and Program Participation (SIPP) is a longitudinal survey that follows individuals for a duration of up to five years, with interviews being held in four-month intervals called waves. In each interview, respondents are asked questions about their income, employment/labor force status, government transfer receipts, and source of health insurance coverage (if any) over the previous four months not including the interview month. Topical modules that are administered on a less frequent basis but contain information on respondents' wealth holdings are also available. I use the SIPP to compute various calibration inputs and other non-targeted moments to validate the model against among which include (1) earnings distribution, (2) employment rates and labor market transitions, (3) access to and participation in employer-provided health insurance, (4) access to and participation in publicly-funded health insurance, and (4) self insurance in the form of asset holdings. Unless otherwise stated, all calibration targets computed from the SIPP are derived from the 2004 panel. Appendix A.1 provides detailed explanation on sample selection and construction of each of these data moments.

1.3.2 Calibration

Demographics A period in the model is one quarter. An individual is born at age 25 and lives to a maximum of 60 years. During the first 40 years of life, agents can participate in the labor market or choose to leave the labor force. Retirement is mandatory at age 65. Agents spend 20 years in retirement and live until age 84 after which death occurs for all agents. This implies that $T = 160$ and $T^R = 80$. The survival probabilities φ_t are chosen to match the average survival probability of males and females as reported in the 2006 Social Security Administration life tables. The survival probabilities imply a ratio of workers to retirees is of 3.0 in the model.

Health and medical expenditures I use the 1999 to 2014 waves of Medical Expenditure Panel Survey (MEPS) data to estimate medical expenditure shocks, following Jeske and Kitao (2009). I restrict the sample to household heads aged 25 to 84. The number of health states is set to three in the model $h \in \{\text{excellent, average, poor}\}$. For each age, I divide the sample into three medical expenditure bins:$\{50\%, 40\%, 10\%\}$. I consider an agent whose medical expenditures is below the median of the distribution of medical expenditures for his age as having excellent health while an agent whose medical expenditures is above the 90th percentile of distribution for his age as having poor health. To obtain the function $m_t(h)$, for each health status bin h, I fit mean medical expenditures over age t with a cubic function of age.[7] It has been well-documented that the MEPS understates medical expenditures when compared data from the National Health Expenditure Account (NHEA) (Selden, 2001). To address this, I scale medical expenses so that the ratio of total medical expenditure to GDP is 13% which is consistent with levels reported in the NHEA in 2006 as in Pashenko and Porapakkarm (2013). The medical expenditure function $m_t(h)$ is plotted in Figure 1.1. The health transition matrix $\Gamma_t^h(h')$ is set by calculating the fraction of agents that are age t with health status h who move into bin h' in the next period.

Health insurance The coinsurance rates $1-\Psi^j$ for insurance type $j \in \{gr, mc, mr\}$ are estimated from the same MEPS sample. I divide the sample into working-age population (age 25 to 64) and a retiree population (age 65 to 84). First, I compute coinsurance rates for the working-age population. To compute coinsurance rates for EPHI, I first restrict the sample to those that report being covered *exclusively* under EPHI for more than six months during the year.[8] Since the MEPS provides a breakdown of total medical expenditures by source of payment, I am able to identify the amount of medical expenditures covered by the respondent's group plan. I find that the fraction of total medical expenditures covered by EPHI is 65 percent in the data. Thus, I set $\Psi^{gr} = 0.65$. I implement the same procedure for Medicaid and set $\Psi^{mc} = 0.80$. The retired population in the U.S. receive substantial health coverage from both the Medicare and Medicaid programs. An average of 30 percent of retiree medical expenditures is covered by Medicaid while 60 percent is covered by Medicare.[9] Hence, I set the retiree health insurance coverage rate in the model to be $\Psi^{mr} = 0.90$. The fraction of health premiums firms passed onto the worker by firms

[7]Medical expenditures are in 2006 dollars and is scaled by the ratio of mean wages (1.49) in the model to mean wages in the MEPS ($35,417).

[8]The resulting estimates are robust to alternative values of this threshold.

[9]In computing for aggregates program statistics, I attribute publicly funded retiree medical expenditures to the Medicaid and Medicare programs accordingly.

Figure 1.1: Lifecycle medical expenditures by health status

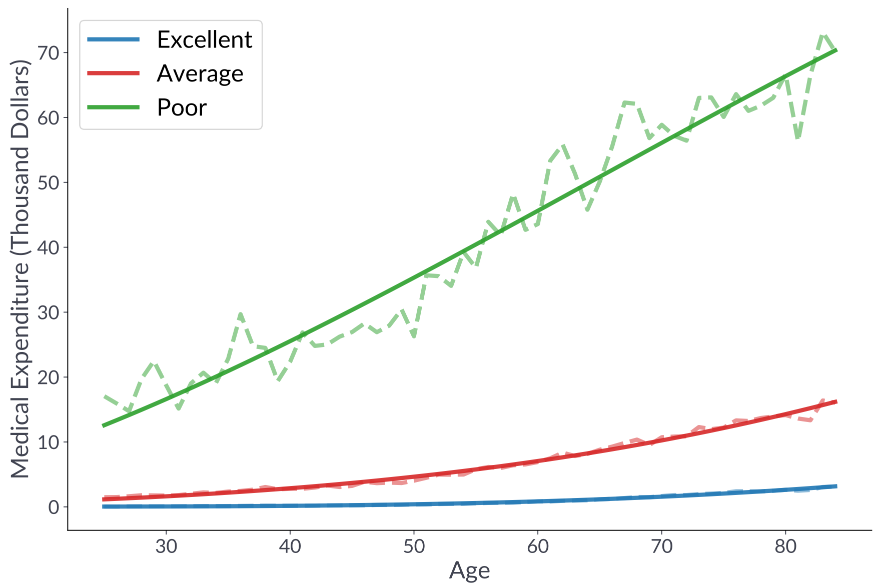

Note: This figure plots the average medical expenditures by health status estimated from the 1999-2004 waves of the MEPS. Dashed lines represent the raw data and solid lines are obtained by fitting the raw data with a cubic function of age. I convert these amounts in the data into model units using the ratio of mean wages in the model and mean wages in the MEPS.

ω_h is set to be 0.24 from Pashchenko and Porapakkarm (2013).

Medicaid recipients are required to meet certain financial eligibility requirements. The income threshold z^{mc} of Medicaid is set to be 65% of the Federal Poverty Line (FPL) which is the median value among all states in 2006. The asset threshold is estimated to match the fraction of working-aged population receiving Medicaid of 9.2% while the probability of meeting non-financial eligibility requirements is set to match the fraction of uninsured households with non-positive wealth which is 32%, both moments being calculated from the SIPP. Finally, the distribution of health insurance administrative cost shocks $\Gamma^{\epsilon}(x)$ is used to target the joint distribution of EPHI coverage and wages from the SIPP.

Preferences I specify preferences to take the following form: $u(c_t) = \frac{c_t^{1-\sigma}}{1-\sigma}$ and $\nu_t(h_t) = \nu_b + \nu_m(h) \times t$. I set risk aversion parameter σ to be equal to 3 which is within the range commonly used in the life-cycle literature. The cost of participating in the labor force (being employed or looking for work) is parameterized to be linear in age with a slope that varies with health. The constant ν_b is used to target an average labor force participation rate (LFPR) of 79% while the slope parameter $\nu_m(h)$ is used to target labor force participation over the lifecycle across individuals with different levels of health. The LFPR and its age profile by health status bin is calculated from the CPS using information on self-reported health status and labor force participation.[10] The parameters $\{\nu_m(h_1), \nu_m(h_2), \nu_m(h_3)\}$ are set to match the ratio of LFPR of young and old households for each health bin. Specifically, I choose these three parameters to match the ratio of LFPR of those aged 25-29 to LPPR of those aged 60-64 for each of the three health bin. Finally, discount factor β is set to match the median value of asset holdings relative to quarterly before-tax labor income distribution calculated from the SIPP 2004 Panel Wave 6 topical module covering 2005.

Government taxes and transfers I parameterize $T^p(z)$ to be piece-wise linear as in McGrattan and Prescott (2017). This is constructed by linearizing $T^p(z)$ on AGI income intervals $\left[\underline{z}_j, \bar{z}_j\right], j = 1, \ldots, I$ as follows:

$$\begin{aligned} T^p(z) &\simeq T_j^{p'}(\bar{z})\, z - \left\{T_j^{p'}(\bar{z}) - T_j^p(\bar{z})/\bar{z}\right\}\bar{z} \\ &= \beta_j^{Tax} z + \alpha_j^{Tax} \end{aligned} \tag{1.11}$$

where $\bar{z}$ is the midpoint in each bracket. The term $T_i^{p'}(\bar{z})$ is simply the local marginal tax rate at $\bar{z}$ and $T_i^p(\bar{z})/\bar{z}$ is simply the local average tax rate at $\bar{z}$. I compute for each bracket's marginal and average tax rates using data from the CPS March Supplement from 2007 with estimates for 2006 and the National Bureau of Economic Research's (NBER)

[10] Respondents report being in excellent, very good, good, fair or poor health. For comparability to the model, I let "excellent" and "very good" correspond to h = excellent, "good" and "fair" correspond to h = average, and poor correspond to h = poor. For the 2007 CPS March Supplement, the fraction of adults age 25-64 reporting excellent/very good health is 63 percent, good/fair health is 25 percent, poor health is 12 percent, with labor force participation rates of 86 percent, 78 percent, and 43 percent respectively. The fact that the bin size for h = excellent and h = average do not correspond exactly with the model has little effects on the calibration because the LFPR between the two groups is very similar. The fraction of individuals who report poor health matches well with the fraction h = poor in the model.

Figure 1.2: Estimated income tax function

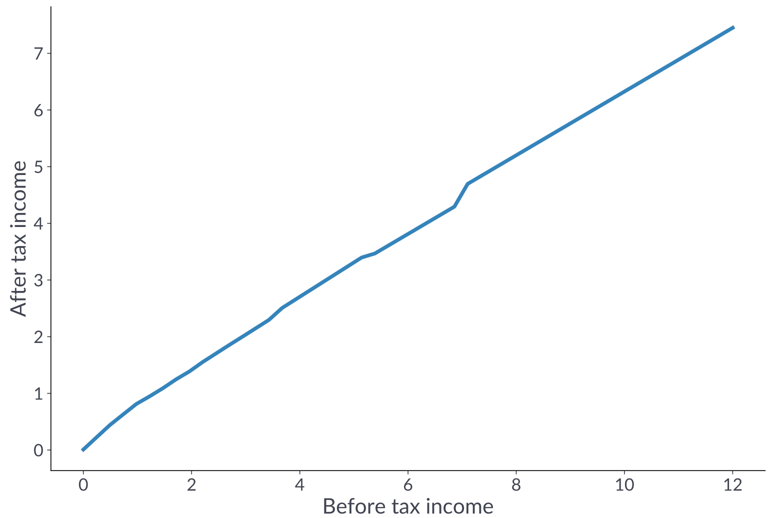

Note: This figure plots after-tax income as a function of before-tax income resulting from the income tax schedule estimated using the 2007 CPS Annual Social and Economic Supplement (ASEC). Details of the procedure can be found in the main text and in Appendix A.1.

TAXSIM program. Figure 1.2 plots the estimated tax function. Details of this calculation can be found in the Appendix A.1.

Consumption tax τ_c is set to match a ratio of government spending to GDP of 19.7 percent as reported in the 2006 National Income and Product Accounts (NIPA). Unemployment insurance benefit b and social security payments b^R are set to match the ratio of aggregate unemployment benefits to GDP of 0.22 percent and 3.93 percent respectively. Finally, the consumption floor $\underline{c}$ is used to target that 29% of the population have with non-positive wealth as calculated from the SIPP 2004 Panel Wave 6 topical module covering 2005.

Firms, production technology and wages I use the following production function

$$F(x, y, k) = A(x, y)\, k^{\psi}$$

where the TFP term $A(x,y)$ is a CES aggregator of worker skill x and firm productivity y:

$$A(x,y) = [\alpha x^{\rho} + (1-\alpha) y^{\rho}]^{\frac{1}{\rho}}$$

I set $\alpha = 0.60$ and the complementarity parameter $\rho = -0.90$.[11] This implies an elasticity of substitution of 0.53, i.e. worker skill and firm productivity are complements in production. Capital is set to depreciate at a quarterly rate of $\delta = 0.015$ which implies an annual depreciation rate of 6 percent. The exponent to capital ψ is used to target a capital output ratio of 3.2 calculated from the NIPA and Flow of Funds (FoF) for the year 2006. Here, I assume that the stock of capital includes (1) private fixed assets, (2) inventories, and (3) real estate/land. I assume that wages are paid as a piece-rate of output ω so that

$$w(x,y) = \omega F(x,y,k)$$

The parameter ω is set to be the ratio of wages and salaries to GDP in 2006 which is 0.44 from the NIPA. Finally, the firm productivity distribution Γ^y is restricted to be an approximation of a Weibull distribution with mean 1, scale parameter λ^y and shape parameter γ^y. The scale parameter λ^y is set to 2, while the shape parameter γ^y is set to match average an average tenure of 18 quarters. The length of tenure is informative of the distribution of firm productivity. A distribution where there is high certainty of drawing high productivity firms would encourage workers to quit more often, thus shortening tenure.

Labor market The labor market matching function is $M(v,S) = \frac{uv}{\left[u^{\zeta}+v^{\zeta}\right]^{1/\zeta}}$ as in den Haan, Ramey, and Watson (2000). This CES functional form of the matching function implies that both the job finding rate $f(\theta) = \theta(1+\theta^{\gamma})^{-1/\gamma}$ and the vacancy filling rate $q(\theta) = (1+\theta^{\gamma})^{-1/\gamma}$ are between 0 and 1. The elasticity of the matching function ζ is set to be 0.95, well within the range of estimates used in the literature.[12]

[11] I assume that worker skill and firm productivity are complements in production. The complementarity parameter is set to be $\rho = -0.90$ which is the value estimated by Lise, Meghir, and Robin (2016) for college graduates.

[12] See Schaal (2017) and Mitman and Rabinovich (2015). Also note that the matching elasticity parameter is typically estimated using data on cyclical movements of the unemployment-employment (UE) rate and aggregate labor market tightness as in Shimer (2005) and Menzio and Shi (2011). Since the model is stationary, I do not attempt to estimate this parameter as in Menzio, Telyukova, and Visschers (2016).

The vacancy cost κ is calibrated to match the average unemployment rate of 4.8 percent, while the probability of exogenous separation γ is used to target the average employment-unemployment (EU) transition of 1.4 percent, both of which are calculated from the 2004 Panel of the SIPP for years 2005-2007.

Worker skill The incremental skill increase while employed and skill decrease while unemployed is set to be $\Delta x = 0.1$. The initial skill accumulation probability π_0^W is used to match the economy-wide p90-p10 wage ratio of 6.30. Skill growth deceleration over the lifecycle is determined by parameter χ. This parameter is used to match the concavity of the lifecycle profile of log wages as in Jung and Kuhn (2018).
The initial distribution of skills that newborn agents draw from is restricted to be an approximation of a Weibull distribution with mean 1, scale parameter λ^x and shape parameter γ^x. These two parameters are used to jointly match (1) the ratio of age 25 mean wage to the economy-wide mean wage of 0.68 and (2) the p90-p10 ratio of age 25 wages of 5.6. All wage moments are calculated from the SIPP 2004 Panel for the year 2006, averaged across all months.

Borrowing Limit $\underline{a}$ Kaplan and Violante (2014) calculate the median value of the credit limit to quarterly labor income ratio for households aged 22 to 59 as 74 percent using Survey of Consumer Finances (SCF) data. I choose the borrowing limit parameter $\underline{a}$ so that the median value of the ratio of $\underline{a}$ to after-tax quarterly labor income in the model is 0.74.

1.3.3 Calibration results

The results of the calibration are summarized below. Internally calibrated parameters are shown in Table 1.1 while externally calibrated parameters are shown in Table 1.2.

1.4 Validation

In this section, I discuss the model's implications for several important but non-targeted data moments. I first compare the model's predicted labor supply elasticity with respect to the generosity of public health insurance with available quasi-experimental evidence. I also validate the model against differential lifecycle job finding rates between individuals with

Table 1.1: Internally calibrated parameters

Parameter	Explanation	Value	Target	Data	Model
			Health insurance		
a^{mc}	Asset threshold	4.8	Medicaid take-up	9.2%	8.3%
γ^{n}	Non-fin elig.	0.89	Uninsured non-pos. wealth	32%	28%
$\Gamma^{\epsilon}(x)$	Admin cost	-	EPHI by wage quintile	$[.5, .75, .85, .9, .9]$	$[.6, .72, .83, .9, .95]$
			Preferences and borrowing constraints		
ν_b	Disutil, level	1.2	Labor force part.	79%	75%
$\nu_m(h)$	Disutil, slope	[.002, .003, .007]	$\frac{\text{LFP}R_{25,29}}{\text{LFP}R_{60,64}}$ by h	$[1.3, 1.4, 2.5]$	$[1.1, 1.2, 1.7]$
β	Discount factor	0.992	Median asset-to-inc	0.78	0.78
$\underline{a}$	Borrowing limit	1.33	Median credit-lim-to-inc	0.74	0.74
			Taxes and transfers		
$\underline{c}$	Cons. floor	0.3	Frac. with non-pos. wealth	0.29	0.31
b	Unemp ben	0.28	UI/GDP ratio	0.22	0.28
b^{R}	Social security	0.8	SS/GDP ratio	3.93	4.26
τ_c	Cons. tax	3.4%	G/GDP ratio	19.07	19.07
			Firms, wages, and production technology		
ψ	Capital exponent	0.34	Cap-output ratio	2.9	3.4
γ^{y}	Firm dist. shape	3.2	Tenure	18	22
			Labor Market		
κ	Vacancy cost	0.82	Unemp rate	4.8%	4.9%
γ	Job destruction	0.012	E-U rate	1.4%	1.5%
			Skill Accumulation		
λ^{x}	Worker dist. scale	1.3	$\frac{\text{Wage}_{25}}{\text{Wage}_{\text{all}}}$	0.68	0.43
γ^{x}	Worker dist. shape	1.8	Wage_{25} p90-p10	5.6	4.8
π_0^{W}	Prob. skill growth	0.2	Wage p90-p10	6.30	4.54

Note: For the explanations of parameter choices and data moment constructions please refer to the main text and Appendix A.1.

Table 1.2: Externally calibrated parameters

Parameter	Explanation	Value	Data
		Health insurance	
Ψ^{mc}, Ψ^{gr}, Ψ^{mr}	Coinsurance rates	[.80, .65, .90]	MEPS
$m_t(h)$	Medical expenditures	-	MEPS
$\Gamma_t^h(h')$	Health transition	-	MEPS
ω_h	Premium share	0.24	Pashchenko and Porapakkarm (2013)
z^{mc}	Medicaid inc threshold	0.8	65% of FPL in 2006
		Preferences	
σ	Risk aversion	3	-
		Taxes and transfers	
$T^p(z)$	Income tax sched.	-	CPS, NBER Taxsim
		Firms, wages, and production technology	
λ^y	Firm dist. scale	2	-
α	Share of skill	0.60	Lise, Meghir, and Robin (2016)
ρ	Complementarity param	-0.90	Lise, Meghir, and Robin (2016)
δ	Capital depreciation	0.015	Annual depreciation 6 pct.
		Labor market	
ζ	Matching func elasticity	0.95	-
		Skill Accumulation	
Δx	Skill gain/loss	0.1	-
χ	Skill growth dep.	0.008	Jung and Kuhn (2018)

Note: For the explanations of parameter choices and data moment constructions please refer to the main text and Appendix A.1.

and without health coverage. These moments are relevant because they are informative of the incentive costs of expanding health insurance coverage which come in the form of higher reservation wages and consequently, lower job-finding rates and employment rates. Second, I present how the model compares to the data on the consequences of job loss which include the likelihood of losing health insurance, the consumption drop, and the persistence of earnings losses. In addition, I also validate the model against other moments of the asset-to-income distribution. The magnitude of the negative consequences of job loss and the amount of self-insurance available to individuals are key to measuring the insurance benefits of delinking health insurance with employment status. For example, if the probability of losing health insurance in the model were to be too large relative to the data, then the insurance benefit of universal healthcare would be overstated. The following sections present the results of these exercises.

1.4.1 Labor supply elasticity to public health insurance generosity

I first validate my model against quasi-experimental evidence obtained by papers that study the effect of public health insurance coverage on the labor market behavior of households. These studies treat sudden expansions or contractions of health insurance by governments as natural experiments to measure how outcome variables such as employment and earnings respond to such changes in the generosity of public health insurance programs. It is important for the model to generate a reasonable labor supply elasticity with respect to health insurance coverage in order to properly measure the incentive costs of the policy reform. Furthermore, the magnitude of household response to policy also affect the extent to which firms will adjust vacancies. Finally, to the extent possible, I also compare the model's predictions on the heterogeneous labor supply responses of households that differ in age, health status, and employment.

Garthwaite, Gross, and Notowidigdo (2014) exploit the unexpected discontinuation of Tennessee's Medicaid (TennCare) expansion program in 2005 to identify the causal effect of public health insurance on labor supply and find large effects. Dague, DeLeire, and Leninger (2017) use the Wisconsin BadgerCare enrollment cap in 2009 and find large but more modest effects, while Baicker, Finkelstein, Song, and Taubman (2014) use the Oregon Medicaid expansion in 2008 and find no effect.

My model predicts large labor supply responses, the magnitude of which lie in between the range of estimates in the literature. Furthermore, the model finds that the composition of the subpopulation targeted in these experiments explain to some degree the dispersion in estimates found in the data. Finally, I also find substantial heterogeneity in labor supply responses to public health insurance generosity that are qualitatively consistent with the empirical findings, where older, less healthy, and unemployed workers are more responsive to policy.

Tennessee disenrollment

Garthwaite, Gross, and Notowidigdo (2014) exploit a reform in Tennessee's Medicaid system (TennCare) to study the labor supply effects of public health insurance. In 2005, TennCare's expansion was abruptly discontinued, leading to the loss of health insurance of approximately 170,000 adults (roughly 4 percent of the state's non-elderly population). The discontinuation was essentially a reversion of eligibility requirements back to that of traditional Medicaid. Thus, TennCare disenrollment affected a subpopulation that was not traditionally eligible for Medicaid – childless adults who also had higher incomes than income eligibility threshold. They find that for the population of childless adults affected by the disenrollment, a 7.3 percentage point decrease in the coverage rate resulted in a 4.6 percentage point increase in employment rate. This is equivalent to around a 60 percentage point increase in employment among disenrollees stemming from the loss of public health coverage.

In order to implement such policy change in my model, I first solve a stationary equilibrium to match Tennessee's employment rate of 64 percent and public health coverage rate of 23 percent for childless adults aged 25-64 in 2004, pre-reform.[13] I match public coverage rates by increasing the Medicaid threshold z^{mc} to 400 percent of FPL and by increasing the non-financial eligibility probability γ^n.[14] Employment rates are matched by adjusting vacancy posting cost κ. To make the exercise comparable to the TennCare reform, I first

[13]Data is obtained from the 2005 CPS March Supplement.

[14]TennCare allowed families with income up to 400 percent of the FPL to enroll and receive premium subsidies (Conover and Davies, 2000). According to Wooldridge et al. (1996), more than 40 percent of TennCare enrollees had incomes above 100 percent of the FPL, while 6.3 percent had incomes between 200 and 400 percent of FPL in 1995.

implement a sudden disenrollment modeled as an unexpected decrease in the Medicaid threshold down to 100 FPL and a decrease of the non-financial eligibility shock γ^n such that public coverage declines by 7.3 percentage points. This mimics a reversion towards traditional Medicaid with strict income requirements and non-financial eligibility rules. Second, I hold taxes and prices fixed after the policy change given that it only affected a small fraction of state population. I then compare the change in employment rate predicted by the model to those found by Garthwaite, Gross, and Notowidigdo (2014). The first row of Table 1.3 shows that the experiment in my model results in a 2.7 percentage points increase in employment rates – large but lower than the 4.6 percentage points increase found in the study.[15] The model implied response equivalent to more than a 30 percent increase in the employment rate among disenrollees.

Garthwaite, Gross, and Notowidigdo (2014) also study the impact of disenrollment on the employment rates of different age- and health- subgroups. They find that the labor supply of older and less healthy adults were more elastic to changes in the public health insurance coverage. As seen in Table 1.3, the model captures these dimensions of heterogeneity qualitatively albeit with less dispersion in elasticities across the subpopulations.[16]

Wisconsin enrollment cap

Dague, DeLaire, and Leininger (2017) study a sudden imposition of an enrollment cap in Wisconsin's BadgerCare Plus Core Plan – a Medicaid expansion program which included non-traditional Medicaid beneficiaries such as households with incomes below 200 FPL, which is a more relaxed threshold. Public enrollment began in July 2009 but was abruptly terminated in October of the same year due to budgetary constraints. Applications received after October 09, 2009 were put on a waitlist but most were never granted coverage in the end. Using a regression discontinuity design, they compare the employment outcomes of applications before and after the cut-off and find that applications who

[15] The sample of disenrollees in Garthwaite, Gross, and Notowidigdo (2014) consist of individuals who were required to actively participate in "re-verification" process to determine if they were eligible for Medicaid. This implies that their sample may have a stronger preference for health insurance – a feature the model abstracts from.

[16] For the subpopulation "Age 21-39", I simulate the model only for agents between 25 and 39 as agents are born at age 25.

Table 1.3: Tennessee disenrollment policy and its effects on employment

Employment rate (pp)	**GGN (2013)**	**Model**
State level	4.6	2.7
Heterogeneity by age		
Age 21-39	1.0	1.5
Age 40-64	6.0	3.1
Heterogeneity by health status		
Excellent health	2.0	2.3
Good/poor health	5.3	3.4

Note: In 2005, Tennessee's Medicaid system's (TennCare) expansion was abruptly discontinued, leading to the loss of health insurance of approximately 170,000 adults (roughly 4 percent of the state's non-elderly population). This table compares the effects of this policy change on the i) state level employment rate and ii) employment rates of different age- and health-subgroups, estimated by Garthwaite, Gross, and Notowidigdo (2014) for childless adults and estimated by my model.

received Medicaid were 5.2 percentage points less likely to be employed relative to the control group of waitlisted applications who eventually did not receive coverage. They also study the heterogeneous effects of Medicaid coverage to a subset of applicants and find a larger magnitude of labor supply responses among older and unemployed households. The first column of Table 1.4 reports findings of Dague, DeLaire, and Leininger (2017) on the percentage point employment rate differences between applicants before and after the enrollment cap.

I implement a similar experiment in my model in the following steps. I first take my sample as uninsured agents with income below 200 FPL. I simulate the enrollment cap by unexpectedly enrolling a fraction of the sample into Medicaid by changing their household state variables from insured by Medicaid to uninsured (treatment group), and leaving the rest as uninsured (control group). I then run a regression of a dummy variable of employment one year after the policy on a dummy variable of Medicaid eligibility using model-simulated data to estimate the difference in employment outcomes between the treatment and control group. I also run the same regression on subpopulations of the sample by age and employment status. The second column of Table 1.4 shows that my model predicts an average labor supply that is twice the estimates of Dague, DeLaire, and Leininger (2017). However, consistent with their findings, the model predicts larger elasticities for older workers and the unemployed. This also suggests that it is the unemployment-employment margin that is largely affected by the generosity of public health insurance programs.

Oregon lottery

In 2008, the state of Oregon implemented a limited expansion of its Medicaid program. Using a lottery, around 30,000 low-income uninsured adults from a roughly 90,000-long waitlist were given the opportunity to apply for coverage. Baicker, Finkelstein, Song, and Taubman (2014) use this lottery to conduct a randomized study on the effects of public health insurance on labor supply. To enroll in Medicaid, lottery winners must have completed the application process and met eligibility criteria such as being aged $19-64$, otherwise ineligible for Medicaid or other public health insurance programs, uninsured for six months, have income below the FPL and own assets below $2000. Using administrative data on those who participated in the lottery, they find that Medicaid enrollment leads to

Table 1.4: Wisconsin enrollment cap and its effects on employment

Emp. rate (Treatment vs. Control)	**DDL (2013)**	**Model**
All Sample	-5.2	-11.8
Heterogeneity by age		
Age ¡35	3.4	-4.9
Age 35-55	-10.2	-13.7
Age ¿55	-17.6	-25.1
Heterogeneity by employment		
Employed	3.8	-8.1
Unemployed	-7.0	-23.7

Note: In 2009, the enrollment to Wisconsin's BadgerCare Plus Core Plan abruptly terminated due to budgetary constraints. This table compares the effects of this policy change on i) the employment rate of the treatment group who received Medicaid relative to the control group who missed the application period (and did not receve Medicaid) and ii) the employment rate of the treatment group relative to the control group when the sample is divided into different age- and employment- subgroups estimated by Dague, DeLaire, and Leininger (2017) and estimated by my model.

a (statistically insignificant) decline in employment of 1.6 percentage points relative to the control group.

Oregon's 2008 expansion essentially relaxed its non-financial eligibility rules since it allow uninsured individuals to apply as long as they met pecuniary income- and asset tests. To conduct a similar experiment in the model, I unexpectedly grant Medicaid coverage to a subset of uninsured households whose income fall below 100 FPL and assets $a < 0.34$ but are excluded from coverage in the model simply because of failing non-financial eligibility requirements.[17] The households in the treatment group are thus comprised of low-income and low-asset households who unexpectedly receive Medicaid coverage. I take as control group low-income, low-asset households who are excluded from Medicaid because of non-financial eligibility requirements. This can be interpreted as individuals in the waitlist who are not drawn from the lottery or lottery winners who are not granted Medicaid.[18] I then run a regression of a dummy variable of employment on a dummy variable of unexpected receipt of Medicaid using model-simulated data to estimate the difference in employment outcomes between the treatment and control group.[19]

I find that the unexpected receipt of Medicaid results in a decline of employment rate by 8.1 percent relative to the control group, which is larger than estimates obtained by Baicker, Finkelstein, Song, and Taubman (2014). This may be possibly because, as Garthwaite, Gross, and Notowidigdo (2014) points out, the policy change in Oregon was implemented during the Great Recession, when the effects of such policy change on labor markets may be mitigated due to lower job finding rates affecting employment outcomes of non-recipients. Finally, notice that the estimated elasticities by the model for the Oregon experiment (8.1%) is much lower than the estimates for the Wisconsin experiment (12%) and for the

[17] I convert the asset threshold in the Oregon experiment into model units by first calculating the ratio of the asset threshold \$2000 to mean quarterly wages in the MEPS for 2006 which is \$8854.25. This ratio is 0.23. I then set the model asset threshold for the Oregon experiment to be the same fraction of mean wages in the model which is equivalent to 0.34.

[18] I experiment with the income composition of the control group to account for the possibiltiy that (1) lottery applicants and (2) lottery winners who did not receive Medicaid may have slightly higher income than lottery winners who receive Medicaid (since they meet income and asset thresholds). Including those below 200 FPL in the control group raises the estimate to 10.5% .

[19] I do not need to implement the same two-stage regression as in Baicker, Finkelstein, Song, and Taubman (2014) since being selected in the lottery is equivalent to Medicaid coverage in the model.

Tennessee experiment (35%). This is because the treatment group in the Oregon experiment are comprised of individuals with very low income (below 100 FPL with negligible wealth) compared to the Wisconsin experiment (below 200 FPL) and Tennessee experiment (below 400 FPL). In the model, individuals with very low income and wealth are less elastic to policy change because for them, the value of employment is very high. Thus, the model partially reconciles the large dispersion in estimates of employment responses to health reform through heterogeneity in the income and wealth of the subpopulation being treated.

1.4.2 Job finding rate differences depending on health insurance

The extent to which job-finding (UE) rates vary depending on an individual's health insurance coverage is informative of the value of a job given that it is a primary source of health insurance. A key mechanism in the model is that uninsured agents will have lower reservation wages and shorter unemployment durations as they are less selective of jobs. I validate the model against data on life-cycle job-finding rates from the SIPP 1996-2008 Panels. I restrict my sample to individuals between the ages of 25 and 64 who do not own a business or derive income from self-employment. I classify the individual as employed (E) if he/she reports having a job and either working or not on layoff, but absent without pay and as unemployed (U) if he/she reports either having no job and actively looking for work or having a job but currently laid off. An individual experiences a UE transition in a given month if he/ is unemployed at the beginning of the quarter and employed at the beginning of next quarter. I then identify the insurance status of the unemployed throughout their unemployment spell. To do so, I consider the respondent uninsured during a quarter if he/she indicates *not* being covered by either Medicare, Medicaid, military-related health insurance, or private health insurance. Individuals are considered to be without coverage during their unemployment spell if they are uninsured for more than 50 percent of their unemployment duration.[20] Figure 1.3 compares the quarterly job finding (UE) rates over the lifecycle of unemployed with and without health insurance coverage both in the model and in the data. It shows that, in the data, the unemployed without health insurance coverage exhibit much higher job finding rates consistently over the lifecycle. Moreover, job

[20]Results are robust to other reasonable values of this threshold.

Figure 1.3: Lifecycle job finding (UE) rates by health insurance coverage

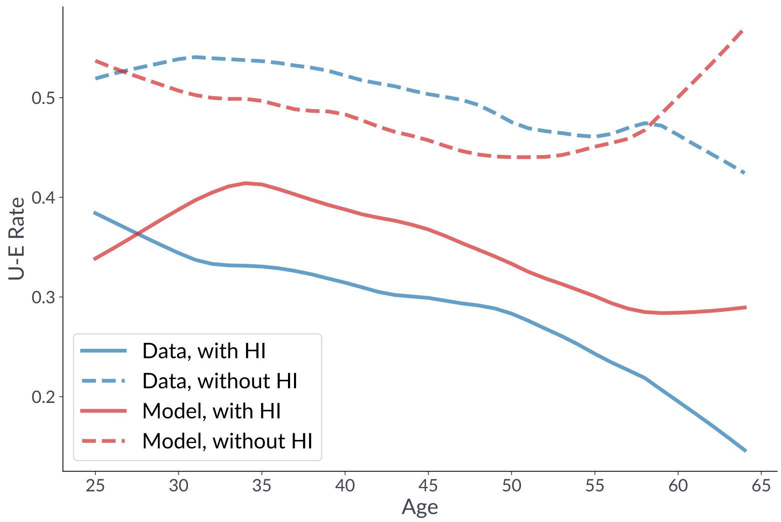

Note: This figure compares the quarterly job finding (UE) rates over the lifecycle of unemployed with and without health insurance coverage both in the model and in the data. The data UE profile is calculated using the SIPP 1996-2008 Panels. Details of the calculation can be found in the main text and in Appendix A.1.

finding rates of unemployed with and without health insurance coverage decline over the lifecycle. The model also generates a similar magnitude of the observed difference in job finding rates between insured and uninsured unemployed. Although the job finding rates of both groups are typically decreasing over the lifecycle in the model, job finding rates of the uninsured increase near retirement since the need for a job with health insurance is stronger given higher medical expenditures.

1.4.3 Effects of job loss on HI, earnings, and consumption

In this section, I validate the model against data on the consequences of job loss on health insurance coverage, earnings, as well as consumption. In order to measure the effects of job loss on these variables, I use the SIPP 2004 Panel to estimate the following distributed lag regression:

$$Y_{it} = \alpha_i + \gamma_t + \beta X_{it} + \sum_{k=-2}^{10} \delta_k D_{it}^{k,} + \epsilon_{it}, \tag{1.12}$$

where Y_{it} is the outcome variables of interest (health insurance coverage, earnings, as well as consumption) for individual i in period t, α_i are coefficients on individual fixed effects, γ_t are coefficients on quarterly fixed effects, X_{it} is a set of explanatory variables that includes education, marital status, number of children below 18, and a quadratic of age, and the error ϵ_{it} represents random factors. The indicator variables D_{it}^k identify all individuals k quarters prior to or after a job loss, where $k = 0$ is the quarter of job loss. For instance, $D_{it}^4 = 1$ for individual i who experiences job loss at time $t - 4$, and zero otherwise. The treatment group consists of individuals who experience at least one job loss during the survey time frame. On the other hand, the control group consists of individuals who never lost their jobs. Thus, $D_{it}^k = 0$ for all quarters t for individuals who belong to the control group. The coefficients $\{\delta_k\}_{k\in\{-2,\ldots,10\}}$ measure the effect of job loss on consumption k quarters prior-to or after the incident relative to individuals who do not experience any job loss.

I once again restrict the sample to individuals between the ages of 25 and 64 who do not own a business or derive income from self-employment. To identify job displacements, I use a question in the SIPP that asks the respondent for the main reason he/she stopped working for an employer. A job displacement is defined to be a transition out of employment caused by a layoff, employer bankruptcy, sale of employer's business, or slack work/business conditions.

I then run the same regression in the model, this time only controlling for a quadratic in age. In order to preserve comparability with the data, I only consider employment to non-employment transitions caused by the exogenous separation. The model implied effects of job loss are then compared with those estimated from the SIPP.

Health insurance and job loss In the SIPP, I classify a respondent as uninsured during a given month if he/she reports not being enrolled in Medicaid, Medicare, a military-related health plan, or a private health plan. Panel A of Figure 1.4 plots the estimated values for $\{\delta_k\}_{k\in\{-2,\ldots,10\}}$ in the regression specified in Equation 1.12 when Y_{it} is an indicator of health insurance coverage. According to the data, upon job loss, there is a 30 percentage points decrease in the probability of having health insurance coverage. This decline in probability of being insured is persistent and remains low even 10 quarters after job displacement. Moreover, displaced individuals have 8 percentage points lower probability of being insured

Figure 1.4: Effects of job loss

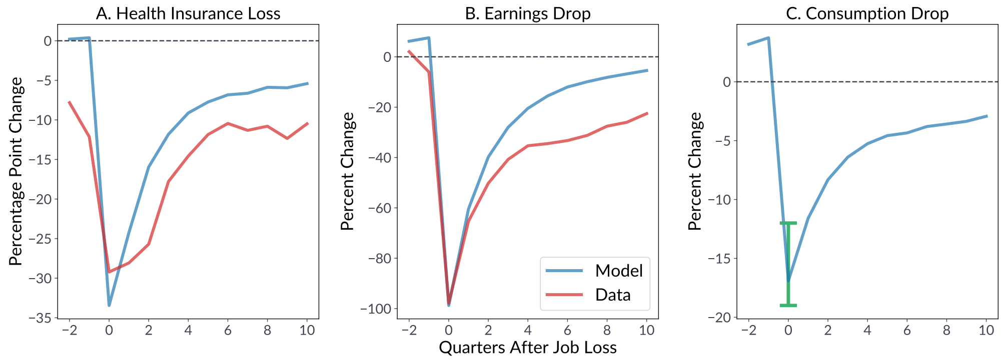

Note: This figure compares the consequences of job loss on health insurance coverage, earnings and consumption in the model and the data. In the data, I estimate the effects of job loss on health insurance loss (Panel A) and earnings loss (Panel B) from SIPP 2004 Panel using a distributed lag regression model. A range of estimates from the literature on consumption drop upon job loss is presented in Panel C. Same distributed lag regression models for each case are estimated using model simulated data.

even before their job displacement relative to the individuals in the control group. As seen in the same figure, the model predicts a slightly larger drop in the probability of being insured upon job displacement. This may be because of the fact that I do not model potential coverage coming from a family member's plan or the purchase of private non-group health insurance markets. I explore these extensions in Section 1.6. Furthermore, while the recovery of the probability is a bit faster in the model, the model also generates persistence in the drop of the probability.

Earnings and job loss To estimate the effect of job loss on earnings in the data, I use SIPP data on respondents' total earnings received from working in a job. Using earnings and employment data, I estimate Equation 1.12 to obtain the level and persistence of earnings drop after job loss. Panel B of Figure 1.4 plots the estimated coefficients as a fraction of mean pre-displacement earnings so that it is expressed percentage terms. It shows that both in the model and the data, earnings drop by as much as 100 percent in the quarter of displacement and do not fully recover even 10 quarters after displacement. The data, however, exhibits a higher degree of persistence in earnings losses over time.

Consumption and job loss Finally, I compare the model-implied value of the average drop in consumption upon experiencing a job loss to available empirical estimates in the literature. I use model-simulated data to estimate Equation 1.12 where the outcome variable Y_{it} is the log of consumption $\log c_{it}$. Panel C of Figure 1.4 plots the estimation results. It shows that in the quarter of job loss, consumption drops 17 percent on average and then slowly recovers over time. However, it still stays around 5 percent lower than its pre-displacement level even 10 quarters after the displacement.

There exists a number of studies that measure the average drop in consumption incurred at the moment of job loss from various data sources. Gruber (1997) finds a decline in food expenditure of 6.8 percent using the Panel Study of Income Dynamics (PSID) from 1968 to 1987. Saporta-Eksten (2014) uses the 1999-2009 biennial waves of the PSID with more detailed consumption data on consumption beyond food expenditures and estimates an 8 percent decline in the year during which job loss occurs. Stephens (2004) measures negative impact of job loss on food expenditure in the Health and Retirement Survey (HRS) and the PSID and finds that it is between 12 percent (PSID) and 15 percent (HRS) when an individual experiences a job loss between interviews. Browning and Crossley (2001) report a 14 percent decline using Canadian Out of Employment Panel (COEP) survey data. Finally, Chodorow-Reich and Karabarbounis (2016) study the effects of job loss on consumption in both the PSID and the Consumer Expenditure Survey (CE) and find that total food expenditures decline by 14 percent (PSID) and 21 percent (CE). Finally, Aguiar and Hurst (2005) use scanner data and report a 19 percent decline in food expenditure among the unemployed.

In conclusion, the model generates an estimate for the average decline in consumption upon job loss that is in line with available empirical estimates in the literature

1.4.4 Asset-income distribution

In addition to monthly data on income and government transfers as well as weekly data on employment status, the SIPP also contains data on respondents' asset holdings. In each SIPP panel, respondents provide information on various types of asset holdings during two or three waves within the panel, usually one year or, equivalently, three waves apart. I use Wave 6 of the 2004 panel of SIPP, which covers interview months October 2005 - January

2006. Sample selection is the same as in previous sections.
In the model, there are two reasons for accumulating assets: i) precautionary savings for employment and health risk, and ii) savings for retirement. Hence, in the data, I consider two sources of savings along these dimensions: i) net liquid wealth because of its immediate availability as a means to smooth consumption upon a negative financial shock, and ii) retirement savings. The net liquid asset holdings of an individual are calculated by adding transaction accounts (checking, saving, money market, call accounts), tradable assets (mutual funds, stocks, bonds), vehicle equity, and then deducting unsecured debt. Here, I follow Koehne and Khun (2015) and include net vehicle equity when calculating net liquid asset holdings. The reason is that, as I have shown in the previous section, income decreases substantially upon unemployment, and some unemployed could resort to liquidating other forms of assets (i.e., the sale of vehicles) to smooth consumption upon job loss. Retirement savings include the market value of saving instruments in the form of a 401k, 403b, thrift plan, Individual Retirement Account (IRA), or KEOGH account.
To normalize wealth and better capture the level of self-insurance, I compute respondents' asset-to-income ratio by dividing the sum of net liquid assets and retirement savings by quarterly before-tax labor income. Table 1.5 shows the computed percentiles of the asset-to-quarterly labor income distribution in the data and the model. The calibrated model comes close to matching the empirical asset distribution. In particular, the model reasonably captures the left tail of the distribution and at the same time closely matches the fraction of the population holding non-positive wealth. Matching the left tail of the distribution matters for my analysis because agents in this region of the distribution are the most affected by changes in public health insurance coverage. Job losers with low wealth have little to no capacity to self-insure or smooth consumption using their own liquid assets and are thus very sensitive to changes in government transfer generosity.

1.5 Policy Reform

The shift to a universal healthcare system is implemented by introducing a government-run public health insurance with no means-test for the working-age population. The coinsurance rate of this system is set to be identical as the current Medicaid system. For the baseline policy experiment, the universal reform is funded by a proportional increases in

Table 1.5: Percentiles of the distribution of asset holdings relative to quarterly before-tax labor income

	Percentiles					Fraction of population
	10th	25th	50th	75th	90th	with non-positive wealth
Data	-0.91	0	0.78	3.97	11.22	0.29
Model	-1.2	0	0.78	2.76	13.94	0.31

Note: This table shows the asset to quarterly labor income distribution both in the data and in the model. The empirical distribution is calculated using the SIPP 2004 Panel, where assets include net liquid assets (defined to be the sum of net financial assets, and equity in vehicles), and the market value of retirement savings (401k, 403b, thrift plan, Individual Retirement Account (IRA), or KEOGH account). Details of the calculation can be found in Appendix A.1.

taxes. This is equivalent to setting the new income tax schedule to be

$$T^p(z) = \left(\beta_j^{Tax} + \xi\right) z + \alpha_j^{Tax}$$

where ξ is used to balance the government's budget under the new policy. Financing the universal healthcare system would require a proportional tax increase of $\xi = 0.043$ under the new steady state.

In order to solve for the transition path, I begin with the stationary distribution of the calibrated economy under the baseline employer-based health insurance system. I then introduce an unexpected and permanent shift towards a universal healthcare system. Along the transition, interest rates and market tightness adjust to clear the capital and labor markets period by period. The increase in taxes are introduced gradually and phased-in. I assume a linear increase in taxes from the moment of policy implementation at time 0 until the end of the transition period.[21]

1.5.1 Macroeconomic effects

To understand the effects of the policy change, I first compare two steady states, one under the baseline economy and another under the universal healthcare system. Table 1.6 shows the impact of the policy on important macroeconomic outcomes. The shift to a universal policy results in a substantial decline in job-finding rates from 45 percent to 31 percent. This is driven by both an increase in individual reservation wages as well as a decrease in firm vacancy posting, as evidenced by lower equilibrium market tightness. Consequently, the unemployment rate and the duration of unemployment rise. Furthermore, quit rates also increase but not substantially. This is due to the fact that quitting becomes less desirable when aggregate job-finding rates are low and it is more difficult to find a new job once an agent leaves employment. Now that health insurance is no longer tied to employment, a large fraction of workers decide to leave the labor force, leading a 6 percentage point increase in non-participation rates. This implies that in the baseline economy, a large number of agents work because of health insurance as they experience the phenomenon of "employment lock".

The provision of generous insurance against medical expenditure risk also weakens precautionary saving motives. This leads to a higher cost of capital and a roughly 8% decline in average capital stock. As a result, output per worker $\frac{Y}{L}$ is lower despite the slight increase in average match TFP $A(x, y)$. Match quality does not rise by much despite better sorting. This is because even if workers become more selective in their choice of jobs, increased labor market frictions decrease incentives to quit from poorly matched jobs.

Given these large effects of the policy change on labor market outcomes, it is also insightful to check the change in lifecycle labor market transition profiles of households. Figure 1.5 plots the change in the age profile of UE and EU rates resulting from the policy reform. I find that while job-finding rates decline across all ages, the response in job finding rates is much larger to older individuals who face higher medical expenditure risk. This makes health insurance much more valuable to this subpopulation, thus explaining why they are more elastic to changes in public health insurance generosity. Meanwhile E-U rates over the lifecycle do not change as drastically due to the small response of quits as mentioned previously.

[21] Details of the computation of transitional dynamics are discussed in Appendix A.2.

Figure 1.5: UE and EU rates: Baseline vs. universal health insurance policy

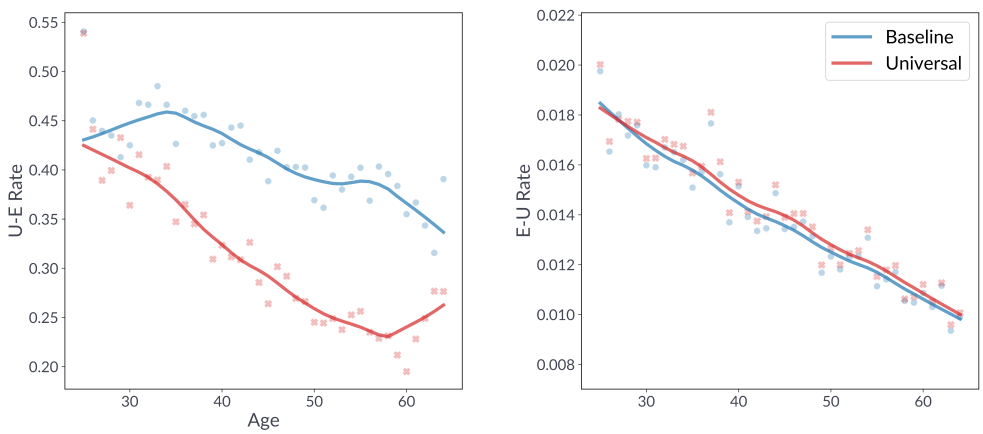

Note: This figure plots the age-profile of unemployment-employment (UE) and employment-unemployment (EU) transitions predicted by the model under the baseline policy and under the universal healthcare policy.

Finally, Figure 1.6 plots the transition path of selected macroeconomic outcomes. Unemployment rate initially declines because firms relinquish the role of providing health insurance and thus no longer pay premiums to insurance providers. The increase in reservation wages are gradual because households first accumulate savings to smooth consumption when unemployment duration increases after they increase reservation wages. Thus, initially, the benefit of lower expenses (no premiums) dominates rising reservation wages. However, when reservation wages finally dominate the decline in premiums, the unemployment rate and job-finding rate begin to increase. Furthermore, the initial increase in savings is only temporary and savings as well as aggregate capital stock eventually decline.

Holding firm vacancy posting fixed In order to understand the importance of incorporating firm responses to household labor supply decisions resulting from the policy change, I consider a shift to a universal heatlhcare system when I hold market tightness θ constant after the policy reform. The last row of Table 1.6 shows these effects. Compared to the case where firms are allowed to endogenously respond to the reform, the decline in job-finding rates and the increase in unemployment rate is more muted given that these

Figure 1.6: Transitional dynamics

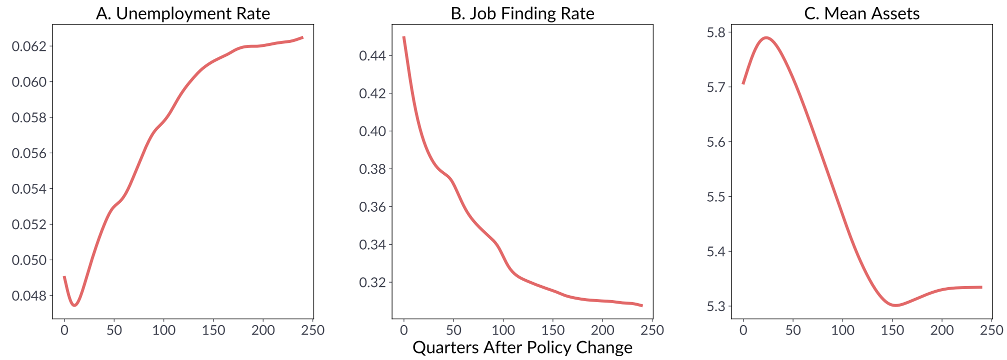

Note: This figure plots the transition path of various macroeconomic outcomes upon introducing the universal healthcare system in the economy at period 0.

are purely driven by labor supply responses. In contrast, the change in non-employment rate is not very different compared to the endogenous firm model. This is because the non-employment margin is less affected by aggregate job-finding rate $p(\theta)$ which is the primary channel through which firm reduction in vacancies affect labor markets. Quit rates, on the other hand exhibit a much larger response when firm decisions are treated as exogenous. Holding aggregate job-finding rates $p(\theta)$ constant makes quitting attractive because it overstates the probability of meeting a firm and finding a better job. When we incorporate the decrease in firm vacancies by allowing market tightness to adjust, quit rates do not increase as much as workers become more reluctant to quit to search for better jobs.

1.5.2 Welfare effects

In the section, I explore the welfare implications of shifting to a universal healthcare system. I pay close attention to the heterogeneity of welfare gains or losses across different types of agents. First, I discuss how I measure welfare gains/losses and then I elaborate on the results.

Welfare Calculation

I measure the welfare effects of any shifting to a universal healthcare system by answering the following question: how much additional lifetime consumption must be endowed to all agents in an economy under the current policy so that average welfare will be equal to an economy where the policy reform is implemented? This welfare measure is defined by the consumption equivalent under the veil of ignorance between an economy under the current policy and the proposed reform.

Let $V_t^P(s, \lambda)$ be the value of an agent with age t and state s in economy P if his consumption is multiplied by $(1+\lambda)$ for the remainder of his lifetime. Formally, this can be written as

$$V_t^P(s, \lambda) = \mathbb{E} \sum_{j=t}^{T+T^R} \beta^{j-t} \prod_{k=j}^{T+T^R} \varphi_k U\left(c_j^P(s)(1+\lambda), l_j^P(s), h_j^P(s)\right)$$

where c_j^P refer to the consumption policy function, and l_j^P and h_j^P refer to employment and health status at age j when $\lambda = 0$. Here, λ is the additional percent consumption given to the agent every period in economy P.

Let b denote the baseline (current) policy and n denote the new (universal policy). We can compute the additional percent lifetime consumption $\bar{\lambda}$ that makes the average welfare equal across these two economies using the following equation:

$$\int_{t,s} V_t^b\left(s, \bar{\lambda}\right) \mu(t, s) = \int_{t,s} V_t^n(s, 0)\, \mu(t, s)$$

where $\mu(t, s)$ refers to the stationary distribution of agents who are alive when the policy is implemented.

We can also compute for conditional (ex-post) welfare for any agent with age t and state s by calculating $\lambda(t, s)$ that solves

$$V_t^b(s, \lambda(t, s)) = V_t^n(s, 0)$$

This measure of welfare asks whether the agent with age t and state s will be better/worse off in an economy under the policy reform and thus allows me to study its the heterogeneous welfare impacts

Welfare Results

I find that the policy reform results in an aggregate welfare gain of $\bar{\lambda} = 0.33$ percent in lifetime consumption equivalents. However, this small welfare gain masks substantial heterogeneity in welfare outcomes across heterogeneous households. Table 1.7 computes the welfare gains/losses from the health insurance policy reform when households are grouped by their asset and employment states based on the stationary distribution before the policy change. I find that welfare gains mostly accrue to wealthier households while welfare losses accrue mostly to poorer and non-employed households. This result is driven by three reasons. First, poorer households suffer the most from worsened labor market conditions. Being close to borrowing constrained, it is precisely for these households for whom an extra quarter of non-employment is very costly given their inability to self-insure against employment risk. Households with sufficient levels of wealth benefit from higher job selectivity and the provision of health insurance because they are able to withstand lower aggregate job finding rates by using their available savings to smooth consumption during non-employment spells. Second, it is typically low-wealth and low-income households who are already covered by Medicaid in the current system. This further reduces the insurance benefits of providing universal healthcare to this population. Finally, a proportional increase in taxes impose a larger utility cost on poorer households for whom the marginal utility of consumption is much higher.

To further corroborate this result, I show the distribution of welfare gainers and losers across various dimensions of heterogeneity in Table 1.8. First, note that welfare losers tend to be younger, healthier, and unskilled. While this subpopulation benefits the least from access to health insurance due to low medical expenditures, they bear the burden of lower aggregate job finding rates. Beyond the effects of longer unemployment spells, the inability to find jobs affects the ability of younger workers to accumulate skills and grow their wages over the lifecycle. The opposite is true for retirees. Welfare losers tend to be older retirees with less savings but face the higher taxes brought about by the policy reform.

1.5.3 Tax burden

Given the substantial welfare losses incurred by low-wealth and non-employed households, a natural question to ask is: what revenue-generating scheme would be welfare-improving?

I answer this question by first introducing a new healthcare tax schedule $T^H(z)$ to fund the additional expenses needed to finance the universal healthcare program while holding other tax rates (income tax $T^P(z)$ and consumption tax τ_c) fixed to their original schedules. To feasibly study this, I parameterize the tax schedule as in Heathcote, Storesletten, & Violante (2014) when implementing the policy reform:

$$T^H(z) = z - \lambda_H z^{(1-\tau_H)}$$

where λ_H defines the level of taxation, τ_H defines the progressivity of the new tax, and z denotes gross income. For any proposed progressivity level τ_H, I adjust λ_H to balance the government's budget. Finally, I solve for the optimal progressivity of the new health tax τ_H^* and its implied λ_H^* that maximizes the ex-ante welfare of the economy.

I find that the optimal revenue scheme is given by $\tau_H^* = 0.33$, implying a highly progressive healthcare financing scheme. This financing scheme raises ex-ante welfare gains to 0.54 percent of lifetime consumption. More importantly, Table 1.9 shows that welfare losses for low-wealth households are ameliorated by shifting the tax-burden of financing universal healthcare towards richer households. This serves as a way to compensate poorer households for the decrease in aggregate job-finding rates caused by the labor supply responses of other agents in the economy. In practice, this result implies that a shift to a universal healthcare system will lead to better welfare outcomes if the additional funding is sourced from higher-income/higher-wealth households as opposed to a proportional increase in payroll taxes or value-added taxes.

1.6 Robustness

Family insurance and COBRA In the baseline model, job losers are partially insured against the loss of health coverage only through either Medicaid or the consumption floor guarantee of government. In reality, these households have access to other forms of insurance such as coverage through a family member/spouse or being eligible for coverage through a former employer as mandated by the Consolidated Omnibus Budget Reconciliation Act (COBRA) of 1985.[22] The prevalence of obtaining insurance either from a

[22] COBRA requires employers with more than 20 employees and provides health insurance to its employees to offer former employees the option of continuing this coverage. The option of COBRA is only given

family member or COBRA will have implications on the insurance benefits of expanding public health insurance programs. If the health risk of job losers are already well insured by inclusion into these mechanisms, then a universal healthcare system will provide small insurance benefits to these group of people.

In order to understand the extent to which family health insurance and COBRA are used to insure against the loss of coverage upon job loss, I use the SIPP 2004 Panel data to measure the likelihood that a job loser will switch to a family member's health plan or participate in a health plan administered a former employer. A respondent is defined to be covered under a family's health insurance plan if he/she reports being (1) insured through a private health insurance plan, (2) covered under someone else's plan, and (3) covered under someone inside the household. Participation in COBRA, on the other hand, is identified in the SIPP by selecting non-employed respondents who report being insured through a private health insurance plan sourced from a former employer. I estimate the distributed-lag regression specified in Equation 1.12 twice: one where the outcome variable of interest is a dummy on health insurance coverage by a family member and another for health insurance coverage through a former employer. Sample selection and the definition of a job displacement is similar to Section 1.4.3. Panel A of Figure **1.7** shows that displaced workers have a modestly higher probability of being covered by a family member relative to the non-displaced group. The relative probability of inclusion into a family health plan rises and peaks to around 4 percent at eight quarters after displacement.[23] One potential reason behind this is positive assortative matching in marriage which would imply that individuals who are predisposed to displacement are more likely to be married with spouses in low-paying jobs with limited access to employer-provided health insurance. In contrast, Panel B of Figure 1.7 shows that the relative probability of obtaining health insurance from a

---

to workers who voluntarily or involuntarily terminate employment for reasons not due to gross misconduct or to workers who experience a cut in hours. However, COBRA participants generally pay for the entire premium themselves as opposed to active employees whose premiums are partially paid for by the employer.

[23]Note that the small increase in the probability of family coverage does not necessarily mean that enrollment in a family member's health plan is not prevalent. It merely indicates that among individuals who are displaced, the incidence of switching to family coverage upon job loss is not widespread. Indeed, there exists workers who to begin with are already under a spouse's plan and when displaced, remains to be under a spouse's plan. Dey and Flinn (2008) find that in two-earner households, health insurance is often only purchased (through their employer) by one of the spouses.

Figure 1.7: Alternative health insurance coverage upon job loss: Family and COBRA

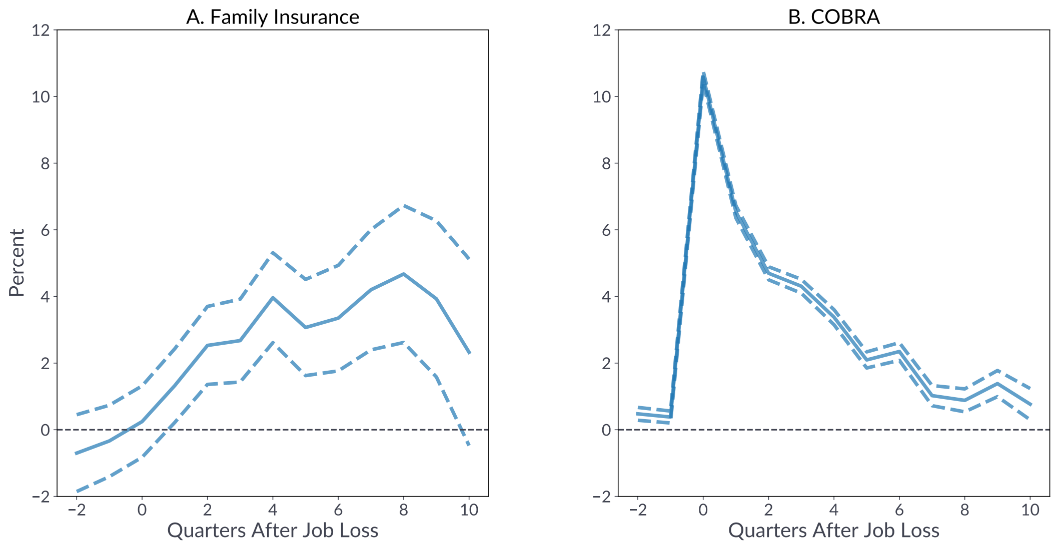

Note: This figure plots the coefficients from a distributed lag regression of coverage under a family member's health plan (Panel A) and undera former employer, i.e. COBRA (Panel B) on job-loss controlling for education, marital status, number of children below 18, age, as well as fixed- and year-effects. Dashed-lines are confidnece intervals. Data is obtained from the SIPP 2004 Panel.

former employer (COBRA) is around 10 percent upon job loss. However, access to COBRA quickly diminishes and becomes close to zero 6 quarters after job loss. This is reflects the fact that continuation coverage through COBRA typically expires after 18 months.

I now explore the implications of extending the model to incorporate family health plans and COBRA as alternative sources of health insurance during non-employment spells. To discipline the extended models, I use the measured take-up of each source of health insurance documented above.

First, I consider an extension of the baseline model that incorporates the possibility of acquiring health insurance through a family member upon job loss. I model family insurance as a random probability γ^f of receiving coverage upon job loss. This probability is set to be constant at 4% and only given to individuals who are not covered by Medicaid to approximate the above finding. The coinsurance rate is set to be the same as EPHI. As a simplification, I assume that family coverage is granted throughout an individual's non-employment spell. The covered medical expenditures of all agents enrolled in family plan is added uniformly to the premium payments paid by all firm-worker pairs that offer health insurance. This captures the fact that the cost of family coverage is borne jointly by both employers and employees.

The first row of Table 1.10 shows the labor market and welfare effects of a universal healthcare system in the model with family health insurance. Given the small utilization of spousal health insurance, it is unsurprising that job-finding rates decline by a similar magnitude under a universal healthcare system when compared with the baseline model specification. Similar effects can also be observed for both unemployment rates and non-employment rates. Aggregate welfare gain declines slightly due to the lower insurance benefits expanded coverage provides while agents in the bottom quintile of the asset distribution still suffer substantial welfare losses as in the baseline results.

In the next exercise, I extend the baseline model by incorporating the option of continued coverage. COBRA allows employees who are separated from their job to purchase health insurance at more affordable group premiums through their former employer. However, this option typically expires roughly six quarters after job loss. To mimic these regulations, I model COBRA as an option to purchase group health insurance by paying the group premium w_h, first made available at the moment of job loss with probability γ^{COBRA}[24]. However, this option expires permanently with probability $e^{COBRA} = \frac{1}{6}$. The second row of Table 1.10 shows that the substantial decline in job-finding rates and increase in unemployment rate is preserved but is slightly muted because the presence of COBRA raises the value of non-employment under the current system (and thus lowers the relative increase in the value of non-employment under the universal healthcare system). The effect on non-participation rates is identical to that of the baseline model because COBRA's expiration provides little value to the long-term non-employed who are comprised mostly of non-participants. Finally, welfare gains are also much lower, again, because of lower insuranc benefits universal healthcare provides given expanded access to alternative health insurance sources such as COBRA in the current set-up.

[24]The group premium w_h is adjusted accordingly to account for COBRA participants who take up the offer. The availability of the continuation coverage option occurs with probability $\gamma^{COBRA} = 0.32$ to match an average take-up rate of 11% at the moment of job loss and capture the fact that some types job separations and employers do not fall within the scope of COBRA.

Segmented markets In the baseline model specification, all workers and firms meet in a single labor market. This implies that the labor market behavior of, say, high-skilled workers affects job-finding probability of low-skilled workers, through its effects on firm vacancies. The extent to which labor markets are segmented can limit such externalities and introduce differential aggregate job-finding rate responses to the policy reform across segmented markets. In order to understand if the labor market and welfare effects are robust to the single-market assumption, I extend the model by segmenting labor markets by skill x. The extended model features two markets: low skilled $x \in [\underline{x}, \hat{x}]$ and high skilled $x \in (\hat{x}, \bar{x}]$. Under this set-up, there are two free entry conditions and thus two equilibrium market tightness θ_L and θ_H which may have differential responses to the policy reform. To map the two markets in the model to the data, I pick the threshold $\hat{x}$ to match the ratio of mean wages of workers with at least a college degree and workers with less than a college degree which is 1.86 for the year 2006 in the SIPP 2004 Panel. The distribution of administrative cost shocks $\Gamma^{\epsilon}(x)$ is set equal to the baseline calibration, implying that the low-skilled market will have less access to EPHI. Table 1.11 shows that job finding rates under both markets decline significantly, but much more so for the low-skilled markets. This is explained by the fact that low-skilled workers comprise of low-income and low-wealth agents and a large fraction of the population who are on the margin (and thus exhibit high elasticity) fall within this market.

1.7 Conclusion

This paper studies the macroeconomic and welfare effects of shifting to a universal healthcare system, paying particular attention to the reform's effects on the labor market. The expansion of public health insurance changes the relative value of non-employment to employment, which in turn determines non-employed individuals' job search and acceptance decisions, and employed individuals' quit decisions. On the other side of the labor market, these changes in the behavior of households affect the value of a vacancy and a filled job for firms. Thus, firms react to this by endogenously adjusting the level of job creation. As a result, while universal healthcare can insure against employment risk and health expenditure risk, these insurance benefits may be offset by potential incentive costs in the labor market.

To study the reform, I build a general equilibrium lifecycle model with incomplete asset markets, employer-provided health insurance and a limited government-funded means-tested health insurance, and frictional labor market with two sided-heterogeneity. The model is calibrated to match the distribution of individuals who are eligible of Medicaid and the joint distribution of wages and access to employer-provided health insurance since these directly affect the insurance benefits and incentive costs of expanding publicly-provided healthcare. Then, I compare the model-implied labor supply elasticities with respect to public health insurance generosity to existing estimates from quasi-experimental studies. I find that the model is capable of partially reconciling the wide-range of estimates found in experiments conducted in Oregon, Wisconsin, and Tennessee and is consistent with heterogeneous elasticities across income, age, health, and employment status.

Using the calibrated model, I conduct a policy experiment by introducing a universal healthcare system that is funded through a proportional increase in taxes. I find that the reform results in a substantial decline in aggregate job finding rates from 45 percent to 31 percent and concomitantly, higher non-employment rates and longer unemployment spell durations, mostly due to lower vacancy posting rates of firms. The universal healthcare system results in a small ex-ante welfare gain of 0.3 percent additional lifetime consumption. However, I find that welfare losses coming from labor market effects accrue mostly to wealth-poor individuals for whom lower aggregate job-finding rates and longer unemployment durations are most costly since these households have severely limited ability to self-insure against consumption fluctuations caused by unemployment risk. Finally, the optimal revenue-generation scheme involves a highly progressive system where high income households bear the brunt of the additional taxes.

These conclusions arise from the labor market effects of health care reform. Naturally, health policy would affect other dimensions of the economy. For example, this study abstracts from endogenous health spending, health goods markets and prices, as well as non-monetary benefits of health, all of which merit a careful study in future research.

Table 1.6: Macroeconomic effects of universal heatlh care

	Baseline	Universal	Universal (Vacancies Fixed)
Labor Markets			
Job-finding rate %	45	31	37
Market tightness $\theta = \frac{v}{u}$	7.6	3.3	7.6
Unemployment rate %	4.9	6.2	5.5
Unemployment spell duration	2.1	2.8	2.5
Quit rate %	0.50	0.53	0.58
Non-participation rate %	25	31	30
Capital and Productivity			
Annual interest rate %	3.4	3.6	3.6
Average capital stock	5.7	5.3	5.4
Average match TFP	2.21	2.24	2.27
Output per worker $\frac{Y}{L}$	3.40	3.35	3.39

Note: This table compares the macroeconomic outcomes between two steady states: i) an economy under the baseline health insurance system vs ii) an economy under the universal health care system.

Table 1.7: Welfare gains/losses by wealth and employment status

	Asset Groups				
Employment	a_1	a_2	a_3	a_4	a_5
Employed	-0.41	-0.13	0.28	0.75	1.4
Non-employed	-2.1	-0.39	-0.01	1.4	1.9

Note: This table presents the welfare gains/losses from the health insurance policy reform across households who are grouped by their asset and employment states, where a_1 to a_5 represents the quintiles of the wealth distribution, based on the stationary distribution before the policy change.

Table 1.8: Distribution of Welfare Gains and Losses

	Gainers	**Losers**
Working Age		
Age	47.1	35.7
Assets	8.3	3.9
Fraction of Healthy %	67	84
Fraction of Medicaid Recipient %	3	14
Skill	1.6	0.7
Fraction of Unemployed %	2.1	6.6
Retired		
Age	69.2	78.1
Assets	19.9	8.5
Fraction of Healthy %	54	49

Note: This table compares the characteristics of households who enjoy welfare gains from the policy reform to characteristics of households who suffer welfare losses from the policy reform.

Table 1.9: Welfare Gains/Losses by Wealth and Employment Status under Progressive Financing

	Asset Groups				
Employment	a_1	a_2	a_3	a_4	a_5
Employed	0.1	-0.05	0.39	0.47	1.1
Non-employed	0.4	-0.07	0.46	0.93	1.5

Note: This table presents the welfare gains/losses from the health insurance policy reform under the optimal progressivity of taxation scheme across households who are grouped by their asset and employment states, where a_1 to a_5 represents the quintiles of the wealth distribution, based on the stationary distribution before the policy change.

Table 1.10: Family health insurance plans and COBRA

	Job-finding rate	Unemp. rate	Non-part. rate	Ex-ante welfare	$a < p20$ welfare
Family health insurance	32%	6.1%	31%	0.32	-1.29
COBRA	34%	5.9%	30%	0.26	-1.82

Table 1.11: Segmented markets

	Baseline	Universal
Job-finding rate (low skill)	0.49	0.32
Job-finding rate (high skill)	0.36	0.29

Chapter 2

How Should Unemployment Insurance Vary Over the Business Cycle?

2.1 Introduction

The sharp increase in unemployment during the Great Recession was associated with dramatic expansions to the unemployment insurance (UI) program. While intended to provide adequate insurance to the large pool of jobless individuals, the question of whether UI policy played a quantitatively significant role in slowing the recovery of employment remains at the center of discussion.[1] Alongside this positive debate, an equally important policy question emerges: how then should UI policy vary over the business cycle? Addressing this question will shed light on how UI policy must adjust to economic fluctuations, especially during economic downturns.

Our main contribution to the growing literature on optimal UI over the business cycle is to study the endogeneous interaction between precautionary savings and changes in UI policy

[1] For example, Hagedorn et al. (2016) find that a generous UI policy during the recession is partly responsible for the drastic and sustained rise in unemployment that followed. On the other hand, Chodorow-Reich and Karabarbounis (2017) show that the extensions have had limited influence on macroeconomic outcomes.

over recessions and expansions, a mechanism that we show is crucial to correctly measure the welfare benefits and costs of any proposed policy. This is because the level of wealth determines not only the insurance value of any public transfer but also its incentive costs, since the labor market behavior of individuals holding different levels of assets responds in varying degrees to changes in the level of generosity of these programs. Moreover, as wealth holdings and the strength of precautionary saving motives vary over the business cycle, they inevitably influence the cyclicality of the insurance benefits and incentive costs of UI payments. It is precisely the cyclicality of the net benefits of UI that will determine how benefit generosity should vary over the business cycle.

We address this question using a heterogeneous agent job search model that incorporates labor productivity driven business cycles and incomplete asset markets. To overcome the computational difficulties encountered in models of this nature, we show that the model's market structure admits a block recursive equilibrium, a subset of recursive equilibria where the endogenous distributions generated by the model are not part of the state space (Menzio and Shi 2010, 2011). This allows us to compute the optimal UI policy in a model with aggregate shocks and saving decisions.

We find that the optimal UI policy is countercyclical. In particular, when the aggregate labor productivity is at its mean, it features a 30 percent replacement rate for 4 quarters. When aggregate labor productivity is depressed by 3.5 percent, however, it offers more generous benefits of a 54 percent replacement rate for 10 quarters, financed by higher labor income taxes. Compared to a UI policy that mimics the policy implemented by the U.S. government during the Great Recession, the optimal policy represents an ex-ante welfare gain of 0.58 percent additional lifetime consumption.

The countercyclicality of the optimal policy is explained by how the insurance benefits of extra UI payments expand during recessions relative to expansions while relative incentive costs contract. Two important insurance benefit channels expand during recessions: (1) consumption insurance against unemployment risk and (2) consumption insurance against aggregate labor productivity risk. First, generous benefits insure against unemployment risk by alleviating the consumption drop experienced by job losers. This is especially important in recessions when unemployment rises and spells are prolonged. Second, it also insures against aggregate risk since it reduces the burden of having to engage in (costly) precautionary savings during economic downturns. Recessions trigger a strong need to accumulate a buffer stock of savings, which in turn entails a concomitant reduction in consumption. In the absence of public insurance, this makes consumption fluctuate heavily with the business cycle. However, this effect is mitigated when individuals are promised more generous payments for future unemployment spells.[2] Remarkably, this results in sizeable welfare gains not only for job losers but also for those who are employed.

Insurance benefits come with a trade-off: generous UI payments during recessions decrease the job finding rates of the unemployed through a decline in job search effort and an increase in the wages that they seek. This results in longer unemployment durations. However, we show that these costs are relatively lower in recessions for two reasons: (1) the value of job search is low during recessions, and (2) borrowing constraints impose discipline on the unemployed's job search behavior. First, the value of job search during recessions is low because, to begin with, jobs are difficult to find and available jobs offer relatively lower wages. Hence, even if generous benefits were to discourage job search during a recession, the forgone search effort would not have been very productive anyway. Second, a reduction in wealth holdings during recessions induces the unemployed to find a job more quickly as they get closer to becoming borrowing constrained. In this sense, the presence of borrowing constraints is a device to discipline the job search behavior of the unemployed. For both of these reasons, the incentive costs associated with generous benefits are partially offset in recessions.[3]

These channels remain active even under a high level of the opportunity cost of employment calibration. In this case, we find that while the mean replacement rate and duration of the optimal policy reduce to a 19 percent replacement rate for one quarter, the degree of countercyclicality remains roughly similar. As fluctuations in consumption are less pronounced under this calibration, the government implements a low replacement rate for short durations when aggregate labor productivity is at its mean value. Still, insurance benefits expand and incentive costs contract in recessions. Thus, the government finds it optimal to transfer funds from expansions toward recessions. The resulting optimal policy

[2] This channel is consistent with Engen and Gruber's (2001) empirical finding that UI payments crowd out private savings.

[3] This result is consistent with Kroft and Notowidigdo (2016), who empirically find that the moral hazard cost of UI is procyclical.

in this case provides ex-ante welfare gains of 0.25 percent lifetime consumption, which is less than half of the welfare gains provided by the optimal policy under the baseline calibration of the opportunity cost of employment.

We quantify various sources of ex-ante welfare gains of the optimal policy and find that most of them are attributable to changes in consumption patterns, whereas the welfare gains from economizing on relatively unproductive search during recessions are negligible. These changes in consumption patterns can potentially increase ex-ante welfare for three reasons: (1) an increase in consumption levels, (2) a decrease in consumption volatility, and (3) a reduction in consumption inequality across individuals. We find large welfare gains due to an increase in the average consumption level along the transition path after the implementation of the optimal UI policy. This is because agents decumulate savings and consume more of their labor income when public insurance is generous, and this effect dominates the increase in labor income taxes. Steady state welfare decomposition reveals that long-run welfare gains are attributable mostly to reduced consumption uncertainty, but at the cost of lower consumption levels. The reduction in the consumption level is due to higher taxes and lower wealth holdings once the economy converges to a new steady state, although this change is not large enough to overturn uncertainty gains. Finally, welfare gains due to a reduction in consumption inequality are small because the optimal policy has two offsetting effects on consumption inequality. On the one hand, the redistribution of labor income from workers to the unemployed creates more equal consumption paths across heterogeneous agents. On the other hand, the optimal policy increases wealth inequality in the stationary distribution. This is because while most of the individuals in the economy under the optimal policy reduce their savings, the response of the agents in the top percentiles of the distribution is very small. The rise in wealth inequality, in turn, increases consumption inequality among heterogeneous agents. We find that these two opposing effects quantitatively cancel each other out and thus result in negligible welfare gains attributable to a decline in consumption inequality.

Next, we analyze the heterogeneous welfare effects of the optimal policy. Unsurprisingly, the unemployed who are eligible for UI benefits gain significantly, with the poor within this group enjoying the largest welfare gains, since each additional dollar of benefit payments is more valuable to them. Workers also enjoy a sizeable welfare gain, albeit to a smaller degree due to two opposing effects. Although they are the primary financers of the increased government expenditures because of the generous policy, they also experience large consumption smoothing benefits over the business cycle. Similarly, gains are also much larger among poor workers for whom a reduction in precautionary savings diverted toward consumption is most beneficial. Finally, the unemployed who are ineligible for UI gain the least because they will enjoy benefits only if they find a job and become eligible through the loss of that job. They are also adversely affected by lower job finding rates during recessions without the insurance that UI provides.

When solving for the optimal UI policy, we follow a large strand of literature that uses calibrated models to study the optimal policy for a restricted class of policy instruments.[4] The model simultaneously matches the liquid asset-to-income distribution and salient features of the labor market prior to the Great Recession. The policy instruments in our welfare analysis are restricted to take the form of the UI replacement rate and UI payment duration as functions of current aggregate labor productivity, and a constant labor income tax used to balance the government's budget for any proposed UI program.

Related Literature Our paper contributes to the growing literature on optimal UI over the business cycle. Recent papers in this literature are Landais et al. (2017), Jung and Kuester (2015), Mitman and Rabinovich (2015). However, in these models, risk-averse agents do not have access to asset markets for self-insurance purposes.[5] This assumption has several important implications for the level and cyclicality of the insurance benefits and incentive costs of any proposed UI policy. First, the insurance value of UI payments for job losers is overstated because public insurance is the only way of smoothing consumption upon job loss. Second, since the elasticity of search effort and the wage choice of the unemployed are both decreasing in wealth holdings, a model that abstracts from self-insurance altogether also overestimates the level of the moral hazard costs associated

[4]See Hansen and Imrohoroğlu (1992), Acemoglu and Shimer (2000), Abdulkadiroğlu et al. (2002), Wang and Williamson (2002), Krusell et al. (2010), Koehne and Kuhn (2015), and Eeckhout and Sepahsalari (2015).

[5]In addition to this difference, there are other important modeling differences between our paper and these papers. For example, Jung and Kuester (2015) and Landais et al. (2017) do not consider UI expiration. See Mitman and Rabinovich (2015) for a discussion on the implications of these assumptions.

with introducing a more generous UI policy. Third, disregarding asset markets completely eliminates the interaction between self-insurance and public insurance. Importantly, the decline in precautionary saving motives as a response to a generous UI policy contributes to the expansion of insurance benefits of UI in recessions because it also provides consumption insurance against aggregate risk. The novelty of our analysis is to study this endogenous response of the asset distribution to changes in UI policy over the business cycle, which is crucial for the true measurement of the cyclicality of insurance benefits of UI. Among these papers, our model is closest to Mitman and Rabinovich (2015) with two differences: our model 1) allows for self-insurance through incomplete asset markets, and 2) features directed search, making the model still tractable due to block recursivity even under the presence of incomplete asset markets, whereas job search is random in their model. In terms of welfare exercise, Mitman and Rabinovich (2015) are able to solve a Ramsey problem to obtain the optimal UI policy as a function of the entire history of past aggregate shocks, whereas we use our calibrated model to study the optimal policy for a restricted class of policy instruments that only depend on the current period realization of the aggregate shock in order to maintain tractability.

Another strand of literature studies the optimal design of UI policy under the presence of asset markets. However, these papers use models that do not incorporate either unemployment risk (Kroft and Notowidigdo 2016) or aggregate risk (Hansen and Imrohoroğlu 1992, Acemoglu and Shimer 2000, Abdulkadiroğlu et al. 2002, Wang and Williamson 2002, Lentz 2009, Krusell et al. 2010, Koehne and Kuhn 2015, and Eeckhout and Sepahsalari 2015) or both features (Shimer and Werning 2008, Chetty 2008).[6] Absent unemployment risk, assets have no role for precautionary savings purposes, and they are simply used for consumption smoothing until the single spell ends and a permanent job is found.[7] Importantly, we show in our model that saving decisions interact with the changes in UI policy because wealth is a substitute for UI payments for self-insurance purposes. The

[6]Although the baseline model in Krusell et al. (2010) incorporates aggregate fluctuations, they study the welfare effects of changes in UI policy in a steady-state experiment. The baseline model in Chetty (2008) has no unemployment risk, but he presents an extension to incorporate it, and he shows that his main results hold under extra assumptions.

[7]Typically, in these models, all agents are initially unemployed, and they decide when to accept a permanent employment offer. These models are often called single-spell models.

changes in saving decisions in turn significantly affect the search effort and wage choices of the unemployed as well as the consumption patterns of everyone in the economy. On the other hand, a model in which aggregate risk is absent makes the insurance value of UI time-invariant. In our framework with aggregate risk, the strength of precautionary saving motives significantly varies with the level of unemployment risk over the business cycle. Incorporating this feature is especially important to understand the optimality of time-varying UI policy.[8]

Finally, other papers investigate the impact of the Great Recession extensions of UI duration on macroeconomic outcomes.[9] Pei and Xie (2016) relax the perfect commitment assumption and analyze the effects of time-consistent policy over the business cycle in a model with search frictions but risk averse agents are not allowed to save or borrow. They find that while benefit extensions resulted in higher unemployment, it provided welfare gains ex post compared to a no-extensions policy. We show that even when government can commit perfectly to its UI policy, the optimal policy is countercyclical when we account for changes in precautionary saving motives over the cycle. Two recent papers study this question in a framework with search frictions and incomplete markets. First, Nakajima (2012) carefully models UI extensions during the Great Recession and its subsequent recovery using a model with business cycle dynamics and then measures the effect of these extensions on the unemployment rate. He does not, however, study the welfare effects of these changes in UI policy. We extend his model to a general equilibrium model in which the government finances the UI benefits and use the model to study how UI policy must vary over the business cycle. Second, Kekre (2017) studies the macroeconomic and welfare effects of UI extensions during the Great Recession in a model with nominal rigidities and

[8] Our paper has other important features compared to these papers in the literature. In terms of modeling, previous papers (except for Krusell et al. 2010 and Eeckhout and Sepahsalari 2015) use partial equlibrium models of the labor market. In these models, the changes in aggregate conditions of the economy or in UI policy do not affect firm hiring decisions and offered wages. In terms of welfare analysis, Shimer and Werning (2008) use an optimal contracting approach to study the optimal variation of UI over the unemployment duration. Chetty (2008) and Kroft and Notowidigdo (2016) find a locally optimal UI policy in a welfare exercise that can be used only to calculate the marginal welfare effects of small changes in the UI benefit level, relative to the observed UI benefit level in the data.

[9] See Hagedorn et al. (2016), Mitman and Rabinovich (2014), and Chodorow-Reich and Karabarbounis (2017), among many others.

constraints on monetary policy but without business cycle dynamics in the real business cycle tradition. In his model, when the unemployed have a higher marginal propensity to consume than the employed, generous UI policy increases the aggregate demand for consumption both in the current period and in the previous period because individuals endogenously reduce precautionary savings when they expect generous public transfers in the future. As a result, he finds that UI extensions reduced the unemployment rate and provided welfare gains during the Great Recession. Rather than only focusing on discretionary UI policy changes during the Great Recession, we solve for the optimal UI policy over the business cycle and find that it should be countercyclical even when business cycles are completely exogenous and that UI policy has no role on smoothing these fluctuations through its impact on aggregate demand. Complementary to his findings, we also show that the endogenous response of precautionary savings to changes in UI generosity is key to understanding the true welfare benefits and costs of UI benefits.

On the theoretical side, our model is a heterogeneous agent general equilibrium directed search model of the labor market with aggregate labor productivity driven business cycles as in Menzio and Shi (2010, 2011). The market structure enables us to overcome the computational difficulties of solving a model of this type by utilizing the block recursive equilibrium. We extend their framework by incorporating asset markets as in Herkenhoff (2017) to study the optimal UI over the business cycle with endogenous wealth distribution. To the best of our knowledge, our model is the first to study this question in a model with endogenous wage determination, search frictions, incomplete markets, and aggregate fluctuations.

This paper is organized as follows. We present our model in Section 2.2. Then, Section 2.3 describes the calibration strategy and model fit. Section 2.4 explains the calculation of the welfare effects of various UI policies. Section 2.5 contains the main results. In Section 2.6, we provide a detailed discussion on our results and conduct robustness checks. Section 2.7 provides preliminary evidence from the micro-data that support the model's main mechanism. Finally, Section 2.8 concludes.

2.2 Model

This section first introduces the environment of the model in Section 2.2.1. We then lay out the problem of the household and firm in Section 2.2.2 and Section 2.2.3, respectively. Next, we explain the government's UI policy in Section 2.2.4. Finally, Section 2.2.5 defines the equilibrium of the model and characterizes the job search behavior of the unemployed.

2.2.1 Environment

Time is discrete and denoted by $t = 0, 1, 2, ...$ Individuals are infinitely lived and ex-ante identical, with preferences given by

$$\mathbb{E}_0 \sum_{t=0}^{\infty} \beta_t \left[u\left(c_t\right) - \mathbf{1}_U \left[\nu\left(s_t\right) \right] \right]$$

where $u\left(\cdot\right)$ is a strictly increasing and strictly concave utility function over consumption level c that satisfies Inada conditions, $\mathbf{1}_U$ is an indicator function that takes the value of one if the agent is unemployed, and $\nu\left(\cdot\right)$ represents the disutility associated with search effort of the unemployed and is a strictly increasing and strictly convex function of search intensity s. Moreover, β_t is a stochastic variable that is idiosyncratic - i.i.d. across agents - and describes the cumulative discounting between period 0 and period t. In particular, $\beta_{t+1} = \tilde{\beta}\beta_t$, where $\tilde{\beta}$ is a five-state, first-order Markov process as in Krusell et al. (2009). The heterogeneity in discount rates allows us to match important features of the empirical asset distribution, as we will discuss in Section 2.3.1.

In the model, individuals are heterogeneous in terms of their labor market status, asset holdings, labor market earnings, and stochastic discount rate. An agent can be classified into one of the following labor market statuses: a worker W, an unemployed individual who is eligible for unemployment insurance benefits UE, or an unemployed individual who is ineligible for unemployment insurance benefits UI.[10]

[10]Farber et al. (2015) find that UI extensions reduced the labor force exits by 20 to 30 percent during 2008-2011 and 2012-2014 respectively. Notice that even if our model does not incorporate a labor force participation margin, we find that the optimal policy is countercyclical. As a result, given that UI genorosity increases labor force participation, the welfare gains from the optimal policy actually constitute a lower bound in our model.

The labor market features directed search. Unemployed individuals direct their search effort $s \in [0,1]$ toward wage submarkets indexed by w. Once matched with a firm within submarket w, the household is paid a fixed wage w every period until the match exogenously dissolves, as in Menzio and Shi (2010).[11] Unemployed individuals who are eligible for UI benefits receive a fraction of the wage they received during their last employment, whereas the unemployed ineligible do not receive any benefits. In order to finance the unemployment insurance program, the worker and unemployed eligible pay a fraction τ of their wages/benefits to the government every period. In addition to labor earnings, all households have access to incomplete asset markets where they can save/borrow at an exogenous interest rate r.[12] On the other side of the labor market, firms decide the wage submarket in which to post a vacancy. Once matched with a worker, the firm-worker pair operates a constant returns to scale technology that converts one indivisible unit of labor into final consumption goods. All firm-worker pairs are assumed to be identical in terms of their production efficiency; that is, the amount of production only depends on aggregate labor productivity.

The timing of the model is as follows. At the beginning of each time period t, aggregate labor productivity p and the idiosyncratic discount rate β for each agent realize. The period labor productivity level p completely determines 1) the UI replacement rate $\phi(p) \in [0,1]$ and the stochastic UI expiration rate $e(p) \in [0,1]$, and 2) the exogenous job separation rate $\delta(p) \in [0,1]$. This implies that $\delta(p)$ fraction of those who were workers in $t-1$ lose their jobs and must spend at least one period being unemployed. Among those who lose their job, $e(p)$ fraction become ineligible for unemployment benefits. After the realization of the exogenous shocks, there are two stages in each time period t where agents make endogenous decisions.

First, in the labor market stage, firms decide the wage submarket in which to post a

[11]In Section 2.6.1, we extend our baseline model to endogenize the quit decisions of workers and explore the quantitative implications of this assumption on our main results.

[12]We could endogenize the interest rate by modeling an asset market in which financial intermediaries post asset returns in different locations and individuals look for saving/borrowing opportunities in these different locations depending on their state variables. This is similar to Herkenhoff (2017). In our baseline model, we abstract from this and assume a constant and exogenous interest rate. In Section 2.6.1, we explore the quantitative implications of this assumption.

vacancy, while the unemployed choose a wage submarket w within which to look for a job. Second, the production and consumption stage of time t open where each firm-worker pair produces p units of consumption goods, wages are paid to workers, UI benefits are paid to eligible unemployed as a fraction $\phi(p)$ of their previous wages, and any unemployed receive the monetized value of non market activities h.[13] The households then make their saving/borrowing decision. Finally, prior to time $t+1$, unemployed households decide the search effort level s they will exert in the labor market stage of time $t+1$ where the utility cost of that search effort is incurred at time t.

It is important to discuss the reasons why this environment is useful in answering our question. Beyond the obvious features of the presence of incomplete markets, a UI program, and equilibrium unemployment, we would like to consider an equilibrium model of the labor market in which firm and household decisions are affected by both aggregate fluctuations and changes in UI policy. This way, we are able to incorporate the moral hazard costs of generous UI policies on the job search intensity and wage choice behavior of the unemployed, as well as changes in the vacancy creation incentives of firms over the business cycle. Moreover, directed search is useful not only because of tractability reasons but also because of its implications for equilibrium efficiency. In particular, under some conditions, the equilibrium is efficient in the directed search model but not in a random search model with Nash bargaining.[14] Hence, in our framework, the government insurance program aims to fix the inefficiencies caused by incomplete asset markets.

2.2.2 Household problem

A household's state vector consists of current employment status $l \in \{W, UE, UI\}$, net asset level $a \in \mathcal{A} \equiv [\underline{a}, \bar{a}] \subseteq \mathbb{R}$, the current wage level $w \in \mathcal{W} \equiv [\underline{w}, \bar{w}] \subseteq \mathbb{R}_+$ if the

[13] The variable h encompasses both the value of leisure/home production and other income such as spousal and family income and other transfers. Our results would be similar if h is a utility value instead of a monetary value.

[14] See Acemoglu and Shimer (1999), Burdett, Shi, and Wright (2001), Shi (2001), and Menzio and Shi (2011) for the efficiency of directed search equilibrium. As discussed by Menzio and Shi (2011), however, the equilibrium of our baseline model does not maximize the joint value of a match (and thus it is not bilaterally efficient) because of the limitations in the contract space. In Section 2.6.1, we extend our baseline model to a model with endogenous quit decisions and show that the effects of inefficiencies present in the labor market of the baseline model on our main results are negligible.

employment status is W or the wage level from the previous job if the employment status is UE, and the current discount rate $\beta \in \mathcal{B} \equiv \left[\underline{\beta}, \bar{\beta}\right] \subset (0,1)$.

The aggregate state is denoted by $\mu = (p, \Gamma)$, where $p \in \mathcal{P} \subseteq \mathbb{R}_{+}$ denotes the current aggregate labor productivity and $\Gamma : \{W, UE, UI\} \times \mathcal{A} \times \mathcal{W} \times \mathcal{B} \to [0,1]$ denotes the distribution of agents across employment status, asset level, current/previous wage level, and discount rate. The law of motion for the aggregate states is given by $\Gamma' = H(\mu, p')$ and $p' \sim F(p' \,|\, p)$.

The recursive problem of the worker is given by

$$\begin{aligned} V^{W}(a, w, \beta; \mu) \quad = \quad \max_{c, a'} \quad & u(c) + \beta \mathbb{E}\Bigg[\delta\left(p'\right)\left(1 - e\left(p'\right)\right) V^{UE}\left(a', w, \beta'; \mu'\right) \\ & + \delta\left(p'\right) e\left(p'\right) V^{UI}\left(a', \beta'; \mu'\right) \\ & + \left(1 - \delta\left(p'\right)\right) V^{W}\left(a', w, \beta'; \mu'\right) \Big| \beta, \mu \Bigg] \end{aligned} \tag{2.1}$$

subject to

$$\begin{aligned} & c + a' \leq (1 + r)\, a + w(1 - \tau) \\ & a' \geq -\underline{a} \\ & \Gamma' = H\left(\mu, p'\right) \quad \text{and} \quad p' \sim F\left(p' \,|\, p\right). \end{aligned}$$

Notice in the above problem that the worker may not qualify for UI benefits with probability e after losing her job due to exogenous job separation, which captures both voluntary and involuntary reasons for job loss in our model. This feature intends to capture the fact that according to the current UI policy in the United States, not all workers transitioning into unemployment qualify for UI benefits. In particular, individuals do not qualify for benefits if they voluntarily quit their job or if they do not meet certain work/earnings requirements.[15]

The unemployed directs her job search effort toward a wage submarket indexed by w with an associated market tightness given by $\theta(w; \mu)$, which is an equilibrium object defined later. Let $f(\theta(w; \mu))$ be the job finding probability for the unemployed who visits submarket w when the aggregate state is μ. Then, we lay out the recursive problem of eligible

[15]The unemployed must meet requirements for wages earned or time worked during an established period of time referred to as the base period. In most states of the United States, this is usually the first four out of the last five completed calendar quarters prior to the time that a UI application is filed.

unemployed as follows:

$$
\begin{aligned}
V^{UE}(a,w,\beta;\mu) \quad = \quad & \max_{c,a',s} \quad u\left(c\right)-\nu\left(s\right)+\beta\mathbb{E}\Bigg[\max_{\tilde{w}}\Big\{sf\left(\theta\left(\tilde{w};\mu'\right)\right)V^{W}\left(a',\tilde{w},\beta';\mu'\right) \\
& +\left(1-sf\left(\theta\left(\tilde{w};\mu'\right)\right)\right)\left(1-e\left(p'\right)\right)V^{UE}\left(a',w,\beta';\mu'\right) \\
& +\left(1-sf\left(\theta\left(\tilde{w};\mu'\right)\right)\right)e\left(p'\right)V^{UI}\left(a',\beta';\mu'\right)\Big]\Big\}\Big|\beta,\mu\Bigg]
\end{aligned} \tag{2.2}
$$

subject to

$$
\begin{aligned}
& c+a'\leq\left(1+r\right)a+h+\phi\left(p\right)w\left(1-\tau\right) \\
& a'\geq-\underline{a} \\
& \Gamma'=H\left(\mu,p'\right) \quad \text{and} \quad p'\sim F\left(p'\,|\,p\right).
\end{aligned}
$$

where the eligible unemployed receives a fraction ϕ of her previous wage as UI benefits and pays τ fraction as labor income tax. As described earlier, she may lose her eligibility with probability e if she is unable to find a job during the labor market stage of the current period. When choosing the wage submarket to search for jobs, the unemployed individual faces the trade-off between the level of surplus from a possible match (i.e., the wage level) and the probability of finding a job because of the lower number of vacancies posted for high-paying jobs.

Finally, the recursive problem of the ineligible unemployed is given by

$$
\begin{aligned}
V^{UI}(a,\beta;\mu) \quad = \quad & \max_{c,a',s} \quad u\left(c\right)-\nu\left(s\right)+\beta\mathbb{E}\Bigg[\max_{\tilde{w}}\Big\{sf\left(\theta\left(\tilde{w};\mu'\right)\right)V^{W}\left(a',\tilde{w},\beta';\mu'\right) \\
& +\left(1-sf\left(\theta\left(\tilde{w};\mu'\right)\right)\right)V^{UI}\left(a',\beta';\mu'\right)\Big\}\Big|\beta,\mu\Bigg]
\end{aligned} \tag{2.3}
$$

subject to

$$
\begin{aligned}
& c+a'\leq\left(1+r\right)a+h \\
& a'\geq-\underline{a} \\
& \Gamma'=H\left(\mu,p'\right) \quad \text{and} \quad p'\sim F\left(p'\,|\,p\right).
\end{aligned}
$$

Notice that in the above problem, the unemployed ineligible is unable to regain eligibility for UI benefits if job search fails. This captures the fact that according to current UI policy in the United States, the unemployed receive UI benefits only for a certain number of weeks - which varies over the business cycle - and once that threshold is reached, the unemployed cannot continue to collect UI benefits.

2.2.3 Firm problem

Firms post vacancies offering fixed wage contracts in certain wage submarkets. The labor market tightness of submarket w is defined as the ratio of vacancies v posted in the submarket to the aggregate search effort S exerted by all the unemployed searching for a job within that submarket. It is denoted as $\theta(w;\mu) = \frac{v(w;\mu)}{S(w;\mu)}$. Let $M(v,u)$ be a constant returns to scale matching function that determines the number of matches in a submarket with S level of aggregate search effort and v vacancies. We can then define $q(w;\mu) = \frac{M(v(w;\mu),S(w;\mu))}{v(w;\mu)}$ to be the vacancy filling rate and $f(w;\mu) = \frac{M(v(w;\mu),S(w;\mu))}{S(w;\mu)}$ to be the job finding rate in submarket w when aggregate state is μ. The constant returns to scale assumption on the matching function guarantees that the equilibrium object θ suffices to determine job finding and vacancy filling rates since $q(\theta) = \frac{M(v,S)}{v} = M\left(1, \frac{1}{\theta}\right)$ while $f(\theta) = \frac{M(v,S)}{S} = M(\theta, 1)$.
First, consider a firm that is matched with a worker in submarket w when the aggregate state is μ. The pair operates under a linear production technology and produces p units of output, and there is no capital in the economy. The worker is paid a fixed wage of w and with some probability $\delta(p)$, the match dissolves. Hence, the value of a matched firm is given by

$$J(w;\mu) = p - w + \frac{1}{1+r}\mathbb{E}\left[\left(1-\delta\left(p'\right)\right) J\left(w;\mu'\right) \Big| \mu\right] \tag{2.4}$$

subject to

$$\Gamma' = H\left(\mu, p'\right) \quad \text{and} \quad p' \sim F\left(p' \,|\, p\right).$$

Meanwhile, the value of a firm that posts a vacancy in submarket w under aggregate state μ is given by

$$V(w;\mu) = -\kappa + q(\theta(w;\mu)) J(w;\mu), \tag{2.5}$$

where κ is a fixed cost of posting a vacancy that is financed by risk-neutral foreign entrepreneurs who own the firms.
When firms decide the submarket in which to post vacancies to maximize profits, they face the trade-off between the probability of filling a vacancy and the level of surplus from a

possible match. This is because if a firm posts a vacancy in a low (high) wage submarket, then the level of the surplus from the match in that submarket will be higher (lower) for the firm, but the probability of filling the vacancy will be lower (higher) as less (more) unemployed individuals visit that submarket to search for a job.

The free entry condition implies that profits are just enough to cover the cost of filling a vacancy in expectation. As a result, the owner of the firm makes zero profits in expectation. Thus, we have $V(w;\mu) = 0$ for any submarket w such that $\theta(w;\mu) > 0$. Then, we impose the free entry condition to Equation (2.5) and obtain the equilibrium market tightness:

$$\theta(w;\mu) = \begin{cases} q^{-1}\left(\frac{\kappa}{J(w;\mu)}\right) & \text{if} \quad w \in \mathcal{W}(\mu) \\ 0 & \text{otherwise.} \end{cases} \tag{2.6}$$

The equilibrium market tightness contains all the relevant information needed by households to evaluate the job finding probabilities at each submarket.

2.2.4 Government policy

The UI policy is characterized by $\{\phi(p), e(p), \tau\}$, where $\phi(p)$ is the replacement rate and $e(p)$ is the expiration rate, both of which may vary with current aggregate labor productivity p.[16] A labor income tax τ is levied on the labor earnings of the worker and on the UI benefits of the eligible unemployed in order to finance the UI program.[17] The benefit expiration rate $e(\cdot)$ is stochastic, as in Fredriksson and Holmlund (2001), Albrecht and Vroman (2005), Faig and Zhang (2012), and Mitman and Rabinovich (2015). This assumption simplifies the solution of the model because we do not need to carry the unemployment duration as another state variable for the eligible unemployed.

[16] We restrict the UI policy to depend on the aggregate state of the economy μ only through the current aggregate labor productivity p and not through the distribution of individuals across states Γ. This restriction allows our model to retain the block recursitivity, which we will explain in Section 2.2.5.

[17] We focus on the optimality of government policies that can be conditioned on the employment status of the individuals so that the government policies provide insurance against unemployment risk. Also, if the government finds it optimal to make transfers (by reducing taxes) during recessions, it can obviously do this by increasing the UI replacement rate and duration. For these reasons, we consider time-invariant income tax schedules in our analysis.

The government balances the following budget constaint in expectation:[18]

$$\sum_{t=0}^{\infty}\sum_{i}\left(\frac{1}{1+r}\right)^{t}\times\left[\mathbf{1}_{\{l_{it}=W\}}\times w_{it}+\mathbf{1}_{\{l_{it}=UE\}}\times w_{it}\phi\left(p_{t}\right)\right]\times\tau$$
$$=\sum_{t=0}^{\infty}\sum_{i}\left(\frac{1}{1+r}\right)^{t}\times w_{it}\phi\left(p_{t}\right)\times\mathbf{1}_{\{l_{it}=UE\}} \tag{2.7}$$

where the left-hand side is the present discounted value of tax revenues collected from the labor income of workers and the unemployed eligible, and the right-hand side is the present discounted value of UI payments to the unemployed eligible.

2.2.5 Equilibrium

Definition of the Recursive Equilibrium: Given a UI policy $\left\{\tau,\phi\left(p\right),e\left(p\right)\right\}_{p\in\mathcal{P}}$, a recursive equilibrium for this economy is a list of household policy functions for assets $\left\{g_a^l\left(a,w,\beta;\mu\right)\right\}_{l=\{W,UE\}}$ and $g_a^{UI}\left(a,\beta;\mu\right)$, wage choices $g_w^{UE}\left(a,w,\beta;\mu\right)$ and $g_w^{UI}\left(a,\beta;\mu\right)$, search effort $g_s^{UE}\left(a,w,\beta;\mu\right)$ and $g_s^{UI}\left(a,\beta;\mu\right)$, a labor market tightness function $\theta\left(w;\mu\right)$, and an aggregate law of motion $\mu'=\left(p',\Gamma'\right)$ such that

1. Given government policy, shock processes, and the aggregate law of motion, the household's policy functions solve their respective dynamic programming problems (2.1), (2.2), and (2.3).

2. The labor market tightness is consistent with the free entry condition (2.6).

3. The government budget constraint (2.7) is satisfied.

4. The law of motion of the aggregate state is consistent with household policy functions.

Notice that in order to solve the recursive equilibrium defined above, one must keep track of an infinite dimensional object Γ in the state space, making the solution of the model infeasible. To address this issue, we utilize the structure of the model and use the notion of block recursive equilibrium developed by Menzio and Shi (2010, 2011).

[18] This assumption is motivated by the fact that according to the current UI system in the United States, states are allowed to borrow from a federal UI trust fund when they meet certain federal requirements, and thus they are allowed to run budget deficits during some periods.

Definition of the Block Recursive Equilibrium (BRE): A BRE for this economy is an equilibrium in which the value functions, policy functions, and labor market tightness depend on the aggregate state of the economy μ, only through the aggregate productivity p, and not through the aggregate distribution of agents across states Γ.

Now, we prove that our model admits block recursitivity.

Proposition 1: *If i) utility function $u(\cdot)$ is strictly increasing, strictly concave, and satisfies Inada conditions; $\nu(\cdot)$ is strictly increasing and strictly convex, ii) choice sets $\mathcal{W}$ and $\mathcal{A}$, and sets of exogenous processes $\mathcal{P}$ and $\mathcal{B}$ are bounded, iii) matching function M exhibits constant returns to scale, and iv) UI policy is restricted to be only a function of current aggregate labor productivity, then there exists a Block Recursive Equilibrium for this economy. If, in addition, $M = min\{v, S\}$, then the Block Recursive Equilibrium is the only recursive equilibrium.*

Proof: See Appendix B.2

Proposition 1 is very useful because it allows us to solve the model numerically without keeping track of the aggregate distribution of agents across states Γ. One should be careful when interpreting this result. Even though we can solve for the policy functions, value functions, and labor market tightness independent of Γ, it does not mean that the distribution of agents is irrelevant for our analysis. Notice that the evolution of macroeconomic aggregates such as the unemployment rate, average spell duration, and wealth distribution of the economy is determined by household decision rules in the labor market and financial market. These decisions, in turn, are functions of individual states whose distribution is determined by Γ. Hence, the evolution of aggregate variables after a change in UI policy will depend on the distribution of agents in the economy at the time of the policy change. Notice that if the UI policy instruments were to depend on the unemployment rate of the economy, then it would break the block recursivity of the model. This is because agents would need to calculate next period's unemployment rate to know the replacement rate and UI duration next period. However, this requires calculating the flows in and out of unemployment, the latter of which depends on the distribution of agents across states Γ. Although the changes in UI policy are triggered by the changes in the unemployment rate according to the current UI program in the United States, the assumption that UI policy depends on aggregate productivity is not too restrictive because of the strong correlation between the unemployment rate and aggregate labor productivity in the model.

Job search decision rules We now characterize the job search behavior of the unemployed. This will supplement our discussions of the main results of the paper in Section 2.5.

Figure 2.1 plots the labor market behavior of the eligible unemployed holding various levels of wealth under a less generous UI policy and a generous UI policy. It shows that the search intensity is decreasing in wealth, whereas the wage choice is increasing in wealth for any UI policy.[19]

Moreover, similar to Krusell et al. (2010), the marginal effect of an increase in assets on wage choice and search effort is relatively higher for the borrowing-constrained unemployed.[20] While this result is unsurprising and intuitive, it highlights the importance of accounting for wealth heterogeneity across agents, since the aggregate search effort and wage levels in the economy now crucially depend on the underlying wealth distribution. An economy where agents are relatively wealthy is likely to exhibit lower levels of aggregate search and higher wages, whereas the opposite is true when wealth levels are low. Since business cycles induce changes in precautionary savings and thus variation in aggregate search effort and wage choices, the optimal design of UI policy over the business cycle must account for this channel. For instance, in a recession where many individuals deplete their existing wealth, this channel exerts an upward pressure on search effort and downward pressure on wage choices as agents seek to find jobs more quickly. This effect dampens the moral hazard costs induced by introducing a more generous UI policy during recessions, since poorer agents tend to ramp up job-finding efforts themselves.

Next, a comparison of the two policy functions across UI policies highlights two important points. First, generous UI payments entail incentive costs because they lead the eligible

[19] Notice that there is little dispersion across wage choices of the unemployed holding different levels of wealth. Hornstein et al. (2011) show that frictional wage dispersion - measured by the mean-min wage ratio - is very small in a directed search model. When calibrated to match the empirical asset distribution and salient features of the labor market prior to the Great Recession in the United States, our model generates a mean-min wage ratio of 1.034, in line with their conclusion (less than 1.05).

[20] These patterns are also present for the ineligible unemployed.

Figure 2.1: Search effort and wage choice policy functions of the eligible unemployed

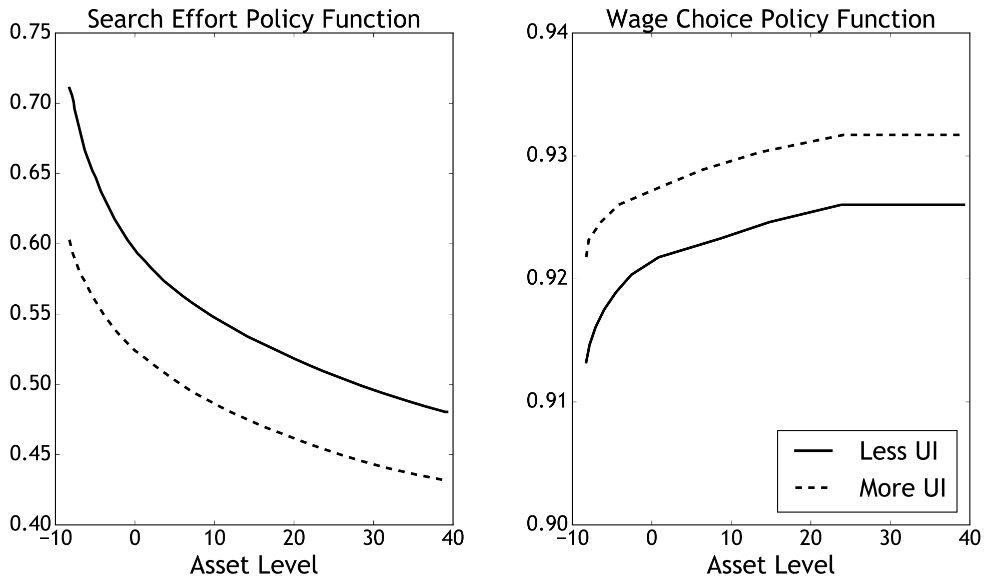

Note: These figures plot the search effort and wage choice policy functions of the eligible unemployed holding different levels of wealth for average levels of labor productivity, discount rate, and previous wage under a less generous and a more generous UI policy.

unemployed to decrease their search effort and increase their wage choices.[21] The combined effect of lower search effort and a shift toward higher-paying jobs, which are more difficult to find, results in a lower aggregate job finding rate and prolonged unemployment spells. Second, the unemployed holding different levels of wealth respond in varying degrees to changes in UI policy. Similar to Chetty (2008), wealthier agents are less responsive to changes in UI policy because the insurance value of a marginal increase in benefits is less important to them. This implies that a model that abstracts from self-insurance altogether overestimates the level of the moral hazard costs of introducing a more generous UI policy. The assumption that agents have no access to asset markets effectively raises the aggregate elasticity of search effort and wage choice to changes in UI policy, since the most responsive agents are precisely those with the least available self-insurance. As a result, it is crucial for the model to match the observed asset distribution in the data in order to generate the

[21]This result is also established in the previous literature. See Shavell and Weiss (1979), Hopenhayn and Nicolini (1997), and Acemoglu and Shimer (1999), among many others.

true magnitude of moral hazard costs in the model.

2.3 Calibration

We calibrate the stochastic steady state of our model to match salient features of the labor market and asset distribution of the U.S. economy prior to the Great Recession. In doing so, we feed into the model a constant replacement rate and expiration rate, which we call the *acyclical/flat* policy.

The model period is taken to be a week. We use the following separable functional form for the period utility function:

$$u(c_t) - \mathbf{1}_U [\nu(s_t)] = \frac{c_t^{1-\sigma}}{1-\sigma} - \mathbf{1}_U \left[\alpha \frac{s_t^{1+\chi}}{1+\chi} \right],$$

which is also used by Chetty (2008) and Nakajima (2012). We restrict the values of discount rates to be symmetric around an average value $\bar{\beta}$ with a difference of η between two adjacent values. Moreover, we allow β to take five different values. In our simulations of the model, we set 40 percent of the population to the middle discount rate value and 10 percent to each extreme point in any time period. The expected duration of being in the extreme discount rate value is set to be 50 years, where transitions can only occur between adjacent values.

The labor market matching function is $M(v, S) = \frac{vS}{[v^\gamma + S^\gamma]^{1/\gamma}}$ as in den Haan et al. (2000). This CES functional form of the matching function implies that both the job finding rate $f(\theta) = \theta (1 + \theta^\gamma)^{-1/\gamma}$ and the vacancy filling rate $q(\theta) = (1 + \theta^\gamma)^{-1/\gamma}$ are between 0 and 1.

Following Shimer (2005), we use a process for the job destruction rate that depends only on labor productivity, $\delta_t = \bar{\delta} \times \exp(\omega(p_t - 1))$, where $\bar{\delta}$ is the average weekly exogenous job destruction rate in the data. These separation shocks can be interpreted as idiosyncratic match quality shocks that drive down the productivity of a match to a low enough level so that the match endogenously finds it optimal to dissolve, as in Lise and Robin (2017). Moreover, the probability of this idiosyncratic event is correlated with the aggregate state of the economy. As a result, this specification allows the model to capture the cyclicality of employment-to-unemployment (E-U) transitions.[22] We then calibrate ω so that the volatility of quarterly E-U transitions in the model matches its data counterpart, which we calculate using E-U transition rates measured by Fujita and Ramey (2009) for the time period 1976:I-2005:IV.[23]

The logarithm of the aggregate labor productivity p_t follows an AR(1) process:

$$\ln p_{t+1} = \rho \ln p_t + \sigma_\epsilon \epsilon_{t+1}.$$

We take p_t as the mean real output per person in the non-farm business sector. Using the quarterly data constructed by the Bureau of Labor Statistics (BLS) for the time period 1951:I-2007:IV, we estimate the above process at a weekly frequency and find that $\rho = 0.9720$ and $\sigma_\epsilon = 0.0025$.

Next, we calibrate the replacement rate and expiration rate of the acyclical/flat policy by matching the long-run empirical averages of U.S. government policy. First, we discuss the calibration of the replacement rate. Chodorow-Reich and Karabarbounis (2016) measure the mean of pretax benefits per recipient as 21.5 percent of pretax marginal product.[24] Under a mean take-up rate of UI benefits among the eligible unemployed of 65 percent, this implies setting the mean of pretax benefits per recipient to 14 percent, since we do not model UI take-up decisions given the completixity of our framework.[25] Second, we take the UI benefit duration as 26 weeks (2 quarters), which is the standard benefit duration

[22] Empirically, Elsby et al. (2009), Fujita and Ramey (2006, 2009), Yashiv (2007), and Fujita (2011a) show that the separation rate into unemployment is countercyclical.

[23] The model-implied Beveridge curve, which plots the relationship between unemployment and vacancies, exhibits a negative slope as in the data. This is because when labor productivity declines, firms cut back on vacancies, which translates to lower job finding rates and higher unemployment. Moreover, the rise in separation shocks further amplifies the increase in unemployment. As a result, unemployment and vacancies move in the opposite direction.

[24] This value is consistent with a replacement rate level that accounts for the difference between wage and total compensation, the difference between compensation and the marginal product, and the gap in productivity and compensation between those receiving UI and the economywide average. In our model, wages are not exactly equal to marginal product because of frictions, but the difference between the two is small.

[25] Estimates in the literature for the fraction of all eligibles who receive UI range from 50 to 77 percent using Current Population Survey (CPS) data for different samples. Fuller, Ravikumar, and Zhang (2013) find that during the Great Recession, only about 50 percent of those eligible collected their benefits. Vroman (1991) uses CPS supplements from 1989 and 1990 and finds 53 percent. Blank and Card (1991) estimate

without extensions. Under the stochastic steady state calibration of our model, these two numbers require us to set $\phi_t = 0.14$ and $e_t = 1/26\ \forall t$ as the acyclical/flat policy.[26] Then, a labor income tax rate of $\tau = 0.36$ percent balances the government budget in equilibrium when the unemployment rate is 4.8 percent.

Having specified functional forms, the law of motion of the productivity process, and UI policy, we now calibrate several parameters outside of our model. We choose a coefficient of relative risk aversion $\sigma = 2$ and set $r = 0.095$ percent, which generates an annual return on assets of around 5 percent. Hagedorn and Manovskii (2008) estimate the combined capital and labor costs of vacancy creation as 58 percent of weekly labor productivity. Following their estimate, we set the cost of vacancy creation as $\kappa = 0.58$.

We measure the average weekly job separation rate $\bar{\delta}$ using data from the Survey of Income and Program Participation (SIPP) for the time period between 2005 and 2007. The SIPP comprises individual level longitudinal data in which each respondent provides information on monthly income and government transfers as well as weekly labor force status. We restrict our sample to individuals between the ages of 24 and 65 who do not own a business or derive income from self-employment. We classify the individual as employed (E) if he/she reports having a job and either working or not on layoff, but absent without pay. We classify the individual as unemployed (U) if he/she reports either having no job and actively looking for work or having a job but currently laid off. We then calculate the average E-U transition rate in the data where we account for seasonality by removing weekly fixed effects and obtain $\bar{\delta} = 0.0022$.

This leaves us eight parameters to be calibrated: i) the average value of discount rates $\bar{\beta}$, ii) the difference between two adjacent discount rates η, iii) the borrowing limit $\underline{a}$, iv) the

---

the take-up rate as 71 percent for the period 1977–1987. Auray, Fuller, and Lkhagvasuren (2013) estimate the average take-up rate as 77 percent from 1989 to 2012 using detailed state-level eligibility criteria. Meanwhile, Anderson and Meyer (1997) use administrative data between the late 1970s and early 1980s and find that the take-up rate is 54 percent for a subsample that represents mainly separations from mass layoffs. In our baseline calibration, we set the take-up rate as 65 percent, which is around the mean of the above estimates in the literature.

[26] In Section 6.1, we also calculate the welfare gains from the optimal policy under 40 percent of the replacement rate (i.e., $\phi = 0.4\ \forall t$), which is the unadjusted replacement rate value calculated by the Department of Labor. We show that the optimal UI policy still yields significant welfare gains relative to the benchmark policy under this alternative high calibration of the replacement rate.

level parameter of the search cost function α, v) the curvature parameter of the search cost function χ, vi) the matching function parameter γ, vii) the separation rate parameter ω, and viii) the monetized value of non-market activity h. We jointly calibrate these parameters to match the following eight data moments, respectively: i) the median value of liquid asset holdings relative to weekly after-tax labor income distribution, ii) fraction of the population with non-positive liquid wealth, iii) the median value of the credit limit to labor income ratio, iv) the average unemployment rate, v) the response of the average unemployment duration to changes in the replacement rate, vi) the standard deviation of the unemployment rate, vii) the standard deviation of the job separation rate, and viii) the level of the opportunity cost of employment.

The first two moments related to the asset-to-income distribution is calculated from SIPP 2004 data and details are given in Section 2.3.1. Kaplan and Violante (2014) calculate the median value of the credit limit to quarterly labor income ratio for households aged 22 to 59 as 74 percent using Survey of Consumer Finances (SCF) data. We choose the borrowing limit parameter $\underline{a}$ so that the median value of the ratio of $\underline{a}$ to after-tax quarterly labor income in the model is 0.74.

The average unemployment rate and its standard deviation are calculated from U.S. data. In our baseline calibration, we choose the curvature parameter of the search cost function χ so that a 10 percentage point increase in the replacement rate generates an increase of 0.5 week in average unemployment duration among the UI eligible, which is within the range of available empirical estimates.[27] Hence, this parameter is important because it controls the magnitude of the incentive costs associated with the increase in UI payments.

Finally, Chodorow-Reich and Karabarbounis (2016) use a complete markets model and estimate the level of the opportunity cost of employment as 47 percent of the marginal product of employment under separable preferences. We choose the monetized value of non-market activity h so that the opportunity cost of employment generated by our model is 0.47. Given the incomplete markets model we have, to make the calibration comparable, we only simulate agents from the top 1 percent of the stationary asset-to-income distribution

[27] See Nakajima (2012) for the summary of empirical estimates. We evaluate the welfare gains from the optimal policy under different values of χ that match other levels of the available estimates in the literature. We find that the welfare gains from the optimal policy remain similar for different values of χ. These results are available upon request.

Table 2.1: Internally calibrated parameters

Parameter	Explanation	Value	Target	Model	Data
$\bar{\beta}$	Average discount rate	0.9986	Median asset-to-income ratio	6.17	6.22
η	Deviation from $\bar{\beta}$	0.0005	Frac. of pop. with non-positive wealth	0.27	0.27
$\underline{a}$	Borrowing limit	-8.25	Median credit-limit-to-income ratio	0.74	0.74
α	Level of search cost	5.02	Average unemployment rate	0.048	0.048
χ	Curvature of search cost	1.49	Response of average unemp. duration to changes in replacement rate	0.5	0.5
γ	Matching function parameter	0.217	Std. dev. of unemployment rate	0.10	0.12
ω	Separation rate parameter	-14.3	Std. dev. of separation rate	0.18	0.16
h	Value of nonmarket activity	0.342	Level of opportunity cost of emp.	0.47	0.47

Note: The average unemployment rate is calculated using monthly data between January 2005 and December 2007 provided by FRED - Federal Reserve Economic Data from the Federal Reserve Bank of St. Louis. The average standard deviation of the unemployment rate is reported in logs as quarterly deviations from an HP-filtered trend with a smoothing parameter of 1600, using quarterly data between 1951:I–2007:IV provided by FRED. The same procedure is applied to obtain the volatility of separation rates using data from Fujita and Ramey (2009) from 1976:I–2005:IV. The rest of the data moments are discussed in the main text.

when calculating the opportunity cost of employment using our model. This is because the behavior of the very rich agents in the incomplete markets model converges to the behavior of agents in the complete markets model. Section 2.3.2 explains how we calculate the opportunity cost of employment in our model. Later in Section 2.6.2, we target 0.955 as an alternative level of the opportunity cost of employment, which is the value calibrated by Hagedorn and Manovskii (2008), and discuss its implications for our main results.
Table 2.1 summarizes these calibrated parameters and compares the model's match to these data moments.

2.3.1 Asset distribution

In addition to monthly data on income and government transfers as well as weekly data on employment status, the SIPP also contains data on respondents' asset holdings. In each SIPP panel, respondents provide information on various types of asset holdings during two or three waves within the panel, usually one year or, equivalently, three waves apart. We use Wave 6 of the 2004 panel of SIPP, which covers interview months October 2005 - January 2006 and is the wave closest to the Great Recession that provides wealth holding information. We restrict our sample to individuals ages 24-65 and to those who neither own a business nor derive income from self-employment.

We use individual net liquid asset holdings as our primary measure of wealth because of its immediate availability as a means to smooth consumption in the event of job loss. The net liquid asset holdings of an individual are calculated by adding transaction accounts (checking, saving, money market, call accounts) and tradable assets (mutual funds, stocks, bonds), and then deducting unsecured debt. We follow Koehne and Khun (2015) and include net vehicle equity when calculating net liquid asset holdings. The reason is that income can decrease substantially upon unemployment, and some unemployed could resort to liquidating other forms of assets (i.e., the sale of vehicles) to smooth consumption upon job loss.

To normalize wealth and better capture the level of self-insurance, we compute respondents' asset-to-income ratio by dividing net liquid assets by weekly after-tax labor income.[28] We determine after-tax income using the statutory income tax codes. Table 2.2 shows the computed quantiles of the asset distribution in the data and model. The calibrated model comes close to matching the empirical asset distribution. In particular, our model reasonably captures the left tail of the distribution and at the same time exactly matches the fraction of the population holding non-positive liquid wealth. Matching the left tail of the distribution matters for our analysis because agents in this region of the distribution are the most affected by changes in UI policy. Job losers with low wealth have little to no

[28] We use weekly employment status information to obtain weekly labor earnings from monthly labor earnings data. We simply divide the monthly labor earnings by the number of weeks with a job for that month to obtain weekly labor earnings. Appendix B.1 provides more details on the calculation of the asset holdings and after-tax labor income.

Table 2.2: Percentiles of the distribution of liquid asset holdings relative to weekly after-tax labor income

	Quantiles					Fraction of population
	10th	25th	50th	75th	90th	with non-positive wealth
Data	−8.59	0.00	6.22	20.23	56.57	0.27
Model	−3.84	−0.85	6.17	33.32	42.46	0.27

Note: This table shows the liquid asset to after-tax labor income distribution in both the data and the model. The empirical distribution is calculated by the authors using the SIPP 2004 Panel. The main text provides the details of the calculation.

capacity to self-insure or smooth consumption using their own liquid assets and are thus very sensitive to changes in unemployment insurance generosity.

2.3.2 Opportunity cost of employment

To calculate the opportunity cost of employment across individual and aggregate states, we first derive surpluses obtained from moving from eligible unemployment to employment, and from ineligible unemployment to employment separately. Let $S^{UE}\left(a, w^{UE}, \beta; p\right)$ be the surplus derived by an unemployed eligible with state $\left(a, w^{UE}, \beta; p\right)$ who transitions into employment in a job that pays her optimal wage choice $\tilde{w}\left(a, w^{UE}, \beta; p\right)$. Similary, let $S^{UI}\left(a, \beta; p\right)$ be the surplus associated with moving from ineligible unemployment with state $(a, \beta; p)$ to a job that pays the optimal wage choice $\tilde{w}\left(a, \beta; p\right)$. We can then write

$$S^{UE}\left(a, w^{UE}, \beta; p\right) = V^{W}\left(a, \tilde{w}\left(a, w^{UE}, \beta; p\right), \beta; p\right) - V^{UE}\left(a, w^{UE}, \beta; p\right) \tag{2.8}$$

and

$$S^{UI}\left(a, \beta; p\right) = V^{W}\left(a, \tilde{w}\left(a, \beta; p\right), \beta; p\right) - V^{UI}\left(a, \beta; p\right). \tag{2.9}$$

Now consider the same individual who loses the aforementioned job that pays $\tilde{w}\left(a, w^{UE}, \beta; p\right)$. We can define the next period surplus of an eligible unemployed as

$$\begin{aligned} S\left(a'^{W}, \tilde{w}\left(a, w^{UE}, \beta; p\right), \beta'; p'\right) = & V^{W}\left(a'^{W}, \tilde{w}\left(a'^{W}, \tilde{w}\left(a, w^{UE}, \beta; p\right), \beta'; p'\right), \beta'; p'\right) \\ & - V^{UE}\left(a'^{W}, \tilde{w}\left(a, w^{UE}, \beta; p\right), \beta'; p'\right), \end{aligned} \tag{2.10}$$

where the right-hand side is the difference in the value of again finding a job that pays optimal wage choice $\tilde{w}\left(a'^{W}, \tilde{w}\left(a, w^{UE}, \beta; p\right), \beta'; p'\right)$ and remaining as unemployed eligible. Similarly, the next period surplus for the ineligible unemployed is given by

$$S^{UI}\left(a'^{W}, \beta'; p'\right) = V^{W}\left(a'^{W}, \tilde{w}\left(a'^{W}, \beta'; p'\right), \beta'; p'\right) - V^{UI}\left(a'^{W}, \beta'; p'\right) \tag{2.11}$$

Evaluating V^{W}, V^{UE}, and V^{UI} at a'^{W} in Equations (2.10) and (2.11) restricts the continuation surpluses to only that part associated with entering next period in the employed state. Next, substituting (2.1), (2.2), and (2.3) into (2.8) and (2.9), we obtain

$$\begin{aligned} \frac{S^{UE}\left(a, w^{UE}, \beta; p\right)}{\lambda^{W}} = & \tilde{w}\left(a, w^{UE}, \beta; p\right) \times (1-\tau) - \underbrace{\left(z^{UE}_{flow} + z^{UE}_{a} + z^{UE}_{w} + z^{UE}_{elg}\right)}_{z^{UE}} \\ & + \beta \mathbb{E}\left[\frac{\lambda'^{W}}{\lambda^{W}} \times \frac{\left(1 - \delta\left(p'\right) - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right) S^{UE}\left(a'^{W}, \tilde{w}\left(\cdot\right), \beta'; p'\right)}{\lambda'^{W}}\right] \end{aligned} \tag{2.12}$$

and

$$\begin{aligned} \frac{S^{UI}\left(a, \beta; p\right)}{\lambda^{W}} = & \tilde{w}\left(a, \beta; p\right) \times (1-\tau) - \underbrace{\left(z^{UI}_{flow} + z^{UI}_{a} + z^{UI}_{w} + z^{UI}_{elg}\right)}_{z^{UI}} \\ & + \beta \mathbb{E}\left[\frac{\lambda'^{W}}{\lambda^{W}} \times \frac{\left(1 - \delta\left(p'\right) - sf\left(\theta\left(\tilde{w}\left(a'^{UI}, \beta'; p'\right); p'\right)\right)\right) S^{UI}\left(a'^{W}, \beta'; p'\right)}{\lambda'^{W}}\right], \end{aligned} \tag{2.13}$$

where λ^{W} is the the marginal utility of consumption for the worker. The opportunity cost of employment z^{l} for each unemployed type $l = \{UE, UI\}$ consists of four components: z^{l}_{flow} is simply the flow utility difference between a worker and an unemployed type l, z^{l}_{a} is the change in value due to differential asset accumulation between the employed and the unemployed type l, z^{l}_{w} measures the change in value due to wage differences that result from the possibility of losing a job the next period and finding another job with a different wage as opposed to keeping the same job, and finally, z^{l}_{elg} represents the difference in value due to changes in the likelihood of ineligibility. Appendix B.2 provides derivations of these terms in detail.[29] This calculation yields the opportunity cost of employment $z^{UE}\left(a, w^{UE}, \beta; p\right)$ and $z^{UI}(a, \beta; p)$ for each state. As discussed above, we then simulate agents from the top 1 percent of the stationary asset-to-income distribution and calculate a weighted average of the opportunity cost of employment among this group. We then choose the monetized value of non-market activity h so that the average opportunity cost of employment for the richest agents in our model is 0.47.

The derivations above show that the opportunity cost of employment in our model is beyond the flow utility difference between the employed and unemployed. Importantly, our calculation takes into account the dynamic effects of one period of additional employment on the opportunity cost of employment. Intuitively, one period of additional employment causes a relative decline in the budget, since the employed typically accumulate more assets. However, entering next period with higher levels of wealth creates an offsetting gain in the continuation value. Moreover, higher wealth holdings encourage the unemployed to search for higher wages and thus increase the possibility of higher labor income. Finally, one extra period of employment decreases the probability of ineligibility because the worker must separate from his job first before being subject to eligibility risk, as opposed to an unemployed eligible who constantly faces the risk of losing benefits. As a result, these dynamic benefits of employment measured respectively by z_a^l, z_w^l, and z_{elg}^l jointly dampen the flow opportunity cost of employment z_{flow}^l.[30]

[29] Our calculation extends the opportunity cost of employment derivation in Chodorow-Reich and Karabarbounis (2016). In addition to the asset differential z_a in the incomplete markets version of their model, we account for the wage differential z_w^l and ineligibility probability differential z_{elg}^l in our opportunity cost of employment formula for each unemployment type l. In addition, z_a varies for each unemployment type l in our setup.

[30] In our model, $z_a^l + z_w^l + z_{elg}^l$ is small for the richest agents, and thus z^l approaches z_{flow}^l in the calibration. This is because the dynamic benefits of one period of extra employment have little value for this group of agents. While disregarding these benefits does not affect the calibration of value of non-market activity h, $z_a^l + z_w^l + z_{elg}^l$ is relatively large for poorer agents. Thus, it is crucial to account for these dynamic benefits when calculating the opportunity cost of employment across different agents in the economy so that the insurance benefits and the incentive costs of any proposed UI policy are correctly measured.

2.3.3 Testable implications

In this section, we discuss our model's implications for several important untargeted moments of the data. First, we measure the economy wide size and cyclicality of marginal propensity to consume (MPC) as well as the average consumption drop upon job loss predicted by the model. These are then compared to available empirical estimates in the literature. It is important for the model to generate a reasonable level and cyclicality of MPCs and average consumption drop in order to properly measure the insurance benefits of any proposed UI policy. For example, if the consumption drop were very low, then the insurance benefit of UI would be understated. Second, we present how the model compares to the data on labor market transitions, survival probabilities into unemployment, and the aggregate impact of UI extensions on the unemployment rate. Generating transition rates and unemployment survival functions, that are in line with the data is crucial to understanding the individual labor market response (incentive costs) of the unemployed to changes in UI policy, and generating a reasonable response of the unemployment rate ensures that the aggregate effects of UI are well accounted for. The following sections present the results of these exercises.

Marginal propensity to consume

Figure 2.2 qualitatively demonstrates the consumption choices of agents across different asset holdings and employment states. The unemployed not only consume less than workers but also exhibit higher marginal propensities to consume. The differences in MPCs between workers and the unemployed is most evident for agents holding little wealth, but this differential eventually diminishes as wealth increases.

In order to quantitatively understand how MPCs differ across heterogeneous agents in the economy, Table 2.3 presents the average quarterly MPC of different asset-to-income and employment groups based on the stationary distribution of the economy. We compute the MPC of an agent by calculating the fraction of an unexpected transfer, scaled such that it is equivalent to \$500, that an agent spends on consumption. As in Kaplan and Violante (2014), we implement a \$500 rebate in order to ensure consistency with available empirical estimates that study the impact of tax rebates on consumption. Noticeably, the poor unemployed ineligible exhibit the highest MPC given the absence of both public and

Figure 2.2: Consumption policy function

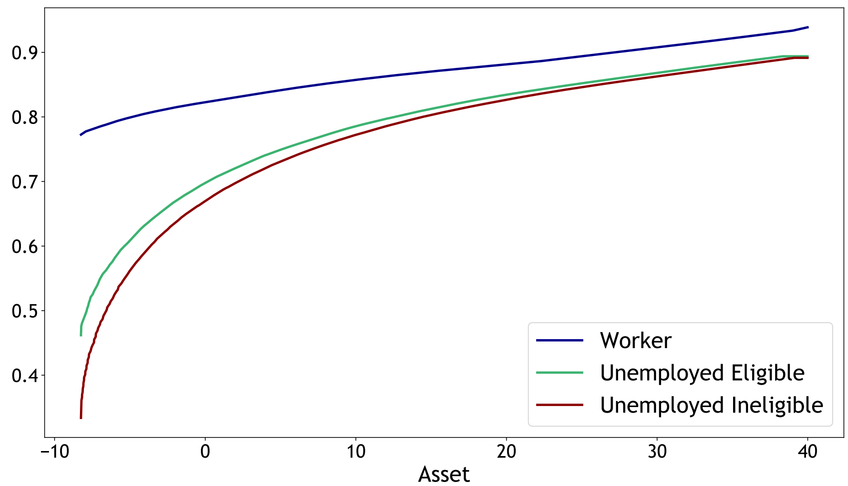

Note: This figure plots the consumption choices of agents with different employment statuses and asset holdings. The wages of workers and the unemployed eligible are set to be the economy's mean wage. Productivity and discount rates are also set to their means.

private insurance. Across employment states, the unemployed have significantly higher MPCs than workers, especially for agents in the lower end of the wealth distribution. Meanwhile, for any given employment status, the MPC is decreasing in wealth holdings.

The empirical literature documents two aggregate MPC data moments that we can use to validate our model. To do so, we calculate two untargeted average quarterly MPC moments in our model using the stationary distribution of agents across states and compare it to these available empirical estimates. Results are summarized in Table 2.4.

First, we find that the average quarterly economy wide MPC is 8 percent in our model. On the empirical side, Parker et al. (2013) measure that households, under different specifications, spend between 12 and 30 percent of unexpected tax rebates in the quarter that they are received. Thus, the fraction of borrowing-constrained individuals who have large MPCs as shown in Table 2.3 is too small to generate a sizeable response in the aggregate in our model.[31] Second, Gross et al. (2016) measure the cyclicality of the

[31] Notice that since our model generates a lower average MPC than its empirical counterpart, the households spend a relatively lower fraction of UI receipt on consumption. However, even if this is the case, we still find that the optimal policy is countercyclical. Thus, welfare gains provided by the optimal policy can

Table 2.3: Heterogeneous MPCs

	Asset-to-Income Groups				
Employment	a_1	a_2	a_3	a_4	a_5
Worker	0.13	0.11	0.07	0.05	0.04
Unemployed Eligible	0.34	0.18	0.12	0.07	0.04
Unemployed Ineligible	0.64	0.20	0.12	0.06	0.04

Note: This table shows the average quarterly MPCs of various type-groups, where columns represent agents with varying asset-to-income ratios and rows represent agents of differing employment statuses. Individual MPCs are calculated by computing the fraction consumed out of an unexpected \$500 transfer. Asset-to-income groups are $a_1 < p(10)$, $a_2 \in [p(10), p(25))$, $a_3 \in [p(25), p(50))$, $a_4 \in [p(50), p(75))$, and $a_5 \geq p(75)$, where percentiles are from the stationary asset-to-income distribution.

Table 2.4: Model fit of average MPCs

	Model	Data
Economy wide MPC	0.08	0.12 − 0.30
MPC difference of borrowing-constrained between 2008 and 2011	0.08	0.08

Note: This table shows the average quarterly economy wide MPC, and the average semiannual MPC of borrowing-constrained individuals between 2008 and 2011 implied by the model's simulations. Individual MPCs are calculated by computing the fraction consumed out of an unexpected \$500 transfer. These model-generated average values are then compared to available empirical estimates in the literature.

MPC by exploiting the unexpected changes in credit card borrowing limits of previously bankrupt individuals and find that the MPC is countercyclical over the Great Recession. In particular, they show that the average semiannual MPC difference of borrowing-constrained individuals between 2008 and 2011 is 8 percent. Using the Great Recession simulation of our model, we calculate the same moment and find that it is also 8 percent.[32] Hence, while the economy wide average MPC in our model is lower than its empirical counterpart, our model replicates the observed variation in the average MPC over the business cycle. This implies that our model successfully generates cyclical variation in the insurance value of

be considered as a lower bound.

[32]Section 2.5.1 explains the details on how we simulate the Great Recession using the model.

Figure 2.3: Model fit

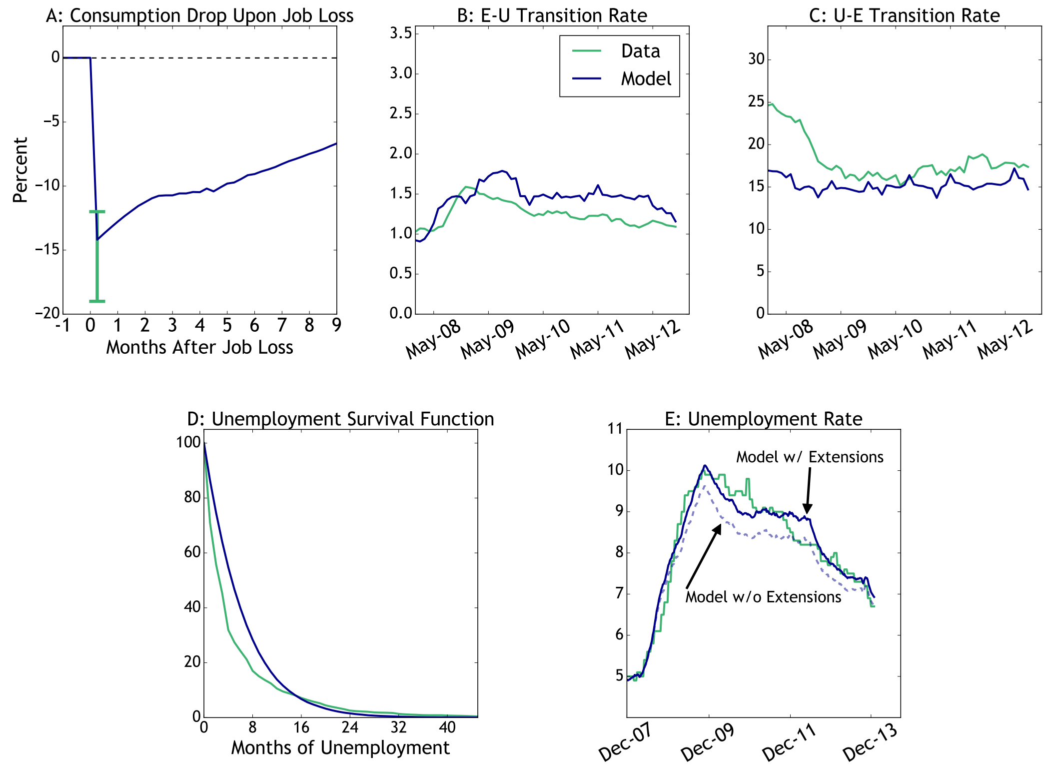

Note: This figure shows our model's implications for several important untargeted moments of the data. The main text discusses the details of this comparison.

UI, which is crucial when studying the optimal design of UI policy over the business cycle.

Average consumption drop upon job loss

First, we compare the model-implied value of the average drop in consumption upon experiencing a job loss to the available empirical estimates in the literature. To do so, we estimate the following distributed-lag regression using the simulation data:

$$\log(c_{it}) = \alpha_i + \gamma_t + \beta a_{it} + \sum_{k=-4}^{36} \delta_k D_{it}^{k,} + \epsilon_{it}, \tag{2.14}$$

where the outcome variable $\log(c_{it})$ is the logarithm of consumption of individual i in week t, α_i are coefficients on individual fixed effects, γ_t are coefficients on week fixed effects, a_{it} is the net asset level of individual i in week t, and the error ϵ_{it} represents random factors. The indicator variables D_{it}^k identify all individuals k weeks prior to or after a job loss, where $k = 0$ is the week of job loss. For instance, $D_{it}^4 = 1$ for individual i who experiences job loss at time $t - 4$, and zero otherwise.

Our treatment group consists of individuals who experience at least one job loss during the simulation period. Thus, the control group consists of individuals who never lost their jobs. Thus, $D_{it}^k = 0$ for all weeks t for individuals who belong to the control group.[33] Thus, the coefficients $\{\delta_k\}_{k\in\{-4,\ldots,36\}}$ measure the effect of job loss on consumption k weeks prior-to or after the incident relative to individuals who do not experience any job loss. Panel A of Figure 2.3 plots the estimated values for $\{\delta_k\}_{k\in\{-4,\ldots,36\}}$. It shows that in the week of job loss, consumption drops 14 percent on average and then slowly recovers over time.

Several papers in the literature estimated the average consumption drop upon job loss from various data sources. Gruber (1997) finds a decline in food expenditure of 6.8 percent using the Panel Study of Income Dynamics (PSID) for the period up to 1987. Saporta-Eksten (2014) uses cross-sectional variation in the PSID and estimates an 8 percent decline in consumption expenditure in the year during which a job loss occurs.[34] Stephens (2004) estimates the average decline in food expenditure upon job loss in the Health and Retirement Survey (HRS) and the PSID and finds that the decline is between 12 percent (PSID) and 15 percent (HRS) when an individual experiences a job loss between interviews. Browning and Crossley (2001) report a 14 percent decline using Canadian Out of Employment Panel (COEP) survey data. Chodorow-Reich and Karabarbounis (2016) conduct an analysis of the effects of job loss on consumption in both the PSID and the Consumer Expenditure Survey (CE) and find that the decline in total food expenditure is between 14 percent (PSID) and 21 percent (CE). Finally, Aguiar and Hurst (2005) report a 19 percent decline in food expenditure among the unemployed using scanner data.

[33]Notice that since the job loss event is exogenous in our model, simulated groups should not exhibit any selection bias.

[34]However, this estimate does not condition on the fraction of the year spent as unemployed. When we assume an average unemployment duration of 17 weeks, this would imply a decline in consumption of around 24 percent.

In summary, our model generates an estimate for the average decline in consumption upon job loss that is in line with available empirical estimates in the literature.[35]

Labor market transitions

We focus on the employment-to-unemployment (E-U) and unemployment-to-employment (U-E) transition rates implied by the model during the Great Recession and how they compare with the data. This way, we are able to evaluate the model's implications for the cyclical patterns of labor market transition rates. Since the timing of SIPP panels misses the rise in the E-U rate and the decline in the U-E rate during the first months of the Great Recession, the transition rates in Panel B and C of Figure 2.3 are taken from Current Population Survey (CPS) data as calculated by Kroft et al. (2016).[36] First, Panel B shows that the model is able to generate the initial rise in the E-U rate due to the rise in exogenous job separations in the model. It is also able to match the observed slow decline throughout the recovery, although the model-implied E-U rate decreases relatively earlier due to the recovery of aggregate labor productivity and the resulting decline in job separation shocks. Second, Panel C reveals that the model generates a smaller decline in job finding rates at the start of the Great Recession relative to the drastic decline in the data, but the levels of the model and the data become similar afterward. This is because in the model, when labor productivity decreases and firms do not post vacancies in submarkets offering high wages, the unemployed optimally direct their search effort toward submarkets offering lower wages where job-finding rates are relatively higher. As a result, the magnitude of the drop in the average job finding rate of the model during economic downturns is relatively smaller than its data counterpart. This, however, does

[35] Notice that the magnitude of the average consumption drop upon job loss in our model is largely controlled by the value of non-market activity h, which is calibrated to match the level of the opportunity cost of employment. Hence, the result that our model generates a similar magnitude of the average consumption drop upon job loss to the data lends support to our baseline calibration of the level of the value of non-market activity h.

[36] Kroft et al. (2016) report that CPS transition rates are not consistent with the stock levels of unemployment, employment, and non-participation. Then, they describe a procedure to adjust these rates so that the transition rates become consistent with observed changes in stocks between months. The data also account for seasonality by residualizing out month fixed effects and are smoothed by taking three-month moving averages.

not mean that the model underestimates the costs of recessions. While not as drastic as the Great Recession, the drop in the job finding rate is still sizeable and is accompanied by a significant decline in offered wages. Furthermore, even if the model generates a smaller drop in job finding rates in response to changes in aggregate productivity, it generates the observed elasticity of average unemployment duration with respect to changes in UI generosity, as this is one of the data moments in our calibration. This is also evident in Figure 2.8, where we show the impact of a countercyclical UI policy on the job finding rate.

Unemployment survival function

In the model, the likelihood of exiting from an unemployment spell depends on the aggregate labor productivity as well as the unemployed agent's choice of search intensity and wage submarket. A useful summary of how long individuals spend unemployed is given by the unemployment survival function, which shows the probability that an agent will remain unemployed beyond a given unemployment duration.

First, we use the SIPP 2008 panel to measure the survival function in the data. We restrict our sample to working-aged individuals age 24 to 65 who do not own a business or derive income from self-employment. As in Rothstein and Valletta (2017), we require at least one quarter of employment prior to the spell in order to focus on individuals who have sufficient attachment to the labor market. Spells that are left-truncated and spells with missing information for which we cannot ascertain the employment status of the respondents are dropped. Finally, we define spells to be uninterrupted months of unemployment and thus do not consider time spent out of the labor force, since we do not model the non-participation margin. Panel D of Figure 2.3 shows that the unemployment survival function generated by the model under the baseline calibration is close to its data counterpart. While survival data exhibit sharp drops during early months, the model survival function decays in a smooth fashion given the probabilistic nature of eligibility and job-finding rates in the model.

Impact of UI extensions on aggregate unemployment

In order to understand the model's predictions about the aggregate effect of benefit extensions on the labor market during the Great Recession, we simulate the model for the Great

Recession period with and without UI benefit extensions and measure the time path of the unemployment rate. Panel E of Figure 2.3 shows that during the depth of the recession, the model-implied unemployment rate would have been 0.6 percentage points lower in the absence of benefit extensions.

The body of work that studies the impact of UI on macroeconomic aggregates has found mixed results. Rothstein (2011) exploits variation in UI benefits across states with similar economic conditions, the behavior of UI ineligible as a control group, and several other strategies to address endogeneity problems in measuring the impact of UI on labor market conditions. Using CPS data, he finds that UI extensions raised the unemployment rate in early 2011 by only about 0.1 to 0.5 percentage points. Consistent with this finding, Chodorow-Reich and Karabarbounis (2017) implement a novel empirical strategy by using exogenous variation coming from measurement error in real-time state unemployment rates and find that benefit extensions increased the unemployment rate by at most 0.3 percentage points. Coglianese (2015) uses a similar strategy and also finds small effects. Meanwhile, Farber and Valletta (2015) use variation in individuals' time-to-exhaustion and find that extended benefits account for an increase of around 0.4 percentage points in the 9 percent unemployment rate in 2010. Valletta and Kuang (2010) find that in the absence of extended benefits, the unemployment rate would have been about 0.4 percentage points lower at the end of 2009, while Marinescu (2017) also finds small effects due to the reduced congestion resulting in a higher job-finding rate of any given job application.

On the other hand, Hagedorn et al. (2016) highlight that benefit extensions lead to higher equilibrium wages and thus lower vacancies. They also emphasize the role of firm expectations on future UI policies when making vacancy or hiring decisions. Accounting for this additional channel, they find that UI generosity increased the unemployment rate by 2.0 to 2.7 percentage points. This result is consistent with the findings of Johnston and Mas (2016), who find large effects of reductions in UI duration on unemployment. Fujita (2011b) also finds that extensions led to a substantial 1.2 percentage points increase in male workers' unemployment rate.

2.4 Welfare Calculation

We measure the welfare effects of any proposed UI policy by answering the following question: how much additional lifetime consumption must be endowed to all agents in an economy where some benchmark policy is being implemented so that average welfare will be equal to an economy where the proposed policy is implemented? In effect, we are evaluating whether an alternate UI policy will be welfare improving when compared to a benchmark policy, a natural choice being the actual UI policy implemented during the recession. Henceforth, we will refer to the UI policy implemented by the U.S. government during the Great Recession as the *benchmark policy*.[37]

Let b denote the benchmark policy and n denote the new/proposed policy. We can compute the additional percent lifetime consumption $\bar{\pi}$ that makes the average welfare equal across these two economies using the following equation:

$$\int_i \left[E_0 \sum_{t=0}^{\infty} \beta_{it} U \left(c_{it}^{b} \left(1+\bar{\pi}\right), s_{it}^{b} \right) \right] d\Gamma_{ss}\left(i\right) = \int_i \left[E_0 \sum_{t=0}^{\infty} \beta_{it} U \left(c_{it}^{n}, s_{it}^{n} \right) \right] d\Gamma_{ss}\left(i\right) \tag{2.15}$$

where c_{it}^{j} and s_{it}^{j} denote the consumption and search effort levels of agent i at time t under UI policy $j \in \{b, n\}$, and Γ_{ss} is the stationary distribution of the economy.

One can interpret the welfare exercise in Equation (2.15) as follows. Consider two countries populated by people with the same type-distribution. The only difference between both countries is that the government of the first country changes UI policy to policy b, while the second changes UI policy to policy n. The question is how much additional lifetime consumption $\bar{\pi}$ should the first government compensate an individual who is behind the veil of ignorance (i.e., does not know her initial type in the stationary distribution) in order to make her indifferent between being part of one of these two countries? Thus, the best UI policy n that the second government can implement is the one that makes the first

[37]During the Great Recession, the U.S. government increased the duration of UI payments to as much as 99 weeks in some states but kept replacement rates almost constant. We set a duration of UI payments that increases from 26 weeks (2 quarters) to up to 90 weeks (7 quarters) over the decline in aggregate productivity p, while the replacement rate of UI payments is kept fixed at its long-run average of 14 percent for all levels of p. This policy closely mimics the UI policy in the United States during the Great Recession and its recovery, assuming the United States as a single state.

government pay the highest compensation $\bar{\pi}_{max}$ to weakly attract this prospective citizen. This policy will be the optimal UI policy.

We restrict the class of candidate UI policies to be linear in current productivity level p such that $\phi(p) = q_\phi + m_\phi p$ and $e(p) = q_e + m_e p$. Under this restriction of UI policy instruments, we search over five UI policy parameters $(q_\phi,\, m_\phi,\, q_e,\, m_e,\, \tau)$ to solve for the optimal UI policy.

In order to obtain ex-ante welfare gains/losses $\bar{\pi}$ for each policy n, we begin from the stationary distribution of our calibrated economy Γ_{ss} where (1) aggregate labor productivity is constant at its mean level and (2) the unemployment benefit policy is fixed at a 14 percent replacement rate and 26 weeks expiration, which we call the *acyclical/flat* UI policy f. In each economy, an unanticipated but permanent policy change toward benchmark policy b and proposed policy n, respectively, is implemented. Given any guess of $\bar{\pi}$, we can now compute for both sides of Equation (2.15) by integrating over the stationary distribution. We then solve for the $\bar{\pi}$ that equates both sides of Equation (2.15) and select the UI policy that yields the highest welfare gain $\bar{\pi}_{max}$ as the optimal UI policy.[38]

2.5 Main Results

We find that $m_\phi = -6.44$, $q_\phi = 6.75$, $m_e = 0.34$, $q_e = -0.32$, and $\tau = 1.06$ percent, implying that the optimal UI policy should be countercyclical in both replacement rate and duration. These values imply that the optimal policy offers a 30 percent replacement rate for 4 quarters when aggregate labor productivity is at its mean value, and a 54 percent replacement rate for 10 quarters when aggregate labor productivity is depressed by 3.5 percent. This means that the optimal policy offers a more generous replacement rate for a longer duration compared to the U.S. government's UI policy during the Great Recession, which provided 14 percent of the replacement rate for around 7 quarters of payments for the same drop in labor productivity. Compared to this benchmark policy, the optimal policy increases welfare by 0.58 percent additional lifetime consumption for all agents. Meanwhile, compared to an acyclical policy that offers 14 percent of the replacement rate for 2 quarters for all levels of aggregate labor productivity, the optimal policy yields a

[38] Given the functional form of the utility function, there is no closed-form solution for $\bar{\pi}$.

welfare gain equivalent to 0.74 percent additional lifetime consumption.[39]

These welfare gains from the optimal UI policy are much larger when compared to welfare gains of eliminating the business cycle obtained by Lucas (1987), who finds that the welfare of an infinitely lived representative agent increases by only 0.008 percent in consumption equivalents for logarithmic preferences if cycles are removed. A more relevant comparison to our model is Krusell et al. (2009), who extend this analysis in an incomplete asset markets model with heterogeneous households and study the welfare effects of eliminating both aggregate risk and its impact on idiosyncratic risk when there is a correlation between these two shocks. They find that the welfare gains of eliminating the cycle and its effect on idiosyncratic risk are as much as 1 percent in consumption equivalents for the same period utility function. Importantly, they show that the effect of business cycles on idiosyncratic risk has great quantitative consequences. Specifically, if one does not correctly integrate out the effect of cycles on idiosyncratic risk, then the welfare gains of eliminating cycles are only slightly larger than those calculated by Lucas (1987). Similar to Krusell et al. (2009), our model features aggregate shocks and incomplete asset markets in which aggregate risk significantly affects the magnitude of idiosyncratic risk, as job finding and job separation rates are functions of aggregate labor productivity. Both models are also similar in that households are heterogeneous in terms of their employment status, discount rates, and wealth holdings. These similarities suggest that welfare results in their study are a useful benchmark against our model. In our model, optimal UI smooths aggregate shocks by introducing cyclicality into the generosity of benefits and also attenuates idiosyncratic unemployment risk by providing higher benefit levels on average. Nonetheless, as the optimal UI policy in our framework can only partially smooth the effect of cycles on consumption, welfare gains are much lower than the upper bound provided by Krusell et al. (2009).

The following discussions elucidate on the sources and distribution of welfare gains brought about by the countercyclical policy. First, we simulate the Great Recession using our model and compare how consumption patterns and labor market aggregates differ between the optimal policy and an acyclical policy. This will provide useful insight about the insurance

[39]To obtain this number, we repeat the welfare calculation procedure in Section 2.4 where we set the benchmark policy b as the acyclical/flat policy f. In this case, the first country continues to implement UI policy f, while the second country changes it to the optimal UI policy.

benefits and incentive costs of the countercyclical optimal policy, especially when a recession triggers more generous benefits. We then proceed to quantitatively decompose ex-ante welfare gains of the optimal policy attributable either to changes in consumption patterns resulting from altered saving and wage choices, or to changes in the search intensity exerted by the unemployed. Finally, we look at ex-post welfare outcomes among heterogeneous agents in order to understand how welfare gains are distributed across agents with different employment statuses and wealth holdings.

2.5.1 Great Recession exercise

We now use the Great Recession as an interesting test case that allows us to understand the insurance benefits and incentive costs associated with the countercyclical optimal policy. In order to discipline this exercise, we take as given the U.S. government's UI extension policy during the Great Recession and then pick the realizations of aggregate labor productivity to match the unemployment rate from December 2007 to December 2013 – the period that spans the beginning of the recession until the time when the Emergency Unemployment Compensation Act of 2008 (EUC08) was no longer renewed. Matching the realized unemployment rate by imposing that government policy mimics benefit extensions during the Great Recession is important, since using the model's aggregate labor productivity alone to match the unemployment rate disregards the fact that more generous UI policies implemented during the recession and recovery may have contributed to the heightened level of unemployment. Thus, in this exercise, the drop in labor productivity triggers lower job finding rates, higher separation rates, and longer benefit durations, all of which contribute to the rise in unemployment. Figure 2.4A shows the realizations of the labor productivity process that we obtain from this procedure, while Figure 2.4B compares the unemployment rate generated by the model to its counterpart in the data.[40]

In this exericse, we consider two economies that both experience the Great Recession between December 2007 and December 2013 but differ in the UI policy that is implemented.

[40]We acknowledge that labor productivity in the data during the Great Recession did not decline in a similar way. However, given that labor productivity in our model is the only source of aggregate fluctuations, we place more emphasis on matching the observed unemployment rate and less on the manner by which we do it. While we call the decline in p "labor productivity shock", it can stand in for other forms of shocks such as TFP, aggregate demand, or financial shocks.

Figure 2.4: Labor productivity and unemployment rate

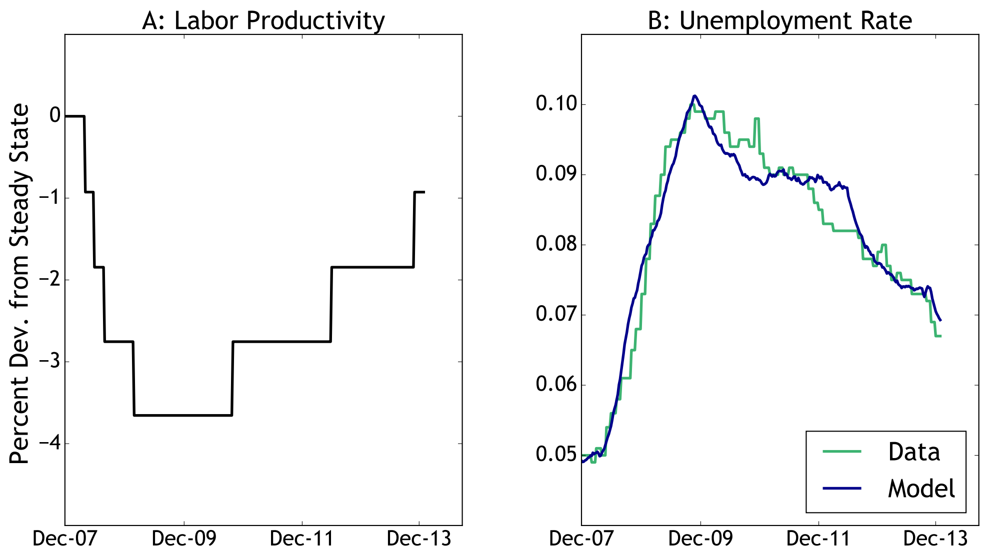

Note: Panel A shows the labor productivity series that we use in our model to generate the observed time series of the unemployment rate in the data during the Great Recession. Panel B compares the paths of the unemployment rate in the data and the model.

In both economies, the simulation begins under the stationary distribution.[41] At $t = 0$, we introduce a recession to both economies by feeding the labor productivity series into Figure 2.4A. It must be noted, however, that agents use the AR(1) process to take expectations on labor productivity. One economy introduces the optimal policy o, and the other maintains the less generous acyclical policy f. In both cases, the policy change is unanticipated by agents. This is a reasonable assumption, as UI extensions during deep recessions (such as the EUC08) are typically beyond the scope of pre-existing triggers that households are aware of. This policy change is permanent and will thus apply the same UI policy to future fluctuations of the same magnitude.

In the following sections, we separately analyze the consumption-smoothing benefits and

[41]We select the number of agents to simulate N to be large enough such that $\bar{\pi}$ does not change with further increases in N. We find that $N = 120,000$ is sufficient for this goal.

incentive costs of the optimal policy if it had been implemented during the Great Recession and compare them to that of the acyclical policy.[42] We place emphasis on how the cyclicality of these benefits and costs rationalizes a countercyclical optimal policy.

Insurance Benefits

Consumption Smoothing Upon Job Loss We first show the effect of the optimal UI policy on the consumption drop experienced upon job loss. We ask what would happen to the consumption profile of agents who experience a job loss in the economy that introduces the optimal policy and the economy that remains under the acyclical policy. The comparison of consumption profiles across these two economies will reveal the welfare benefits of the generous optimal policy coming from smoothing consumption between E-U transitions. Using model-generated data, we run the same distributed-lag regression in Equation (2.14) for each economy.

Figure 2.5 compares the consumption drop upon job loss between an acyclical policy and the generous optimal policy. On average, the consumption drop upon job loss is 15 percent under the acyclical policy, and 9 percent under the optimal policy, implying that the decline is 6 percentage points less under the optimal policy. This simply demonstrates the role of UI in dampening large fluctuations in consumption when job loss occurs, an insurance benefit on which the literature has traditionally focused. Moreover, this lower drop in consumption upon job loss is enjoyed by a larger number of agents in a recession due to the higher incidence of unemployment and longer spells during which wealth is depleted. As a result, the insurance value of UI payments in smoothing consumption upon job loss is larger in recessions.

Note that the reduction in the consumption drop is the net effect of two opposing forces: a more generous UI policy (1) directly increases consumption upon job loss due to higher benefits but also (2) indirectly crowds out precautionary savings. The first channel raises public insurance and thus decreases the consumption drop, while the second channel decreases self-insurance and thus increases the consumption drop as individuals enter unemployment

[42] Comparing the optimal and acyclical policies makes the illustration of the idea clear, as the acyclical policy offers the same replacement rate and duration across different realizations of the aggregate state. The intuition provided by the exercise also holds qualitatively when comparing the optimal policy with the benchmark policy.

Figure 2.5: Average consumption drop upon job loss

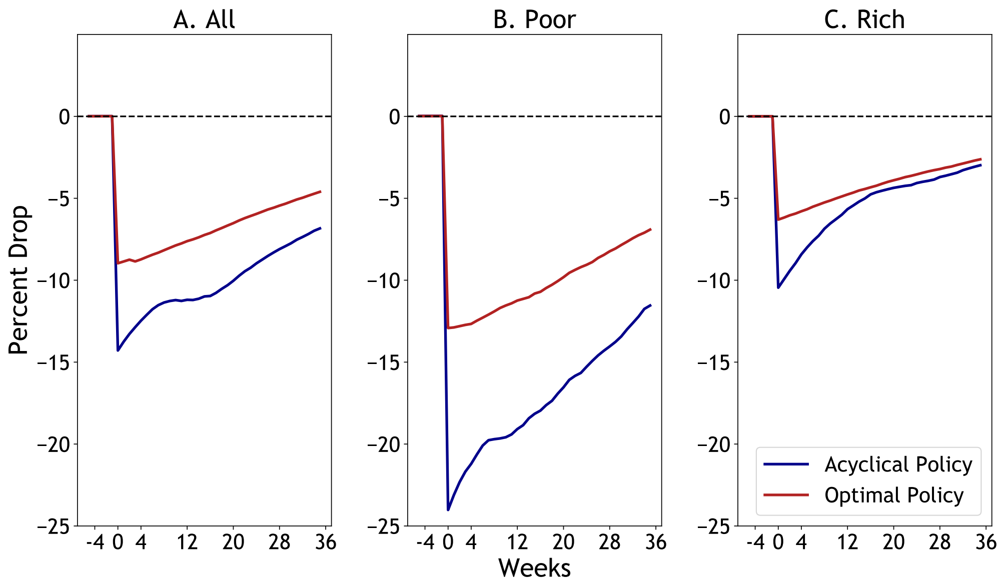

Note: Panel A plots the path of the average consumption drop upon job loss between 4 weeks prior to job loss and 36 weeks after the job loss. Two different consumption profiles are obtained from a distribution-lag regression in Equation (2.14) using model-generated data under the acyclical policy and the optimal policy. Panels B and C repeat this exercise for poor and rich agents experiencing job loss. "Poor" refers to agents who enter unemployment with an asset-to-income ratio below the 75th percentile of the stationary asset-to-income distribution, while "Rich" refers to those above the threshold.

with less wealth. In addition, notice that the recovery of consumption is slightly faster under the acyclical policy given how agents are forced to find jobs more quickly compared to an economy where the optimal policy is implemented.

It is also insightful to understand the effect of the optimal policy on the consumption drop upon job loss among rich and poor households. In Figure 2.5, we group individuals based on their asset-to-income ratio at the moment of job loss when the acyclical policy is implemented and then plot their consumption profiles. The first group consists of those who enter unemployment with an asset-to-income ratio below the 75th percentile of the stationary asset-to-income distribution, while the second group consists of those above that threshold. Using the same grouping (and the same job destruction shocks), we calculate the consumption drop that individuals would have experienced had the optimal policy been implemented instead. Panels B and C of Figure 2.5 demonstrate substantial heterogeneity in the consumption-smoothing benefits agents derive from the optimal policy. Among the poor, the consumption drop is reduced by around 12 percent, while for the rich, it is only around 4 percent. This result highlights the need to carefully calibrate the model's wealth distribution to match the data in order to correctly evaluate the true magnitude of any proposed policy's insurance benefits.

In summary, the optimal policy provides substantial insurance against E-U transitions, the magnitude of which varies significantly across the wealth distribution. More importantly, these benefits are larger during recessions simply because more agents experience job loss and remain unemployed for longer durations during which wealth declines.

Consumption Smoothing over the Business Cycle Although the insurance benefits of UI are traditionally seen to accrue mostly to job losers, we show in this section that in the presence of aggregate shocks and incomplete asset markets, UI also provides consumption-smoothing benefits even to those who do not lose their jobs. Under this framework, UI policy plays an important role in smoothing consumption over the business cycle. In order to demonstrate this channel, consider for the moment a worker in an economy that does not have a UI program. When a recession occurs, the worker anticipates that there is a higher risk of losing her job and that the unemployment spell is likely to be prolonged given the persistence of negative shocks. In the absence of government insurance, the worker self-insures by cutting back on consumption in order to build a buffer stock of savings that she could use to attenuate the impact of potential job loss. This means that consumption fluctuates heavily with aggregate fluctuations even if job loss does not actually occur. This reaction is simply a manifestation of the inefficiencies resulting from over-saving in an incomplete markets model, relative to its first best. The government then uses its UI program to reduce the excessive precautionary saving behavior of workers by promising higher public insurance during times when the unemployment risk is large in order to bring the economy closer to the efficient allocation. When UI is generous during recessions, individuals are relieved of the burden to reduce consumption in order to build savings, since UI makes the prospect of losing one's job less painful. This further contributes to the expansion of insurance benefits during recessions because it is precisely during this time when excessive precautionary saving behavior is triggered. While this channel is also present in previous models with incomplete markets, the literature on the optimal design of UI over the business cycle has not quantified the effect, possibly because of computational difficulties, which we are able to overcome.

Panel A of Figure 2.6 demonstrates this channel by comparing the average consumption of the economy during the Great Recession under the optimal and acyclical UI policies. It reveals that average consumption is much smoother under the optimal policy. The large drop in consumption at the onset of the recession when UI is acyclical is caused precisely by agents diverting consumption toward savings. This is corroborated by Panel A of Figure 2.7 which plots the average wealth of job losers during the first week of entering unemployment. At the start of the recession when labor productivity starts declining, it is clear that workers in the economy under the acyclical policy engage in precautionary savings due to the higher risk of losing a job and staying unemployed for longer durations. Thus, we see that average asset holdings upon entering unemployment rise during this period and only begin to decline during the recovery. In the case of the optimal policy, however, the need for precautionary saving is offset by the generous UI payments, implying that agents enter their unemployment spell with less self-insurance compared to their counterparts under the acyclical policy. The same idea is also apparent in Panel B of Figure 2.7, which plots the evolution of various percentiles of the asset distribution when a recession hits both economies. It shows that the level of precautionary savings under the generous optimal policy is markedly muted. Furthermore, similar to the consumption-smoothing benefits upon job loss, consumption smoothing through this channel is also cyclical. It is stronger during recessions precisely because it is during this time when precautionary saving motives are strong and thus significant cuts in consumption occur.

Next, we analyze the consumption-smoothing benefits of the optimal policy over the business cycle for agents with varying wealth levels. To do this, we again group agents based on their asset-to-income level at the start of the Great Recession. The first group consists of agents whose asset-to-income level at the start of this period is below the 75th percentile of the stationary asset-to-income distribution, while the second group comprises of those above this threshold. Panels B and C of Figure 2.6 then plot the average consumption paths of these two groups over the Great Recession. Comparing average consumption paths under the acyclical and optimal policies shows intuitively that the consumption-smoothing benefits of the optimal policy over the business cycle are largely different for poor and rich agents. While the optimal policy improves consumption smoothing for the poor, it does not for the rich, as they are already well insured.

Incentive Costs

While the optimal policy provides consumption-smoothing benefits to a large fraction of agents in the economy, it also induces certain moral hazard costs. This section discusses the magnitude of these costs associated with introducing the optimal UI policy vis-a-vis the acyclical UI policy. First, we look at how these costs manifest through lower job-finding probabilities and thus longer durations in unemployment. Second, we discuss how the magnitude of these moral hazard costs varies over the business cycle.

When a more generous UI policy is implemented, the unemployed eligible reduce their search effort and ask for higher wages because of an increase in the opportunity cost of employment. To provide a useful summary of the combined effects of both margins, in Figure 2.8, we look at how job finding rates and survival in unemployment change between the two economies. Panel A demonstrates that job finding rates during the recession shift downward when the optimal policy is introduced. Meanwhile, Panel B plots the Kaplan-Meier estimates of the unemployment survival function under both policies, as described in Section 2.3.3. The lower job finding rates result in the outward shift of the unemployment survival function under the optimal policy when compared to that of the less generous acyclical policy. This simply means that the likelihood that a duration will last beyond t months is always higher in the economy under the optimal policy. For instance, the probability that an unemployment spell will last beyond one month is around 40 percent under the acyclical policy, whereas it goes up to 60 percent under the optimal policy.

It is now evident that the optimal UI policy induces nontrivial costs through lower job-finding rates and thus longer unemployment durations. However, what is key to determining the optimal policy over the business cycle is the cyclicality of the size of these moral hazard costs, that is, how they expand and contract over the business cycle.

First, the value of job search is cyclical. A forgone unit of search during a recession is less costly than a forgone unit of search during a boom because jobs are difficult to find during a recession and conditional on finding a job, wages are likely to be lower as well. This means that while an extra dollar of benefits received during a recession induces the unemployed to search less, this reduction in search effort is not as costly compared to when the same dollar is received in an expansion during which firms are posting a lot more vacancies at higher wages. The cyclicality of the value of search effort is evident in Panel A of Figure 2.9 which shows that the consumption value of a unit of search effort is markedly lower during a recession compared to a boom. The same message is conveyed in Panel B, which shows that the average value of job search drops during the Great Recession and rises during the recovery for both eligible and ineligible unemployed, although the change is larger for the eligible unemployed, as they are the direct recipient of UI payments.

Second, wealth effects that discipline job search are more likely to manifest during recessions. For any given UI policy, recessions generally lead to prolonged unemployment spells during which agents draw down their assets to supplement consumption. Getting closer to their borrowing constraints, the unemployed have a higher incentive to ramp up their job finding efforts through a combination of higher search intensity and lower wage choices, as they seek to find work more quickly. This is evident in the household decision rules in Figure 2.1, which shows that for every unit of the decline in asset holdings at the time of unemployment, there is a disproportionate increase in search effort and decline in wage choices as the unemployed get closer to becoming borrowing constrained. Simply put, the presence of borrowing constraints acts like a self-disciplining device for job search efforts of the unemployed during recessions. As a result, the moral hazard costs are dampened by the fact that agents are more ill-prepared in terms of their own private savings during recessions.

In summary, while a generous UI policy decreases the job finding rate and increases the average unemployment spell duration, these moral hazard costs are partially offset in recessions because the consumption value of job search is low during recessions, and the decline in asset holdings in recessions incentivizes the unemployed to ramp up their job search. This result is consistent with Kroft and Notowidigdo (2016), who empirically find that the moral hazard cost of UI is procyclical.

2.5.2 Welfare decomposition

The Great Recession exercise in the previous section demonstrates the qualitative effects of the optimal policy on individual decision rules as well as the aggregate outcomes. We now proceed to quantitatively decompose the welfare contribution of the aforementioned changes. The ex-ante welfare gains of the optimal policy can be decomposed into either its effects on consumption coming from changes in savings and wage choices or its effects on the search intensity exerted by the unemployed. In order to isolate the welfare gains attributable to changes in consumption from those attributable to search effort, we shut down endogenous search decisions in the model.[43] This version of the model is then re-calibrated and used to evaluate the welfare gains coming from the countercyclical optimal policy. When policy has no effects on search intensity, welfare increases by 0.56 percent of additional lifetime consumption for all agents relative to the benchmark policy. Thus, the welfare gains of the optimal policy attributable to changes in search effort are negligible.[44] As a result, we conclude that the welfare gains come largely from changes in consumption patterns.

Having isolated the welfare gains attributable to search, we then want to understand how the optimal policy changes consumption patterns in the model without endogenous search effort. Our first step is to disentangle welfare gains along the transition from the long-run (steady state) gains. To do this, we make a slight but important modification in Equation (2.15). In particular, we change Γ_{ss} to Γ_b (where b denotes the benchmark policy) on the left-hand side, and Γ_{ss} to Γ_n on the right-hand side, where n is set to be the optimal UI policy. This implies that the first economy has implemented the benchmark policy, while the second economy has implemented the optimal policy for a very long time so that these two economies are in their respective steady states. We then ask an unborn agent who does not know her type within the respective stationary distributions which economy she prefers to live in. The ex-ante steady state welfare gains/losses from the optimal policy π_{ss} are then given by the percentage of additional lifetime consumption that the first government

[43] We do this by assuming that the unemployed searches for a job full-time, (i.e., $s = 1$), without incurring a disutility from search effort, (i.e., $\alpha = 0$).

[44] The result that changes in UI policy have small effects on the job search intensity of the unemployed is consistent with previous empirical evidence in the literature. For example, Ashenfelter et al. (2005) find that low job search effort is not a significant source of UI overpayments using evidence from randomized trials in four U.S. sites. Recently, Hagedorn et al. (2016) carefully analyze the effect of changes in UI policies on both the search intensity of unemployed workers (the micro effect), and the aggregate job finding rate per unit of search effort through vacancy posting decisions of the firms (the macro effect). They also find a small micro effect.

should compensate this agent in order to make her indifferent between being part of one of these two economies. We find that $\pi_{ss} = 0.18$, which is smaller than the welfare gain with a transition of 0.56. This result suggests sizeable welfare gains along the transition from the economy under the acyclical policy to the economy under the optimal policy. We know from our earlier analysis that the optimal policy reduces the precautionary saving motives, as agents substitute away from self-insurance to public insurance for consumption-smoothing purposes. As a result, agents decumulate savings and consume more of their labor income along the transition path. This increase in consumption is enough to overcome any rise in taxes brought about by the policy change, thus providing large welfare gains along the transition.

Next, we decompose the steady state welfare gains of the optimal policy. In particular, under a utilitarian equally weighted social welfare function as in Equation (2.15), the optimal policy can increase steady state welfare for three reasons: (1) an increase in the average consumption of the economy (the level effect), (2) a decline in the volatility of individual consumption paths (the volatility effect), and (3) a decline in inequality across individual consumption paths (the egalitarian effect). Following Floden (2001), the welfare gain from the optimal policy under the steady state comparison can be decomposed approximately into (1), (2), and (3):

$$\pi_{ss} = (1 + \pi_{lev})(1 + \pi_{vol})(1 + \pi_{egal}) - 1. \tag{2.16}$$

Comparing the average consumption level of economies under the optimal and benchmark UI policies, we find that average consumption is 0.18 percent *lower* in the steady state of the optimal policy, (i.e., $\pi_{lev} = -0.18$). This is because once the economy converges to a new steady state with lower wealth holdings and higher taxes, consumption levels decrease. On the other hand, we find that the optimal policy significantly reduces the volatility of average consumption and that there are sizeable welfare gains because of this channel. On average, we find that $\pi_{vol} = 0.35$, which implies that uncertainty gains overcome any reduction in long-run consumption levels. This is again due to the endogenous response of saving decisions to changes in UI policy over the business cycle. Recall from our Great Recession exercise in the previous section that the government uses its UI program to reduce the excessive precautionary saving behavior of workers by implementing a generous UI during times when unemployment risk is large in order to bring the economy closer to the efficient allocation. Therefore, the impact of fluctuations in aggregate labor productivity on the consumption path of individuals is lower under the optimal policy relative to that under the benchmark policy. This smoother consumption profile over the business cycle provides significant welfare gains. Finally, we find that $\pi_{egal} = 0.01$, implying that there are negligible welfare gains from the optimal policy due to equalizing the consumption paths across heterogeneous agents. However, this result masks the two opposing effects of the optimal policy on the inequality across individual consumption paths. On the one hand, generous UI payments to the unemployed and higher income tax rates create more equal consumption paths across heterogeneous agents and thus increase π_{egal}. On the other hand, the steady state asset distribution under the optimal policy is more unequal than its counterpart under the benchmark policy. This is because while most of the individuals in the economy under the optimal policy save less, the response of the agents in the top percentiles of the distribution is very small. As a result, the Gini coefficient of the asset distribution increases from 0.68 under the benchmark policy to 0.91 under the optimal policy. This rise in the inequality of the steady-state wealth distribution in fact reduces π_{egal}, as it makes individual consumption paths across heterogeneous agents more unequal. We find that these two opposing effects quantitatively cancel each other out, and thus on average π_{egal} is small.[45]

2.5.3 Heterogeneous welfare effects

While the previous section decomposes the average ex-ante welfare gains into various mechanisms at work in our model, it is also insightful in understanding which type of agents stand to gain or lose from the optimal policy compared to the benchmark policy. In order to measure the ex-post heterogeneous welfare gains/losses from the optimal UI policy, we group agents by their employment status and asset level based on the stationary distribution. We then calculate $\bar{\pi}$ from Equation (2.15) for each group by only integrating over agents that belong to each group.

[45]The welfare decomposition exercise presented here can be modified to incorporate the effects of transition on π_{lev}, π_{unc}, and π_{egal}. The reason why we decompose the welfare gains across two different steady states is to isolate the long-run effects of the optimal policy as the policy change is permanent. However, we also did this exercise with transition and find that the level gains in consumption from the optimal policy are large because of the decline in savings along the transition.

Table 2.5: Heterogeneous welfare impacts of optimal policy

	Asset Groups				
Employment	a_1	a_2	a_3	a_4	a_5
Worker	0.73	0.67	0.58	0.53	0.52
Unemployed Eligible	1.89	1.55	1.28	0.96	0.84
Unemployed Ineligible	0.61	0.58	0.55	0.50	0.51

Note: This table shows the heterogeneous welfare gains from the optimal policy on various type-groups, where columns represent agents holding various levels of assets and rows represent agents of differing employment statuses. Welfare numbers are in percent lifetime equivalent consumption terms. Asset groups are $a_1 < p(10)$, $a_2 \in [p(10), p(25))$, $a_3 \in [p(25), p(50))$, $a_4 \in [p(50), p(75))$, and $a_5 \geq p(75)$, where percentiles are from the stationary asset distribution. Gains are calculated relative to the benchmark policy.

Table 2.5 shows the heterogeneous welfare impacts of the optimal policy on various type-groups, where columns represent agents holding various levels of assets (set to be the different ranges in the asset distribution) and rows represent agents of differing employment statuses.

It is clear that the unemployed eligible stand to gain the most from the optimal policy. This result is unsurprising, since the unemployed eligible are the direct beneficiaries of more generous payments and durations, and thus enjoy the largest consumption-smoothing gains. Intuitively, among the unemployed eligible, poorer individuals also enjoy larger welfare gains compared to their richer counterparts, given how each additional dollar of benefit payment is more valuable to them.

Consistent with our earlier discussion, the unemployment eligible are not the sole beneficiaries of the optimal UI policy. Workers also enjoy a sizeable welfare gain, albeit to a smaller degree, because of two opposing effects. On the one hand, workers maintain smoother consumption over the business cycle given the weaker need to engage in precautionary savings afforded to them by optimal UI benefits; on the other hand, they are the primary financers of the optimal UI policy and would thus face higher taxes and lower consumption levels. Nonetheless, the ability to maintain smoother consumption during economic fluctuations dominates the financing effect. Note that if we had not accounted for this benefit, then we would expect workers to experience welfare losses, as they would be paying taxes without enjoying the benefit of being able to smooth consumption over fluctuations in aggregate labor productivity. Unsurprisingly, welfare gains are also much larger among poor workers for whom savings (and the corresponding forgone consumption) is most costly.

Meanwhile, the unemployed ineligible only receive the generous UI payments in the event that they find a job, lose it, and become eligible, which is a small probability. While they do not contribute to financing the optimal UI policy, they incur costs because of lower job-finding rates resulting from the generous UI payments. Having to spend longer weeks without benefits and being forced to exert more effort in finding a job results in this group experiencing the lowest gains from the optimal policy.

2.6 Robustness

2.6.1 Welfare under different specifications

In this section, we compute the welfare gains or losses from the optimal policy relative to the benchmark policy under different specifications of the baseline model. In these exercises, whenever a change in parametrization is necessary, the model is recalibrated to match the moments found in Section 2.3 and tax rates are adjusted under each UI policy so that the government budget constraint holds in equilibrium. The nature of the first three exercises in this section requires us to simulate a recession in order to compute the welfare gains. To preserve consistency within this section, we report the welfare gains of the remaining specifications under a scenario in which a recession occurs initially as well. The results are summarized in Table 2.6.

First, in order to quantify how welfare gains change depending on the timing of the policy change, we evaluate the welfare gains from the optimal policy when the policy change is introduced at the onset of a recession. This exercise follows the Great Recession simulation discussed in Section 2.5.1 where an unanticipated UI policy change is implemented. The only difference here is that for the first economy, the benchmark policy b is introduced at $t = 0$, whereas in the alternate economy, the optimal UI policy is implemented.

We modify the welfare criterion in Section 2.4 slightly, as we now require a simulation-based welfare calculation. Additional details regarding the computational procedure are provided in Appendix B.3. We compute for $\bar{\pi}$ in Equation (2.15) modified to account for the recession that occurs right at the same time the policy change is made and find that the optimal policy increases ex-ante welfare by 1.25 percent additional lifetime consumption relative to the benchmark policy. The welfare gains of the optimal policy are unsurprisingly much higher when the policy is implemented right before a sharp drop in aggregate labor productivity, since there is a frontloading of gains coming from large net insurance benefits provided during the recession. At the onset of a recession, stronger precautionary motives cause larger drops in consumption, and a larger pool of unemployed individuals experiences the consumption drop upon job loss. This is in contrast to welfare gains of 0.58 when we do not take a stance on the realizations of aggregate productivity.[46]

The second exercise we perform considers how welfare gains change if the policy were temporary. While we study permanent changes in the UI benefit schedule, our framework is also useful to assess the welfare effects of discretionary fiscal policies such as the one implemented during the Great Recession. We now assume that the optimal policy is only implemented during the period of the Great Recession, and it unexpectedly reverts back to the acyclical policy f at the end of this period. This is to closely pattern the simulation of the model to the events that occurred during the Great Recession where the EUC08 was completely terminated in December 2013 and UI policy returned to what it had been prerecession. We find that the welfare gains from the optimal policy become 0.83 percent additional lifetime consumption relative to the benchmark policy. The difference between this value and welfare gains when the policy change is permanent (1.25 percent) reveals that around 35 percent of the total welfare gains are attributable to the expectation of generous UI payments during future economic downturns.

Third, we test the quantitative effects of assuming a time-invariant interest rate r on the welfare gains from the optimal policy. In this exericse, we consider an interest rate that varies with the state of the economy such that it is procyclical and closely mimics its data counterpart during the Great Recession.[47] Under this exercise, we find that the optimal

[46]Furthermore, when the government implements the optimal policy right before a boom, we find that it increases ex-ante welfare by 0.23 percent additional lifetime consumption relative to the benchmark policy. Given that the optimal policy raises welfare more when implemented right before a recessiom compared to that of a boom provides strong evidence about the countercyclicality of insurance benefits net of incentive costs.

[47]The weekly real interest rate reduces from its baseline value of 0.00095 to -0.0003 at the depth of the Great Recession. This way, we are able to measure the quantitative effects of significant changes in the real

Table 2.6: Welfare gains under different specifications

Exercise	Welfare gains (%) from the optimal policy
Great Recession simulation	1.25
Temporary policy change	0.83
Procyclical interest rates	0.64
Endogenous quit decisions	1.10
Replacement rate $\phi = 0.4$	0.77
UI eligibility requirements	0.94
Permanent discount factor	1.24

Note: This table shows welfare gains from the optimal policy under different specifications of the baseline model. Welfare numbers are in percent lifetime equivalent consumption terms. Gains are calculated relative to the benchmark policy under a labor productivity series that generates the observed unemployment rate time path during the Great Recession.

policy yields a welfare gain equivalent to 0.64 percent additional lifetime consumption relative to the benchmark policy. The reason for the reduction in welfare gains from 1.25 percent to 0.64 percent is that the decline in the real interest rate reduces precautionary saving motives during recessions, making agents' consumption profiles relatively smoother over the business cycle even under a less generous benchmark policy. This reduces the welfare gains from the optimal policy. While the welfare gains under a recession are reduced to half their original value, the countercyclical optimal policy still provides substantial gains. Moreover, given that interest rate fluctuates drastically with the state of the economy in this exercise, this result places an upper bound on the likely effects of endogenizing interest rates.

Fourth, we address the feature of the baseline model where a matched worker receives the same wage throughout her tenure within a firm. These fixed-wage contracts introduce "job lock" since an unemployed individual who is desperate for work may land a low-paying job during a recession but be unable to switch to a higher-paying job unless the match exogenously dissolves. This feature of the model may be a source inefficiency that the optimal policy is trying to correct, since generous benefits during recessions can nudge agents toward looking for higher-paying jobs. Hence, generous benefits during recessions not only may be providing consumption insurance but also may serve as a means of convincing the unemployed to look for jobs that will be paying higher even after the recession ends. In order to understand whether the optimal UI policy is also correcting inefficiencies introduced by the fixed-wage contract assumption of the baseline model, we solve for the welfare gains of the optimal policy in an extended model that allows for endogenous quits.[48] In this extended model, workers can choose to quit their jobs in order to begin searching for a new job. Under this setup, the artificial job lock problem is eliminated, as workers who place a higher value on the option of becoming unemployed and looking for a higher-paying job can actually leave their employer. The model details and a modified computational algorithm can be found in Appendices B.4 and B.5, respectively. The welfare gains under the model with endogenous quits is given by 1.10 percent when the optimal policy is implemented at the onset of the Great Recession. Introducing endogenous quit decisions into the model has

interest rate.

[48]Without the fixed-wage contract assumption, solving for the optimal policy will be computationally burdensome, as firms would now need to keep track of household decisions.

a small effect because the option of quitting is not widely used by workers, given that the value of becoming unemployed ineligible is very low. As a result, inefficiencies created by fixed wage contracts in the baseline model have a small quantitative impact on the welfare gains from the optimal policy.

The fifth robustness exercise considers the calibration of the replacement rate of the benchmark UI policy. Recall that our benchmark replacement rate of 14 percent takes into account the effect of partial take-up among those eligible for benefits and adjusts for differences between wages and total compensation. To understand the effects of this adjustment, we calculate the welfare gains from the optimal policy when the benchmark policy replacement rate is set to 40 percent, (i.e., $\phi(p) = 0.4 \quad \forall p$), the (unadjusted) value calculated by the Department of Labor. The goal of this exercise is to understand whether the countercyclical optimal policy would still be welfare improving when compared to a new benchmark policy that has a significantly higher but time-invariant level replacement rate. We find that the optimal policy increases welfare by 0.77 percent relative to the new benchmark. This result implies that there are still sizeable welfare gains when the government transfers funds from booms to recessions, as the insurance value of UI payments expands and incentive costs contract during recessions. This also emphasizes that welfare gains are not mostly driven by more generous benefits levels but by the introduction of cyclical generosity.

The sixth exercise considers eligibility rules for workers at the moment of job loss. According to the UI program in the United States, workers have to satisfy some monetary and nonmonetary requirements to be eligible for UI benefits.[49] Under these requirements, on average, around 75 percent of the workers are in fact eligible for UI benefits upon job loss.[50] When studying the optimal design of UI program, it will be interesting to consider the welfare implications of treating these eligibility requirements as another policy instrument. In our baseline setup, eligibility requirements upon job loss are controlled by the UI expiration rate e. Thus, an extension of UI duration also implies a relaxation of UI eligibility requirements for workers in our baseline model. In order to understand the effects

[49] For example, in terms of monetary requirements, workers must receive enough wages during the base period to establish a claim. In terms of nonmonetary requirements, the reason for the workers' job loss must be through no fault of their own, and they must be actively looking for work while unemployed.

[50] See Chodorow-Reich and Karabarbounis (2016).

of this relationship, we change the problem of the worker in Equation (2.1) such that the probability of being eligible upon job loss is fixed at 75 percent rather than controlled by changes in e. We then evaluate the welfare gains from the optimal policy, and find that it yields 0.94 percent additional lifetime consumption relative to the benchmark policy. Since a lower fraction of workers are now eligible for UI benefits upon job loss relative to the baseline model, the welfare gains from the optimal policy are slightly reduced under this exercise.

Finally, we explore the implications of time-varying discount factors β_t. The stochasticity of discount factors introduces another idiosyncratic shock to households, and so one might be concerned about the presence of an unintended role of UI payments as providing insurance against the discount factor risk. In order to quantify this effect, we set discount factors to be permanent and use an equally weighted social welfare function in computing the welfare gains. We find that the optimal policy yields 1.24 percent additional lifetime consumption relative to the benchmark policy, implying that the effect of time-varying discount factors on welfare gains of the optimal policy is negligible. This result is expected given that discount factors are calibrated to be highly persistent in our baseline calibration.

2.6.2 High level of opportunity cost of employment

We now explore the features of the optimal policy under a high level of opportunity cost of employment calibration. In particular, we recalibrate our baseline economy so that the model matches the same labor market and asset-to-income distribution moments as in our baseline calibration, but the level of opportunity cost of employment is set to be 0.955, as calibrated by Hagedorn and Manovskii (2008). Next, we evaluate the welfare gains/losses of the same set of linear policies and obtain the optimal policy for this case under the welfare criterion in Section 2.4.

We find that the optimal policy is still countercyclical even under a high level of opportunity cost of employment. Specifically, it features a 19 percent replacement rate for one quarter when aggregate labor productivity is at its mean value, and a 59 percent replacement rate for 4 quarters when aggregate labor productivity is depressed by 3.5 percent. Compared to the U.S. government's UI policy during the Great Recession (the benchmark policy), this optimal policy increases welfare by 0.25 percent additional lifetime consumption for all agents. Relative to the optimal policy under the baseline calibration of opportunity cost of employment, the optimal policy in this case offers a lower replacement rate for a much shorter duration when labor productivity is at its mean, while the cyclicality of the optimal policy remains roughly the same. This result is intuitive because when the value of unemployment is close to the value of employment because of a high opportunity cost of employment, the consumption drop upon job loss becomes less pronounced. Thus, the government implements a low replacement rate for short durations under the mean level of aggregate labor productivity. Moreover, consumption still fluctuates because of changes in the saving behavior of agents as a response to fluctuations in aggregate labor productivity. Hence, the government still finds it optimal to transfer funds from expansions to recessions. However, the magnitude of these fluctuations in consumption is relatively smaller, as the precautionary saving motives are not as strong under a high level of opportunity cost of employment. For this reason, the welfare gains from the optimal policy in this case are less than half of the welfare gains provided by the optimal policy under the baseline calibration of opportunity cost of employment.

This exercise is also useful to compare our result to the findings of the previous literature. As we discussed in Section 2.1, Mitman and Rabinovich (2015) also study the optimal cyclicality of UI replacement rate and duration in an equilibrium search model in which agents are not allowed to save/borrow. In their baseline calibration, the summation of UI benefits b and the value of nonmarket activity h is equal to 0.984, implying that the flow opportunity cost of employment is high. In this setup, they find that the optimal UI policy is procyclical. Then, in Section 5.4 of their paper, they discuss the implications of relaxing the no saving/borrowing assumption on their results. In this discussion, they also acknowledge that when agents are allowed to save/borrow, fluctuations in agents' wealth holdings over the business cycle may create a force that has a potential to reverse the cyclicality of their optimal UI policy. In our model, we allow agents to save/borrow through incomplete asset markets and indeed show that this channel is strong enough to rationalize the countercyclicality of the optimal policy even under a high level of opportunity cost of employment.

2.7 Evidence on the Mechanism: A First Pass

In this section, we empirically test the interaction between UI generosity and savings decisions in order to check whether our main mechanism is also observed in the micro data. This exercise builds on Engen and Gruber (2001), who find that UI benefits tend to crowd out individual savings.[51] We focus on the Great Recession period to understand the impact of drastic changes in UI policy on the saving decisions of individuals. Using the SIPP panel 2008 core data, we obtain household employment, labor income, and state of residence information. Wealth data are once again obtained from the topical data of the same panel, which is typically released on a yearly basis as opposed to the monthly frequency of the core data. State-level UI duration data during the Great Recession consist of maximum potential duration by adding up standard weeks, Extended Benefits (EB), and EUC tiers 1-4 (when applicable).[52] Meanwhile, the state-level replacement rate is defined as either (1) the weighted average of the ratio of the weekly benefit amount and average claimants' wage or (2) the ratio of the weighted average of the weekly benefit amount and the weighted average of claimants' wage.[53] To obtain the expected benefit receipt of a worker, we compute the average weekly wage of the respondent for one quarter prior to the wealth observation and multiply it by the replacement rate offered by her state of residence during that time.

Our sample includes workers ages 24 to 65 who report not owning any business in part or in full and has worked for at least one quarter prior to the first observation and are always working in between observations. This more or less guarantees eligibility for UI if the observed worker is displaced in the future. Moreover, focusing only on employed individuals between observations eliminates other reasons for changes in asset holdings, such as experiencing unemployment. We organize the data into person-time information (where $t = \{2009, 2010\}$) and run the following regression:

$$a_{it} = \gamma_{ben} benefit_{it} + \gamma_{dur} dur_{st} + \beta X_{it} + \alpha_i + \alpha_s + \alpha_t + \epsilon_{ist}$$

[51]While Engen and Gruber (2001) study the effect of the UI replacement rate on saving decisions, we also include time- and state-varying UI duration in order to account for the effect of expected length of UI receipts on wealth for the period of the Great Recession.

[52]We thank A. Yusuf Mercan for kindly sharing this dataset with us.

[53]See https://oui.doleta.gov/unemploy/ui_replacement_rates.asp for more details.

where a_{it} is the asset-to-income ratio of individual i at time t, $benefit_{it}$ is the expected weekly benefit receipt of individual i at time t, dur_{st} is the maximum potential duration of UI in state s during time time t, X_{it} is a set of controls which include education, martial status, and age, and $\alpha_{j\in\{i,s,t\}}$ are individual, time, and state fixed effects.[54] The coefficients of interest are the impact of the unemployment benefit level and duration on the asset-to-income ratio given by γ_{ben} and γ_{dur}. Note that a selection problem arises if there is a systematic movement of a certain type of worker to states with high levels of UI generosity. In order to control for this, we also expand the original regression to control for individuals moving from one state to another.

Given that isolating the causal effect of benefit generosity on self-insurance is beyond the scope of this exercise because of endogeneity, our intention is simply to provide correlational evidence on this relationship. Table 2.7 shows that expected benefit receipt has a negative and statistically significant impact on self-insurance. While the length of UI duration has a negative coefficient, it is not statistically significant. For example, $\gamma_{ben} = -.0135$ implies that a \$100 increase in the expected benefit amount received each week should unemployment occur results in a decrease in the asset-to-income ratio that is equivalent to 1.35 weeks' worth of insurance. Alternatively, this would also imply a reduction in savings by around \$1124 for a worker earning the median weekly wage of \$833. This relationship is consistent with the crowding-out effect of UI on precautionary savings documented by Engen and Gruber (2001). This result lends evidentiary support to the idea that the insurance benefits of a generous UI policy during recessions are partially attributable to the relief UI benefits provide workers who no longer need to experience sudden drops in consumption in order to build a buffer stock of savings. Results in the second and fourth columns also indicate that the issue of selection caused by state-to-state moves is not consequential. Finally, comparing the first two columns with the last two reveals that the relationship is robust to the manner by which replacement rates are calculated.

Motivated by the above empirical evidence, we revisit our welfare analysis in the model to understand if the replacement rate is a more important instrument in providing welfare gains relative to the UI duration. We find that a UI policy that consists of an optimal

[54] Notice that $benefit_{it}$ is affected by the replacement rate offered by the state s that individual i resides in during time t.

Table 2.7: Regression results

	Benefit Calculation 1		Benefit Calculation 2	
	(1)	(2)	(1)	(2)
benefit	−.0135***	−.0135***	−.0153***	−.0153***
	(.0010)	(.0010)	(.0012)	(.0012)
dur	−.0122	−.0123	−.0119	−.0119
	(.0164)	(.0164)	(.0164)	(.0164)
moving		−.0924		−.0933
		(.2255)		(.2254)
individual fixed effects	Yes	Yes	Yes	Yes
state fixed effects	Yes	Yes	Yes	Yes
time fixed effects	Yes	Yes	Yes	Yes
observations	33,012	33,012	33,012	33,012

Note: The dependent variable is the asset-to-income ratio of individuals. “Benefit Calculation 1” uses a replacement ratio calculated as the weighted average of the following ratio: weekly benefit amount (WBA) / weekly wage. “Benefit Calculation 2” uses a replacement rate ratio calculated as the ratio of the weighted average of WBA and the weighted average of the weekly wage. *, **, *** denote statistical significance at the 10%, 5%, and 1% level, respectively.

replacement rate but a UI duration of the benchmark policy, together with the tax rate that balances the government's budget constraint for this hybrid policy, yields an average welfare gain that is equivalent to 0.46 percent additional lifetime consumption relative to the benchmark policy. This implies that around 80 percent of the welfare gains from the optimal policy are attributable to the optimality of the UI replacement rate, and the remaining 20 percent of the gains come from the optimality of UI duration. This is consistent with the above empirical result that the changes in replacement rates significantly affect the self-insurance decisions of individuals.

2.8 Conclusion

We study optimal UI over the business cycle using a tractable heterogeneous agent job search model that features labor productivity driven business cycles and incomplete asset markets. We find that the optimal UI policy is countercyclical. In particular, when aggregate labor productivity is at its mean, it features a 30 percent replacement rate for 4 quarters, but when aggregate labor productivity is depressed by 3.5 percent, it offers more generous benefits of a 54 percent replacement rate for a duration of 10 quarters financed by higher labor income taxes. Compared to a UI policy that mimics the policy implemented during the Great Recession by the United States government, the optimal policy represents an average welfare increase of 0.58 percent additional lifetime consumption. We show that incorporating the response of individual saving behavior to changes in UI policy is quantitatively important in measuring the welfare benefits and costs of UI policy.

Insurance benefits are larger in recessions relative to expansions, while incentive costs exhibit the opposite pattern. Insurance benefits expand during recessions because (1) consumption insurance upon job loss is provided for a larger pool of unemployed and long jobless spells, and (2) it attenuates the need to engage in precautionary savings by cutting back on consumption at the onset of a recession. Meanwhile, incentive costs are also relatively smaller in recessions because (1) jobs are difficult to find and forgone search is not as worthwhile, and (2) borrowing constraints impose discipline on individual job search behavior because of a wealth effect. As a result, the optimal policy is countercyclical.

A quantitative decomposition of ex-ante welfare gains reveals that in the long run, the optimal policy provides a substantial reduction in consumption uncertainty at the cost of lower consumption levels. Along the transition, however, large consumption level gains are enjoyed by agents as they decumulate savings in response to more generous public insurance during recessions. Meanwhile, gains from reduced inequality and lower search effort are present but limited. In addition, ex-post welfare gains are shown to be heterogeneous across different types of agents. The unemployed eligible gain the most, but the employed remarkably enjoy large gains as well because of the reduced precautionary motives during recessions. Unsurprisingly, gains are largest for the poor across all employment types.

Our contribution to the existing literature lies in carefully accounting for the welfare effects of endogenous interaction between savings and UI policy over the business cycle. The natural extension of our analysis is to analyze how other sources of private insurance (such as family labor supply) react to changes in UI policy and how this interaction would affect the optimal policy. Another avenue for future research is to incorporate capital accumulation in order to account for the effect of government programs on aggregate capital stock. However, given the complexity of our current model, we leave these extensions to future work.

Figure 2.6: Average consumption over the business cycle

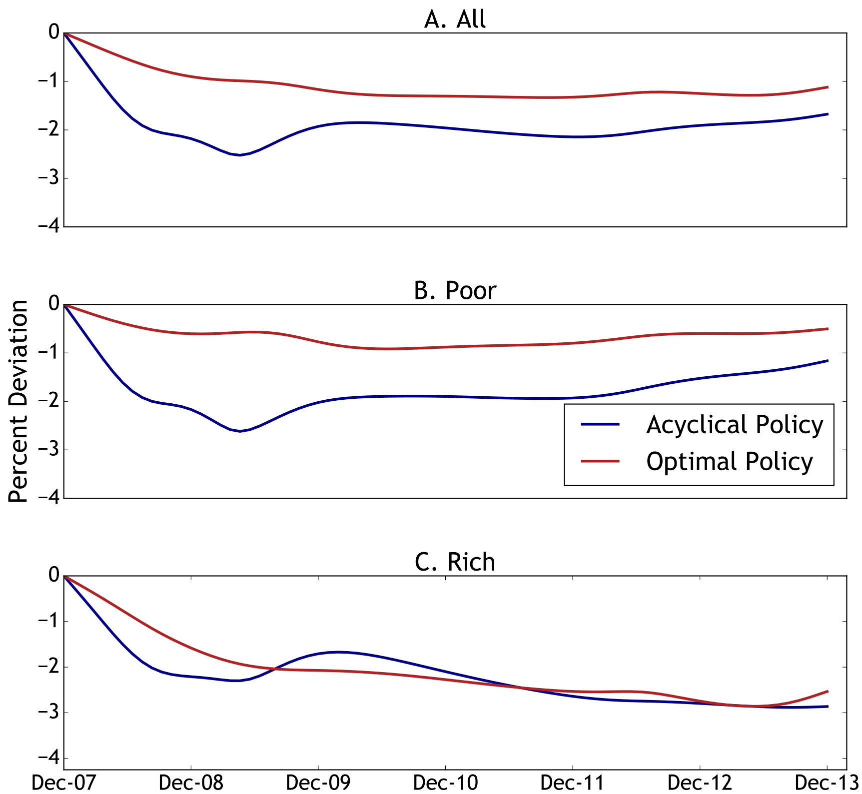

Note: Panel A plots the percent deviation of average consumption's trend during the Great Recession from its steady-state level at the start of this period under the acyclical and optimal UI policies. Panels B and C repeat this exercise for poor and rich agents. "Poor" refers to agents who enter unemployment with an asset-to-income ratio below the 75th percentile of the stationary asset-to-income distribution, while "Rich" refers to those above the threshold.

Figure 2.7: Average assets upon job loss over the business cycle

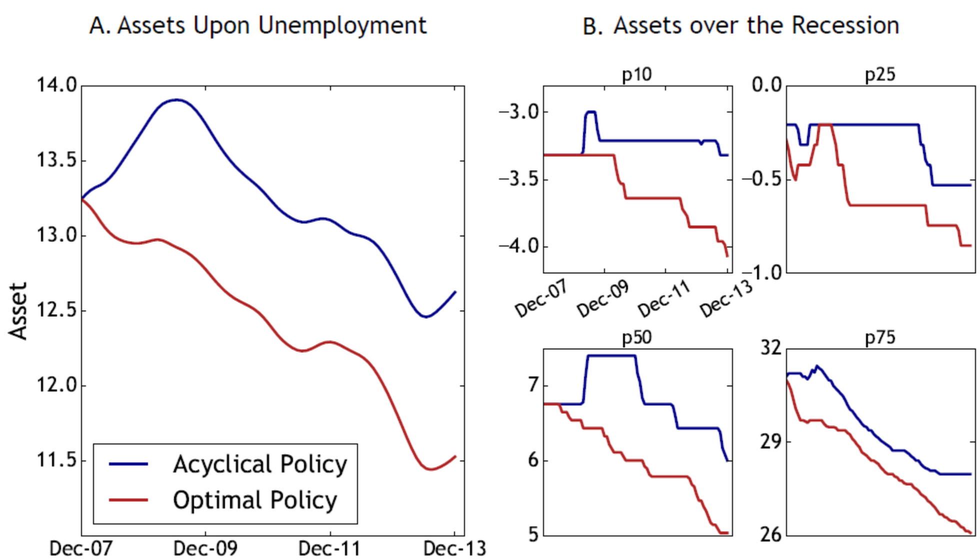

Note: Panel A plots the trend of average asset holdings of the unemployed during the first week of entering unemployment over the Great Recession under the acyclical and optimal UI policies. Panel B shows the evolution of various percentiles of the asset distribution over the Great Recession under these two policies.

Figure 2.8: Incentive costs of the optimal policy

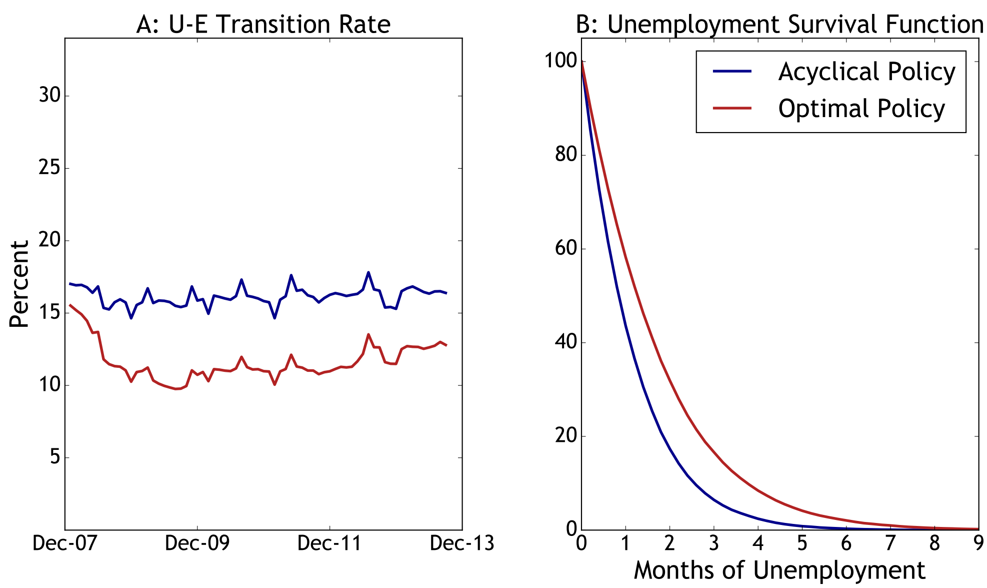

Note: Panel A shows the average job finding rates during the Great Recession under the acyclical and optimal UI policies. Panel B plots the unemployment survival function under these two policies, which shows the probability that an unemployment spell will last beyond t months.

Figure 2.9: Value of job search

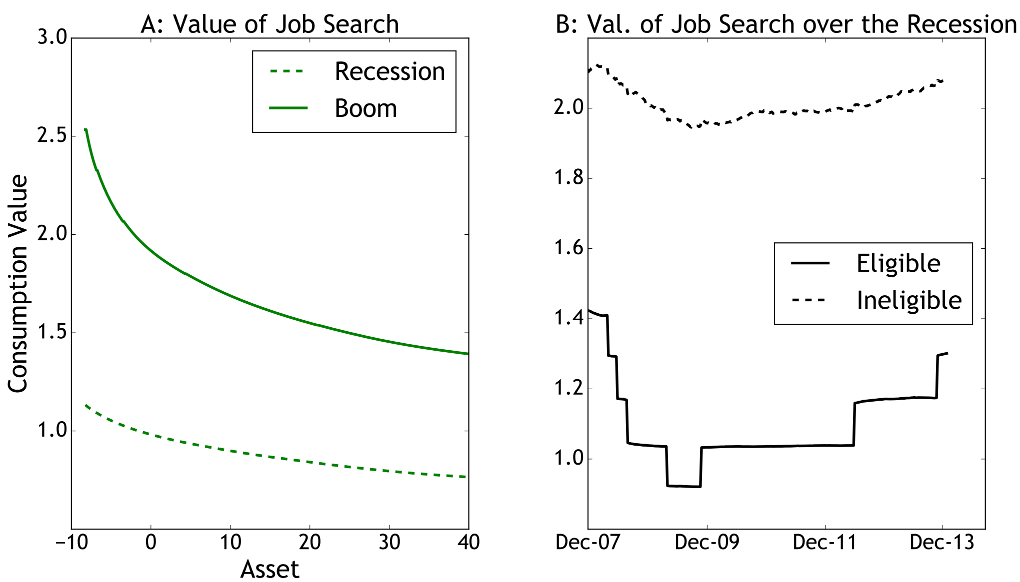

Note: Panel A shows the consumption value of job search across different levels of asset holdings in recessions and booms. Panel B plots the average consumption value of job search during the Great Recession for unemployed eligible and ineligible.

Chapter 3

What Do Survey Data Tell Us about U.S. Businesses?

3.1 Introduction

Representative surveys of households and firms have become an important source of data on business owners and their activities, and are now used extensively in studies of wealth inequality and entrepreneurial choice. This paper examines the reliability of these data for research on U.S. businesses, including pass-through entities and subchapter C corporations.[1] Pass-through businesses account for roughly half of business net income in the United States and have been a focus of recent tax reforms and debates about income inequality.[2] Subchapter C corporations account for the remaining half and include all publicly traded firms. We document issues arising from nonrepresentative samples and measurement errors in survey data and discuss the implications of the errors for economic research.

[1] For tax purposes, pass-through entities classify themselves as sole proprietorships, S corporations, or partnerships. They are called "pass-through" because the income earned by such businesses is taxed under the owners' individual income tax. In contrast, C corporations pay corporate taxes on income earned, and individual shareholders pay dividend or capital gains taxes when profits are distributed.

[2] Smith et al. (2017) use tax audit data to conclude that rising business income accounts for all of the growth in the top 1 percent income share since 2000. Furthermore, the majority of rising top business income resulted from rising income of pass-through businesses.

We start by examining the reliability of measures of business incomes, receipts, and valuations in the Federal Reserve's Survey of Consumer Finances (SCF), which is a publicly available and widely used triennial cross-sectional survey of U.S. households. Households with actively managed businesses are asked to report business receipts and net income from specific lines on their tax forms.[3] This aspect of the survey design makes it easy for us to compare the household responses with administrative data from the IRS *Statistics of Income* (SOI).[4] Averaging across survey years, we find that the SCF overstates pass-through business income per tax return by 400 percent and business receipts per return by 169 percent. For C corporations, net income and receipts are on average understated in the SCF by 26 and 21 percent, respectively, but the SCF does not include publicly traded corporations, whereas the IRS does. Since publicly traded corporations have much higher receipts and net incomes per tax return than private corporations, including them would result in a significant overstatement for all business entities.

The overstatement of incomes and receipts in the SCF varies in the cross section and year by year, making it difficult, if not impossible, to systematically correct for the errors. To demonstrate this, we provide evidence of both sampling and measurement errors. Sampling errors arise from an understatement or overstatement of the reported number of tax returns in the SCF relative to the IRS. We find that the SCF significantly understates the number of sole proprietorships, S corporations, and C corporations and significantly overstates the number of partnerships, with the degree of under- or overstatement varying across the income distribution. Looking at data in the cross section, our findings suggest a significant underrepresentation of low-income businesses driving the overstatement of business incomes. This may be attributable to measurement error resulting from how the questions are framed. For example, there are many IRS businesses with net losses but few in the SCF, possibly because the respondents answered that they had no net income rather

[3] Sole proprietors are asked to report business gross receipts and net income from Form 1040 Schedule C (lines 1 and 31), shareholders of partnerships from Form 1065 (lines 1c and 22), shareholders of S corporations from Form 1120S (lines 1c and 21), and shareholders of C corporations from Form 1120 (lines 1c and 30).

[4] The IRS compiles the SOI based on stratified probability samples of income tax returns and other forms. We use information from the SOI Corporation Income Tax Returns, Individual Income Tax Returns, and Partnership Returns that are available in the historical data tables from www.irs.gov/statistics.

than a negative net income. Another measurement issue we document is that the frequency of referencing supporting documents is strikingly low. For example, if we condition on all business owners in the SCF, we find that 75 percent never referenced any tax document.

The SCF survey has also been used extensively to study the level and dispersion of business wealth. Households with actively managed businesses are asked to assess the value of their business, net of all loans, if they were to sell it. Since there are no measures of total valuations for ongoing businesses other than publicly traded C corporations, we construct net income-to-value ratios in the SCF and compare them to available income yields from brokered private business sales recorded by Pratt's Stats and publicly traded companies, both small and large, recorded by the Center for Research in Security Prices (CRSP) (merged with Compustat). For virtually all subsamples and all years, the SCF income yields are significantly higher than comparable measures from the other datasets. The overstatement in yields is even greater than for incomes, which is suggestive of an understatement in business valuations. For example, the SCF average value-weighted income yield is 19 percent, much higher than the Pratt's Stats estimates of 2 percent or the CRSP estimates of 7 percent for all businesses, and −9 percent for those in the bottom quintile when firms are ranked by total assets. We also find that the SCF distributions are more right-skewed than those based on Pratt's Stats or CRSP-Compustat data.

For unincorporated businesses, we can compare the SCF estimates of business incomes per owner and, if available, income yields to those of three other widely used surveys: the Current Population Survey (CPS), the Panel Study of Income Dynamics (PSID), and the Survey of Income and Program Participation (SIPP).[5] We show that all four surveys overstate incomes per owner relative to the IRS data, but by differing amounts. Averaging across survey years, business income per owner for unincorporated businesses is overestimated by 586 percent in the SCF, 179 percent in the CPS, 185 percent in the PSID, and 34 percent in the SIPP. Average value-weighted income yields calculated for the PSID and SIPP are high relative to Pratt's Stats and CRSP data in all cases but are not very different from those in the SCF. The main differences in yields across surveys are found

[5]We also compare the SCF results to the results of Gurley-Calvez et al. (2016), who match responses of new businesses in the Kauffman Firm Survey (KFS) with IRS tax data and find that these firms understate business incomes. For some other surveys of businesses, such as the Panel Study of Entrepreneurial Dynamics (PSED), we find the response rates of business-related questions to be too low to be reliable.

when we compare the distributional statistics.

An overstatement of business incomes in the survey data relative to the IRS could be the result of misreporting in the IRS or miscategorization of incomes by survey respondents. To check misreporting, we use auxiliary tax audit data to adjust the IRS data but still find a significant mismatch with the survey data. For example, for pass-through businesses, the SCF overstates the average income per return by 178 percent relative to the tax audit data. To check miscategorizations, we use a broader definition of business income. For example, business owners might confuse business incomes on Schedule C, E, and F, overstating one category and understating another. When we combine these categories into a broader concept of business income, we still find incomes to be significantly overstated. Respondents are not miscategorizing the incomes but rather are often overstating all categories of business income. Other adjustments, such as correcting for within-survey inconsistencies regarding business ownership and income and correcting for the fact that the survey only includes individual owners, do not alleviate the measurement issues.

The inconsistencies across surveys and the conceptual measurement issues that we highlight suggest that the "facts" about business income and wealth based on current survey data should be treated with caution. Measurement problems related to business incomes are surmountable given that respondents are asked about specific lines on the tax forms. Measurement problems related to business valuations and returns may be insurmountable without data on actual business sales transactions. First, it is difficult for owners to estimate business valuations when one considers that businesses are heavily invested in intangible assets.[6] Second, survey business owners answer questions separately about income and valuations. For example, if the net incomes derive from both capital and labor inputs, while the business valuations are based on fixed assets owned by the business, then the estimated income yields from surveys may not be comparable across owners who interpret the question differently. Interpreting survey-based measures of business returns or valuations requires a consistent framework for true returns, stocks, and valuations. Given current measurement issues, such interpretations may not be possible.

[6]McGrattan and Prescott (2010a,b) and Bhandari and McGrattan (2018) both find estimates of the value for intangible assets to be close to estimates of tangible assets used by businesses. Intangible assets come in the form of research and development, software, advertising, brands, and investments in customer lists, goodwill, and other forms of investments in building organizations.

This paper is organized as follows. Section 3.2 discusses the related literature and the implications of our findings for theoretical and applied research on businesses. In Section 3.3, we document that business incomes and receipts measured from the SCF survey data are largely inconsistent, and we discuss problems regarding nonrepresentativeness and measurement errors. Section 3.4 studies business valuations and rates of return. Section 3.5 presents robustness checks. Section 3.6 compares the SCF results with other widely used surveys. Finally, Section 3.7 concludes.

3.2 Related Literature

Our findings have implications for three active areas of economic research. The first area is the empirical literature that documents levels and trends in the dispersion of income and wealth and emphasizes the role of entrepreneurs in wealth accumulation. The second area, which is motivated by and builds upon the first, is the theoretical literature developing models of entrepreneurial choice, which are specifically designed to fit the "stylized facts" of the empirical literature. The third area includes quantitative policy analyses that use the empirical findings and theoretical developments of the first two literatures as their laboratory for the study of counterfactual policies. Our findings cast doubt on the facts that have been uncovered in the empirical literature—specifically documenting that survey data are unreliable for business-related statistics—and thus raise issues concerning the theoretical developments and policy analyses that have been designed around them.

A large and burgeoning empirical literature documents trends in income and wealth and has been particularly focused on increased dispersion over time. Greater dispersion is attributed to top earners, and therefore researchers work primarily with survey data from the SCF or administrative tax data from the IRS. For example, Kuhn and Rios-Rull (2016) provide a nearly exhaustive summary of distributional facts about U.S. earnings, income, and wealth based on the SCF. A starting point for several papers in this literature is the observation that, for broad income categories, aggregated SCF responses match up well to the aggregated IRS data. In Figure 3.1, we replicate the time series plot for adjusted gross income (AGI) from the SCF and plot it against the corresponding data from the IRS. We see that the SCF tracks the level and cyclical trends for AGI in the IRS.

Our focus is on measures that relate to business activity. Of particular relevance are the findings in Kuhn and Rios-Rull (2016) that business income is one of the main contributors to income inequality and that business equity is one of the main contributors to wealth inequality, which they document for the history of the SCF surveys between 1989 and 2013. Bricker et al. (2016) also use the SCF to document the rise of the top share of wealth holdings over time and find that the share of wealth attributable to the top 1 percent rose from 30 percent in 1992 to 36 percent in 2013. Our paper exploits the fact that SCF answers can be compared to administrative data from the IRS and finds that respondents are not reliably or consistently answering questions about their business income or business equity, and therefore we cannot trust the SCF distributions.

Given issues with measuring business incomes, researchers might be tempted to combine all nonwage income into a residual "capital income" category, since SCF aggregates match up well with aggregated IRS data. Here, we argue that this capital income measure would not be appropriate for either research on U.S. businesses or research on U.S. capital. For research on businesses, the residual income measure would be inappropriate because significant nonbusiness income is included with interest payments, capital gains, pensions and annuities, alimony, trusts, and government transfers. Furthermore, as we noted earlier, there is evidence that owners are not miscategorizing income categories, and therefore using broader categories of income would do little to ameliorate the measurement issues. For research on capital, the nonwage income in AGI would be inappropriate because a significant fraction of capital income is untaxed and the corresponding assets are held by fiduciaries. Furthermore, as we show later, the majority of respondents do not reference financial documents, making it nearly impossible to have reliable estimates of their total capital income or wealth.

Saez and Zucman (2016) document trends in wealth dispersion by capitalizing incomes from administrative tax data. They compare their results to the SCF and find similar levels and trends for wealth in the top 10 percent of the distribution but differences for the top 1 percent.[7] The Saez and Zucman (2016) capitalization method is inappropriate for estimating wealth in business for several reasons. First, there is no way to validate the procedure except by comparing to survey data, which we find are unreliable. Second,

[7] They also compare results to estate taxes and foundation records, but these data are not informative about most businesses in the United States.

the U.S. flow of funds aggregates used to compute capitalization factors include ad hoc imputations made by the Federal Reserve for private businesses.[8] Third, the assumption that returns are the same for everyone and constant over time is hard to reconcile with the fact that there is significant entry into and exit out of business (see Bhandari and McGrattan (2018)).

Because of the problems with data from the SCF and the capitalized IRS incomes, the main message of our findings for the theoretical literature is a cautionary one, namely, that these data tell us little about business valuations or returns, and therefore theorists should not insist on models that replicate "stylized facts," which are not actually facts. The most popular stylized facts are that entrepreneurs, as a group, own a substantial share of household wealth and income, with shares increasing throughout the distribution, and that entrepreneurs have high savings rates relative to the population, implying much more dispersion in wealth than in income (see De Nardi, Doctor, and Karen (2007) and Gentry and Hubbard (2004)). These findings have led researchers to model entrepreneurs as overcoming significant market frictions to run highly risky businesses with the expectation of earning high returns and amassing significant wealth (see, for example, Quadrini (2000), Cagetti and De Nardi (2006), and Buera (2009)). Furthermore, the theoretical frameworks parameterized to match the survey data have been used as a laboratory for policy work, especially when considering tax policy reform (see, for example, Meh (2005), Kitao (2008), Bohacek and Zubricky (2012), and Scheuer (2013)). Our results cast doubt on the survey data underlying the models of financial frictions and, hence, the subsequent policy recommendations.

Our paper is also related to a second strand of the empirical literature, which reaches

[8]For example, when constructing capitalization factors, Saez and Zucman (2016) use aggregate flow of funds wealth measures for closely held corporations (both subchapter C and S) and unincorporated businesses. These businesses are not publicly traded and thus have no market valuations. The Federal Reserve imputes market values for closely held corporations by taking a ratio of market value to revenues for publicly traded companies and then applying that ratio to private businesses with similar industry, employment, and revenue profiles—after arbitrarily adjusting the estimate downward by 25 percent to reflect the lack of liquidity of closely held shares. Valuations for unincorporated businesses are based on balance sheet data reported to the IRS, which are historical-cost accounting measures, not market valuations.

different conclusions about the payoffs to entrepreneurial activity.[9] Hamilton (2000) uses survey data from the 1984 SIPP and finds that self-employed individuals—who could be running an incorporated or unincorporated business—have lower median earnings than similar individuals in paid employment. Moskowitz and Vissing-Jorgensen (2002) extend his analysis and work with SCF data, allowing for a more comprehensive treatment of equity returns and including adjustments for firm entry and exit. They find that returns to private businesses are no higher than returns to publicly traded firms and thus puzzlingly low given the risks entrepreneurs face.[10] Using PSED data, Hurst and Pugsley (2011) report that more than 50 percent of new business owners cite flexible hours and being one's own boss as the primary reason for starting their own business. These findings have led researchers to conclude that the nonpecuniary benefits of self-employment play an important role in occupational choice.

We document that survey data *overstate* business incomes, and the overstatement leads to income yields for private businesses that are significantly higher than those for publicly traded companies. In other words, we find that the private returns computed with survey data are puzzlingly high, not puzzlingly low. When comparing our results to Moskowitz and Vissing-Jorgensen (2002), we find that the main difference is the concept of return: they add an imputation for capital gains that drive their results. Regrettably, neither the income yield nor the capital gain imputation are reliable estimates, leaving us with little to say about whether private returns are low or high relative to public returns. This finding is relevant for policy discussions related to business taxation and subsidization. For example, Hurst and Pugsley (2017) followed up on the work of Hamilton (2000) and Moskowitz and Vissing-Jorgensen (2002) by incorporating nonpecuniary benefits in a model of entrepreneurship and then analyzed the impact of small business subsidies. Our results cast doubt on SIPP and SCF survey data and hence on policy recommendations that arise

[9] Later, we document inconsistencies across surveys and across time that could lead to different empirical insights.

[10] Moskowitz and Vissing-Jorgensen (2002) used samples over the period 1989–1998. Kartashova (2014) extended their analysis to 2010 and documented that for the longer sample, the private equity premium is about 6 percent as compared to about 0 percent in the shorter sample. Since the difference in means is swamped by the variability of the CRSP index returns, which has a standard deviation of 20 percent over the period 1988–2015, we would argue that their estimated private returns are not significantly different.

from frameworks that incorporate nonpecuniary benefits of business entrepreneurs.

3.3 Business Incomes

In this section, we compare SCF pretax business incomes that correspond to specific lines on U.S. tax forms with incomes reported to the IRS. We first document that the SCF responses—both in the aggregate and across the distribution—are statistically different and argue that the magnitudes of the differences are economically important. We then explore reasons for the mismatch.[11]

3.3.1 Aggregates

We start with aggregated business incomes and show that, while the SCF does well in matching the IRS total AGI (as shown in Figure 3.1), there are large discrepancies between the survey and tax data for businesses. For pass-through businesses, business income per return is significantly and consistently overstated in the SCF relative to the IRS. For C corporations, the average per-return business income in the SCF is not very different from the IRS but should be much smaller given that the survey excludes publicly traded companies.

To demonstrate that there is a significant discrepancy between SCF and IRS data, we start by defining *business income* as gross receipts from sales minus expenses (including depreciation) incurred in running the business. Information on business incomes is obtained from the respective business tax forms: Form 1040, Schedule C (line 31) for sole proprietors, Form 1065 (line 22) for partnerships, Form 1120S (line 21) for S corporations, and Form 1120 (line 30) for C corporations. In each survey year, we use the SCF sampling weights and ownership information to compute the aggregate business income and the aggregate number of business tax returns by legal form of the business.

Figure 3.2 plots aggregated business income divided by the number of business tax returns using the SCF and the data actually reported to the IRS for tax years between 1988

[11] In Appendix C, cited henceforth as Bhandari, et al. (2019), we provide a comprehensive collection of statistics for all survey years and subsamples.

and 2015.[12] All data are reported in thousands of current dollars. The shaded region in Figure 3.2 shows the 90 percent confidence interval of the average per-return business incomes.[13] Panel A of Figure 3.2 shows the results for sole proprietorships. For this group, average incomes reported in the IRS are about \$8,000 per return in 1988 and rise gradually to \$12,000 per return by 2015. Average incomes per return reported in the SCF are significantly higher, rising from \$32,000 in 1988 to \$40,000 by 2015, and display large year-to-year variation. If we construct a percentage error (that is, 100(SCF−IRS)/IRS) in each year, we find the average error is 289 percent and ranges from 158 percent to 384 percent across all years. We see a similar result in Panel B, which shows the data for S corporations. The SCF S-corporation incomes per return are significantly higher than the IRS analogues for the entire sample period. The errors in this case average 273 percent and range from 142 percent to 387 percent.

In the case of partnerships, a sampling issue implies that the discrepancy between the SCF and IRS income could be even larger than that shown in Panel C of Figure 3.2. The SCF only surveys owners of partnerships who are individuals, whereas the IRS includes information on partnerships owned by individuals and other legal entities such as corporations.[14] The exclusion of corporate partners in the SCF should lead to an understatement of aggregate business income but, in principle, should not affect the business income per

[12] In the SCF, we assume that a business owner who owns multiple sole proprietorships files one return. This assumption is made to be consistent with IRS statistics that state: "For purposes of the statistics, if a proprietor owned more than one business, the statistics for each business were combined with those of the proprietor's dominant business and included in the industrial group for that business activity" (see Dungan (2017, p. 2)). For partnerships, S corporations, and C corporations, we calculate the number of returns taking into account the ownership share of the family from each reported business.

[13] Standard errors are computed using a bootstrap method. For each survey year, the SCF provides a set of 1,000 replicate weights. We use these weights to compute 1,000 values for the relevant statistic, for example, business income per return, and then we compute the confidence intervals using the 5th and 95th percentiles of these 1,000 draws.

[14] For the year 2011, Cooper et al. (2016) estimate that 32 percent of total business income from partnerships is generated by individual partners, who account for 73 percent of all partners.

return.[15] However, we find that both aggregate income and per-return income are overstated for partnerships. Aggregate business income from partnerships in the SCF is *higher* than the IRS by 305 percent on average, with the range of errors between 52 percent and 889 percent across survey years. Per-return income is also overstated by 300 percent on average, with errors between 31 percent and 837 percent, as shown in Panel C of Figure 3.2.

If we include all pass-through businesses in one category (using data in Panels A through C of Figure 3.2), we find that the SCF error is 400 percent on average and ranges from 230 percent to 568 percent for business income per return. Contrast this with business incomes per return for C corporations, shown in Panel D of Figure 3.2. For these businesses, we find that in most years, the average SCF business income per return is actually understated by about 26 percent as compared to the IRS data. However, the IRS data include publicly traded corporations, whereas the SCF data do not. Publicly traded C corporations are typically much larger than their private counterparts. If we could include the incomes from these publicly traded corporations with the SCF estimates, we would find that the SCF total incomes would be significantly higher than the IRS estimates, as is the case for the pass-through businesses. Despite these measurement issues, we compare business incomes of C corporations in the SCF with the IRS and interpret the results with the understanding that the SCF will underrepresent large businesses.

While incomes per return are overstated in the SCF relative to the IRS, the number of returns filed by businesses are significantly understated for sole proprietors and corporations in all years. Figure 3.3 plots the number of business returns filed as reported by the IRS and the SCF, over time and by legal entity, with shading marking the 90 percent confidence interval. In the case of sole proprietors and S corporations shown in Panels A and B, the understatement has worsened over time as the number of IRS filings has grown and the number reported in the SCF has not. In Panel C, we see that the number of partnership returns in the SCF is undercounted in only a few years and not by as much as in the case of the other business types. However, as mentioned before, the SCF data only include partners who are individuals, implying that the SCF significantly *overstates* the number

[15] Here, we are operating under the assumption that the SCF is representative and partnerships with all corporate partners (which are entirely missed by the SCF) either are small in terms of their share of aggregate business income or else are not systematically different from the rest of the partnerships.

of returns for partnerships owned by individuals. The number of C-corporate returns is shown in Panel D of Figure 3.3. In this case, part of the difference between the IRS and SCF results is the exclusion of publicly traded companies in the SCF, but these businesses only account for about 5,000 out of the roughly 1.6 million C corporations and therefore cannot account for the large understatement of returns shown in the figure.

One possible reason for understated returns is that the SCF data may not include owners that earn very little business income (for example, part-time Uber drivers or AER referees), while the IRS includes all business tax filers. If this were indeed the case, then the aggregate business incomes—found by multiplying values in Figure 3.3 by values in Figure 3.2—would be similar for the IRS and SCF. We find that this is not the case: aggregate business incomes are significantly overstated in the SCF. For example, in the case of pass-through businesses, we find average overstatements of 34, 137, and 305 percent for sole proprietors, S corporations, and partnerships, respectively, with a large range in the errors over time. (See Bhandari et al. (2019) for details.) The large overstatement of aggregate incomes, especially in S corporations and partnerships, is clearly inconsistent with the hypothesis that nonactive business owners explain the differences between the SCF and IRS results.

Finally, we extend the analysis to business receipts and again find large and variable errors in the SCF responses when compared to the IRS counterparts. (Full details are provided in Bhandari et al. (2019).) For example, in the case of pass-through businesses, we find that the average error in business receipts per return over the period 1988–2015 is 169 percent, with errors over the period in the range of 89 percent to 367 percent. Thus, our main finding is an overstatement of aggregated business incomes and receipts in the SCF across all legal forms, with large variation in the discrepancies across survey years.

3.3.2 Distributions

Next, we show that the overstatement of business income per return documented in the previous section varies in the cross section and year by year, making it impossible to systematically correct the SCF responses. The discrepancies between the SCF and IRS statistics are so large and variable as to make the cross-sectional moments based on the survey data unusable for applied work.

To demonstrate this, we compute percentage errors by grouping businesses in two different ways. First, we group those with positive business income (profits) separately from those with negative business income (losses). For pass-through businesses, the SCF overstates business income per return for those that have profits by 277 percent on average, with the year by year errors in the range of 151 percent to 446 percent. The SCF understates the business income per return for those that have losses by about 82 percent, with the errors in the range of 66 percent to 94 percent. Both the overstatement of profits and the understatement of losses affect the errors in cross-sectional statistics in quantitatively important ways. We demonstrate this in Table 3.1 for pass-through businesses by decomposing the percentage errors in the aggregate business income into the fraction that arises due to overstatement of profits and the fraction that arises due to understatement of losses. For example, in tax year 1988, the overstatement of profits and understatement of losses each account for about 50 percent of the aggregate error. In subsequent years, these fractions vary but are nontrivial in both the overstatement of profits and understatement of losses. (See Bhandari et al. (2019) for results across all legal forms and survey years.)

A second method of grouping businesses is by ranking them according to total income of the owners. Here, we do this in the case of sole proprietorships since we have comparable IRS data in all SCF survey years. (See Bhandari et al. (2019) for a more limited analysis of S corporations.) Specifically, we rank sole proprietors in the SCF by their AGI and then assign them to income brackets using the same bins as the IRS. In Figure 3.4, we plot the fractions of business income for owners with below-median AGI and for those with AGIs in the top 1st percentile. For most years, the SCF income shares for these two groups are understated and display large year by year variation. For example, the share for those with below-median AGI is nearly doubled or halved from one survey to the next. Since the fractions sum to 100 percent across all AGI groups, the SCF must necessarily overstate incomes for some bins. We find the largest overstatement of shares for those with AGIs between the 50th and 75th percentile.

In Figure 3.5, we see that the overstatement of business income per return in the SCF data also varies a lot across years and across AGI bins, with no systematic pattern. The panels of this figure can be compared to the aggregate data for sole proprietorships in Panel A of Figure 3.2. In contrast, the incomes per return in the IRS data show little variation over time and vary similarly across AGI bins.

3.3.3 Nonrepresentativeness and measurement error

We now investigate the reasons behind the discrepancies in business incomes between the SCF and IRS and provide evidence for two types of errors in the SCF: nonrepresentativeness of business owners and misreporting of business income by business owners. The evidence of both types of errors again demonstrates that there is no easy correction for the survey data.

To fix ideas, we decompose the difference between a survey aggregate, X^S, and an IRS aggregate, X^I, for some measure into three terms as follows:

$$X^S - X^I = \left(\frac{\sum_i(\omega_i^S - \omega_i^I)}{\sum_i \omega_i^I}\right) X^I + \sum_i w_i^I X_i^I \left(\frac{\omega_i^S}{\omega_i^I} - \frac{\sum_i \omega_i^S}{\sum_i \omega_i^I}\right) + \sum_i \omega_i^S \left(X_i^S - X_i^I\right), \quad (3.1)$$

where sums are taken over household types indexed by i. The average X for type i is denoted by X_i^s, and the population weight for type i is denoted by ω_i^s, with $s \in \{S, I\}$. The first two terms capture differences in weights, $\omega_i^S \neq \omega_i^I$, which we refer to as "sampling errors," and the last term captures differences in averages, $X_i^S \neq X_i^I$, which we refer to as "measurement errors."

With linked survey-IRS data, one can fully decompose the difference on the left-hand side of (3.1) into the sampling and measurement error components. For the SCF, we do not have access to such linked data but can provide evidence that is strongly indicative that both sampling and measurement errors are nontrivial. We start with evidence on the sampling errors. Earlier, we provided evidence based on the total number of business returns that the first term in (3.1) is significantly different from zero. Recall that we found a significant underrepresentation of sole proprietors and corporations and an overrepresentation of partnerships, indicating severe sampling issues.

To shed light on the second term in (3.1), we analyze how the returns are distributed in the cross section. We define the groups of businesses as we did previously in Section 3.3.2, first on the basis of whether they earned profits or losses and second by ranking them according to their owners' AGI. We then compare ratios of population weights, ω_i^S/ω_i^I, across types and across time. For example, if we compare these ratios for pass-through businesses that have profits with those that have losses, we find significant differences in most years, with the highest difference being 35 percentage points. Similarly, if we compare ratios for sole proprietors in different AGI bins, we find significant differences across AGI bins. Figure

3.6 shows this in the case of sole proprietors with AGIs per return below and above the median. For businesses that have owners with below-median AGIs, the number of IRS returns (ω_i^I) has risen from about 5 million in 1988 to over 12 million in 2015, but the SCF estimate (ω_i^S) has remained at roughly 2 million for the entire period. For businesses with above-median AGIs, the number of IRS returns has risen from a little over 8 million to above 12 million, but the SCF estimate has hovered around 5 million. Comparing these data to the full sample in Figure 3.6, we find that the share of sole proprietorship returns below the median AGI is around 25 percent in the SCF and 43 percent in the IRS. These findings suggest a significant underrepresentation of low-income businesses, which leads to an overstatement of business incomes if business income is positively correlated with AGI.

Finally, we provide evidence on the last term in (3.1), which summarizes the measurement error. As mentioned earlier, without linked records, it is impossible to directly validate measurement errors, but we document several aspects of the SCF survey results that suggest they are nontrivial. The first is related to the aforementioned observation that a significant part of the overstatement of income is due to an understatement of losses. This could arise from the framing of the question "What is your net income?" which could be misinterpreted as being a question about positive net income. For instance, consider the distribution of losses by AGI bins for tax year 2015 as shown in Table 3.2. We see that 10 out of 19 bins, which account for 23 percent of the total number of returns and 26 percent of the total losses in the IRS, have an aggregate zero (that is, all respondents in those income brackets reported a zero net income) in the SCF data.

A second reason to be suspicious about misreported incomes in the SCF is that a very small fraction of respondents refer to their tax documents when responding to questions about the specific line items on tax forms. At the end of the survey, SCF interviewers note how frequently respondents accessed particular documents while answering questions and the type of documents they referenced, if any. Using this information, we calculate the frequency with which business owners referenced either tax or other financial documents in tax year 2015.[16] These tabulations are shown in Table 3.3. The first row shows that 75 percent of business owners in the SCF never referenced tax documents, 2 percent rarely

[16] Other financial documents include account statements, investment and business records, loan documents, and pension documents. If any of these documents are referenced, we assume all are.

did, 9 percent sometimes did, and 14 percent frequently did. The second row shows that 64 percent never referenced any other financial documents, 6 percent rarely did, 15 percent sometimes did, and 15 percent frequently did.

To provide further evidence on measurement errors, we show that the SCF fails a simple consistency check by comparing answers to two closely related questions. In the case of sole proprietors, respondents are asked to report incomes listed on lines 12 and 18 of their Form 1040, which correspond to Schedule C and F incomes, respectively. Separately, they are asked about business income from a sole proprietorship and told it is listed on line 31 of Schedule C.[17] By design, the difference in responses to these two questions must be farm income from Schedule F. In Figure 3.7, we see that the differences across the two answers vary between $17,000 and $40,000 per return, considerably more than could be attributable to farm incomes. In a typical year, only 4 percent of business profits listed on Form 1040 are farm income, and farm losses exceed profits in many of the years of our sample.

A related exercise is to check whether there are SCF respondents who report that they do not own or actively manage a sole proprietorship but still report nonzero income from a sole proprietorship or farm. For example, in 2015, of the 16 million who reported a nonzero income (coded as X5704), only 6 million reported that they actively manage a sole proprietorship (coded as X3119, X3219, or X3319), while 10 million reported that they did not. More importantly, the fraction of misreported income is significant. According to SCF data, 65 percent of the business income from Schedule C and F was earned by those reporting that they did not actively manage a sole proprietorship.

We turn next to measures of business valuations and rates of return, which are key for measuring wealth inequality as well as disciplining theories of entrepreneurial activity.

3.4 Business Valuations and Rates of Return

A challenge in estimating valuations and returns for privately held businesses is that they are not frequently traded, and for this reason, most researchers use the SCF to study the

[17] The first answer is coded as X5704 and the second as X3132, X3232, and X3332, combined with the response to legal status of the actively managed business with codes X3119, X3219, and X3319.

distributional aspects of business wealth and returns.[18] SCF business valuations are self-reported evaluations of the current net worth of the business if sold. We can use these estimates along with net incomes to construct SCF income yields that are comparable to available yields from brokered private business sales and publicly traded firms, both small and large. We find that for virtually all subsamples and all years, the SCF yields are significantly higher and more right-skewed than comparable measures from the brokered sales and public firms. We relate this finding to a previous empirical literature that has documented a puzzlingly small private equity premium and show that the earlier estimates are driven by an inconsistent imputation of capital gains on private businesses.

We start by describing the measurement of SCF income yields, which will be compared to yields based on broker data from Pratt's Stats and public firm data from the CRSP-Compustat merged database. The SCF income yield, which is computed for each business, is the ratio of total pretax net income from businesses divided by the self-reported total net worth of businesses. Let $\{\omega_{i,t}\}$ be the SCF population weights for survey year t. We compute an equally weighted and value-weighted mean yield for t, denoted as R_t^{ew} and R_t^{vw}, respectively:

$$R_t^{ew} = \sum_i \omega_{i,t} \left(\frac{NI_{i,t}}{V_{i,t}} \right), \quad R_t^{vw} = \sum_i \left(\frac{\omega_{i,t} V_{i,t}}{\sum_i \omega_{i,t} V_{i,t}} \right) \left(\frac{NI_{i,t}}{V_{i,t}} \right), \tag{3.2}$$

where NI is total pretax net income and V is the self-reported total business value. In Figure 3.8, we plot time series of yields for all businesses by legal form across years. Across SCF survey years 1989–2016, the average equally weighted yield, R_t^{ew}, is 102 percent for all businesses, 104 percent for pass-through businesses, and 57 percent for C corporations. The average value-weighted yield, R_t^{vw}, is 19 percent for all businesses, 20 percent for pass-through businesses, and 17 percent for C corporations. Yields vary significantly across surveys. For example, in the case of C corporations, R_t^{ew} is in the range of 14 to 102 percent, and R_t^{vw} is in the range of 11 to 28 percent.

Next, we compare the SCF income yields to comparable statistics from Pratt's Stats and show that the SCF yields are much higher and more right-skewed. The Pratt's Stats

[18] Some studies use aggregate business valuations from the Federal Reserve's flow of funds, which imputes values for privately held businesses using estimates of private business revenues, publicly traded business revenue-to-value ratios, and an estimate of the liquidity premium on public versus private business.

database has transaction-level data on sales of private and public businesses over the period 1988–2017. The dataset includes financial information about the target business and other attributes of the sale including payment terms, purchase price allocations, and employment agreements. We compute the income yield in Pratt's Stats by dividing the pretax income earned by the business in the year before the sale by the sale price. The advantage of Pratt's Stats is that it records the price at which the business was actually sold, thus, it is conceptually close to the ideal answer to the SCF question on business valuation. The results of this comparison are shown in Table 3.4. In the first two rows and columns, we report equally weighted and value-weighted mean yields for all businesses in the SCF dataset and the Pratt's Stats dataset. The differences are dramatic: the Pratt's Stats equally weighted yield is 27 percent as compared to 102 percent for the SCF, and the Pratt's Stats value-weighted yield is 2 percent as compared to 19 percent for the SCF. The fact that there is a larger discrepancy in the equally weighted yield than in the value-weighted yield suggests that there are also discrepancies in the distribution of yields. In the last three rows of Table 3.4, we report percentiles of the income yield distribution across data sources. Here, we see that the 75th percentile yield in the SCF is substantially higher than the counterpart in Pratt's Stats. This result suggests that the SCF overestimates the right-skewness of the cross-sectional distribution of business returns.

An obvious concern about the broker data is its representativeness. There may be a bias in business returns arising from the comparison of ongoing concerns (in the SCF) and a possibly selected subset of businesses that were sold. For instance, if businesses with higher rates of return also have a higher probability of finding potential buyers, then Pratt's Stats returns will be biased upward. On the other hand, there could be factors that lead to a downward bias in the Pratt's Stats returns. Sales triggered by distress, say because of health-related issues facing the owner, would imply a lower yield. Similarly, sales of technology- or research-intensive businesses would imply lower yields because of the significant expensing done by these firms. We can partially correct for the downward bias by ignoring transactions in which the target company is in technology- and research-intensive sectors (that is, with NAICS codes 51, 5415, or 5417) and those for which the stated reason for the sale was health related. In the third column of Table 3.4, we report the data for this subset of firms and find the results are nearly the same as for all businesses.

Next, we compare SCF income yields to those of publicly traded firms in the CRSP-Compustat database. This comparison has the advantage that we have a sample of businesses that are active, but it has the disadvantage that the typical firm in CRSP is much larger than the typical firm in the SCF. To address this issue, we restrict our attention to larger firms in the SCF (S and C corporations) and compare them to both the full sample of CRSP-Compustat firms and a subset of small firms. Our baseline definition of "small" is a firm that belongs to the bottom quintile of firms ranked by the book value of assets.[19] The income yield for a firm in CRSP-Compustat is computed by dividing business income (before taxes and retained earnings) by the firms' market capitalization. The equally weighted and value-weighted means are computed as in (3.2) where we use the end-of-year market capitalization as a measure of business value. In Table 3.5, we see that the equally weighted income yield is negative for both the full sample (-9 percent) and the subsample of small firms (-27 percent) in the CRSP-Compustat database, whereas the yield is large and positive for both C corporations (57 percent) and S corporations (76 percent) in the SCF. Similarly, SCF value-weighted yields are significantly higher than CRSP yields, although less dramatically different. Considering the distribution, we again find that the SCF yields are more right-skewed than those in CRSP-Compustat. For example, at the 75th percentile, the SCF C-corporate yield is 36 percent, while the CRSP yield for all businesses is 10 percent.

Our results thus far would appear to be inconsistent with Moskowitz and Vissing-Jorgensen (2002), who constructed private business returns using SCF data and concluded that they were surprisingly low when compared to those of publicly traded firms. We use a longer sample than they do, but we know from Kartashova (2014) that this would account for only about a 6 percentage point difference in the SCF estimates. The more important difference for the quantitative results is the concept of return. The earlier results are based on a measure of return equal to the sum of a value-weighted income yield and an imputed capital gain. In theory, one would need a panel of firm valuations to compute a value-weighted capital gain, namely,

$$R_{t+1}^{cg} = \sum_i \left(\frac{\omega_{i,t} V_{i,t}}{\sum_i \omega_{i,t} V_{i,t}} \right) \left(\frac{V_{i,t+1}}{V_{i,t}} \right), \tag{3.3}$$

[19] In Bhandari et al. (2019), we also report results separately for each survey year and for different definitions of "small," for example, based on gross sales or market capitalization.

using survey weights $\{\omega_{i,t}\}$ and valuations $\{V_{i,t}\}$ for each firm i in year t. Given that the SCF survey is triennial with virtually no panel aspect (other than two surveys), there is no way to compute $V_{i,t+1}/V_{i,t}$ firm by firm. Moskowitz and Vissing-Jorgensen (2002) instead compute their capital gains measure using the following annualized index:

$$\tilde{R}^{cg}_{t+3} = \left(\frac{\sum_i \omega_{i,t+3} V_{i,t+3}}{\sum_i \omega_{i,t} V_{i,t}} \right)^{\frac{1}{3}} - 1. \tag{3.4}$$

Their concept of rate of return is given by $R^{vw}_t + \tilde{R}^{cg}_t$, where R^{vw}_t is defined in (3.2). They adjust the SCF net income by subtracting imputed measures of taxes and retained earnings and compare their measure of return to the value weighted mean holding period return on the CRSP index portfolio.[20] This procedure generates private returns that are similar in magnitude to the CRSP returns. Considering the higher risk for private businesses, Moskowitz and Vissing-Jorgensen (2002) conclude that there is a puzzle as to why individuals become entrepreneurs. Their preferred explanation is that there are non-pecuniary benefits of running a business.[21]

Replicating the exercise of Moskowitz and Vissing-Jorgensen (2002) for our full sample with income yields and capital gains compared separately, we find that the capital gain imputation drives the differences between our findings and theirs. The results are summarized in Table 3.6. The first two columns show estimates of SCF and CRSP-Compustat income yields, R^{vw}_t, in all SCF survey years. The last three columns show estimates of $\tilde{R}^{cg}_t$ for SCF and both R^{cg}_t and $\tilde{R}^{cg}_t$ for the CRSP-Compustat sample. The table reveals several noteworthy patterns. First, consistent with our findings for the average income yields, the yearly SCF yields are substantially higher than the CRSP-Compustat counterparts for all survey years. Second, the annualized SCF capital gains vary substantially less than those for firms in the CRSP-Compustat gains R^{cg}_t over the sample, which is not surprising given the conceptual differences in the measures and the long interval between survey years.[22] If

[20] Since the assumptions underlying the imputations of taxes and retained earnings are ad hoc, we measure R^{vw}_t using pretax income in both the SCF and CRSP samples.

[21] See also Hamilton (2000) and Hurst and Pugsley (2011), who reach a similar conclusion using data from the SIPP and the PSED, respectively.

[22] Incidentally, the time variation in the capital gains components explains why Moskowitz and Vissing-Jorgensen (2002) and Kartashova (2014) estimate different average returns for the different sample periods they study.

we were to add R_t^{vw} plus $\tilde{R}_t^{cg}$ for SCF and R_t^{vw} plus R_t^{cg} for CRSP-Compustat firms, we would confound two discrepancies and conclude that the private and public returns are not very different on average: 26 percent for SCF versus 21 percent for CRSP-Compustat. If we were to restrict attention to comparable measures, either R_t^{vw} or $R_t^{vw} + \tilde{R}_t^{cg}$, we would instead conclude that the private business yields and the imputed total returns are relatively high for private businesses when compared to public returns, not low as previously thought.

As we discussed in Section 3.3.3, we find that SCF returns are relatively high because there are significant measurement and sampling errors. In the case of self-reported business valuations, respondents would understate values of intangible assets or nontransferable human capital (for example, reputation of the owner). Given that the SCF incomes are verifiably overstated, an understatement of business value would bias the SCF returns upward even more. Furthermore, a rate of return is the ratio of two terms, both of which have measurement errors. This injects a much larger error in the ratio and can cause large discrepancies in the distribution of measured returns with no obvious correction. A constructive way to deal with the measurement issues in the SCF and estimate aggregate and distributional statistics for business valuations is to rely more heavily on a theory that is disciplined by the flows measured from the IRS and business sales data such as Pratt's Stats. A theory featuring business sales would take a stand on the selection bias and also provide a way to impute the valuations for ongoing concerns.[23]

3.5 Robustness

In this section, we show that the overstatement of business incomes in the SCF is robust to potential misreporting in the IRS and to potential miscategorization by SCF respondents across closely related categories of business income.

3.5.1 Adjusting for misreporting in the IRS

One explanation for the overstated business incomes in the SCF is that individuals might report true incomes in the surveys but underreport their incomes to the tax authorities. In

[23] An attempt in this direction is some ongoing work in Bhandari and McGrattan (2018).

principle, this should not be a concern for our analysis because the SCF specifically asks them to report what they *wrote* on their tax forms. Nevertheless, we analyze data from several sources on the extent of misreporting on tax forms to evaluate this hypothesis and conclude that tax misreporting is far from sufficient to close the gap between the SCF and IRS business incomes.

The Bureau of Economic Analysis (BEA) estimates tax misreporting for noncorporate income (that is, incomes from sole proprietorships and partnerships) to reconcile the national accounts with the data from tax audits. For the years 1988–2015, the BEA estimates that reported noncorporate tax incomes are lower by roughly 50 percent because of misreporting. These BEA estimates are in line with studies such as Johns and Slemrod (2010), who used tax audit data for the year 2001 and document that Schedule C income is underreported by 54 percent. For S- and C-corporate business incomes, measures of tax misreporting are hard to obtain. Johns and Slemrod (2010) document underreporting of 18 percent for Schedule E income, which includes all supplemental income from S corporations, partnerships, rental real estate, royalties, estates, trusts, and farm rentals. Based on reviews of documents stemming from the National Research Program at the IRS, the Government Accountability Office (GAO) estimates net misreporting margins for S-corporate incomes on the order of 15 to 20 percent. We construct a measure of adjusted IRS pass-through income by adding back the BEA estimates of misreporting for noncorporate incomes, along with an adjustment of 18 percent for income from S corporations based on the study of Johns and Slemrod (2010) and the reports of the GAO. In Figure 3.9, we compare the SCF business incomes per return to the adjusted IRS incomes per return and find that they are still significantly higher. Computing the SCF errors as before, we find that the average error with the tax audit adjustment is 178 percent, with a range of errors of 98 percent to 274 percent over the sample.

3.5.2 Adjusting for categorization of business income

Another source of measurement error in the SCF is the respondent's possible confusion about closely related categories of business income. For example, when asked about income from a sole proprietorship appearing on line 31 of Schedule C, business owners might also include income appearing on Schedule E, which includes income from real estate, royalties, partnerships, S corporations, estates, and trusts. From our previous analysis, we know that business incomes from Schedule C are overstated in the SCF. If the overstatement was due to classification errors, we should see an understatement in categories of income corresponding to Schedule E. In Table 3.7, we report the percentage errors for Schedule E income and document that they are overstated for most survey years.

Johnson and Moore (2008) conducted a similar exercise but constructed an even broader category of business income by including capital gains and losses (lines 13 and 14 of Form 1040) to the incomes individuals report on Schedules C, E, and F. Neither the IRS nor SCF data allow us to isolate the capital gains for business owners or for business-related assets. Hence, we did not include these data in our baseline analysis of business income. In Bhandari et al. (2019), we replicate the findings of Johnson and Moore (2008) and extend them to all survey years. We find that, although the capital gains in the SCF are lower when compared to the IRS gains, the Johnson and Moore (2008) measure of broader business income in the SCF is still larger in all years than its counterpart in the IRS. The average error is 47 percent, with a range of 18 percent to 115 percent across survey years.

3.6 Other Surveys

In this section, we review evidence from other surveys, namely, the CPS, PSID, SIPP, KFS, and PSED. These surveys contain information about businesses and have been widely used by researchers. When comparing business incomes and valuations across these surveys and with the SCF, we find that there are significant inconsistencies but similar concerns related to sampling and measurement.

We start with the CPS, PSID, and SIPP. All three surveys contain questions about business incomes and organizational forms (that is, whether they are incorporated or unincorporated). The PSID and SIPP additionally contain self-reported estimates of business valuations.[24] Unlike the SCF, the surveys have less detailed information on the legal form of the businesses. For example, these surveys do not distinguish among types of pass-through businesses, and the questionnaires do not specifically connect responses to line items on tax forms. In order to compare across surveys, we focus on business income per owner and

[24] See Bhandari et al. (2019) for more details.

income yields for unincorporated businesses.

In Figure 3.10, we plot incomes per owner for four surveys (SCF, CPS, PSID, and SIPP) and the IRS in Panel A and the number of owners for all surveys in Panel B.[25] As with the SCF, the CPS, PSID, and SIPP have higher business income per owner than is reported by the IRS, but the magnitudes are statistically different across surveys. The SCF is highest with estimates in the range of \$29,000–\$100,000, the PSID is next with a range of \$15,000–\$55,000, the CPS after that with a range of \$15,000–\$35,000, and the SIPP is lowest with a range of \$13,000–\$18,000. All are higher than the IRS, which has a range of \$5,000–\$15,000.[26] The inconsistencies between surveys are driven primarily by differences in aggregate business incomes. The number of owners across these surveys are not significantly different from each other—on the order of 10 to 13 million and stable across years—but are far lower than the IRS, which reports roughly 35 million owners in 1988 and over 50 million by 2015.[27]

Next, we use the responses on self-reported business valuations to compute income yields, as we did for the SCF in Section 3.4. In Table 3.8, we see that value weighted income yields in the PSID and SIPP are comparable to the SCF even though business income per owner is lower than that in the SCF by a factor of two or three. This implies that average business values are even lower in these other surveys. However, if we compare yields across the distribution, we see large differences, especially in the right tail. These observations point to the lack of representativeness in the PSID and SIPP for the universe of unincorporated businesses as well as their lack of comparability to the SCF.

[25] Our sample in the PSID starts in 1992 and provides annual data until 1996 and biennially after that until 2014. The SIPP reports business incomes every four months for the years 2004–2006 and 2009–2012, and valuations are reported once a year for 2004, 2005, and 2009–2011 depending on when the "topical" modules are available.

[26] Hurst, Li, and Pugsley (2014) combine spending data from the Consumer Expenditure (CE) survey with the PSID and estimate that self-employed individuals underreport income by about 25 percent relative to an imputed measure of true income. The imputation relies on estimating the relationship between expenditures and incomes for wage and salary workers and using it along with food expenditures for the self-employed to infer "true" income of the self-employed. We instead compare survey responses directly to IRS data.

[27] As in the SCF, these surveys only account for partners who are individuals. However, as we mentioned before, using estimates from Cooper et al. (2016), this fact alone does not help to account for the massive understatement in the number of owners.

For KFS, Gurley-Calvez et al. (2016) compare responses about receipts, expenses, and profits with matched tax forms for an eight-year panel of new businesses beginning in 2004. They match responses from Form 1040, Schedule C for sole proprietorships, Form 1065 for partnerships, and Form 1120S or 1120 for corporations. Eighty percent of firms are matched to tax files, and the matched data file includes 3,940 firms. They find that the businesses in the survey overstate receipts and overstate expenses by even more, implying that the businesses understate profits across the distribution. These findings are for the most part in contrast to the SCF and IRS comparison, as the SCF overstates business income, while the KFS firms understate business income. We report estimates from their study in Table 3.9 for ease of comparison.

The PSED provides information about business start-ups using a nationally representative sample. However, from the perspective of our study, the PSED suffers from a critical measurement issue: the response rates for business-related questions is very low in all years of the survey. For example, among the 1,214 entrepreneurs in the 2005 panel, only 115 (that is, 9 percent) responded to the question that asks about calculated profits and losses during a follow-up interview for tax year 2006. Thus, we would caution against any use of statistics for quantitative research on entrepreneurial activity from this survey given the small sample size.

In summary, we find severe measurement issues with other surveys currently being used to study U.S. businesses. Key statistics drawn from these surveys are inconsistent with administrative data from the IRS and are inconsistent with each other.

3.7 Conclusion

This paper has examined the reliability of widely used survey data for studies of U.S. businesses. We compared key statistics for net incomes and receipts to counterparts in administrative data from the IRS and found large sampling and measurement errors. In all surveys examined, we found that incomes are significantly overstated relative to IRS data, even when respondents are asked to provide incomes from specific lines on their tax forms. The errors we document are large on average and vary wildly across years and across surveys. We provide evidence that the overstatements of income may be due to the nonrepresentativeness of business owners with lower incomes and to the fact that the majority of respondents do not reference any tax or financial documents. We also consider the implications for key statistics used in economic research, such as the level and dispersion of wealth and the return on businesses.

We hope and expect that our analysis will lead to improved measurement in future surveys. Studies of wealth inequality, entrepreneurial choice, and business taxation are using current surveys as predictive tests for economic theory. Our findings suggest that the current data should be treated with great caution, but we hope improvements in sampling will lead to improvements in quantitative predictions in the future. Attempts should be made to link responses to administrative data where possible. Questions should be limited to queries that are verifiable. In the case of businesses, care should be taken to ensure representative samples of all types of legal organizations.

Table 3.1: Decomposition of SCF-IRS pass-through business income gap

		Percentage of Gap	
Tax Year	SCF−IRS \$ Bill.	Overstatement of Profit	Understatement of Loss
1988	159	50	50
1991	307	64	36
1994	624	83	17
1997	386	63	37
2000	635	68	32
2003	785	71	29
2006	1,096	77	23
2009	750	41	59
2012	218	−56	156
Mean	551	51	49

Note: This table shows the difference (gap) between aggregated business income of all pass-through businesses in the SCF and the IRS. The gap is then decomposed into the fraction attributable to an overstatement of profits or that attributable to an understatement of losses.

Table 3.2: Sole proprietorships with net losses in the IRS and SCF by AGI bins, 2015

	IRS		SCF	
AGI	Returns	Losses	Returns	Losses
Bins	’000	$ Bil.	’000	$ Bil.
No adjusted gross income	426.0	12.2	91.4	0.2
$1 under $5,000	138.3	0.9	39.7	0.2
$5,000 under $10,000	185.7	1.5	33.3	0.0
$10,000 under $15,000	270.8	2.4	10.6	0.0
$15,000 under $20,000	344.3	3.5	47.9	0.0
$20,000 under $25,000	351.4	3.1	60.0	0.2
$25,000 under $30,000	316.8	3.0	77.5	0.2
$30,000 under $40,000	533.0	3.9	102.2	0.6
$40,000 under $50,000	469.3	3.4	62.8	0.0
$50,000 under $75,000	833.7	5.8	159.3	0.1
$75,000 under $100,000	626.4	4.3	199.5	0.8
$100,000 under $200,000	1047.9	7.7	216.2	0.8
$200,000 under $500,000	312.4	3.7	71.6	0.4
$500,000 under $1,000,000	50.4	1.3	0.0	0.0
$1,000,000 under $1,500,000	11.6	0.6	0.6	0.0
$1,500,000 under $2,000,000	5.3	0.4	0.0	0.0
$2,000,000 under $5,000,000	8.4	1.0	0.1	0.0
$5,000,000 under $10,000,000	2.3	0.5	0.7	0.0
$10,000,000 or more	1.8	1.3	36.6	0.0

Note: This table shows the number of business returns that report a net loss and the corresponding amount of these net losses across various AGI bins for tax year 2015.

Table 3.3: Percentage of respondents checking documents in SCF 2016

	Never	Rarely	Sometimes	Frequently
Income tax document	75	2	9	14
Other financial documents	64	6	15	15

Note: This table shows the fraction of business owners that refer to their income tax documents or other relevant financial documents in varying frequency. A respondent who referred to account statements, investment/business records, or loan documents is considered to have checked other financial documents.

Table 3.4: Net income yields in the SCF and Pratt's Stats

	SCF	Pratt's Stats	
	All	All	Non-tech &
Moments	businesses	businesses	nondistressed
Equally weighted mean	102.5	27.4	29.3
Value weighted mean	19.1	1.9	3.5
p25	0.9	3.8	5.0
p50	17.6	21.7	23.1
p75	63.0	46.8	48.3

Note: This table shows moments of the net income yield distribution from the SCF and Pratt's Stats. For Pratt's Stats, we also consider income yields for a subset of businesses that excludes those in technology- and research-intensive sectors (NAICS codes 51, 5415, or 5417) and those for which the stated reason for the sale was health related.

Table 3.5: Net income yields in the SCF and CRSP-Compustat

	SCF		CRSP-Compustat	
	C Corps	S Corps	All businesses	Small businesses
Equally weighted mean	56.8	76.4	−9.2	−26.6
Value weighted mean	16.9	15.2	7.3	−8.5
p25	1.3	2.2	−5.5	−29.0
p50	10.6	14.2	5.4	−7.7
p75	36.2	50.5	10.4	4.0

Note: This table shows moments of the net income yield distribution from the SCF and the CRSP-Compustat database. For the CRSP-Compustat sample, small businesses refer to publicly traded firms in the CRSP database that belong to the bottom 20 percent when ranked by total assets.

Table 3.6: Net income yields and capital gains

Tax Year	Net Income Yields		Capital Gains		
	SCF	CRSP	SCF	CRSP-Compustat	
				$(t-1) \to t$	$(t-3) \to t$
1988	16.6	12.4	—	—	—
1991	20.7	6.2	0.2	26.9	13.2
1994	31.5	9.8	5.3	-3.2	8.5
1997	20.6	6.2	11.4	30.2	29.7
2000	22.6	4.6	11.7	3.7	13.8
2003	17.7	6.2	6.6	28.6	-4.8
2006	18.1	8.0	15.9	10.3	8.9
2009	14.8	5.7	−7.9	21.6	-8.6
2012	14.1	8.0	2.9	12.0	9.6
2015	14.6	5.4	12.8	-3.0	10.7
Mean	19.1	7.3	6.6	14.6	9.0

Note: This table shows estimates of income yields and capital gains for businesses in the SCF and CRSP-Compustat firms. For the SCF, capital gains are computed using Equation 3.4 found in the main text, as in Moskowitz and Vissing-Jorgensen (2002). For the CRSP-Compustat firms, we report two measures of capital gains. The column $(t-1) \to t$ measures the realized capital gains using Equation 3.3 for year t where t corresponds to the fiscal year for which income is reported in the SCF. The column $(t-3) \to t$ measures a geometric mean of the capital gains for the index over the past three periods using equation 3.4.

Table 3.7: Schedule E income comparison

Tax Year	IRS	SCF	Error
	\$ Bill.	\$ Bill.	%
1988	57.3	116.1	102.6
1991	69.9	129.6	85.6
1994	133.0	121.8	−8.4
1997	195.3	147.0	−24.7
2000	249.0	180.3	−27.6
2003	292.7	427.1	45.9
2006	463.1	805.6	74.0
2009	380.8	720.7	89.3
2012	613.3	949.3	54.8
2015	713.2	1142.1	60.1

Note: This table shows aggregated Schedule E income from the IRS and respondents' reported Schedule E income in the SCF. Dollar amounts are in billions.

Table 3.8: Income yield distribution of noncorporate businesses in the SCF, SIPP, and PSID

	SCF	SIPP	PSID
Value-weighted mean	22.6	17.7	14.9
p25	0.8	2.2	3.2
p50	19.6	33.2	27.0
p75	70.6	230.1	114.9

Note: This table shows moments of the income yield distribution for noncorporate businesses for the SCF, SIPP, and PSID. The columns average the income yields for all years that the data are available. The SCF is available triennially between 1989 and 2016, the SIPP for the years 2004–2005 and 2009–2011, and the PSID biennially between 1988 and 2014.

Table 3.9: Comparison of KFS and IRS business tax data, 2004–2011

	Receipts			Expenses			Profit		
	KFS	IRS	Error	KFS	IRS	Error	KFS	IRS	Error
Statistic	'000	'000	%	'000	'000	%	'000	'000	%
Mean	552	417	32	369	188	96	30	169	−82
Median	92	66	29	57	36	57	5	24	−79
p25	21	11	74	1	12	-1,400	-3	1	−700
p75	350	281	25	236	152	55	31	142	−78
p99	11,500	7,434	55	7,450	2,680	178	810	2,478	−67

Note: The source of statistics is Gurley-Calvez et al. (2016).

Figure 3.1: Adjusted gross incomes: SCF vs. IRS

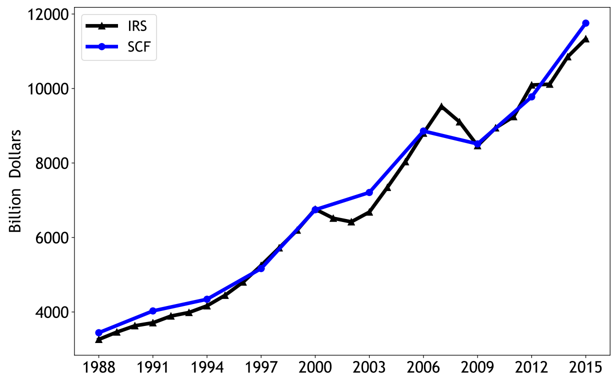

Note: For the IRS, adjusted gross income is obtained from Form 1040. For the SCF, if AGI is not available, we construct it by adding the appropriate income categories.

Figure 3.2: Business income per return by legal entity: SCF vs. IRS

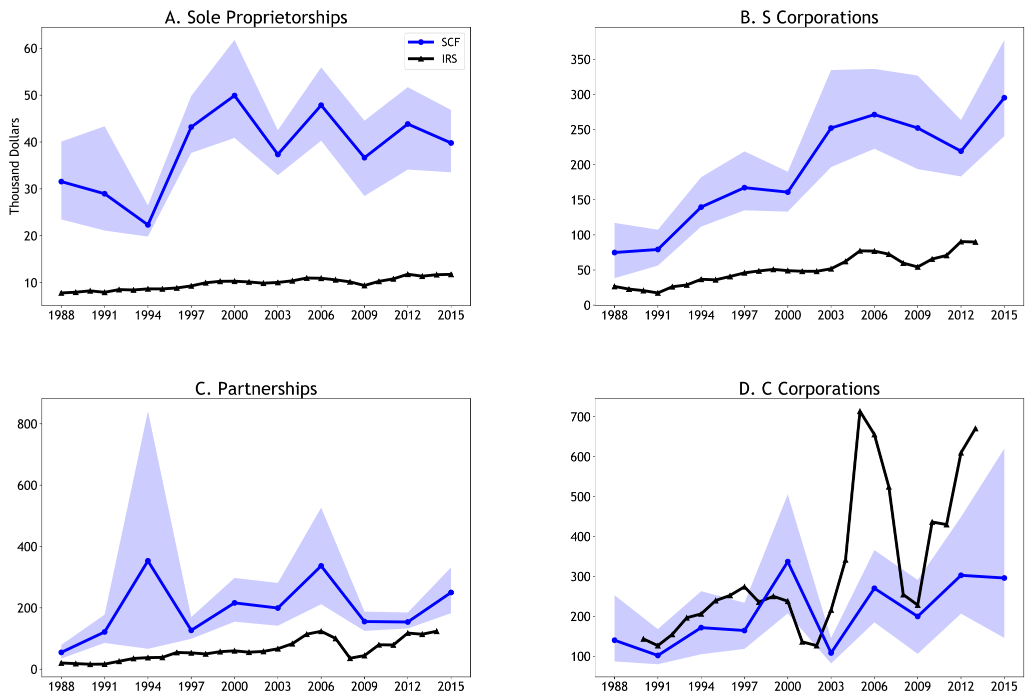

Note: This figure plots business income per business tax return in the IRS and the SCF as reported on Form 1040 Schedule C for sole proprietorships, Form 1120S for S corporations, Form 1065 for partnerships, and Form 1120 for C corporations. IRS data for partnerships, S corporations, and C corporations are available only until 2013. IRS data for C corporations exclude data for those filing 1120A, 1120F, 1120L, 1120PC, 1120REIT, 1120RIC. Prior to 1990, only consolidated information is available and thus is not comparable to the series plotted here. The shaded region for the SCF shows the 90 percent confidence interval.

Figure 3.3: Number of returns by legal entity: SCF vs. IRS

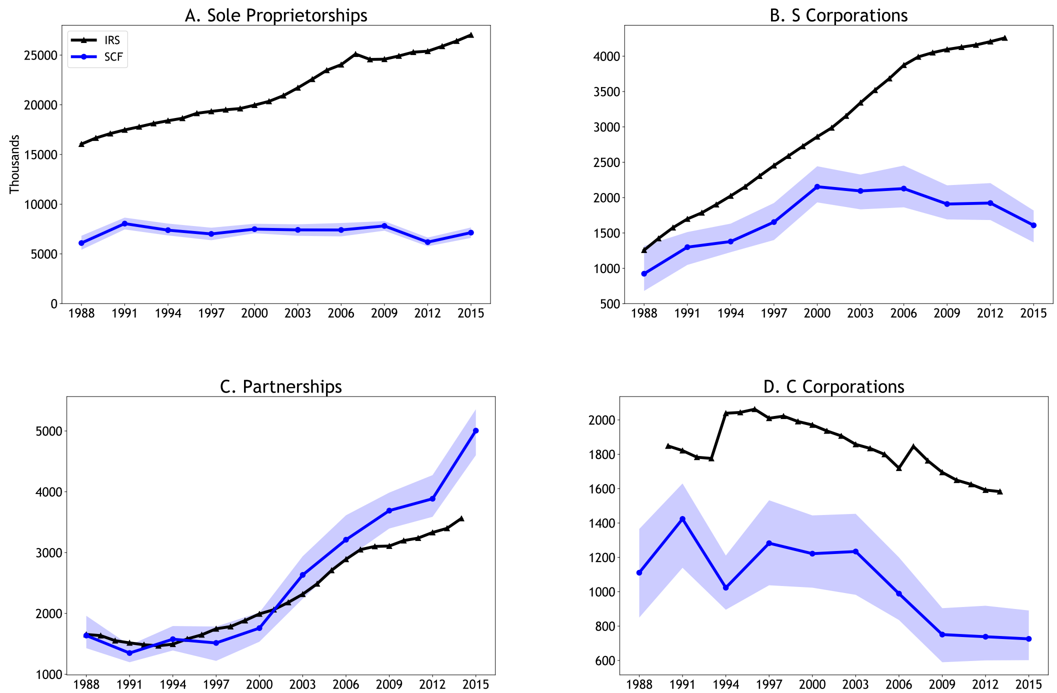

Note: This figure plots the number of business returns of sole proprietorships, S corporations, partnerships, and C corporations over time in the IRS and the SCF. IRS data for partnerships, S corporations, and C corporations are available only until 2013, and C-corporation data are unavailable prior to 1990 when only consolidated information is available and thus is not comparable to the series plotted here. The shaded region for the SCF shows the 90 percent confidence interval.

Figure 3.4: Proprietor income shares: SCF vs. IRS

Note: This figure plots the fraction of business income from sole proprietorships attributable to returns with AGI below the median and above the 99th percentile.

Figure 3.5: Income per return, proprietors with below- and above-median AGI: SCF vs. IRS

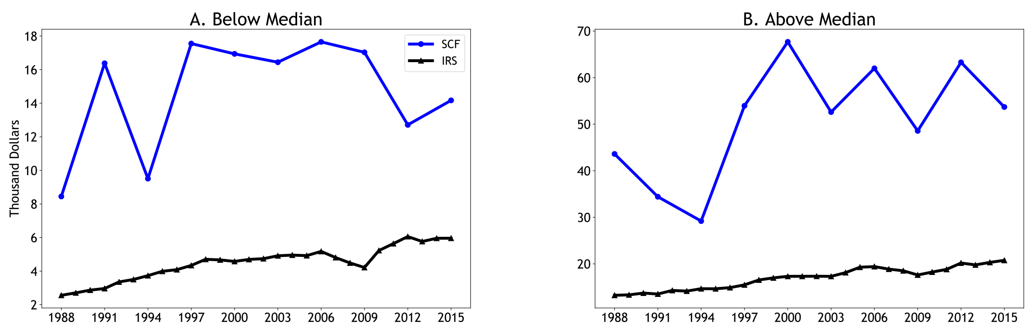

Note: This figure plots sole proprietorship business income per return for those with below- and above-median AGI.

Figure 3.7: Comparing proprietors' individual and business incomes, SCF

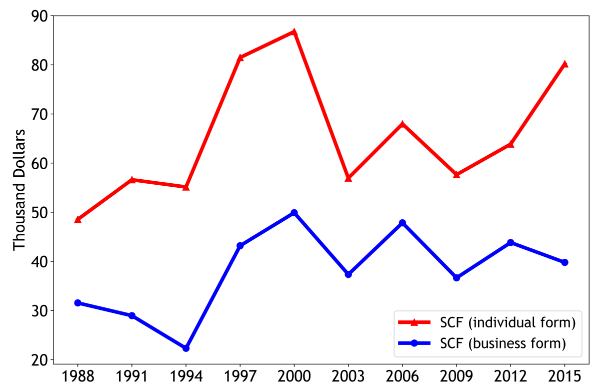

Note: This figure plots business income per return in the SCF for questions that ask respondents to report individual incomes listed on Form 1040, lines 12 plus 18, and business income on Schedule C of 1040, line 31.

Figure 3.6: Number of returns, proprietors with below- and above-median AGI: SCF vs. IRS

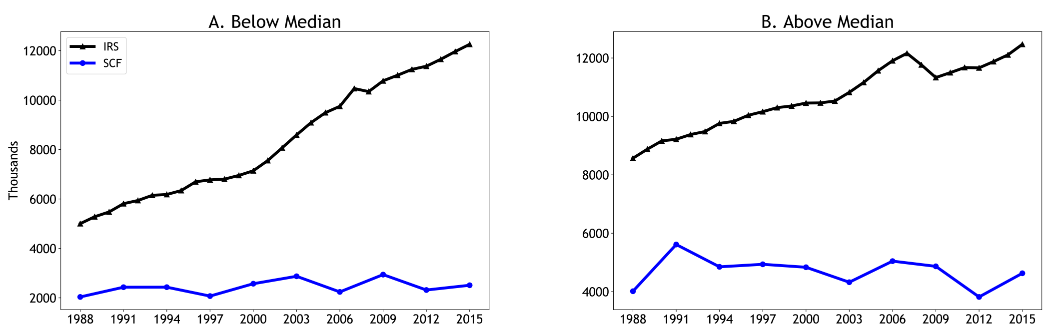

Note: This figure plots the number of sole proprietorship returns (Form 1040 Schedule C) filed by business owners with below- and above median AGI.

Figure 3.8: Equally and value-weighted average net income yields, SCF

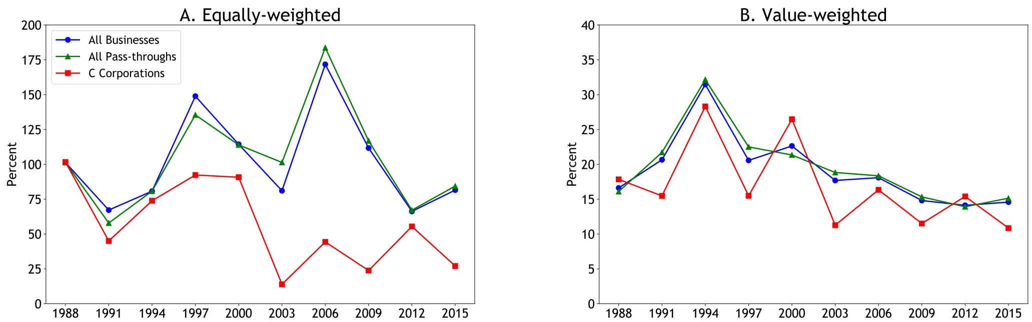

Note: This figure plots equally weighted and value-weighted average net income yields. The SCF sample includes businesses with positive net worth and excludes the bottom 1st percentile of these businesses. The business income of each business that the family members own in the SCF is obtained from SCF variables that correspond to information on business tax forms.

Figure 3.9: Pass-through business income per return with tax misreporting adjustments, SCF vs. IRS

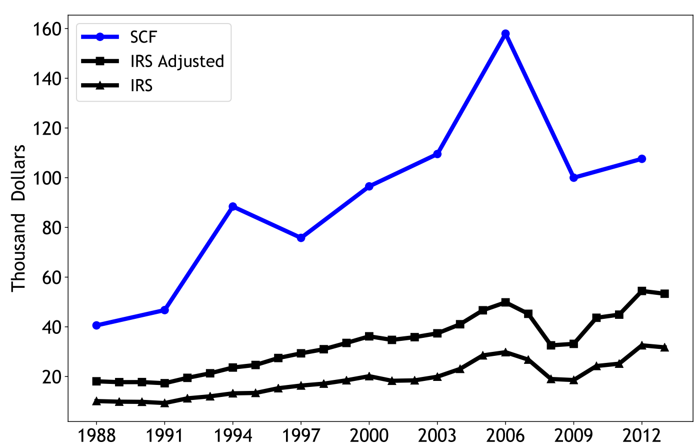

Note: In this figure, we use BEA estimates for misreporting of pass-through business incomes and reports from the Government Accountability Office (GAO) on misreporting of S-corporation business incomes to adjust IRS pass-through business income per return. We add these yearly adjustments to the sum of pass-through income in the IRS, calculate total business income per tax return, and compare it with estimates from the SCF.

Figure 3.10: Unincorporated business income per owner and number of owners

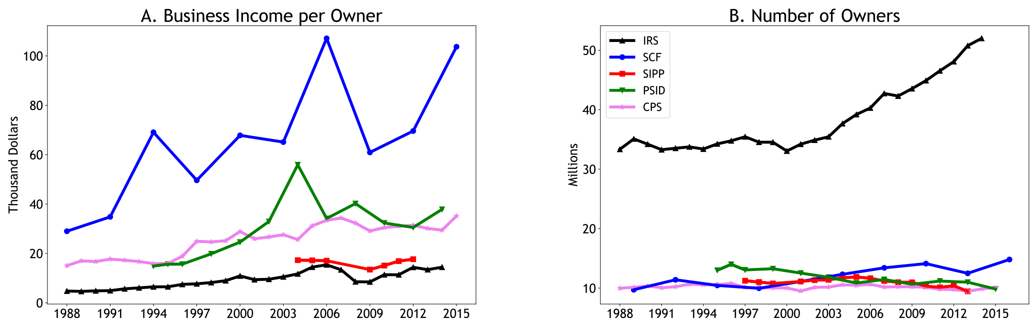

Note: This figure plots the total business income per owner of unincorporated businesses (Panel A) and total number of unincorporated business owners (Panel B) in the SCF, CPS, PSID, SIPP, and the IRS. Before 2004, the SIPP does not provide information about an individual's own share of business income from an unincorporated business. Instead, it contains information about the total income of the business, which is not enough information to calculate the total business income of unincorporated businesses.

References

Abdulkadiroğlu, A., Kuruşçu, B., & Şahin, A. (2002). Unemployment insurance and the role of self-insurance. *Review of Economic Dynamics*, 5(3), 681-703.

Acemoglu, D., & Shimer, R. (1999). Efficient unemployment insurance. *Journal of Political Economy*, 107(5), 893-928.

Acemoglu, D., & Shimer, R. (2000). Productivity gains from unemployment insurance. *European Economic Review*, 44, 1195-1224.

Aguiar, M., & Hurst, E. (2005). Consumption versus expenditure. *Journal of Political Economy*, 113(5), 919-948.

Aizawa, N. (2017) Labor market sorting and health insurance system design. Working Paper.

Aizawa, N., & Fang, H. (2015). Equilibrium labor market search and health insurance reform. NBER Working Paper 18698.

Albrecht, J., & Vroman, S. (2005). Equilibrium search with time-varying unemployment benefits. *The Economic Journal*, 115(505), 631-648.

Anderson, P. M., & Meyer, B. D. (1997). Unemployment insurance takeup rates and the after-tax value of benefits. *The Quarterly Journal of Economics*, 112(3), 913-937.

Ashenfelter, O., Ashmore, D., & Deschênes, O. (2005). Do unemployment insurance recipients actively seek work? Evidence from randomized trials in four US states. *Journal of Econometrics*, 125(1), 53-75.

Attanasio, O., Kitao, S., & Violante, G. L. (2010). Financing Medicare: A general equilibrium analysis. In Demography and the Economy (pp. 333-366). University of Chicago Press.

Auray, S., Fuller, D., & Lkhagvasuren, D. (2013). Unemployment insurance take-up rates in an equilibrium search model. Working Paper 13001, Concordia University.

Baicker, K., Finkelstein, A., Song, J., & Taubman, S. (2014). The impact of Medicaid on labor market activity and program participation: evidence from the Oregon Health Insurance Experiment. *American Economic Review, Papers and Proceedings*, 104(5), 322-28.

Bhandari, A., Birinci, S., McGrattan, E. R., See, K. (2019). Data appendix: What do survey data tell us about U.S. businesses? Federal Reserve Bank of Minneapolis, Staff Report 578.

Bhandari, A., & McGrattan, E. R. (2018). Sweat equity in U.S. private business. Federal Reserve Bank of Minneapolis, Staff Report 560.

Blank, R. M., & Card, D. E. (1991). Recent trends in insured and uninsured unemployment: is there an explanation?. *The Quarterly Journal of Economics*, 106(4), 1157-1189.

Bohacek, R., & Zubricky, J. (2012). A flat tax reform in an economy with occupational choice and financial frictions. *Economic Journal*, 122(565): 1313–1345.

Bricker, J., Henriques, A., Krimmel, J., & Sabelhaus, J. (2016). Measuring income and wealth at the top using administrative and survey data. Brookings Papers on Economic Activity, 2016(1): 261–331.

Browning, M., & Crossley, T. F. (2001). Unemployment insurance benefit levels and consumption changes. *Journal of Public Economics*, 80(1), 1-23.

Buera, F. J. (2009). A dynamic model of entrepreneurship with borrowing constraints: Theory and evidence. *Annals of Finance*, 5(3–4): 443–464.

Burdett, K., Shi, S., & Wright, R. (2001). Pricing and matching with frictions. *Journal of Political Economy*, 109(5), 1060-1085.

Cagetti, M., & De Nardi M. (2006). Entrepreneurship, frictions, and wealth. *Journal of Political Economy*, 114(5): 835–870.

Chetty, R. (2008). Moral hazard versus liquidity and optimal unemployment insurance. *Journal of Political Economy*, 116(2), 173-234.

Chivers, D., Feng, Z., & Villamil, A. (2017). Employment-based health insurance and misallocation: Implications for the macroeconomy. *Review of Economic Dynamics*, 23, 125-149.

Chodorow-Reich, G., & Karabarbounis, L. (2016). The cyclicality of the opportunity cost of employment. *Journal of Political Economy*, 124(6), 1563-1618.

Chodorow-Reich, G., & Karabarbounis, L. (2017). The limited macroeconomic effects of unemployment benefit extensions. Working Paper 386166, Harvard University.

Coglianese, J. (2015). Do unemployment insurance extensions reduce employment?. Working Paper, Harvard University.

Conesa, J. C., Costa, D., Kamali, P., Kehoe, T. J., Nygard, V. M., Raveendranathan, G., & Saxena, A. (2018). Macroeconomic effects of Medicare. *The Journal of the Economics of Ageing*, 11, 27-40.

Conover, C. J., & Davies, H. H.. (2000). The role of TennCare in health policy for low-income people in Tennessee (Occasional Paper No. 33). Washington, D.C.: The Urban Institute.

Cooper, M., McClelland, J., Pearce, J., Prisinzano, R., Sullivan, J., Yagan, D., Zidar, O., & Zwick, E. (2016). Business in the United States: Who owns it, and how much tax do they pay?. *Tax Policy and the Economy*, 30(1): 91–128.

Dague, L., DeLeire, T., & Leininger, L. (2017). The effect of public insurance coverage for childless adults on labor supply. *American Economic Journal: Economic Policy*, 9(2), 124-54.

De Nardi, M., Doctor, P., & Karen, S. D. (2007). Evidence on entrepreneurs in the United States: Data from the 1989–2004 Survey of Consumer Finances. Federal Reserve Bank of Chicago Economic Perspectives, 4:18–36.

De Nardi, M., French, E., & Jones, J. B. (2016). Medicaid insurance in old age. *American Economic Review*, 106(11), 3480-3520.

den Haan, W. J., Ramey, G., & Watson, J. (2000). Job destruction and propagation of shocks. *American Economic Review*, 90(3), 482-498.

Dey, M. S., & Flinn, C. J. (2005). An equilibrium model of health insurance provision and wage determination. *Econometrica*, 73(2), 571-627.

Dey, M. S., & Flinn, C. J. (2008). Household search and health insurance coverage. *Journal of Econometrics*, 145(1-2), 43-63.

Dungan, A. (2017). Sole proprietorship returns, tax year 2015. Statistics of Income Bulletin, 37(2): 2–24.

Eeckhout, J., & Sepahsalari, A. (2015). Unemployment risk and the distribution of assets. Working Paper, University College London.

Elsby, M. W., Michaels, R., & Solon, G. (2009). The ins and outs of cyclical unemployment. *American Economic Journal: Macroeconomics*, 1(1), 84-110.

Engen, E. M., & Gruber, J. (2001). Unemployment insurance and precautionary saving. *Journal of Monetary Economics*, 47(3), 545-579.

Faig, M., & Zhang, M. & (2012). Labor market cycles, unemployment insurance eligibility, and moral hazard. *Review of Economic Dynamics*, 15(1), 41-56.

Farber, H. S., Rothstein, J., & Valletta, R. G. (2015). The effect of extended unemployment insurance benefits: Evidence from the 2012–2013 phase-out. *American Economic Review P&P*, 105(5), 171-176.

Farber, H. S., & Valletta, R. G. (2015). Do extended unemployment benefits lengthen unemployment spells? Evidence from recent cycles in the US labor market. *Journal of Human Resources*, 50(4), 873-909.

Floden, M. (2001). The effectiveness of government debt and transfers as insurance. *Journal of Monetary Economics*, 48(1), 81-108.

Fredriksson, P., & Holmlund, B. (2001). Optimal unemployment insurance in search equilibrium. *Journal of Labor Economics*, 19(2), 370-399.

French, E., & Jones, J. B. (2011). The effects of health insurance and self-insurance on retirement behavior. *Econometrica*, 79(3), 693-732.

Fujita, S. (2011a). Dynamics of worker flows and vacancies: evidence from the sign restriction approach. *Journal of Applied Econometrics*, 26(1), 89-121.

Fujita, S. (2011b). Effects of extended unemployment insurance benefits: evidence from the monthly CPS. Federal Reserve Bank of Philadelphia Working Paper 10-35/R.

Fujita, S., & Ramey, G. (2006). The cyclicality of job loss and hiring. Federal Reserve Bank of Philadelphia Working Paper 06-17.

Fujita, S., & Ramey, G. (2009). The cyclicality of separation and job finding rates. *International Economic Review*, 50(2), 415-430.

Fuller, D. L., Ravikumar, B., & Zhang, Y. (2013). Unemployment benefits: how much money goes unclaimed? Federal Reserve Bank of St. Louis Economic Synopses No. 2013-18.

Garthwaite, C., Gross, T., & Notowidigdo, M. J. (2014). Public health insurance, labor supply, and employment lock. *The Quarterly Journal of Economics*, 129(2), 653-696.

Gentry, W. M., & Hubbard, R. G. (2004). Entrepreneurship and household saving. Advances in Economic Analysis and Policy, 4(1), Article 1.

Gross, T., Notowidigdo, M.J., & Wang, J. (2016). The marginal propensity to consume over the business cycle. Working Paper, Northwestern University.

Gruber, J. (1997). The consumption smoothing benefits of unemployment insurance. *American Economic Review*, 87(1), 192-205.

Gruber, J., & Madrian, B. C. (1994). Health insurance and job mobility: The effects of public policy on job-lock. *ILR Review*, 48(1), 86-102.

Gurley-Calvez, T., Bruce, D., Reedy, E. J., & Russell, J. (2016). Comparing survey data and tax data: Differences in reporting across businesses. Working Paper, Statistics of Income.

Hagedorn, M., Karahan, F., Manovskii, I., & Mitman, K. (2016). Unemployment benefits and unemployment in the Great Recession: The role of macro effects. NBER Working Paper No. 19499.

Hagedorn, M., & Manovskii, I. (2008). The cyclical behavior of equilibrium unemployment and vacancies revisited. *American Economic Review*, 98(4), 1692-1706.

Hamilton, B. (2000). Does entrepreneurship pay? An empirical analysis of the returns to self-employment. *Journal of Political Economy*, 108(3): 604–631.

Hansen, G. D., Hsu, M., & Lee, J. (2014). Health insurance reform: The impact of a medicare buy-in. *Journal of Economic Dynamics and Control*, 45, 315-329.

Hansen, G. D., & Imrohoroğlu, A. (1992). The role of unemployment insurance in an economy with liquidity constraints and moral hazard. *Journal of Political Economy*, 100(1), 118-142.

Heathcote, J., Storesletten, K., & Violante, G. L. (2014). Consumption and labor supply with partial insurance: An analytical framework. *American Economic Review*, 104(7), 2075-2126.

Herkenhoff, K., (2017). The impact of consumer credit access on unemployment. Working Paper, University of Minnesota.

Herkenhoff, K., Phillips, G., and Cohen-Cole, E. (2017) How credit constraints impact job finding rates, sorting, and aggregate output. NBER Working Paper 22274.

Hopenhayn, H. A., & Nicolini, J. P. (1997). Optimal unemployment insurance. *Journal of Political Economy*, 105(2), 412-438.

Hornstein, A., Krusell, P., & Violante, G. L. (2011). Frictional wage dispersion in search models: A quantitative assessment. *American Economic Review*, 101(7), 2873-2898.

Hurst, E., Li, G., & Pugsley, B. W. (2014). Are household surveys like tax forms? Evidence from income underreporting of the self-employed. *The Review of Economics and Statistics*, 96(1): 19–33.

Hurst, E., & Pugsley, B. W. (2011). What do small businesses do?. Brookings Papers on Economic Activity, 43(2): 73–142.

Hurst, E., & Pugsley, B. W. (2017). Wealth, tastes, and entrepreneurial choice. In Measuring Entrepreneurial Businesses: Current Knowledge and Challenges, ed. J. Haltiwanger, E. Hurst, J. Miranda, and A. Schoar. Chicago: University of Chicago Press.

Jeske, K., & Kitao, S. (2009). US tax policy and health insurance demand: Can a regressive policy improve welfare?. *Journal of Monetary Economics*, 56(2), 210-221.

Johns, A., & Slemrod, J. (2010). The distribution of income tax noncompliance. *National Tax Journal*, 63(3): 397–418.

Johnston, A. C., & Mas, A. (2016). Potential unemployment insurance duration and labor supply: The individual and market-level response to a benefit cut. NBER Working Paper No. 22411.

Johnson, B. W., & Moore, K. (2008). Differences in income estimates derived from survey and tax data. In Proceedings of the Joint Statistical Meetings, Survey Research Methods Section. Alexandria: American Statistical Association.

Jung, J., & Tran, C. (2016a). Market inefficiency, insurance mandate and welfare: US health care reform 2010. *Review of Economic Dynamics*, 20, 132-159.

Jung, J., & Tran, C. (2016b). Social Health Insurance: A Quantitative Exploration. Working Paper.

Jung, P., & Kuester, K. (2015). Optimal labor-market policy in recessions. *American Economic Journal: Macroeconomics*, 7(2), 124-156.

Jung, P., & Kuhn, M. (2018). Earnings losses and labor mobility over the life cycle. *Journal of the European Economic Association*, doi: 10.1093/jeea/jvy014

Kaplan, G., & Violante, G. L. (2014). A model of the consumption response to fiscal stimulus payments. *Econometrica*, 82(4), 1199-1239.

Karahan, F., & Rhee, S. (2013). Geographical reallocation and unemployment during the great recession: The role of the housing bust.. New York Fed Staff Reports No. 605.

Kartashova, K. (2014). Private equity premium puzzle revisited. *American Economic Review*, 104(10): 3297–3334.

Kekre, R. (2017). Unemployment insurance in macroeconomic stabilization. Working Paper, University of Chicago.

Kitao, S. (2008). Entrepreneurship, taxation, and capital investment. *Review of Economic Dynamics*, 11(1): 44–69.

Koehne, S., & Kuhn, M. (2015). Should unemployment insurance be asset tested?. *Review of Economic Dynamics*, 18(3), 575-592.

Kroft, K., Lange, F., Notowidigdo, M. J., & Katz, L. F. (2016). Long-term unemployment and the Great Recession: the role of composition, duration dependence, and nonparticipation. *Journal of Labor Economics*, 34(S1), S7-S54.

Kroft, K., & Notowidigdo, M. J. (2016). Should unemployment insurance vary with the unemployment rate? Theory and evidence. *Review of Economic Studies*, 83(3), 1092-1124.

Krusell, P., Mukoyama, T., & Şahin, A. (2010). Labour-market matching with precautionary savings and aggregate fluctuations. *The Review of Economic Studies*, 77(4), 1477-1507.

Krusell, P., Mukoyama, T., Şahin, A., & Smith Jr., A. A. (2009). Revisiting the welfare effects of eliminating business cycles. *Review of Economic Dynamics*, 12(3), 393-404.

Kuhn, M. & Rios-Rull, J. V. (2016). 2013 update on the U.S. earnings, income, and wealth distributional facts: A view from macroeconomics. Federal Reserve Bank of Minneapolis Quarterly Review, 37(1): 2–73.

Landais, C., Michaillat, P., & Saez, E. (2018). A macroeconomic approach to optimal unemployment insurance: Applications. *American Economic Journal: Economic Policy*, 10(2), 182-216.

Lentz, R. (2009). Optimal unemployment insurance in an estimated job search model with savings. *Review of Economic Dynamics*, 12(1), 37-57.57

Lise, J., Meghir, C., & Robin, J. M. (2016). Matching, sorting and wages. *Review of Economic Dynamics*, 19, 63-87.

Lise, J., & Robin, J.-M. (2017). The macrodynamics of sorting between workers and firms. *American Economic Review*, 107(4), 1104-1135.

Ljungqvist, L., & Sargent, T. J. (1998). The European unemployment dilemma. *Journal of Political Economy*, 106(3), 514-550.

Lucas, Jr., R. E. (1987). Models of business cycles. Basil Blackwell, New York.

Madrian, B. C. (1994). Employment-based health insurance and job mobility: Is there evidence of job-lock?. *The Quarterly Journal of Economics*, 109(1), 27-54.

Marinescu, I. (2017). The general equilibrium impacts of unemployment insurance: Evidence from a large online job board. *Journal of Public Economics*, 150, 14-29.

McGrattan, E. R., & Prescott, E. C. (2010a). Unmeasured investment and the puzzling U.S. boom in the 1990s. *American Economic Journal: Macroeconomics*, 2(4): 88–123.

McGrattan, E. R., & Prescott, E. C. (2010b). Technology capital and the U.S. current account, *American Economic Review*, 100(4): 1493–1522.

McGrattan, E. R., & Prescott, E. C. (2017). On financing retirement with an aging population. *Quantitative Economics*, 8(1), 75-115.

Meh, C. (2005), Entrepreneurship, wealth inequality, and taxation. *Review of Economic Dynamics*, 8(3): 688–719.

Menzio, G., & Shi, S. (2010). Block recursive equilibria for stochastic models of search on the job. *Journal of Economic Theory*, 145(4), 1453-1494.

Menzio, G., & Shi, S. (2011). Efficient search on the job and the business cycle. *Journal of Political Economy*, 119(3), 468-510.

Menzio, G., Telyukova, I. A., & Visschers, L. (2016). Directed search over the life cycle. *Review of Economic Dynamics*, 19, 38-62.

Mitman, K., & Rabinovich, S. (2014). Do unemployment benefits explain the emergence of jobless recoveries? Working Paper, Stockholm University.

Mitman, K., & Rabinovich, S. (2015). Optimal unemployment insurance in an equilibrium business-cycle model. *Journal of Monetary Economics*, 71, 99-118.

Moskowitz, T. J., & Vissing-Jorgensen, A. (2002). The returns to entrepreneurial investment: A private equity premium puzzle?. *American Economic Review*, 92(4): 745–778.

Nakajima, M. (2012). A quantitative analysis of unemployment benefit extensions. *Journal of Monetary Economics*, 59(7), 686-702.

Nakajima, M., & Tuzemen, D. (2016). Health-care reform or labor market reform? A quantitative analysis of the affordable care act. Working Paper.

Parker, J.A., Souleles, N.S., Johnson, D.S, & McClelland, R. (2013). Consumer spending and the economic stimulus payments of 2008. *American Economic Review*, 103(6), 2530-2553.

Pashchenko, S., & Porapakkarm, P. (2013). Quantitative analysis of health insurance reform: Separating regulation from redistribution. *Review of Economic Dynamics*, 16(3), 383-404.

Pei, Y., & Xie, Z. (2016). A quantitative theory of time-consistent unemployment insurance. Federal Reserve Bank of Atlanta Working Paper 2016-11.

Quadrini, V. (2000). Entrepreneurship, saving, and social mobility. *Review of Economic Dynamics*, 3(1): 1–40.

Rothstein, J. (2011). Unemployment insurance and job search in the Great Recession. *Brookings Papers on Economic Activity*, 143-214.

Rothstein, J., and Valletta, R.G. (2017). Scraping by: Income and program participation after the loss of extended unemployment benefits. NBER Working Paper No. 23528.

Saez, E. & Zucman, G. (2016). Wealth inequality in the United States since 1913: Evidence from capitalized income tax data. *The Quarterly Journal of Economics*, 131(2): 519–578.

Saporta-Eksten, I. (2014). Job loss, consumption, and unemployment insurance. Working Paper, Stanford University.

Scheuer, F. (2013). Entrepreneurial taxation with endogenous entry. NBER Working Paper 19235.

Selden, T. M., Levit, K. R., Cohen, J. W., Zuvekas, S. H., Moeller, J. F., McKusick, D., & Arnett III, R. H. (2001). Reconciling medical expenditure estimates from the MEPS and the NHA, 1996. *Health Care Financing Review*, 23(1), 161.

Shavell, S., & Weiss, L. (1979). The optimal payment of unemployment insurance benefits over time. *Journal of Political Economy*, 87(6), 1347-1362.

Shi, S. (2001). Frictional assignment. I. Efficiency. *Journal of Economic Theory*, 98(2), 232-260.

Shimer, R. (2005). The cyclical behavior of equilibrium unemployment and vacancies. *American Economic Review*, 95(1), 25-49.

Shimer, R., & Werning, I. (2008). Liquidity and insurance for the unemployed. *American Economic Review*, 98(5), 1922-1942.

Smith, M., Yagan, D., Zidar, O., & Zwick, E. (2017). Capitalists in the twenty-first century. Working Paper, U.S. Department of the Treasury.

Stephens Jr., M. (2004). Job loss expectations, realizations, and household consumption behavior. *Review of Economics and Statistics*, 86(1), 253-269.

Tsujiyama, H. (2015). Implications of health care reform for inequality and welfare, Goethe University Working Paper.

Valletta, R., & Kuang, K. (2010). Extended unemployment and UI benefits. Federal Reserve Bank of San Francisco Economic Letter No. 2010-12.

Vroman, W. (1991). The decline in unemployment insurance claims activity in the 1980s. Unemployment Insurance Occasional Paper 91–2, Washington, DC: U.S. Department of Labor-ETA.

Wang, C., & Williamson, S. D. (2002). Moral hazard, optimal unemployment insurance, and experience rating. *Journal of Monetary Economics*, 49(7), 1337-1371.

Wooldridge, J., Ku, L., Coughlin, T., Dubay, L., Ellwood, M., Rajan, S., & Hoag, S. (1996). Implementing state health care reform: What have we learned from the first year? The first annual report of the evaluation of health reform in five states. Mathematica Policy Research.

Yashiv, E. (2007). US labor market dynamics revisited. *Scandinavian Journal of Economics*, 109(4), 779-806.

Appendix A

Appendix to Chapter 1

A.1 Data

SIPP

I use the U.S. Census Bureau's Survey of Income and Program Participation (SIPP) to document information on wealth holdings, income, health insurance coverage, as well as employment transitions. The SIPP is a longitudinal survey that follows individuals for a duration of up to five years, with interviews being held in four-month intervals called waves. Each respondent is then assigned to one of four rotation groups. The rotation group determines which month within a wave a respondent is interviewed. Each interview covers information about the four months (reference months) preceding the interview month. For example, when a new SIPP panel starts and Wave 1 (the first four months of the new panel) commences, the first rotation group is interviewed in the first month of Wave 1, the second rotation group is interviewed in the second month of Wave 1, and so on. Once all four rotation groups are interviewed at the end of the fourth month of Wave 1, Wave 2 begins with the second interview of the first rotation group. This way, all four rotation groups, and thus all respondents, will have been interviewed at the end of each wave.

In each interview, respondents are asked questions about their income, labor force status, health insurance information, as well as government transfer receipts over the previous four months not including the interview month. The SIPP also contains information on the asset holdings of the respondent. In each SIPP panel, respondents provide information on various types of asset holdings at two or three waves of the panel, usually one year or, equivalently, three waves apart. For documenting wealth holdings, I use the 2004 panel of the SIPP, which contains 12 waves covering information between January 2004 and December 2007. This particular panel allows us to observe data on asset holdings at Waves 3 and 6.

Sample selection I restrict my sample to individuals between the ages of 25 and 64 who do not own a business or derive income from self-employment. Business ownership is identified using variable EBUSCNTR which asks about the number of businesses owned during the reference period. Self-employment earnings is identified using variable TBM-SUM which asks respondents to report the amount of income received from any business during the reference period.

Labor market transitions by health status In this section, I discuss the details of calculating the age profile of quarterly job-finding rates by health insurance coverage presented in the validation section using the SIPP 1996-2008 Panels.

The SIPP collects weekly information on employment status. Since I am interested in quarterly job-finding rates, I convert the data into a quarterly frequency, with the employment status for any given quarter set as the employment status reported during the first week of that quarter. A respondent is considered employed if he/she reports having a job and either working or not on layoff, but absent without pay and as unemployed (U) if he/she reports either having no job and actively looking for work or having a job but currently laid off. An individual experiences a UE transition in a given month if he/ is unemployed at the beginning of the month and employed at the beginning of next month.

I then classify the unemployed as either with health insurance or without health insurance. Doing so requires keeping track of unemployment spells. I restrict the sample of unemployed to those whom I observe both their entry into unemployment, exit back into employment, as well as all periods of unemployment in between. I then further restrict unemployment spells to those that commence because of firm distress (i.e. caused by a layoff, employer bankruptcy, sale of employer's business, or slack work/business conditions). The SIPP also contains information on respondents' monthly insurance status as well as the source of coverage.[1] I consider the respondent uninsured during the reference period if he/she indicates *not* being covered by either Medicare, Medicaid, military-related health insurance, or any type of private health insurance. Individuals are considered to be without coverage during their unemployment spell if the are uninsured for more than 50 percent of their unemployment duration.

The job finding rate for any given age is then calculated by calculating the fraction of age t unemployed individuals who transition into employment in the next quarter.

Asset to income distribution In this section, I describe in detail the calculation of specific moments of the asset-to-income distribution. These moments are used as both calibration targets (median asset to quarterly labor income, fraction with non-positive wealth) and well as moments to validate the model against (percentiles).

I focus on assets that can be used to both insure against income risk and fund retirement consumption. More formally, this encompasses net financial assets, net equity in vehicles, as well as retirement accounts.

The SIPP contains individual level data on financial liquid assets such as interest-earning financial assets in banking and other instutitions, amount in non-interest-earning checking accounts, equity in stocks and mutual funds, and face value of U.S. savings bonds. Moreover, for married individuals, the survey asks about the amount of these assets in joint accounts. Only one spouse is asked about joint accounts; the response is then divided by two, and the divided amount is copied to both spouses' records. The SIPP also contains information about revolving debt on credit card balances at the individual level for both single and joint accounts in the same fashion. The summation of the amounts in liquid asset accounts net of revolving debt gives us the net financial asset holdings of the individual. The SIPP also provides data on equity in cars at the household level. I split that amount between the members of the household who are age 16 or older, and record that value as the amount of equity in cars for each individual within the household. Finally, the SIPP also collects information on the market value of retirement savings instruments in the form of a 401k, 403b, thrift plan, Individual Retirement Account (IRA), or KEOGH account. Respondents are asked to provide information about plans which are only under

[1] EHIMTH asks respondents if they have private health insurnace coverage during the month; ECDMTH, Medicaid; ECRMTH, Medicare; and RCHAMPM, military-related health insurance.

their name. My measure of wealth is thus the sum of net financial asset holdings, net equity in vehicles, and the market value of retirement accounts.[2]

The SIPP also provides information about the monthly gross job earnings for each individual. I use this information to determine the monthly gross labor income of the individual. If the individual is unemployed during the interview month, I use his/her gross labor income associated with the last employment from earlier waves. The ratio of these wealth and income thus measure the number of months an indivual's wealth holdings can replace earnings in the event of job loss. The asset-to-income distribution reported in Table 2.2 reports this measure in terms of quarters.

Joint distribution of wages and employer provided health insurance (EPHI)

The model is calibrated to match the joint distribution of wages and take-up of EPHI calculated using the SIPP 2004 Panel. To calculate this, I first restrict the sample to employed respondents and subsequently group them into wage quintiles. For each wage quintile, I calculate the fraction of employed workers who report being under a private health insurance plan and specifying that it is provided by a current employer.[3] This distribution is reported in Table 1.1 as the average across all months in the SIPP 2004 Panel.

---

[2] Asset holdings are calculated as follows by using the following variables in SIPP data: Net financial assets = TALICHA+TALJCHA+TALSBV+TIMIA+TIMJA+TIAITA+TIAJTA+ESMIV+ESMJV-(EALIDAB+EALJDAB) where TALICHA (TALJCHA) is the amount of non-interest-earning checking accounts in own name (joint account), TALSBV is the face value of U.S. savings bonds, TIMIA (TIMJA) is amount of bonds/securities in own name (joint account), TIAITA (TIAJTA) is the amount in interest earning account in own name (joint account), ESMIV (ESMJV) value of stocks/funds in own name (joint account), and EALIDAB (EALJDAB) amount owed for store bills/credit cards in own name (joint account). Then, net equity in vehicles of the household is given by THHVEHCL. I divide this value among the members of the household above age 16. Thus, I get the net liquid asset holdings of the individual as follows: Net liquid assets = Net financial assets +THHVEHCL /number of persons within the household age 16 and above. Finally, Retirement assets= TALTB + TALRB + TALKB

[3] EHIMTH asks respondents if they have private health insurnace coverage during the month. Respondents who indicate having private health insurnace in any given month during a reference period are then asked the source of health insurance by EHEMPLY.

CPS and NBER Taxsim

I estimate the parameters of the piecewise linear tax function in Equation 1.11 using the 2007 CPS March Supplement and the NBER's tax simulation program (TAXSIM). TAXSIM requires users to input cross-sectional data on marital status, dependents, state of residence, and various sources of income among others. It then provides estimates of average and marginal tax rates.

I use information on the filing status of each respondents in the CPS to determine whether a person or married couple is filing separately or filing jointly. In cases when household members file separately, I assume that the head of the household claims all dependents for personal exemptions.

Taxable household income comprises of wages and salaries; income from self employment (business or farm); income from interest, dividends, rent, and capital gains; income from alimony; taxable pensions, social security benefits, and unemployment insurance receipts. There are four main deductible expenses in the U.S. tax code, namely: medical expenses, state taxes paid, mortgage interest, and charitable contributions. I use CPS data on medical expenses and state tax liability (after deductions) for the first two categories; due to data limitations, I set mortgage income and charitable contributions to zero.

After obtaining average and marginal tax rates for each household in the CPS sample using TAXSIM, I compute bracket-specific income weighted average- and marginal tax rates.

A.2 Computation of Transitional Dynamics

Preliminaries

Before discussing the details of the computational algorithm, I will first outline some notation. Let $P \in \{b, n\}$ denote the policy being implemented, where b is the baseline policy and n is the new (universal) policy. The transition begins at period $j = 0$ and ends at period $j = J^{tr}$. Let $V_t^P(s)$ and $d_t^P(s)$ denote the value functions and policy functions of an agent with age t and state s when the economy is in a steady state under policy P.[4] Likewise, denote $V_{t,j}^P(s)$ and $d_{t,j}^P(s)$ denote the value functions and policy functions of an

[4]For simplicity, the value function V and policy function d is used as a stand-in for the several value functions and policy functions of households and firms alike in the model.

agent when the economy is under policy P during period j of the transition.[5] Let $\{r_j, \theta_j\}$ denote the interest rate and market tightness of the economy in transition period j.

The economy begins with stationary distribution μ^b under the baseline economy at transition period $j = 0$. For transition period $j > 0$, a universal healthcare system is introduced in the following manner. First, the Medicaid asset-test a^{mc} and income-test z^{mc} are abolished while non-financial eligibility probability γ^n is set to 1. This implies that households in the economy are unexpectedly and permanently enrolled in the government health insurance program. Second, the firms no longer offer health insurance and $w_h = 0$. Finally, the additional payroll tax is phased-in, following a linear slope. In particular, $T_j^p(z) = \left(\beta_i^{Tax} + \xi^n \times \frac{j}{J^{tr}}\right) z + \alpha_i^{Tax}$ for all brackets i, where J^{tr} is the length of the transition and ξ^n is the additional payroll tax required to finance the reform.

After having solved the stationary equilibrium under the baseline economy and the economy with universal healthcare, the computational algorithm outlined below is used to solve for transitional dynamics.

Computational algorithm

1. Guess a path of interest rates $\left\{r_j^{old}\right\}_{j=0}^{j=J^{tr}}$ and impose that $r_0 = r^b$ and $r_{J^{tr}} = r^n$. Similarly, guess a path of market tightness $\left\{\theta_j^{old}\right\}_{j=0}^{j=J^{tr}}$ and impose that $\theta_0 = \theta^b$ and $\theta_{J^{tr}} = \theta^n$.

2. Starting from value functions $V_{t,J^{tr}}^n(s) = V_t^n(s)$, compute for the sequence of value functions and policy functions $\left\{V_{t,j}^n(s), d_{t,j}^n(s)\right\}_{j=1}^{j=J^{tr}-1}$ taking as given the path of prices and government taxes.

3. Using the sequence of policy functions $\left\{d_{t,j}^n(s)\right\}_{j=1}^{j=J^{tr}-1}$ and distribution of agents across states $\left\{\mu_j^n\right\}_{j=1}^{j=J^{tr}-1}$, calculate capital supply $K_{s,j}$ and demand $K_{d,j}$ for $j \in \left\{1, \ldots, J^{tr} - 1\right\}$. Let

$$\epsilon_{r,j} = K_{s,j} - K_{d,j}$$

[5] The same notation is used for the value of a posting a vacancy and the value of a firm offering a or matched with a worker.

4. Using the sequence of job acceptance decisions rules, firm value functions, and distribution of agents across states for all $j \in \left\{1, \ldots, J^{tr}-1\right\}$, check if the free entry condition in Equation 1.9 holds $\forall j$. Let

$$\epsilon_{\theta,j} = \int\limits_{y} \left[q\left(\theta_j\right) \frac{1}{1+r} \int\limits_{s_U} d_{a,t,j}\left(s_W\right) \bar{J}_{t,j}\left(s_U, y\right) d\mu_j^U\left(s_U\right) \right] d\mu^J\left(y\right) - \kappa$$

5. (Updating interest rates) To update the guess of the path of interest rates, for each period j, compute $K_j^{new} = \chi_k K_{d,j} + \left(1-\chi_k\right) K_{s,j}$, where $0 < \chi_k < 1$ is an updating weight that aids in convergence. Then, set the new guess for interest rate to be r_j^{new} such that

$$K_j^{new} = \int k^*\left(x, y; r_j^{new}\right) \mu_j^J\left(x, y\right)$$

is satisfied, where $k^*\left(x, y; r_j^{new}\right) = \left[\frac{(1-\omega)\psi A(x,y)}{r_j^{new}}\right]^{\frac{1}{1-\psi}}$ satisfies firm first-order conditions, μ_j^J is the distribution of firm-worker pairs (indexed by their skill x and productivity y) in transition period j . This procedure yields an update of interest rates $\left\{r_j^{new}\right\}_{j=0}^{j=J^{tr}}$

6. (Updating market tightness) Set $\theta_j^{new} = \chi_\theta \hat{\theta}_j + \left(1-\chi_\theta\right) \theta_j^{old}$, where $0 < \chi_\theta < 1$. Here, $\hat{\theta}_j$ is such that the value of a vacancy in period j given by Equation 1.9 is 0. This procedure yields an update of market tightness $\left\{\theta_j^{new}\right\}_{j=0}^{j=J^{tr}}$.

7. Using convergence criterion $\underline{\epsilon}$, check if $\max_{1\leq j\leq J^{tr}} \left|\epsilon_{\theta,j}\right| < \underline{\epsilon}$ and $\max_{1\leq j\leq J^{tr}} \left|\epsilon_{r,j}\right| < \underline{\epsilon}$. Otherwise, update prices $\left\{r_j^{old}, \theta_j^{old}\right\}_{j=0}^{j=J^{tr}} = \left\{r_j^{new}, \theta_j^{new}\right\}_{j=0}^{j=J^{tr}}$ and iterate until asset markets clear and the free entry condition is satisfied for all periods in the transition.

Appendix B

Appendix to Chapter 2

B.1 Data

SIPP data

We use the U.S. Census Bureau's Survey of Income and Program Participation (SIPP) to document the liquid asset holdings of individuals. The SIPP is a longitudinal survey that follows individuals for a duration of up to five years, with interviews being held in four-month intervals called waves. Each respondent is then assigned to one of four rotation groups. The rotation group determines which month within a wave a respondent is interviewed. Each interview covers information about the four months (reference months) preceding the interview month. For example, when a new SIPP panel starts and Wave 1 (the first four months of the new panel) commences, the first rotation group is interviewed in the first month of Wave 1, the second rotation group is interviewed in the second month of Wave 1, and so on. Once all four rotation groups are interviewed at the end of the fourth month of Wave 1, Wave 2 begins with the second interview of the first rotation group. This way, all four rotation groups, and thus all respondents, will have been interviewed at the end of each wave.

In each interview, respondents are asked questions about their income, labor force status and government transfer receipts over the previous four months not including the interview month. In the end, the SIPP provides monthly data on income and government transfers and weekly data on labor force status. Most importantly, the SIPP also contains data on the asset holdings of the respondent. In each SIPP panel, respondents provide information on various types of asset holdings at two or three waves of the panel, usually one year or, equivalently, three waves apart. We use the 2004 panel of the SIPP, which contains 12 waves covering information between January 2004 and December 2007. This particular panel allows us to observe data on asset holdings at Waves 3 and 6. Since it is the closest date to the Great Recession, we calculate the asset distribution using Wave 6.

Asset distribution

We focus on the liquid asset holdings of individuals. The SIPP contains individual level data on financial liquid assets such as interest-earning financial assets in banking and other instutitions, amount in non-interest-earning checking accounts, equity in stocks and mutual funds, and face value of U.S. savings bonds. Moreover, for married individuals, the survey asks about the amount of these assets in joint accounts. Only one spouse is asked about joint accounts; the response is then divided by two, and the divided amount is copied to both spouses' records. The SIPP also contains information about revolving debt on credit card balances at the individual level for both single and joint accounts in the same fashion. The summation of the amounts in liquid asset accounts net of revolving debt gives us the net financial asset holdings of the individual. Finally, the SIPP provides data on equity in cars at the household level. We split that amount between the members of the household who are age 16 or older, and record that value as the amount of equity in cars for each individual within the household. Adding this value to net financial asset holdings of the individual gives us the measure of liquid asset holdings for each individual.[1]

[1] Net financial asset holdings are calculated as follows by using the following variables in SIPP data: Net financial assets = TALICHA+TALJCHA+TALSBV+TIMIA+TIMJA+TIAITA+TIAJTA+ESMIV+ESMJV-(EALIDAB+EALJDAB) where TALICHA (TALJCHA) is the amount of non-interest-earning checking accounts in own name (joint account), TALSBV is the face value of U.S. savings bonds, TIMIA (TIMJA) is amount of bonds/securities in own name (joint account), TIAITA (TIAJTA) is the amount in interest earning account in own name (joint account), ESMIV (ESMJV) value of stocks/funds in own name (joint account), and EALIDAB (EALJDAB) amount owed for store bills/credit cards in own name (joint account). Then, net equity in vehicles of the household is given by THHVEHCL. We divide this value among the members of the household above age 16. Thus, we get the net liquid asset holdings of the individual as follows: Net liquid assets = Net financial assets +THHVEHCL /number of persons within

The SIPP also provides information about the monthly gross job earnings for each individual. We use this information to determine the monthly gross labor income of the individual. If the individual is unemployed during the interview month, we use her gross labor income associated with the last employment from earlier waves. Next, using the weekly employment status of the individual for that month, we calculate the weekly gross labor income of the individual by dividing monthly gross labor income by the number of weeks with a job during the interview month.

We then calculate annual income and payroll tax rates using the statutory U.S. income tax codes in the following steps. First, we calculate the annual income of each individual. Annual income includes labor income, capital income, and all kinds of government transfers including UI received in the fiscal year. Next, we apply the year-specific federal income tax schedule to the annual income net of year-specific personal exemptions and deductions to obtain the total annual income tax for each respondent. After that, we calculate the total annual payroll tax (Social Security and Medicare tax) for each individual. We obtain the total annual payroll tax for each individual by applying the year-specific Social Security and Medicare tax schedule to the total annual labor income of the individual for the time period. [2] Then, our measure for the tax rate is

$$\tau = \frac{\text{Share of labor income} \times \text{Annual income tax} + \text{Annual payroll tax}}{\text{Annual labor income}},$$

where the share of labor income is the ratio of annual labor income to annual income. We then apply the tax rate τ for each individual in our sample and obtain weekly after-tax labor income. Last, dividing the liquid asset holdings measure to weekly after-tax labor income gives us the asset-to-income ratio for each individual.

B.2 Proofs

Opportunity cost of employment

In this section, we show the derivations of Equations (2.12) and (2.13) in the main text.

the household age 16 and above.

[2] We also consider the fact that there is a maximum taxable annual labor income for Social Security tax, while Medicare tax does not have such a limit. As a result, we get total annual tax as the sum of total annual income and payroll taxes.

First, substituting (2.1) and (2.2) into (2.8), we have

$$
\begin{aligned}
S^{UE}\left(a, w^{UE}, \beta; p\right) =& V^{W}\left(a, \tilde{w}\left(\cdot\right), \beta; p\right) - V^{UE}\left(a, w^{UE}, \beta; p\right) \\
=& u\left(c^{W}\right) - u\left(c^{UE}\right) + \nu\left(s\right) \\
&+\beta\mathbb{E}\left[\delta\left(p'\right)\left(1 - e\left(p'\right)\right)V^{UE}\left(a'^{W}, \tilde{w}\left(\cdot\right), \beta'; p'\right)\right. \\
&\left.+\,\delta\left(p'\right)e\left(p'\right)V^{UI}\left(a'^{W}, \beta'; p'\right)\right] \\
&+\beta\mathbb{E}\left[\left(1 - \delta\left(p'\right)\right)V^{W}\left(a'^{W}, \tilde{w}\left(\cdot\right), \beta'; p'\right)\right] \\
&-\beta\mathbb{E}\left[sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)V^{W}\left(a'^{UE}, \tilde{w}\left(\cdot\right), \beta'; p'\right)\right] \\
&-\beta\mathbb{E}\left[\left(1 - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right)\left(1 - e\left(p'\right)\right)V^{UE}\left(a'^{UE}, w^{UE}, \beta'; p'\right)\right] \\
&-\beta\mathbb{E}\left[\left(1 - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right)e\left(p'\right)V^{UI}\left(a'^{UE}, \beta'; p'\right)\right]
\end{aligned}
$$

In order to obtain (2.12), we add and subtract terms, rearrange them, then use (2.10), and divide both sides by λ^{W}. This yields

$$
\begin{aligned}
\frac{S^{UE}\left(a, w^{UE}, \beta; p\right)}{\lambda^{W}} =& \frac{u\left(c^{W}\right) - u\left(c^{UE}\right) + \nu\left(s\right)}{\lambda^{W}} \\
&+\frac{\beta}{\lambda^{W}}\mathbb{E}\Big[sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right) \\
&\times\left(V^{W}\left(a'^{W}, \tilde{w}\left(\cdot\right), \beta'; p'\right) - V^{W}\left(a'^{UE}, \tilde{w}\left(\cdot\right), \beta'; p'\right)\right)\Big] \\
&+\frac{\beta}{\lambda^{W}}\mathbb{E}\Big[\left(1 - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right) - \delta\left(p'\right)e\left(p'\right)\right) \\
&\times\left(V^{UE}\left(a'^{W}, \tilde{w}\left(\cdot\right), \beta'; p'\right) - V^{UE}\left(a'^{UE}, w^{UE}, \beta'; p'\right)\right)\Big] \\
&+\frac{\beta}{\lambda^{W}}\mathbb{E}\Big[\left(1 - \delta\left(p'\right) - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right) \\
&\times\left(V^{W}\left(a'^{W}, \tilde{w}\left(\cdot\right), \beta'; p'\right) - V^{W}\left(a'^{W}, \tilde{w}\left(a'^{W}, \tilde{w}\left(\cdot\right), \beta'; p'\right), \beta'; p'\right)\right)\Big] \\
&+\frac{\beta}{\lambda^{W}}\mathbb{E}\Big[\left(1 - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right) - \delta\left(p'\right)\right)e\left(p'\right) \\
&\times\left(V^{UE}\left(a'^{UE}, w^{UE}, \beta'; p'\right) - V^{UI}\left(a'^{UE}, \beta'; p'\right)\right)\Big] \\
&+\beta\mathbb{E}\left[\frac{\lambda'^{W}}{\lambda^{W}}\frac{\left(1 - \delta\left(p'\right) - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right)S\left(a'^{W}, \tilde{w}\left(\cdot\right), \beta'; p'\right)}{\lambda'^{W}}\right]
\end{aligned}
$$

where the summation of the second and third terms on the right-hand side is $-z_{a}^{UE}$, the fourth term is $-z_{w}^{UE}$, and the fifth term is $-z_{elg}^{UE}$. Given the form of the utility function, we cannot isolate z^{UE}_{flow} from the first term on the right-hand side. However, since we know that the flow value of employment is $\tilde{w}\left(a, w^{UE}, \beta; p\right) \times (1-\tau)$, we can numerically calculate z^{UE}_{flow} using the above equation as follows:

$$\begin{aligned} z^{UE}_{flow} = & \frac{S^{UE}\left(a, w^{UE}, \beta; p\right)}{\lambda^W} - \tilde{w}\left(a, w^{UE}, \beta; p\right) \times (1-\tau) + z^{UE}_a + z^{UE}_w + z^{UE}_{elg} \\ & - \beta \mathbb{E}\left[\frac{\lambda'^W}{\lambda^W} \frac{\left(1-\delta\left(p'\right) - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right) S\left(a'^W, \tilde{w}\left(\cdot\right), \beta'; p'\right)}{\lambda'^W}\right]. \end{aligned}$$

This gives us the opportunity cost of employment for the eligible unemployed $z^{UE} = z^{UE}_{flow} + z^{UE}_a + z^{UE}_w + z^{UE}_{elg}$ for each state $\left(a, w^{UE}, \beta; p\right)$.
Second, substituting (2.1) and (2.3) into (2.9), we have

$$\begin{aligned} S^{UI}\left(a, \beta; p\right) & = V^W\left(a, \tilde{w}\left(\cdot\right), \beta; p\right) - V^{UI}\left(a, \beta; p\right) \\ & = u\left(c^W\right) - u\left(c^{UI}\right) + \nu\left(s\right) \\ & + \beta \mathbb{E}\left[\delta\left(p'\right)\left[\left(1-e\left(p'\right)\right) V^{UE}\left(a'^W, \tilde{w}\left(\cdot\right), \beta'; p'\right) + e\left(p'\right) V^{UI}\left(a'^W, \beta'; p'\right)\right]\right] \\ & + \beta \mathbb{E}\left[\left(1-\delta\left(p'\right)\right) V^W\left(a'^W, \tilde{w}\left(\cdot\right), \beta'; p'\right)\right] \\ & - \beta \mathbb{E}\left[sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right) V^W\left(a'^{UI}, \tilde{w}\left(\cdot\right), \beta'; p'\right)\right] \\ & - \beta \mathbb{E}\left[\left(1-sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right) V^{UI}\left(a'^{UI}, \beta'; p'\right)\right] \end{aligned}$$

Similarly, in order to obtain (2.13), we again add and subtract terms, rearrange them, then use (2.11), and divide both sides by λ^W. This yields

$$\begin{aligned} \frac{S^{UI}\left(a, \beta; p\right)}{\lambda^W} & = \frac{u\left(c^W\right) - u\left(c^{UI}\right) + \nu\left(s\right)}{\lambda^W} \\ & + \frac{\beta}{\lambda^W} \mathbb{E}\left[sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\left[V^W\left(a'^W, \tilde{w}\left(\cdot\right), \beta'; p'\right) - V^W\left(a'^{UI}, \tilde{w}\left(\cdot\right), \beta'; p'\right)\right]\right] \\ & + \frac{\beta}{\lambda^W} \mathbb{E}\left[\left(1-sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right)\left[V^{UI}\left(a'^W, \beta'; p'\right) - V^{UI}\left(a'^{UI}, \beta'; p'\right)\right]\right] \\ & + \frac{\beta}{\lambda^W} \mathbb{E}\Big[\left(1-\delta\left(p'\right) - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right) \\ & \times \left(V^W\left(a'^W, \tilde{w}\left(\cdot\right), \beta'; p'\right) - V^W\left(a'^W, \tilde{w}\left(\cdot\right), \beta'; p'\right)\right)\Big] \\ & + \frac{\beta}{\lambda^W} \mathbb{E}\left[\delta\left(p'\right)\left(1-e\left(p'\right)\right)\left[V^{UE}\left(a'^W, \tilde{w}\left(\cdot\right), \beta'; p'\right) - V^{UI}\left(a'^W, \beta'; p'\right)\right]\right] \\ & + \beta \mathbb{E}\left[\frac{\lambda'^W}{\lambda^W} \frac{\left(1-\delta\left(p'\right) - sf\left(\theta\left(\tilde{w}\left(\cdot\right); p'\right)\right)\right) S^{UI}\left(a'^W, \beta'; p'\right)}{\lambda'^W}\right], \end{aligned}$$

where the summation of the second and third terms on the right-hand side is $-z_a^{UI}$, the fourth term is $-z_w^{UI}$, and the fifth term is $-z_{elg}^{UI}$. Similarly, we numerically calculate z_{flow}^{UI} as follows:

$$\begin{aligned} z_{flow}^{UI} = & \frac{S^{UI}(a,\beta;p)}{\lambda^W} - \tilde{w}(a,\beta;p) \times (1-\tau) + z_a^{UI} + z_w^{UI} + z_{elg}^{UI} \\ & - \beta \mathbb{E}\left[\frac{\lambda'^W}{\lambda^W} \frac{\left(1-\delta(p') - sf\left(\theta\left(\tilde{w}\left(a'^{UI},\beta';p'\right);p'\right)\right)\right) S^{UI}\left(a'^W,\beta';p'\right)}{\lambda'^W} \right]. \end{aligned}$$

This gives us the opportunity cost of employment for the eligible unemployed $z^{UI} = z_{flow}^{UI} + z_a^{UI} + z_w^{UI} + z_{elg}^{UI}$ for each state $(a,\beta;p)$.[3]

Block recursive equilibrium

Proposition 1: *If i) utility function $u(\cdot)$ is strictly increasing, strictly concave, and satisfies Inada conditions; $\nu(\cdot)$ is strictly increasing and strictly convex, ii) choice sets $\mathcal{W}$ and $\mathcal{A}$, and sets of exogenous processes $\mathcal{P}$ and $\mathcal{B}$ are bounded, iii) matching function M exhibits constant returns to scale, and iv) UI policy is restricted to be only a function of current aggregate labor productivity, then there exists a Block Recursive Equilibrium for this economy. If, in addition, $M = min\{v, S\}$, then the Block Recursive Equilibrium is the only recursive equilibrium.*

Proof: The proof presented here follows from Karahan and Rhee (2013) and Herkenhoff (2017), which are extensions of Menzio and Shi (2010, 2011). We extend the former's proof to a model in which government finances the time-varying UI benefits and show that the model still admits block recursivity. We then use the model to study how UI policy must vary over the business cycle. In doing so, the additional assumption we make here is to restrict the class of UI policies to be a function of current aggregate labor productivity.

Existence: We prove the existence of the BRE in two steps. In the first step, we show that the firm value functions and the corresponding labor market tightness depend on the aggregate state of the economy only through the current aggregate labor productivity. Then, in the second step, given that UI policy instruments are restricted to be a function

[3] In this numerical calculation, we calculate the opportunity cost under fixed wages and disregard z_w^{UE} and z_w^{UI}.

of the current aggregate labor productivity, we show that the household value functions do not depend on the aggregate distribution of agents across states. As a result, we show that given the UI policy, the solution of the household's problem together with the solution of the firm's problem and labor market tightness, constitute a block recursive equilibrium.
Let $\mathcal{J}(\mathcal{W},\mathcal{P})$ be the set of bounded and continuous functions J such that $J : \mathcal{W}\times\mathcal{P}\rightarrow\mathbb{R}$ and let T_J be an operator associated with (2.4) such that $T_J : \mathcal{J}\rightarrow\mathcal{J}$. Then, using Blackwell's sufficiency conditions for a contraction and the assumptions of the boundedness of sets of exogenous processes $\mathcal{P}$ and $\mathcal{B}$, and choice sets $\mathcal{W}$ and $\mathcal{A}$, we can show that T_J is a contraction and has a unique fixed point $J^*\in\mathcal{J}$. Thus, the firm value function satisfying (2.4) depends on the aggregate state of the economy μ only through the aggregate labor productivity p. This means that the set of wages posted by the firms in equilibrium $\mathcal{W}$ is determined by the aggregate labor productivity p as well. Then, plugging J^* into (2.6) yields

$$\theta^*(w;p)=\begin{cases} q^{-1}\left(\frac{\kappa}{J^*(w;p)}\right) & \text{if} \quad w\in\mathcal{W}(p)\\ 0 & \text{otherwise}\end{cases}$$

Notice that, as explained in the main text, the constant-returns-to-scale property of the matching function M is crucial here so that we can write the job finding rate and vacancy filling rate as a function of θ only.[4] Hence, we show that equilibrium market tightness does not depend on the distribution of agents across states as well.
Next, using this result and the assumption that the UI policy only depends on p, we show that the household value functions do not depend on the aggregate distribution of agents across states. To do so, we first collapse the problem of households into one functional equation and show that it is a contraction. Then, we show that the functional equation maps the set of functions that depend on the aggregate state μ only through p.
Let Ω denote the possible realizations of the aggregate state μ and define a value function $R : \{0,1\}\times\{0,1\}\times\mathcal{A}\times\mathcal{W}\times\mathcal{B}\times\Omega\rightarrow\mathbb{R}$ such that

$$\begin{aligned} R(l=1,d=0,a,w,\beta;\mu) &= V^W(a,w,\beta;\mu)\\ R(l=0,d=1,a,w,\beta;\mu) &= V^{UE}(a,w,\beta;\mu)\\ R(l=0,d=0,a,w,\beta;\mu) &= V^{UI}(a,\beta;\mu)\end{aligned}$$

[4]The free entry condition (2.6) is also important to pin down market tightness.

Then, we define the set of functions $\mathcal{R} : \{0,1\} \times \{0,1\} \times \mathcal{A} \times \mathcal{W} \times \mathcal{B} \times \mathcal{P} \to \mathbb{R}$ and let T_R be an operator such that

$$\begin{aligned}
(T_R R)(l, d, a, w, \beta; p) =& l \Bigg[\max_{c,a'} u(c) + \beta \mathbb{E} \Big[\delta(p') \left(1 - e(p')\right) R(l = 0, d = 1, \cdot) \\
& + \delta(p') e(p') R(l = 0, d = 0, \cdot) + \left(1 - \delta(p')\right) R(l = 1, d = 0, \cdot) \Big] \Bigg] \\
& + (1 - l) d \Bigg[\max_{c,a',s} u(c) - \nu(s) + \beta \mathbb{E} \Big[\max_{\tilde{w}} \Big\{ sf(\theta(\cdot)) R(l = 1, d = 0, \cdot) \\
& + (1 - sf(\theta(\cdot))) \left(1 - e(p')\right) R(l = 0, d = 1, \cdot) \\
& + (1 - sf(\theta(\cdot))) e(p') R(l = 0, d = 0, \cdot) \Big\} \Big] \Bigg] \\
& + (1 - l)(1 - d) \Bigg[\max_{c,a',s} u(c) - \nu(s) \\
& \quad + \beta \mathbb{E} \Big[\max_{\tilde{w}} \Big\{ sf(\theta(\cdot)) R(l = 1, d = 0, \cdot) \\
& + (1 - sf(\theta(\cdot))) R(l = 0, d = 0, \cdot) \Big\} \Big] \Bigg]
\end{aligned}$$

subject to

$$\begin{aligned}
& c + a' \leq (1 + r) a + lw(1 - \tau) + (1 - l) d \left[\phi(p) w (1 - \tau) + h\right] \\
& \qquad + (1 - l)(1 - d) h \\
& a' \geq -\underline{a} \\
& p' \sim F\left(p' \,|\, p\right)
\end{aligned}$$

where we use the result from above that market tightness does depend on Γ.

Assuming the utility function is bounded and continuous, $\mathcal{R}$ is the set of continuous and bounded functions. Then, we can show that the operator T_R maps a function from $\mathcal{R}$ into $\mathcal{R}$ (i.e., $T_R : \mathcal{R} \to \mathcal{R}$). Then, using Blackwell's sufficiency conditions for a contraction and the assumptions of boundedness of sets of exogenous processes $\mathcal{P}$ and $\mathcal{B}$, and choice sets $\mathcal{W}$ and $\mathcal{A}$, we can show that T_R is a contraction and has a unique fixed point $R^* \in \mathcal{R}$. Thus, the solution to the household problem does depend on Γ. This constitutes a BRE along with the solution to the firm's problem and the implied labor market tightness that does not depend on Γ, given that the UI policy is a function of p only.

Uniqueness: We know that policy functions of the household do not depend on Γ. Now, we prove the uniqueness of the policy functions for assets $\left\{g_a^l(a,w,\beta;p)\right\}_{l=\{W,UE\}}$, and $g_a^{UI}(a,\beta;p)$, wage choice $g_w^{UE}(a,w,\beta;p)$ and $g_w^{UI}(a,\beta;p)$, and search effort $g_s^{UE}(a,w,\beta;p)$ and $g_s^{UI}(a,\beta;p)$.

Wage policy function: Under the assumptions on $u(\cdot)$ and $\nu(\cdot)$ together with the assumptions of boundedness of sets of exogenous processes $\mathcal{P}$ and $\mathcal{B}$, and choice sets $\mathcal{W}$ and $\mathcal{A}$, value functions V^l are strictly concave in w for $l=\{W,UE\}$ and V^{UI} is constant in w. For simplicity, assume that p is non-stochastic and $\delta(p)=\delta$. We then obtain the equilibrium value of a matched firm using Equation (2.4) as follows:[5]

$$J^*(w;p)=\frac{p-w}{r+\delta}(1+r)$$

Then, we can write the equilibrium labor market tightness as

$$f(\theta^*(w;p))=\theta^*(w;p)=\frac{J^*(w;p)}{\kappa}$$

where the first equality uses the assumption that $M=\min\{v,S\}$, and the second equality uses the free entry condition. Using the expression for $J^*(w;p)$ gives

$$f(\theta^*(w;p))=\frac{1+r}{\kappa(r+\delta)}[p-w]>0.$$

This implies that the job finding rate $f(\cdot)$ is linear and decreasing in w. Then, rewriting the objective function for the wage choice of eligible unemployed, we have

$$\begin{aligned}\max_{\tilde{w}}\ & sf(\theta(\tilde{w};p))V^W\left(a',\tilde{w},\beta';p\right)+(1-sf(\theta(\tilde{w};p)))\\ &\times\left[(1-e(p))V^{UE}\left(a',w,\beta';p\right)+e(p)V^{UI}\left(a',\beta';p\right)\right]\end{aligned}$$

Using the result that V^l is strictly concave in w for $l=\{W,UE\}$ and V^{UI} is constant in w, and that $f(\cdot)$ is linear and decreasing in w, it is easy to show that the objective function

[5] The following results can be obtained under N state Markov process assumption for p and no restrictions on the job destruction rate.

above is strictly concave in w. This implies that the wage policy function $g_w^{UE}(a, w, \beta; p)$ is unique.

Similarly, rewriting the objective function for the wage choice of ineligible unemployed yields

$$\max_{\tilde{w}} sf(\theta(\tilde{w}; p)) V^W \left(a', \tilde{w}, \beta'; p\right) + (1 - sf(\theta(\tilde{w}; p))) V^{UI} \left(a', \beta'; p\right),$$

and using the same reasoning implies that the wage policy function $g_w^{UI}(a, \beta; p)$ is also unique.

Asset policy function: Under the assumptions on the utility functions $u(\cdot)$ and $\nu(\cdot)$ and choice sets $\mathcal{A}$, $\mathcal{W}$ and exogenous proceses $\mathcal{B}$, $\mathcal{P}$, value functions V^l are strictly concave in assets. This implies that the objective functions for the asset choice of each employment status are strictly concave in a', and thus asset policy functions $g_a^l(a, w, \beta; p)$ are unique for $l = \{W, UE, UI\}$

Search effort policy function: Using the same reasoning, objective functions for search effort choice of eligible and ineligible unemployed are strictly concave in s. This implies that the search effort policy functions $g_s^{UE}(a, w, \beta; p)$ and $g_s^{UI}(a, \beta; p)$ are unique.

B.3 Welfare Calculation for Great Recession Simulation

First, we focus on individual i. Let $t = 0$ be December 2007 and let T be December 2013. For ease of exposition, we discuss the calculation of welfare in two separate parts: let period (A) include any time $t \in [0, \ldots T]$ during the Great Recession and recovery and (B) represent the terminal time period post-December 2013 $t > T$.

Let $c_i^j(\mathbf{x}_t, p_t)$ and $s_i^j(\mathbf{x}_t, p_t)$ denote the consumption and search effort policy functions of individual i with individual state $\mathbf{x}_t$ at time t when aggregate productivity is p_t and UI policy is $j \in \{b, f, n\}$, where b denotes the benchmark policy, f denotes the flat policy, and n denotes the new/alternative policy.[6] To evaluate the welfare gains from the optimal policy in this exercise, we set policy n to be the optimal policy.

[6]Notice here that we are using the result that policy functions of the agents in our economy depend on the aggregate state of economy only through p as a result of block recursivity.

First consider welfare in period (A). Under the benchmark policy b, the utility of individual i during period (A) when endowed with additional $\bar{\pi}$ percent of consumption for her lifetime is given by

$$\sum_{t=0}^{T} (\beta_{i,t})^t U\left(c_i^b(\mathbf{x}_t, p_t)(1+\bar{\pi}), s_i^b(\mathbf{x}_t, p_t)\right),$$

where $U\left(c_i^b(\mathbf{x}_t, p_t)(1+\bar{\pi}), s_i^b\right) = \frac{\left[c_i^b(\mathbf{x}_t,p_t)(1+\bar{\pi})\right]^{1-\sigma}}{1-\sigma} - \mathbf{1}_U\left[\alpha \frac{s_i^b(\mathbf{x}_t,p_t)^{1+\chi}}{1+\chi}\right]$. Note that in the above expression, $\{p_t\}_{t=0}^T$ represents the labor productivity that is fed into the model during the recession, while $\{\beta_{i,t}\}_{t=0}^T$ represents the realized values of discount factor β. Agents, however, take expectations on aggregate labor productivity using the AR(1) process.

Now consider period (B). The continuation value of the individual post-December 2013 is given by

$$\mathbb{E}_{T+1} \sum_{t=T+1}^{\infty} (\beta_{i,t})^t U\left(c_i^j(\mathbf{x}_t, p_t)(1+\bar{\pi}), s_i^j(\mathbf{x}_t, p_t)\right),$$

which recursively can be written as $(\beta_{i,T+1})^{T+1} V_{\bar{\pi}}^{l_i,j}(a_{T+1}, w_{T+1}, \beta_{T+1}, p_{T+1})$ where $V_{\bar{\pi}}^{l_i,j}$ denotes the value of individual i with labor force status $l_i \in \{W, UE, UI\}$ under the policy j when consumption is multiplied by $1+\bar{\pi}$ every period from $t = T+1$ onward. Computationally, we can find $V_{\bar{\pi}}^{l_i}$ once we have obtained the policy functions associated with the underlying value function V^{l_i}. We do this recursively by policy function iteration with the difference being that consumption is multiplied by $(1+\bar{\pi})$ at every iteration. Under the original exercise where the policy is permanent, we set $j = b$, while under the exercise when the policy is discretionary/temporary, the government reverts back to the flat policy postrecession and thus $j = f$.

Hence, the welfare of agent i who is endowed with an additional $\bar{\pi}$ percent of lifetime consumption over periods (A) and (B) under the baseline policy b can be written as

$$\sum_{t=0}^{T}\left[(\beta_{i,t})^t U\left(c_i^b(\mathbf{x}_t, p_t)(1+\bar{\pi}), s_i^b(\mathbf{x}_t, p_t)\right)\right] + (\beta_{i,T+1})^{T+1} V_{\bar{\pi}}^{l_i,j}(a_{T+1}, w_{T+1}, \beta_{T+1}, p_{T+1}).$$

Now, aggregating across individuals at each point in time, we can write the left-hand side of Equation (2.15) as

$$\sum_{t=0}^{T}\left[\int_{i}\left(\beta_{i,t}\right)^{t}U\left(c_{i}^{b}\left(\mathbf{x}_{t},p_{t}\right)\left(1+\bar{\pi}\right),s_{i}^{b}\left(\mathbf{x}_{t},p_{t}\right)\right)d\Gamma_{t}^{b}\left(i\right)\right]$$
$$+\int_{i}\left(\beta_{i,T+1}\right)^{T+1}V_{\bar{\pi}}^{l_{i},j}\left(a_{T+1},w_{T+1},\beta_{T+1},p_{T+1}\right)d\Gamma_{T+1}^{b}\left(i\right), \tag{B.1}$$

where Γ_t^b is the distribution of the economy at time t under policy b.
Similarly, the right-hand-side of equation (2.15) is computed by solving

$$\sum_{t=0}^{T}\left[\int_{i}\left(\beta_{i,t}\right)^{t}U\left(c_{i}^{n}\left(\mathbf{x}_{t},p_{t}\right),s_{i}^{n}\left(\mathbf{x}_{t},p_{t}\right)\right)d\Gamma_{t}^{n}\left(i\right)\right]$$
$$+\int_{i}\left(\beta_{i,T+1}\right)^{T+1}V^{l_{i},j}\left(a_{T+1},w_{T+1},\beta_{T+1},p_{T+1}\right)d\Gamma_{T+1}^{n}\left(i\right), \tag{B.2}$$

where Γ_t^n is the corresponding distribution under policy n and the superscript j of the value function in (B) depends on whether the policy is permanent ($j = n$) or temporary ($j = f$).
Under a temporary policy, we emphasize that even if the policy reverts to the flat policy f after December 2013, the terminal value will be different for policy b and n because the distribution of each economy at $t = T + 1$ is going to be different from each other (i.e., $\Gamma_{T+1}^b \neq \Gamma_{T+1}^n$).
Then, we simply use a zero-finder to find $\bar{\pi}$ that makes equations (B.1) and (B.2) the same.[7]

B.4 Model with Endogenous Quits

In this section, we present the extended model that incorporates the endogenous quit decisions of workers.

Worker's problem

Under the model with quits, workers matched with a firm can decide to leave employment. After the separation shock realizes, a firm-worker pair that is not dissolved exogenously

[7]Note that there is no closed-form solution for $\bar{\pi}$ given the functional form of the utility function.

may endogenously be separated if the worker chooses to quit. The worker's problem is now given by

$$\begin{aligned}
V^{W}(a,w,\beta;\mu) &= \max_{c,a'} \quad u(c)+\beta\mathbb{E}\Big[\delta(p')(1-e(p'))V^{UE}(a',w,\beta';\mu') \\
&\quad +\delta(p')e(p')V^{UI}(a',\beta';\mu') \\
&\quad +(1-\delta(p'))\max_{d\in\{0,1\}}\Big\{dV^{UI}(a',\beta';\mu')+(1-d)V^{W}(a',w,\beta';\mu')\Big\}\Big] \\
&\text{subject to} \\
&\quad c+a' \leq (1+r)a+w(1-\tau) \\
&\quad a' \geq -\underline{a} \\
&\quad \Gamma' = H(\mu,p') \quad \text{and} \quad p' \sim F(p'\,|\,p).
\end{aligned}$$

Firm's problem

The value of a matched firm is modified to account for the possibility of a quit. Even if a match is not dissolved by the exogenous shock δ, it can be dissolved if the worker's decision to quit is $g_d = 1$:

$$\begin{aligned}
J(a,w,\beta;\mu) &= p-w \\
&\times \frac{1}{1+r}\mathbb{E}\Big[(1-\delta(p'))(1-g_d(a',w,\beta',\mu'))J(a',w,\beta';\mu')\Big|\beta,\mu\Big] \\
&\text{subject to} \\
&\quad \Gamma' = H(\mu,p') \quad \text{and} \quad p' \sim F(p'\,|\,p),
\end{aligned} \tag{B.3}$$

where $a' = g_a(a,w,\beta,\mu)$.
So, the value of posting a vacancy is given by

$$V(a,w,\beta;\mu) = -\kappa + q(\theta(a,w,\beta;\mu))J(a,w,\beta;\mu) \tag{B.4}$$

and market tightness can be obtained by solving

$$\theta(a,w,\beta;\mu) = \begin{cases} q^{-1}\left(\frac{\kappa}{J(a,w,\beta;\mu)}\right) & \text{if} \quad w \in \mathcal{W}(\mu) \\ 0 & \text{otherwise.} \end{cases} \tag{B.5}$$

Notice that the value of a firm J depends on individual states $\mathbf{x} = (a, w, \beta, b)$ because heterogeneous workers will have different quit thresholds. This would then imply that market tightness θ is also a function of these states.

Unemployed's problem

The unemployed's problem remains unchanged, except that market tightness is now a function of other individual states $\mathbf{x} = (a, w, \beta, b)$ for reasons stated in the firm's problem.

B.5 Computational Algorithm

Solving the baseline model

The model is solved using the following steps:

1. Solve for the value function of the firm $J(w, p)$.

2. Using the free-entry condition $0 = -\kappa + q(\theta(w, p)) J(w, p)$ and the functional form of $q(\theta)$, we can solve for market tightness for any given wage submarket w and aggregate productivity p:
$$\theta(w, p) = q^{-1}\left(\frac{\kappa}{J(w, p)}\right),$$
where we set $\theta(w, p) = 0$ when the market is inactive.

3. Given the function θ, we can then solve for the household value functions V^W, V^{UE}, and V^{UI} using standard value function iteration. In order to decrease computation time, we implement Howard's improvement algorithm (policy-function iteration).

4. Once household policy functions are obtained, we are able to simulate aggregate dynamics of the model.

Extended model with endogenous quits

Solving the model will require modifying the baseline algorithm above as follows:

1. Guess a market tightness function $\theta_0(a', w, \beta', p')$.

2. Taking θ_0 as given, solve for the household's problem.

3. Using the household's policy function $g_d^W(\theta_0)$ and $g_a^W(\theta_0)$, solve for the firm's problem.

4. After obtaining $J_0(a, w, \beta, p)$, use Equation B.5 to back out the implied market tightness $\theta_1(a', w, \beta', p')$.

5. If convergence criterion $\| \theta_1 - \theta_0 \| \leq \epsilon_\theta$ is not satisfied, use θ_1 as a guess and repeat the steps outlined above.

Appendix C

Appendix to Chapter 3

In this appendix, we provide details on the data sources and construction of variables for our analysis in "What Do Survey Data Tell Us about U.S. Businesses?" We also include the auxiliary tables and figures omitted from the main text.

C.1 Data Sources

The main data sources are:

- Statistics of Income of the Internal Revenue Service (SOI);
- Survey of Consumer Finances of the Board of Governors of the Federal Reserve System (SCF);
- Survey of Income and Program Participation of the U.S. Census Bureau in the Department of Commerce (SIPP);
- Panel Study of Income Dynamics of the Survey Research Center, Institute for Social Research, University of Michigan (PSID);
- Current Population Survey at the Bureau of Labor Statistics (CPS);
- Center for Research in Security Prices and Compustat (CRSP);
- Pratt's Stats (now renamed as DealStats) from Business Valuation Resources.

Besides the main data sources listed above, we also use information from the national income and product accounts and fixed asset tables of the Bureau of Economic Analysis; financial accounts of the Board of Governors of the Federal Reserve System; Panel Study of Entrepreneurial Dynamics of the Survey Research Center, Institute for Social Research, University of Michigan; and the Kauffman Firm Survey of the Kauffman Foundation.

We construct business incomes and numbers of returns and owners by using variables from the IRS, SCF, SIPP, PSID, and CPS. In addition to these variables, we use BEA estimates of income misreporting by noncorporate businesses and General Accountability Office (GAO) estimates of income misreporting by S corporations to adjust IRS pass-through business income. BEA estimates of income misreporting over time are obtained from NIPA Table 7.14 (line 2). The GAO estimates are taken from reports GAO 14-453 and 10-195, which summarize the progress of the tax compliance studies conducted by the IRS through the National Research Program.

To verify the consistency of sole proprietorship income in the SCF, we compared two measures of net income from a sole proprietorship or a farm. Results of the comparison were reported in the main text. The first uses pretax net income variables (coded by X3132, X3232, and X3332) from Form 1040, Schedule C (line 31). The second is X5704, which measures net annual income from a sole proprietorship or a farm before deductions. Specifically, this variable codes responses to the question asking for combined incomes appearing on lines 12 and 18 of IRS Form 1040.

Finally, to verify whether respondents in the SCF check documents, we use variable X6536, which provides information on the frequency of checking any documents when answering interview questions. Variable X7451 informs us about whether the respondent referred to income tax documents, and variables X7452 through X7455 inform us about whether the respondent referred to other financial documents, namely, pension documents, account statements, investment or business records, and loan documents, respectively. If a respondent says that he or she checked the income tax document (X7451=1), we use his or her answers to X6536 to obtain the frequency of checking this document. The respondent did not check the income tax document if either (X7451=5 or X7451=0 or X7451=-7) or (X6536=4). We use the same steps to check referencing of other financial documents by using X7452–X7455 instead of X7451. We classify a respondent who checks at least one of these four documents as someone who refers to any other tax documents. We then obtain the weighted fraction of the group of respondents who check these two types of documents frequently, sometimes, rarely, or never. Roughly 4 percent of all respondents have nonapplicable responses (NaN). We adjust for this nonresponse rate in the results of the main text so that our fractions sum to 100 percent.

C.2 Additional Results

Next, we report on our auxiliary tables and figures that relate to our findings on business incomes, receipts, and returns.

Business income

Aggregate

In Section 3.3.1 of the main text, we discussed business income per return and the number of returns across years and legal forms. In Figure C.1, we report aggregate business incomes and show that they are overstated for all pass-through businesses. In Figure C.2, we compare the aggregate business income from the SCF with other surveys, namely, the SIPP, PSID, and CPS, and extend the analysis from Section 3.6 of the main text.

Distribution

In Section 3.3.2, we discussed the distribution of business income by splitting pass-through businesses into two categories: those that make profits and those that make losses (or no income). In Figures C.3 and C.4, we plot business income per return by legal status for those making profits and losses, respectively. In Figures C.5 and C.6, we plot the number of returns for the same sets of businesses. In Table C.1, we extend the analysis of decomposing the total percentage error into the overstatement of profits and understatement of losses. In Figure C.7, we report the distributional statistics for S corporations. As we noted in the main text, the data for S corporations are only available for limited years, namely 2003–2012, but these data show similar inconsistencies between SCF and IRS data, as was found with sole proprietorships.

Broad business income

In Figure C.8, we extend the analysis of Section 3.5.2 in the main text by replicating the analysis of Johnson and Moore (2008) for all years. As we noted in the main text, the SCF estimates are still larger in all years than the IRS counterpart even with the broader concept of income.

Business receipts

In Section 3.3.1 of the main text, we reported that business receipts per owner are overstated. In Figures C.9 and C.10, we corroborate that finding by showing aggregate business receipts and business receipts per return across legal forms and across years.

Business returns

In this section, we provide additional details for the comparison of the income yields in SCF to CRSP-Compustat, Pratt's Stats, and other surveys to augment the analysis in Section 3.4 of the main text.

In Table C.2, we provide several additional moments for the distribution of income yields in the SCF. In the main text, we showed evidence that the SCF income yields are high when compared to CRSP-Compustat or Pratt's Stats. The additional moments show that this is true regardless of year or legal structure.

In the main text, we compared the income yields for S and C corporations in the SCF to small firms in CRSP where we defined "small" as corporations that are in the bottom quintile of the size distribution as measured by the book value of total assets. In Table C.3, we extend the analysis to two alternative definitions of "small": (i) those in the bottom quintile by market value and (ii) those in the bottom quintile by gross sales. Although there are some differences in the magnitudes, the equally weighted and value-weighted yields are negative in all years, regardless of how we classify the small firms.

In Table C.4, we report income yields from Pratt's Stats for all legal forms. We see that sole proprietors have higher yields than other pass-throughs and C corporations. However, since these businesses have much smaller valuations, the value-weighted yield for all businesses is relatively low when compared to SCF data.

Finally, in Tables C.5 and C.6, we report the income yields in PSID and SIPP for all years that the data are available. As we noted in the main text, the average yields are comparable across the SCF, PSID, and SIPP, while the distributions are not. These tables more clearly demonstrate this finding.

Table C.1: SCF-IRS business income gap by legal structure

Tax Year	SCF-IRS Gap ($)	Percentage of Gap from Overstatement of Profits (%)	Percentage of Gap from Understatement of Losses (%)
		Sole Proprietorship	
1988	67.09	58	42
1994	5.44	-515	615
2000	168.09	75	25
2006	91.66	29	71
2012	-28.22	359	-259
		Partnership	
1988	56.28	37	63
1994	500.59	92	8
2000	261.03	56	44
2006	724.62	83	17
2012	205.51	0	100
		S Corporation	
1988	35.78	57	43
1994	118.07	74	26
2000	206.06	78	22
2006	279.35	77	23
2012	41.06	-53	153
		C Corporation	
1994	-244.42	148	-48
2000	-57.00	670	-570
2006	-859.87	123	-23
2012	-747.36	138	-38

Note: This table shows the difference (gap) between aggregated business income by legal structure in the SCF and IRS. The gap is then decomposed into the fraction attributable to an overstatement of profits or an understatement of losses. Dollar amounts are in billions. The table shows results for every six years. See my website for the complete results.

Table C.2: Net income yields in the SCF by legal structure

	Sole Proprietorship					Partnership				
Tax Year	Value-Weighted Mean	Equally Weighted Mean	p25	p50	p75	Value-Weighted Mean	Equally Weighted Mean	p25	p50	p75
1988	19.9	105.0	3.2	20.0	80.0	13.6	111.4	0.0	8.0	50.0
1994	19.1	97.8	2.0	24.0	74.0	74.1	49.1	0.3	10.7	42.3
2000	26.6	89.8	0.9	25.5	75.0	24.5	203.1	0.1	11.9	40.0
2006	25.0	254.8	2.3	32.0	100.0	18.8	84.4	0.1	10.0	40.0
2012	24.7	87.4	0.0	23.2	82.4	11.5	36.8	0.0	5.4	33.7
	S Corporation					C Corporation				
1988	12.7	23.5	0.5	6.0	37.5	17.8	101.7	3.2	16.7	30.5
1994	14.3	38.1	0.9	11.7	40.0	28.3	73.9	0.4	8.0	41.1
2000	16.1	120.7	4.4	18.4	40.0	26.5	90.8	2.9	15.8	46.0
2006	15.4	75.1	3.8	16.7	80.0	16.3	44.4	0.0	7.5	36.0
2012	14.4	57.6	2.7	15.2	52.2	15.4	55.4	0.0	9.0	41.3
	All Pass-throughs					All Businesses				
1988	16.1	101.3	1.2	13.3	62.5	16.6	101.3	1.3	14.3	57.0
1994	32.2	80.8	1.1	20.0	64.0	31.5	80.8	1.1	19.0	62.9
2000	21.3	113.9	1.3	21.0	62.9	22.6	114.4	1.6	20.0	62.3
2006	18.4	183.7	2.0	22.0	80.0	18.1	171.7	1.6	20.0	73.3
2012	13.9	67.1	0.0	15.0	60.0	14.1	66.2	0.0	15.0	60.0

Note: This table shows moments of the net income yield distribution of businesses in the SCF by legal structure. The sample includes businesses with positive net worth and excludes the bottom 1st percentile of these businesses. The business income of each business that the family members own in the SCF is obtained from SCF variables that correspond to information on business tax forms. The table shows results for every six years. See my website for the complete results.

Table C.3: Income yield for small firms in CRSP

	by Market Capitalization					by Sales				
Tax Year	EW	VW	p25	p50	p75	EW	VW	p25	p50	p75
1988	−43.6	−27.0	−52.3	−14.3	6.1	−27.2	−8.8	−26.3	−8.6	1.1
1991	−72.9	−49.0	−72.4	−15.9	5.1	−31.7	−6.0	−23.3	−5.6	1.5
1994	−23.3	−14.2	−34.1	−4.1	9.3	−18.1	−9.2	−24.8	−6.6	4.0
1997	−29.9	−19.2	−43.2	−8.5	7.1	−21.1	−8.5	−25.4	−8.0	2.7
2000	−104.1	−71.8	−103.4	−16.4	10.4	−52.8	−12.4	−42.2	−10.7	2.2
2003	−14.2	−9.2	−21.0	−0.9	7.8	−9.5	−7.2	−15.2	−3.3	5.5
2006	−12.1	−8.1	−20.8	−0.2	7.6	−11.9	−8.6	−18.6	−5.1	4.7
2009	−65.0	−47.3	−72.4	−22.5	4.7	−32.6	−11.0	−34.6	−10.8	3.0
2012	−22.7	−12.6	−35.6	−3.8	10.4	−17.1	−5.7	−22.7	−5.4	6.7
2015	−59.6	−35.6	−55.4	−11.5	6.3	−37.6	−11.5	−35.8	−11.9	1.9
Mean	−44.7	−29.4	−51.1	−9.8	7.5	−25.9	−8.9	−26.9	−7.6	3.3

Note: This table shows estimates of income yields for small businesses in CRSP-Compustat firms.The column "EW" reports the equally weighted average, the column "VW" reports the value-weighted average, the column "p25" reports the 25th percentile, the column "p50" reports the 50th percentile, and the column "p75" reports the 75th percentile.

Table C.4: Income yield from Pratt's Stats

Legal Form	EW	VW	p25	p50	p75
Sole Proprietorship	41.3	31.6	13.3	36.7	61.5
Partnership	26.6	4.8	2.7	20.5	48.8
S Corporation	30.3	6.9	6.5	23.3	47.8
C Corporation	6.8	-2.1	-2.3	6.5	29.8

Note: This table shows estimates of income yields from the Pratt's Stats database.The column "EW" reports the equally weighted average, the column "VW" reports the value-weighted average, the column "p25" reports the 25th percentile, the column "p50" reports the 50th percentile, and the column "p75" reports the 75th percentile.

Table C.5: Net income yields of unincorporated businesses in the PSID

Tax Year	Value-Weighted Mean	Equally Weighted Mean	p25	p50	p75
1998	5.2	136.4	0.0	12.5	75.0
2000	21.7	182.4	0.0	7.5	73.3
2002	21.8	187.0	0.0	33.3	139.5
2004	22.2	287.7	3.9	36.9	140.0
2006	20.6	630.1	10.0	42.5	222.2
2008	10.9	175.8	2.7	28.8	125.0
2010	13.9	110.3	3.9	25.0	75.9
2012	10.7	90.8	3.3	23.0	83.3
2014	6.9	182.9	4.8	33.3	100.0
Mean	14.9	220.4	3.2	27.0	114.9

Note: This table shows moments of the net income yield distribution of unincorporated businesses in the PSID. The sample includes businesses with positive net worth and excludes the bottom 1st percentile of these businesses.

Table C.6: Net income yields in the SIPP

Tax Year	Value-Weighted Mean	Equally Weighted Mean	p25	p50	p75
		Sole Proprietorship			
2004	20.2	545.0	6.8	44.8	240.0
2005	19.4	727.7	4.5	41.2	240.0
2009	13.0	3043.1	0.2	24.0	203.3
2010	15.8	5916.6	0.2	31.0	240.0
2011	14.9	8878.2	0.5	29.2	188.0
Mean	16.7	3822.1	2.4	34.0	222.3
		Partnership			
2004	25.1	605.9	0.6	29.2	220.0
2005	19.9	1271.2	0.3	22.6	189.1
2009	17.4	853.4	0.0	7.4	108.0
2010	21.3	2128.0	0.0	22.5	204.0
2011	18.9	1551.7	0.0	11.8	190.7
Mean	20.5	1282.0	0.2	18.7	182.4
		Unincorporated			
2004	22.0	2936.2	6.4	45.7	260.0
2005	19.8	12590.7	4.0	40.4	250.0
2009	14.0	15353.1	0.1	22.5	202.5
2010	17.2	38737.5	0.1	30.8	240.0
2011	15.3	7971.4	0.3	26.7	197.8
Mean	17.6	15517.8	2.2	33.2	230.1

Note: This table shows moments of the net income yield distribution of sole proprietorships, partnerships, and unincorporated businesses in the SIPP 2004 and 2008 panels. Statistics are calculated for years where asset topical modules are available. The sample includes businesses with positive net worth and excludes the bottom 1st percentile of these businesses.

Figure C.1: Business income by legal status, SCF vs. IRS

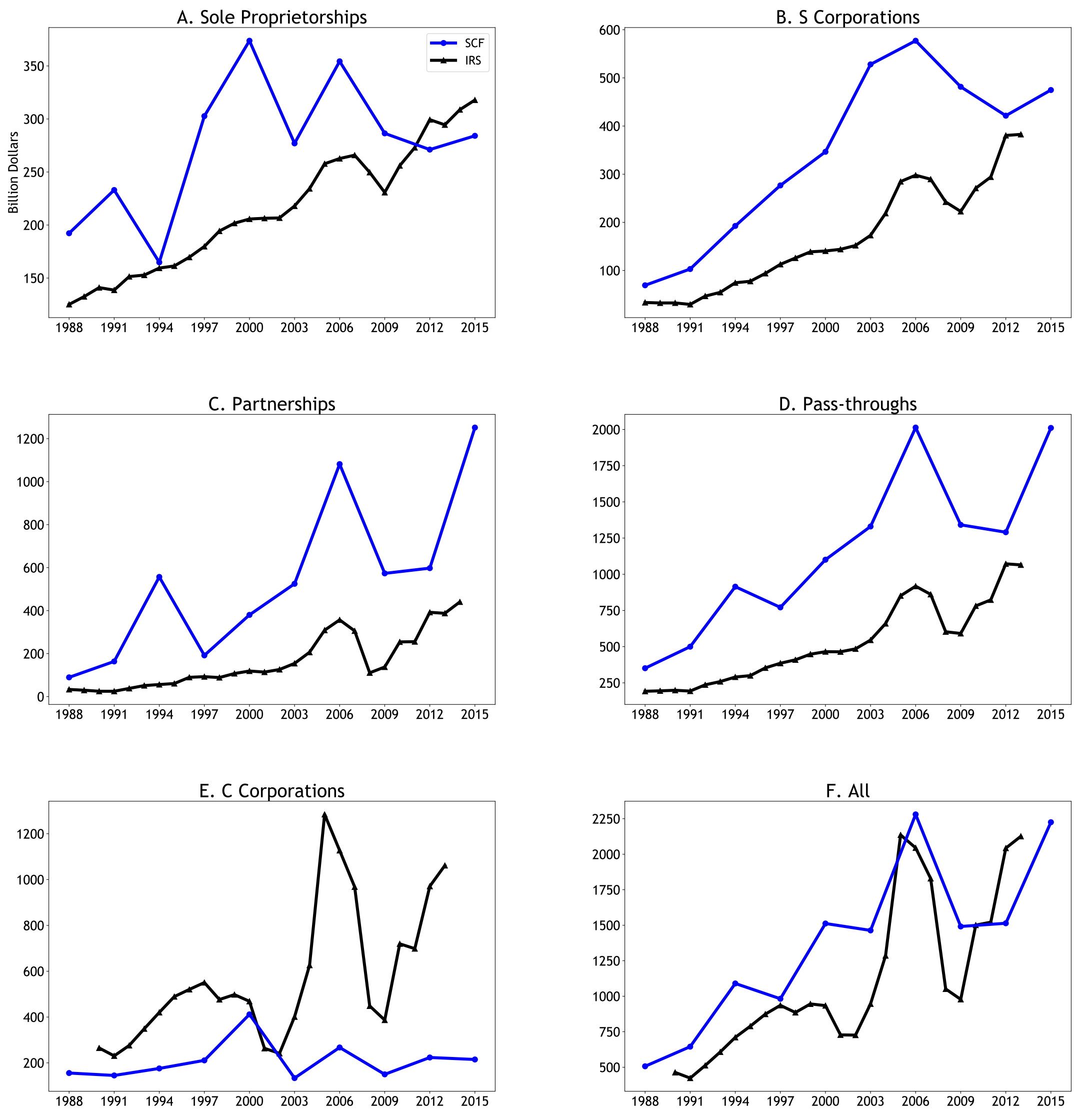

Note: This figure plots the total business income by legal status in the SCF and the IRS. Business income refers to income reported on Form 1040 Schedule C for sole proprietorships, Form 1065 for partnerships, Form 1120S for S corporations, and Form 1120 for C corporations. IRS data for partnerships, S corporations, and C corporations are available only until 2013, and C-corporation data start from 1990 because data for Form 1120 are not available for 1988 and 1989.

Figure C.2: Total unincorporated business income in SCF, SIPP, PSID, and CPS vs. IRS

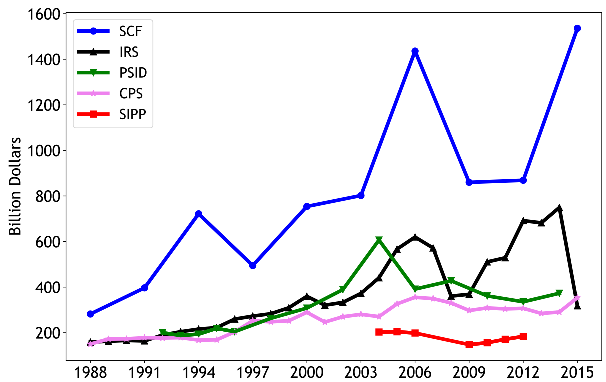

Note: This figure plots the total business income of unincorporated businesses in the SCF, SIPP, PSID, CPS, and IRS. Before 2004, the SIPP does not provide information about an individual's own share of business income from an unincorporated business. Instead, it contains information about the total income of the business, which is not enough information to calculate the total business income of unincorporated businesses.

Figure C.3: Business income per tax return by legal status for businesses with net income, SCF vs. IRS

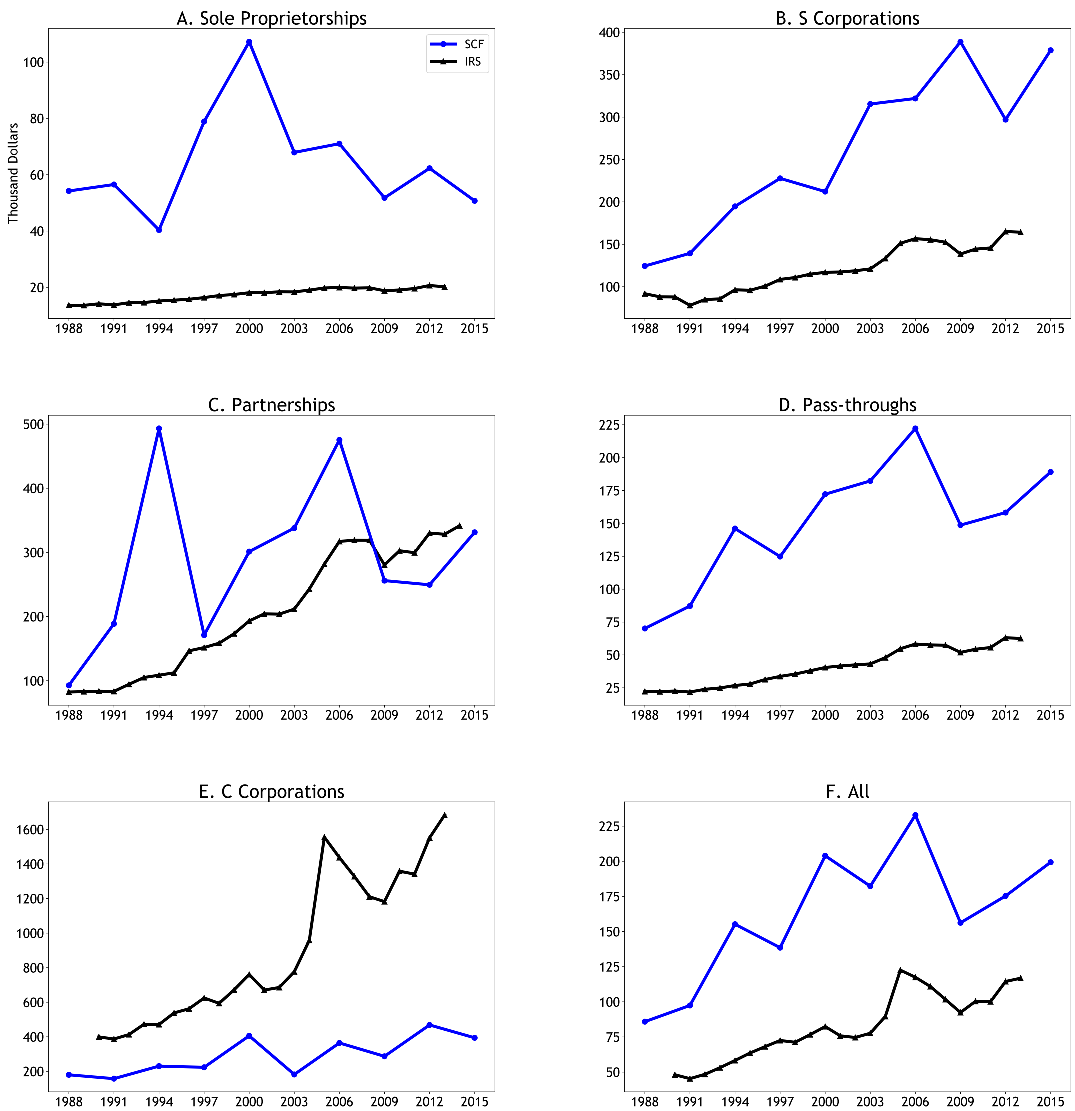

Note: This figure plots the business income per tax return by legal status for businesses with net income in the SCF and IRS. Business income refers to income reported on Form 1040 Schedule C for sole proprietorships, Form 1065 for partnerships, Form 1120S for S corporations, and Form 1120 for C corporations. IRS data for sole proprietorships, partnerships, S corporations, and C corporations are available only until 2013, and C-corporation data start from 1990 because data for Form 1120 are not available for 1988 and 1989.

Figure C.4: Business income per tax return by legal status for businesses with net loss, SCF vs. IRS

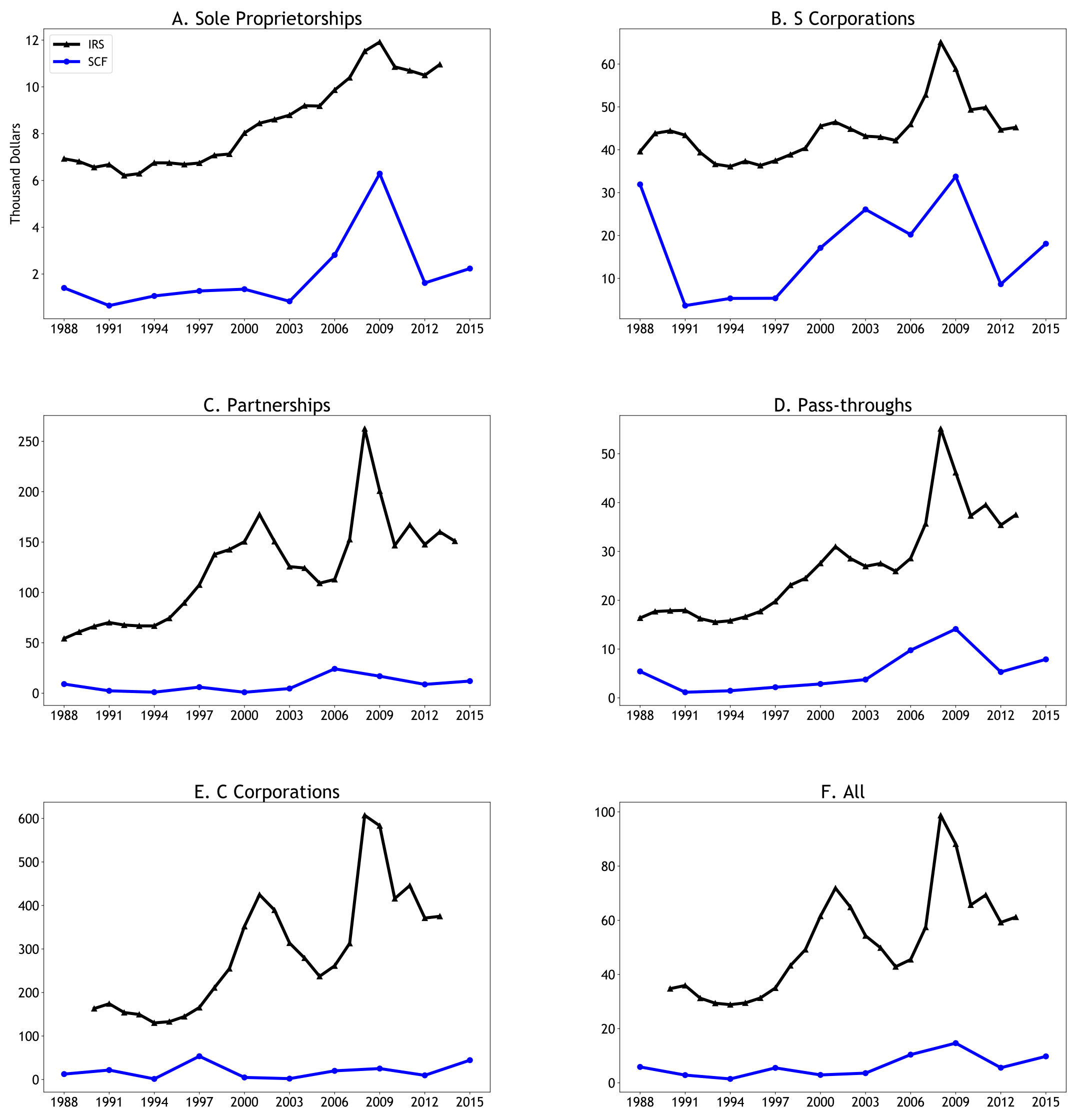

Note: This figure plots the business income per tax return by legal status for businesses with net loss in the SCF and IRS. Business income refers to income reported on Form 1040 Schedule C for sole proprietorships, Form 1065 for partnerships, Form 1120S for S corporations, and Form 1120 for C corporations. IRS data for sole proprietorships, partnerships, S corporations, and C corporations are available only until 2013, and C-corporation data start from 1990 because data for Form 1120 are not available for 1988 and 1989. Businesses with zero net income are included with those that have net losses.

Figure C.5: Number of returns by legal status for businesses with net income, SCF vs. IRS

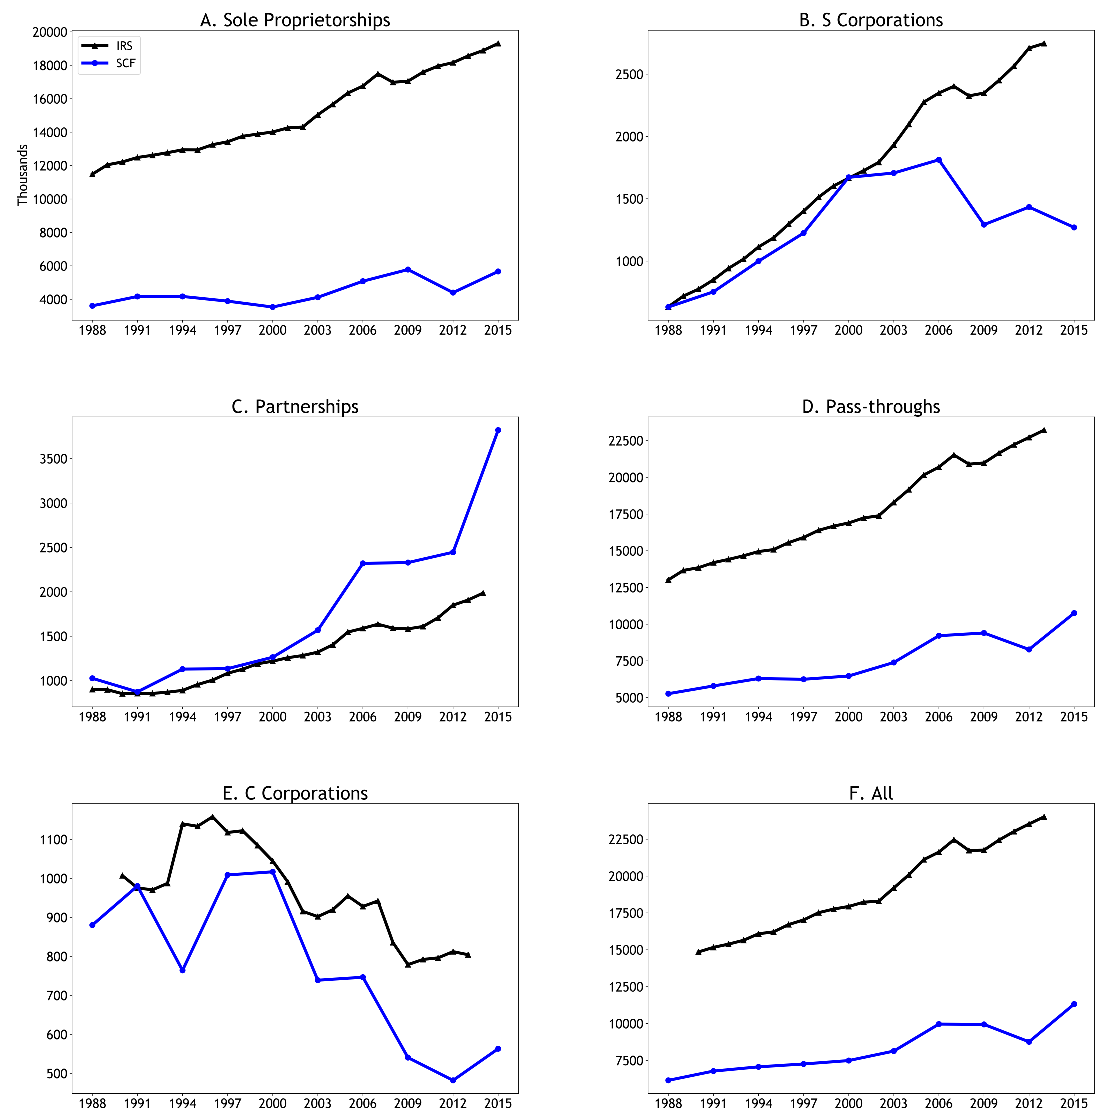

Note: This figure plots the number of business tax returns by legal status for business with net income in the SCF and the IRS. Business income refers to income reported on Form 1040 Schedule C for sole proprietorships, Form 1065 for partnerships, Form 1120S for S corporations, and Form 1120 for C corporations. IRS data for sole proprietorships, partnerships, S corporations, and C corporations are available only until 2013, and C corporations data starts from 1990 because data for Form 1120 is not available for 1988 and 1989.

Figure C.6: Number of returns by legal status for businesses with net loss, SCF vs. IRS

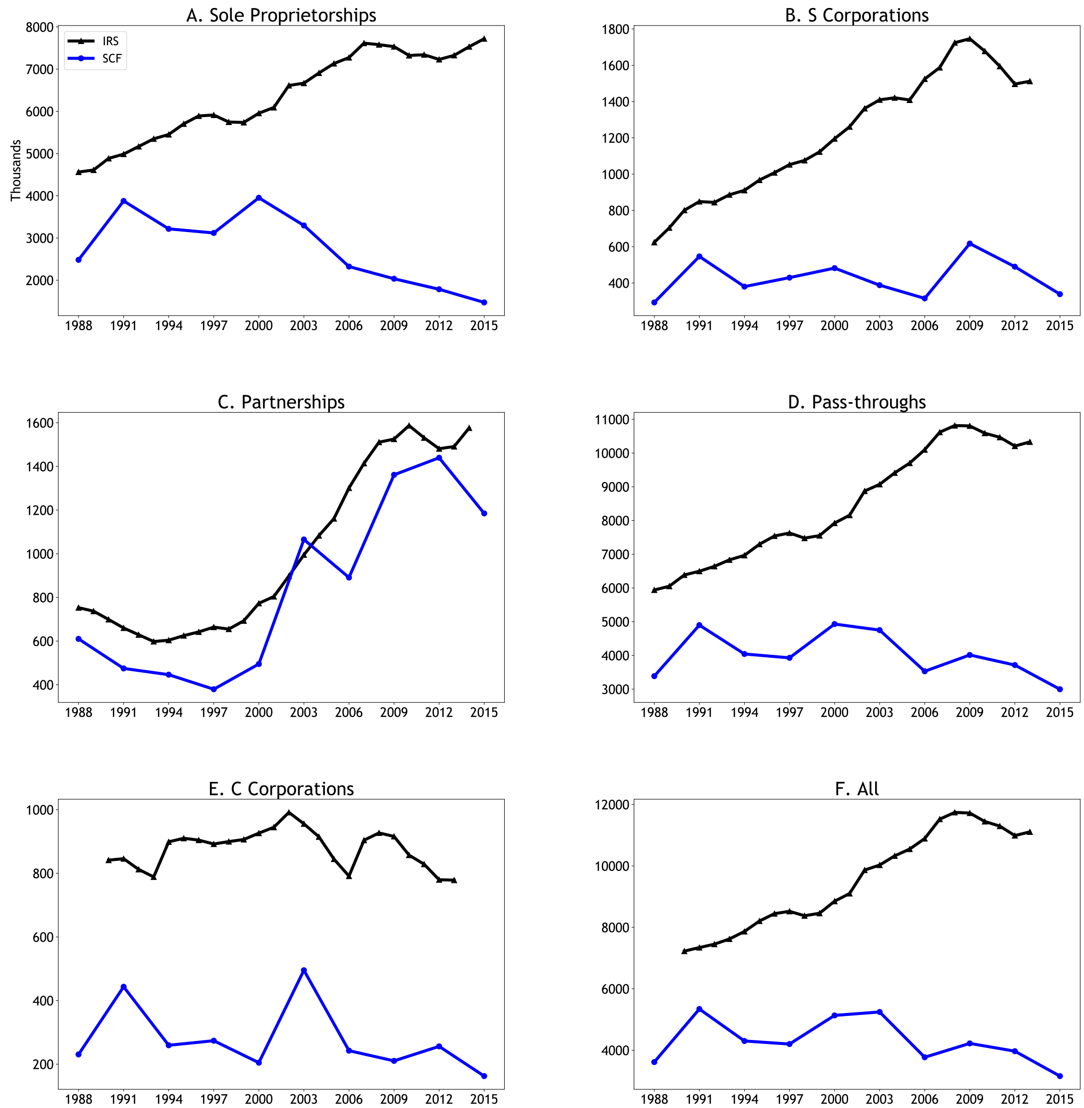

Note: This figure plots the number of business tax returns by legal status for businesses with net loss in the SCF and IRS. Business income refers to income reported on Form 1040 Schedule C for sole proprietorships, Form 1065 for partnerships, Form 1120S for S corporations, and Form 1120 for C corporations. IRS data for sole proprietorships, partnerships, S corporations, and C corporations are available only until 2013, and C- corporation data start from 1990 because data for Form 1120 are not available for 1988 and 1989. Businesses with zero net income are included with those that have net losses.

Figure C.7: Distribution of S-Corporation business income per return, SCF vs. IRS

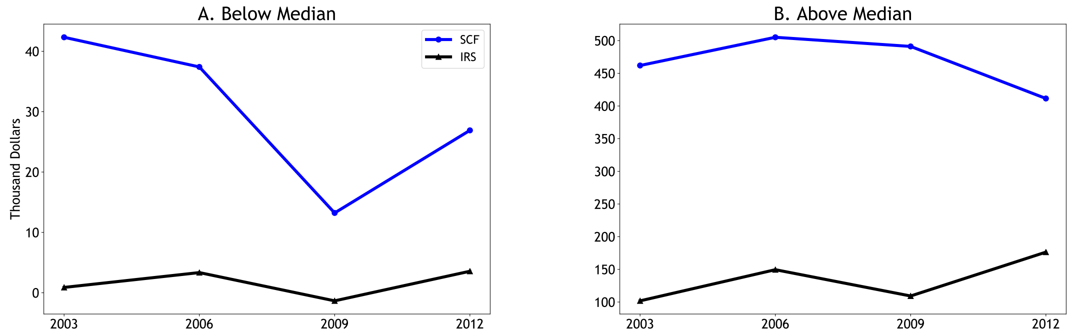

Note: This figure plots S-corporation business income per return for those with below- and above-median business receipts.

Figure C.8: Broad business income, SCF vs. IRS

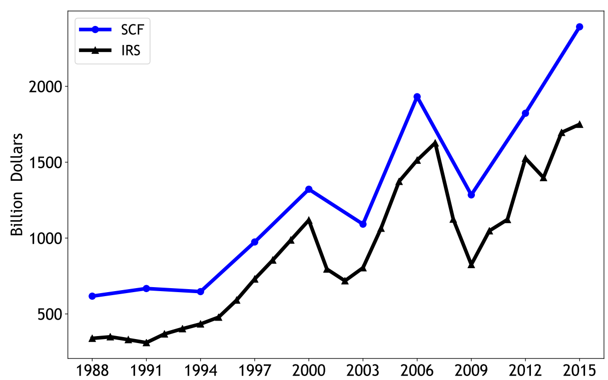

Note: This figure compares a broader measure of business income in the SCF and IRS. Broad business income is defined to be income derived from a business or profession (Form 1040 Schedule C) or farm (Form 1040 Schedule F); income from rental real estate, royalties, partnerships, S corporations, estates, trusts (Form 1040 Schedule E); and income from gains from the sale of capital and other property (Form 1040, lines 13 and 14).

Figure C.9: Business receipts by legal status, SCF vs. IRS

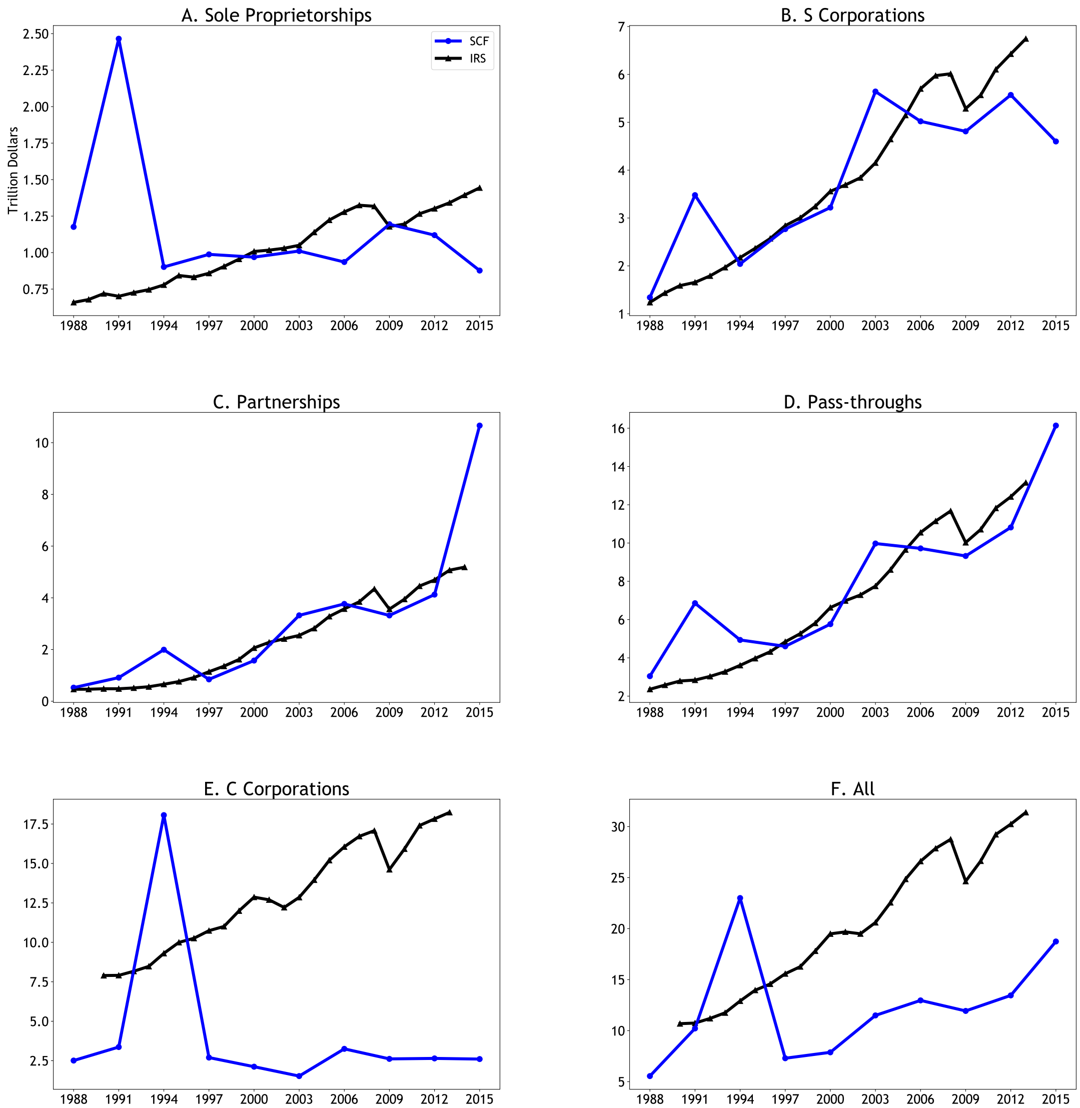

Note: This figure plots the total business receipts by legal status in the SCF and IRS. Business receipts refers to gross sales reported on Form 1040 Schedule C for sole proprietorships, Form 1065 for partnerships, Form 1120S for S corporations, and Form 1120 for C corporations. IRS data for partnerships, S corporations, and C corporations are available only until 2013, and C-corporation data start from 1990 because data for Form 1120 are not available for 1988 and 1989.

Figure C.10: Business receipts per tax return by legal status, SCF vs. IRS

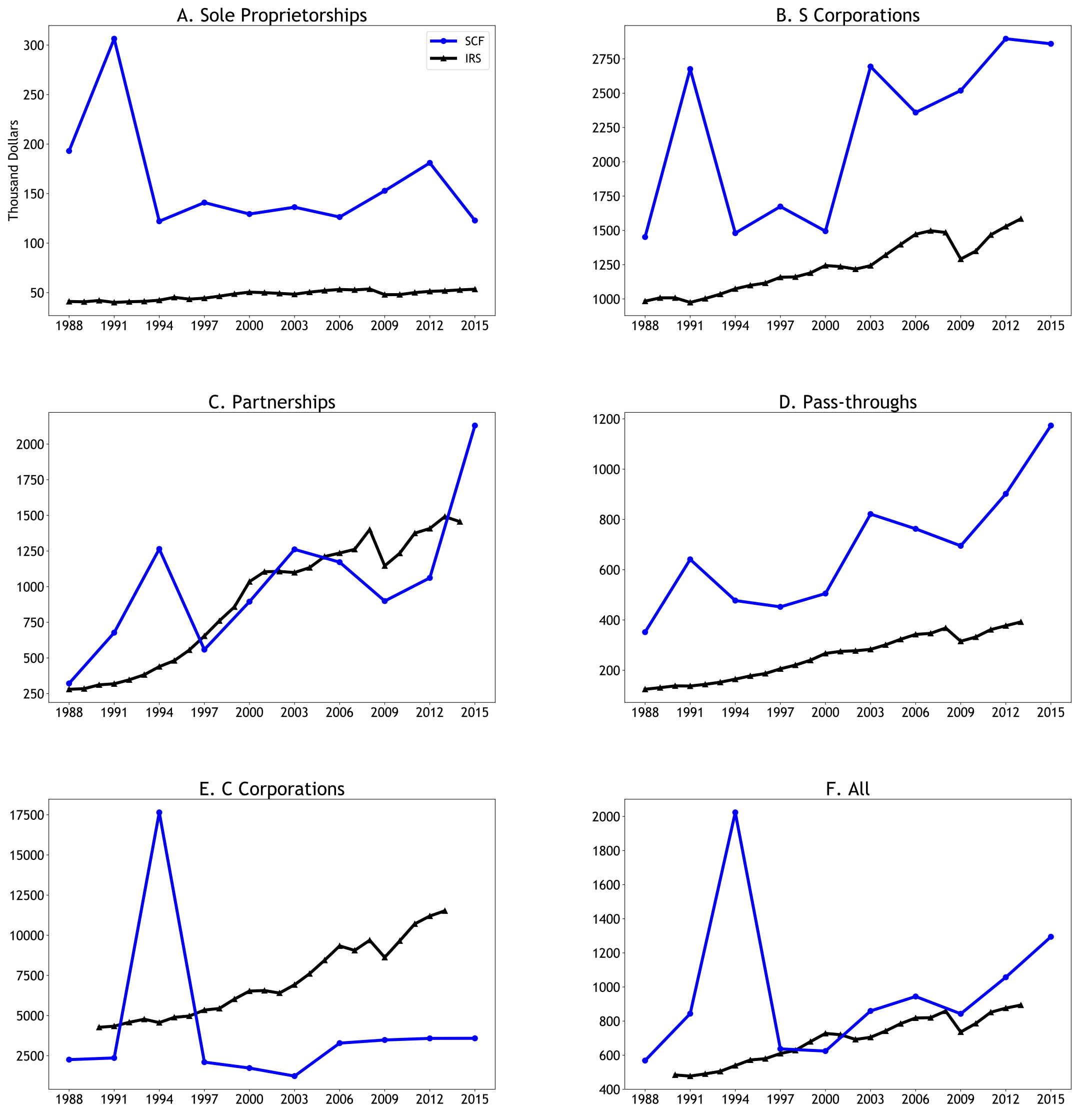

Note: This figure plots the business receipts per tax return by legal status in the SCF and IRS. Business receipts refers to gross sales reported on Form 1040 Schedule C for sole proprietorships, Form 1065 for partnerships, Form 1120S for S corporations, and Form 1120 for C corporations. IRS data for partnerships, S corporations, and C corporations are available only until 2013, and C-corporation data start from 1990 because data for Form 1120 are not available for 1988 and 1989.