

Job Search, Wage Dispersion, and Universal Healthcare*

Dan Beemon[†]

University of Wisconsin - Madison

September 30, 2021

Abstract

I design and estimate a labor market search model with employer sponsored health insurance (ESHI), worker and firm heterogeneity, and wage dispersion arising from firm market power and job transition frictions. The model is used to examine the impact of ESHI on the wage distribution, and to counterfactually predict the effect of removing health insurance from the labor market via automatic enrollment in universal health-care. I estimate this model by the simulated method of moments (SMM) approach using data from the Survey of Income and Program Participation and the Kaiser Family Foundation Employer Health Benefits Survey. The fitted model parameters suggest that the utility of health insurance is valued at roughly a 12 percent increase in post-tax wages. I estimate that shutting down ESHI as a bundled non-wage benefit increases pre-tax wages by an average of \$470 per month, equivalent to a 10.5% increase in post-tax wages, or 73% of the estimated cost to employers of providing health insurance coverage. However, by simulating the transition from the baseline equilibrium, I find that while over 50% of the total expected wage gains accrue to workers within the first 2 years, it takes almost 10 years for the wage distribution to reach the new steady state.

JEL classification: I11, I13, J31, J32, J64

Keywords: health insurance; wage distribution; non-wage benefits; non-wage labor costs; on-the-job search

*I would like to thank my advisors Chris Taber and Naoki Aizawa for their invaluable guidance and support on this paper. I would also like to thank Carter Braxton, Chao Fu, Jesse Gregory, Victoria Gregory, John Kennan, Rasmus Lentz, Corina Mommaerts, Joseph Mullins, Elena Pastorino, Jeff Smith and seminar participants at the University of Wisconsin-Madison for numerous helpful comments and suggestions. I am grateful to Garrett Anstreicher for his teaching on optimization in Julia language. Any remaining errors are my own.

[†]Email: beemon@wisc.edu Website: www.danbeemon.com

1 Introduction

The United States health insurance market has long been a source of research interest given its status as the only industrialized country without a nationalized healthcare system. As of 2019, 54.63% of insured individuals get their health insurance through their employer,¹ and as such, health insurance is typically the primary non-wage benefit considered in labor supply decisions.

In recent years, there has been a growing policy interest in reforming the US healthcare market. In 2010, the passage of the Affordable Care Act represented the first significant attempt to expand health insurance coverage in over 40 years.² More recently, the COVID-19 pandemic, which saw mass unemployment accompanying an increased need for access to medical treatment and public health resources, has highlighted many of the negative effects of tying healthcare to employment. The result has been an increased call for the implementation of some form of universal healthcare, ranging from an affordable public option to a single payer “Medicare-for-all” policy. The primary focus of such policy proposals and the bulk of the accompanying research is on the expansion of access to healthcare and its impact on health and healthcare costs, and there is relatively less work examining the impact that employer sponsored health insurance (ESHI) has on the labor market.³ The goal of this paper is to study the impact of a potential universal healthcare policy on the labor market, and in particular on the distribution of wages. I also aim to estimate the speed with which the transition to any new labor market equilibrium would take place.

Much of the pre-eminent literature in this area suggests that if health insurance (or any other non-wage benefit) were removed from the labor market, wages must almost instantaneously increase so as to keep the total value of compensation the same. However, in the presence of substantial labor market power on the part of firms, the degree of wage adjustment is not so clear. In Berger et al. (2019), the authors posit that most labor markets more closely align with an oligopsony model, and estimate that the welfare impact of a substantial amount of firm-sided labor market power amounts to a roughly 5% decrease in lifetime consumption for workers. Similarly complicating matters is the Monheit et al. (1985) result that actually finds a positive relationship between the cost of health insurance and wages, though most contemporary design-based research finds contrary results.⁴ This raises the question of just how much wages are likely to adjust following an overhaul to the US health insurance system like those mentioned above.

As a result, in order to adequately assess the potential impact of universal healthcare on wages, I develop a model of on-the-job search with worker and firm heterogeneity, employer-sponsored and privately purchased health insurance, and wage dispersion arising from firm market power and job transition frictions. In this model, similar to Dey and Flinn (2005),

¹According to Kaiser Family Foundation 2019 “Health Insurance of the Total Population” estimates.

²Medicare, which provides health insurance only for individuals over the age of 65, and at significant private cost for all but the most basic coverage, was established in 1965.

³Bernie Sanders, chair of the Senate Budget Committee at the time of this writing, is a notable exception. Speaking to a crowd of hospitality workers in Nevada during the 2020 election cycle, he asserted that a benefit of a medicare for all policy was “Your employer will not have to pay \$15,000 a year for your healthcare, your employer will pay \$3,000. That’s a \$12,000 differential. You know who gets that 12,000? You get that 12,000.” Source: <https://www.nytimes.com/2020/02/21/podcasts/the-daily/bernie-sanders-nevada.html>

⁴see Currie and Madrian (1999) for an overview, or Anand (2017) for a more recent example.

workers of different ability levels search for jobs and match with firms of differing productivity levels. They receive an offer in the form of a wage bundled with a binary offer of health insurance (i.e. firms offer health insurance or do not) and accept it if they are at least indifferent between the firm’s offer and their outside option. Unemployed workers receive a wage that leaves them indifferent between their unemployment benefit and employment at the matched firm. Employed workers experience wage increases only in the event of matching with a new employer, as the two firms engage in Bertrand competition for the worker, and the individual may or may not switch jobs.

This model has several key mechanisms through which I might answer my research question. First, the model incorporates a utility parameter for health insurance that allows the value of ESHI to an individual to differ from the firm’s cost of providing health insurance. In particular, I assume diminishing marginal returns to the utility of consumption, which effectively results in heterogeneous preferences for health insurance across worker ability types, despite the homogeneous utility of health insurance. Second, wage determination through bidding for workers results in wage adjustments following a health insurance policy shock only in the result of a new competitive match, and the total level adjustment depends job matching rates. Third, I assume multi-worker firms such that they must commit to their health insurance provision decision prior to observing their worker match. This serves to approximate the reality that employers typically must offer health insurance to all or none of their employees, at least within a specific job category, and allows for another potential wedge between the cost of health insurance and the value to workers of providing it. The development of this novel search model with a commitment mechanism for health insurance, and the application of such a model to a Medicare-for-all type policy proposal are the primary contributions of this paper.

I characterize the solution to this model as the cross-sectional steady state of employment, the wage distribution and health insurance provision probability, similar to Postel-Vinay and Robin (2002). I estimate the parameters of this steady state model using data from the 2008 panel of the Survey of Income and Program Participation (SIPP), as well as years 2009-12 of the Kaiser Family Foundation Employer Health Benefits Survey (KFF). The former is a monthly representative panel of individuals over up to 48 months between 2008 and 2012, and includes information on wages, jobs held and job transitions, and health insurance coverage and its source. The latter is an annual survey of employers containing data on health insurance premiums and other firm-level statistics. Because these datasets do not contain matched worker-firm observations, and due to the complexity of the model to which an analytical solution is not feasible, given the health insurance provision, I estimate the model using the simulated method of moments (SMM). This method uses random draws to generate employment histories for a number of simulated workers, and the simulated data is then calibrated to fit the observed wage distributions and health insurance provision rates using a generalized method of moments (GMM) framework.

The estimated model provides a good approximation of the overall wage distribution and health insurance provision rate, as well as the individual distributions of wages with and without ESHI. The parameter value for the utility of health insurance suggests that the average worker value ESHI at the equivalent of roughly a 10.5% increase in post-tax wages. However, the full value of being employed at a job with health insurance is substantially higher, as both the wage and ESHI provision impact the stream of future wages.

I use the calibrated model to conduct a simple counterfactual experiment where all workers are automatically enrolled in free public health insurance. Computing the steady state equilibrium under the new policy, I find an increase in the average monthly wage of \$469.36, representing an 11.6% average wage gain. This raise, however, comprises only 73.4% of the benchmark model’s cost to firms of providing health insurance, suggesting that firms do not pass through the entirety of cost savings to workers, and so see a significant reduction in labor costs even after wages adjust. Analysis of the counterfactual wages by income group also suggests that lower income individuals disproportionately benefit from the wage adjustment, although absolute wage gains are higher for higher income individuals. The result is a 2 percentage point decrease in wage inequality in terms of Gini coefficient. Furthermore, despite mechanisms within the model that would allow for substantial inefficiencies, I find no evidence of any significant impact of ESHI on unemployment or job mobility.

I also simulate the transition over time to the new equilibrium under universal healthcare by using the last simulated year of the benchmark model as the initial conditions of the counterfactual simulation. I find that there is an initially sharp rise in wages as unemployment becomes more valuable and as the mass of competitive firms for each worker increases in the absence of ESHI costs, and so over 50% of the total expected wage gains accrue by the end of the first two years. However, wage growth quickly diminishes, and the full wage distribution takes almost 10 years to reach the new equilibrium. This suggests that any policy hoping to implement universal healthcare would need to explicitly take into account the long delay in the transfer of welfare gains from employers to workers in designing any revenue-neutral tax proposal.

The rest of this paper will proceed as follows. In section 2, I present the relevant literature in this area; section 3 presents the full structural model; section 4 describes the data used for estimation; section 5 details the estimation procedure and section 6 describes the results of that estimation and their model implications; section 7 presents the counterfactual experiments and their results; and section 8 concludes the paper and discusses potential avenues for future research.

2 Literature

This paper is related to three primary areas of literature. First, and most closely related is the research examining the relationship between health insurance and the labor market. Within this area, most work has been design-based,⁵ but there is a smaller strand focused on estimation of structural models of the labor market and health insurance. Fang and Shephard (2019) estimate a search model incorporating spousal insurance and the Affordable Care Act. Aizawa and Fang (2020) similarly estimate an equilibrium search model that incorporates firm size and healthcare expenditure to evaluate the impact of the ACA and its various individual components. And, similar to this paper though employing a different class of search model, See (2019) develops an equilibrium search model and studies the overall welfare effects of a single-payer healthcare system. See finds that welfare effects would be relatively

⁵See Currie and Madrian (1999) for an overview of reduced-form work on health insurance and the labor market, and Gruber and Madrian (2004) for a review of the literature on the impact of health insurance on labor supply and job mobility.

small, but positive, though employment would decline as reservation wages increase. The most closely related model is that of Dey and Flinn (2005), wherein health insurance enters utility as a binary normal good and within-firm wage dynamics result in considerable wage dispersion among identical workers. This paper expands on this model by introducing multi-worker firms and a greater degree of heterogeneity in order to build a richer model of wage determination in the presence of ESHI. Specifically, modeling multi-worker firms imposes that workers at the same firm must have the same health insurance benefits offered, which this paper posits as an importance potential source of inefficiency in the current labor market structure.

The second strand of literature to which this paper relates is that concerning the evaluation of health insurance policy reform. Of particular prevalence are those examining the impact of the ACA, such as the aforementioned studies by Fang and Shephard (2019) and Aizawa and Fang (2020), as well as Aizawa (2019), which studies alternative specifications for the ACA health insurance exchange (HIX), and Aizawa and Fu (2021) which considers cross-subsidization and risk pooling between ESHI and the HIX. This literature also includes a large body of work evaluating the impact of the Massachusetts health insurance reform that preceded it. Notably Kolstad and Kowalski (2012), Hackmann et al. (2012), Hackmann et al. (2015), and Kolstad and Kowalski (2016). However, working papers by See (2019) and Capatina et al. (2020), whose research focus is on medical expenditure shocks and their effects on human capital accumulation over the life cycle, are some of the only examples of papers that have touched on the potential impact of Medicare-for-all type policies. Furthermore, studies of the Massachusetts health insurance reform and the ACA, which are primarily changes in the intensive margin of employer sponsored health insurance provision, are not readily applicable to a policy which acts on the extensive margin, eliminating ESHI completely. This paper joins this very small body of work studying such a policy proposal and its relationship to the labor market.

Third, this paper relates to the more general area of literature estimating structural labor search models. Burdett and Mortensen (1998) and Postel-Vinay and Robin (2002) are the most prominent papers in this strand of the literature, and form the basis for Aizawa and Fang (2020) and See (2019), and Dey and Flinn (2005) respectively. In addition to these papers modeling on-the-job search, wage dispersion and (in some of them) health insurance provision, this paper relates to the class of models more generally modeling non-wage benefits. Hwang et al. (1998) examine the impact on hedonic wage theory of search models and in doing so develop a model of labor search where workers have heterogeneous preferences over non-wage benefits. In this paper, I contribute to this body of research by developing a search model of wage determination and health insurance provision that can readily be applied to any other (discrete) non-wage benefit.

3 Model

Here I introduce a model of job search with worker and firm heterogeneity, employer-sponsored health insurance and ESHI pre-commitment, and wage dispersion arising from firm market power and job transition frictions. The aim of this model, based on Postel-Vinay and Robin (2002) and Dey and Flinn (2005), is to properly characterize the rela-

relationship between the wage distribution and health insurance provision by estimating the wage determination process in a labor market where firms make a bundled $\{\text{wage, ESHI}\}$ take-it-or-leave-it job offer having previously committed to the ESHI offer at the time of job posting. This functions essentially as a model of wage determination with bundled non-wage benefits, specialized to fit imputed demand for health insurance. I begin by specifying the worker’s job search process and utility function. I then detail the firm’s objective function. Following this, I detail the timing of the model and the wage determination process given the model primitives, and finally, qualitatively characterize the model equilibrium.

3.1 Worker behavior

Workers search for jobs in a discrete time environment, and are heterogeneous only with respect to their ability level, ϵ , which is drawn from some continuous distribution with cdf F , and which is assumed to derive from some unobserved function of their human capital. This is fixed when they enter the labor market, and their ability level determines their unemployment “outside offer.” This is defined as the wage they could earn at some other profession or employment opportunity (i.e. minimum wage or self employment), or from unemployment insurance. This is modeled as a fixed offer b , and I assume an unemployed worker receives a wage of ϵb .

Unemployed workers are randomly matched with a firm with probability λ_u in each period, and assuming it is profitable to hire the specific worker, they will receive an offer $\{w, \eta\}$ where w is the wage and $\eta \in \{0, 1\}$ is ESHI. This model abstracts from individual insurance premiums, and so if ESHI is provided, the worker is assumed to pay nothing for health insurance. Employed workers match with a new firm with independent probability λ_e , and the worker receives a new $\{w, \eta\}$ offer that depends on the competition between the two firms, detailed further in section 3.2.

Workers also have the option to purchase private health insurance $H \in \{0, 1\}$ at price P , if it is not provided by their job, or they are unemployed. This price is exogenously determined by the current structure of the health insurance market and does not depend on the worker’s income.⁶ This price is assumed to be time invariant by workers. This is to say that any change to health insurance policy, such as in my counterfactual experiments, is assumed to enter the market as an unexpected shock.

Workers have a utility function that depends only on their income and access to health insurance. To produce heterogeneous demand for health insurance without increasing the parameter space, I assume diminishing marginal returns to income, $u(Y)$ with $u''(Y) < 0$, with a separable, constant marginal utility of health insurance. This results in an increasing marginal rate of substitution of health insurance for income, with the effect that stochastically, higher ability workers will forgo more wages in exchange for health insurance than low ability workers. Workers with wage w , health insurance offer η and individual health

⁶In the baseline model, I do not allow for individuals to receive health insurance through Medicaid or Medicare, and those individuals are excluded from the estimation sample. Note however, that only persons age 65 or above qualify for Medicare, and very few employed individuals qualified for Medicaid prior to the ACA, and so excluding these cases should not significantly bias my results.

insurance purchase decision H then have flow utility

$$\begin{aligned} v(Y, H, \eta) &= u(Y) + \alpha(\eta + (1 - \eta)H) \\ \text{s.t. } Y &= T(w) - (1 - \eta)PH \end{aligned}$$

where income Y is equal to the after-tax⁷ wage $T(w)$ minus the price of health insurance in the event it is not provided and the individual chooses to buy it, and α is the constant utility of health insurance, which is the same whether purchased privately or provided by the employer. This is the sole manner in which health insurance enters worker utility, as I abstract from any direct health implications of access to healthcare for the sake of tractability. Since the primary goal of this paper is to understand the firm's wage offer response to policy changes, the structure of individual health insurance preferences is not so important, only that, broadly speaking, workers like to have it, and so firms can pay them less in exchange for providing it.

Of course, introducing health insurance as a simple dummy variable that enters utility as a normal good perhaps requires additional justification. Most of the literature treats health insurance as, fittingly, insurance, where individuals face some random distribution of healthcare cost shocks and having health insurance mitigates or eliminates that risk. However, underlying these models is the assumption that the value of health insurance is solely a financial one, whereas a large body of empirical literature has found that the uninsured receive fundamentally different healthcare.⁸ So, it is both easier to estimate and perhaps equally representative of the true value of health insurance to workers to estimate an indirect utility parameter for having it or not.

Returning to worker behavior, employed workers continue to match with firms at the same rate, and, if a new offer gives higher utility, can choose to switch firms at no cost. Searching for jobs is assumed to be costless to the worker as well. Finally, employed workers face an exogenous job destruction rate δ .

3.2 Firm behavior

Similar to workers, firms are heterogeneous only in their productivity parameter, ρ , and receive revenue from hiring a worker of ability level ϵ equal to $\epsilon\rho$. This means that firms value higher ability workers, and higher productivity firms value an identical worker more than lower productivity firms. I assume a continuum of firms with some distribution G of productivity levels. Since all firms must be able to at least match the unemployment benefit in order to hire any worker, I assume a lower bound $\underline{\rho} = b$, as is standard in this class of models.

Firms post jobs at no cost, and hire a continuum of workers at some rate determined by workers' match probabilities. The exact matching rate for each firm is irrelevant to the model equilibrium as I do not track firm size and I assume profit is additively separable

⁷This tax function is relatively inconsequential as it only serves to change the interpretation of α from pre-tax to post-tax relative wage increase. However, as income tax is a realistic feature of the market and easy to implement, I estimate a tax function following Heathcote et al. (2017). Details for this tax function are included in Appendix A.

⁸See for example Monheit et al. (1985) and Kolstad and Kowalski (2012).

across workers. The only meaningful capacity in which firms are multi-worker is that, prior to matching, each firm must decide whether or not to offer health insurance to all potential workers. Each match results in an optimal take-it-or-leave-it wage, ESHI offer upon meeting the worker that depends on the worker's productivity level and current employment status. This commitment mechanism for health insurance is designed to approximate the reality that wages are determined on an individual level upon interviewing job candidates, but health insurance must typically be offered to all workers with analogous job descriptions. Modeling multi-worker firms in this way is a simplifying assumption to avoid more explicitly incorporating firm size, while introducing a realistic inefficiency in firms' ability to adjust offer bundles to fit individual workers.

If firms provide health insurance, they pay a fixed rate ψ per worker per period. So a firm with productivity level ρ employing a worker of ability ϵ with health insurance offer $\eta \in \{0, 1\}$ has flow profit

$$\pi = \epsilon\rho - w - \psi\eta$$

The firm will offer the worker the lowest possible wage such that this worker accepts the offer. However, if this wage is such that profit is negative, given the pre-committed choice to provide ESHI or not, the firm will be unable to hire this worker and profit is equal to zero.

3.3 Timing and wage determination

The timing of this infinite horizon model is as follows. In the pre period, firms observe the unemployment rate and the distribution of worker ability, and each firm makes a choice of whether to offer health insurance in their job postings. These firms then enter the market randomly according to workers matching rates. Each worker starts unemployed at the time they originally enter the market. The timing at which each worker enters the market is assumed to be random, and it is assumed that the distribution of worker abilities entering the market is time-invariant. I assume the market is in a steady state⁹ and workers enter each period with their employment and ESHI status determined from the previous period. Workers then choose whether to buy health insurance if it has not been provided, given their wage. Workers' flow utility is then realized. After realizing their flow utility, some employed workers exogenously lose their jobs with probability δ and enter the next period unemployed.

Workers who entered the period unemployed match with a firm with probability λ_u . Each firm observes the employment status and ability level of their matched worker, and, if profit (conditional on their ESHI decision) is non-negative, offers the worker a $\{w, \eta\}$ bundle. Firms offer unemployed workers the lowest wage that makes them indifferent between remaining unemployed and working at the matched firm. When indifferent, the unemployed worker is assumed to accept the offer. If this wage is such that profit is negative, no offer is made and the worker remains unemployed.

⁹There is an inherent initial conditions problem in simulating this class of models as without explicitly modeling age we cannot observe how long workers have been in the market prior to the sample period. To overcome this, I simulate an initial period in which workers are randomly assigned to unemployment or a firm of random productivity level according to the unemployment rate and distribution of ρ , and I assume that employed workers in this period receive their unemployment indifferent wage. I assume the data are in a steady state, and so I then simulate a 48 months of this labor market and use the final month as the initial conditions for my estimation period.

The remaining employed workers match with a new firm with probability λ_e . A new match is made, the new firm and the current firm engage in Bertrand competition for the worker, given their respective productivities and ESHI policies. Whichever firm is able to offer the combination of wage and ESHI that gives the worker the highest utility will hire the worker, and the worker will receive a wage that makes them indifferent between the old firm and new firm. If both firms have the same ESHI policies, the worker will choose the firm with the highest productivity level. These workers enter the next period employed at the winning firm and with the wage and ESHI status determined from this competition.

Employed workers who do not match with a new firm, do not lose their job, or match with a firm that cannot make a competitive offer given their current wage and insurance status, enter the next period with their wage and insurance status unchanged.

Each period, it is assumed that the labor market demographics are in a steady state. Of course, each period some workers permanently leave the market (i.e. retirement or death), and others enter the labor force, but because I assume ability and match productivity distributions are independent of age or market time period, the entry and exit of workers from the labor force is observationally equivalent to infinitely lived agents who simply lose their jobs randomly and start over in the wage distribution. In this way I abstract from any life cycle wage growth, and all wage growth is instead assumed to come from repeated on the job matching with higher productivity firms. This is done primarily for computational efficiency, as it reduces the number of states that must be kept track of, but also because the data available tracks individuals for only four years, and so a full life-cycle implementation is infeasible.

3.3.1 Value functions

Given these timing assumptions, I can write down the worker's value functions. For a worker with given ability level ϵ the value of being unemployed in any period t is equal to

$$U(\epsilon) = \max_{H \in \{0,1\}} u(T(\epsilon b) - PH) + \alpha H + \beta \left\{ (1 - \lambda_u)U(\epsilon) + \lambda_u \mathbb{E}_{\rho_{t+1}} \mathbb{E}_{\eta_{t+1}(\rho_{t+1})} V(\epsilon, w_{t+1}, \rho_{t+1}, \eta_{t+1}) \right\} \quad (1)$$

with discount factor $\beta \in (0, 1)$, taking expectation over possible firm productivity ρ_{t+1} draws and the probability of being offered health insurance η_{t+1} , which will depend on ρ_{t+1} . The worker has perfect information regarding the determination of wages, which will depend on the worker's time-invariant ability level, current employment status, and the productivity of the firm with which she is matched, and the ESHI probabilities, which depend only on the matched ρ_{t+1} productivity firm. Note that flow utility for unemployed workers is the same in every period, and the expectations over ρ and η will be the same for all unemployed workers, and so $U(\epsilon)$ is time invariant and only depends on the workers ability level.

Similarly, the value to a worker of being employed by a firm with productivity level ρ_t ,

at wage w_t , and with ESHI provision η_t is equal to

$$V_t(\epsilon, w_t, \rho_t, \eta_t) = \max_{H \in \{0,1\}} u(T(w) - (1 - \eta)PH) + \alpha(\eta + (1 - \eta)H) + \beta \left\{ \delta U(\epsilon) + (1 - \delta) \left\{ \lambda_e \mathbb{E}_{\rho'} \mathbb{E}_{\eta_{t+1}(\rho_t, \rho')} V_{t+1}(\epsilon, w_{t+1}(\rho_t, \rho'), \rho_{t+1}(\rho_t, \rho'), \eta_{t+1}(\rho_t, \rho')) + (1 - \lambda_e) V_{t+1}(\epsilon, w_t, \rho_t, \eta_t) \right\} \right\} \quad (2)$$

where again the wage next period, in the event of a match, depends on the productivity of the current firm and the new firm productivity draw ρ' . Next period's firm productivity ρ_{t+1} will depend on the relative productivities of current firm ρ_t the matched productivity ρ' , and their respective ESHI offers. Otherwise, if the worker experiences job destruction with probability δ , they become unemployed next period, whose value depends only on ϵ . If the worker does not lose her job and receives no alternative offer, she receives the same wage and ESHI offer from the current ρ_t firm next period.

As stated above, wage determination follows a very particular process. Workers have no bargaining power, in the manner of Postel-Vinay and Robin (2002). Instead, firms bid against each other, and the worker gets her indifference wage. The result of this feature of the model, in addition to the desired wage growth and dispersion effects, is that wages are deterministic given current productivity ρ , new matched productivity ρ' , and the two firms' pre-committed ESHI offers η and η' respectively. Borrowing notation from Postel-Vinay and Robin, I represent the unemployed and employed indifference wage, respectively, that workers can expect by $\phi_0(\epsilon, \rho, \eta)$ and $\phi(\epsilon, \rho, \rho', \eta, \eta')$. Since the firm will offer the lowest wage possible to incentivize the worker to accept the offer, for unemployed workers, this wage satisfies the following:

$$\phi_0(\epsilon, \rho, \eta) = w^* \text{ s.t. } U(\epsilon) = V(\epsilon, w^*, \rho, \eta) \quad (3)$$

Similarly, for employed workers, the wage that incentivizes the worker of ability level ϵ to switch from a firm of productivity ρ to a firm of productivity ρ' satisfies

$$\phi(\epsilon, \rho, \rho', \eta, \eta') = w^* \text{ s.t. } V(\epsilon, \bar{w}(\rho, \eta), \rho, \eta) = V(\epsilon, w^*, \rho', \eta') \quad (4)$$

where $\bar{w}(\rho, \eta)$ is the maximum wage offer the losing firm can make given its choice of η . Importantly, the job transition will depend on the respective productivities and health insurance policies. In section 3.2, I mention that firms will make an offer only if profit is non-negative. The result is that though it is probable workers will choose to accept the indifference wage offer from the higher productivity firm, cases may arise where the relative productivity gains of the more efficient firm are not sufficient to cover the cost of outbidding the lower productivity firm, if the higher productivity firm has committed to offering health insurance and the lower has not. This can produce productivity-inefficient job transitions in some cases. I discuss this explicitly in section 3.4.

Observing these indifference wages and anticipating the ESHI probabilities of competing firms, a firm with productivity level ρ , committed to ESHI provision choice η has a value of

matching with an unemployed, productivity level ϵ worker of

$$J_t^u(\epsilon, \rho, \eta) = \mathbb{1}_{\pi > 0} \left\{ \pi(\epsilon, \rho, \phi_0(\epsilon, \rho, \eta), \eta) + \beta(1 - \delta) \left\{ (1 - \lambda_e)J_{t+1}^u(\epsilon, \rho, \eta) + \lambda_e \mathbb{E}_{\rho'} J_{t+1}^e(\epsilon, \rho, \rho'_{t+1}, \eta, \eta'_{t+1}) \right\} \right\} \quad (5)$$

where $J^e(\epsilon, \rho, \rho', \eta, \eta')$, the value of matching with a worker currently employed at another firm with productivity ρ' and ESHI provision choice η' , is equal to

$$J_t^e(\epsilon, \rho, \rho'_t, \eta, \eta'_t) = \mathbb{1}_{\pi > 0} \left\{ \pi(\epsilon, \rho, \phi(\epsilon, \rho'_t, \rho, \eta'_t, \eta), \eta) + \beta(1 - \delta) \left\{ (1 - \lambda_e)J_{t+1}^e(\epsilon, \rho, \rho'_t, \eta, \eta'_t) + \lambda_e \mathbb{E}_{\rho'_{t+1}} J_{t+1}^e(\epsilon, \rho, \rho'_{t+1}, \eta, \eta'_{t+1}) \right\} \right\} \quad (6)$$

If the matched worker is unemployed, provided the firm can afford to hire her ($\pi > 0$), the firm realizes flow profits. The next period, the firm expects that, provided the job is not exogenously destroyed, this worker will match with another firm with probability λ_e . If the worker does not match with another firm, this firm gets the same value J^u in the next period. If the worker does match with another firm, this firm must compete with the new firm and will get the employed match value J^e of this worker, taking the expectation over all possible productivity levels of the competing firm.

This employed value functions similarly, but will depend on $\phi(\epsilon, \rho', \rho, \eta', \eta)$, i.e. the wage the firm needs to offer the worker to make her indifferent between staying at the current firm and switching. If this wage results in negative profits, the firm loses the worker and the job is destroyed. If the firm hires the worker, then with probability $(1 - \lambda_e)$ they keep the worker at the same wage next period, and with probability λ_e the worker matches with a new firm and this firm receives the value of competing with the new firm.

Given these value functions, in the initial period (period zero) the firm's probability of offering health insurance is determined by solving

$$\max_{\eta \in \{0,1\}} \mathbb{E}_\epsilon \left(u(\epsilon) J^u(\epsilon, \rho, \eta) + (1 - u(\epsilon)) \mathbb{E}_{\rho'|\epsilon} \mathbb{E}_{\eta'|\rho'} J^e(\epsilon, \rho, \rho', \eta, \eta') \right) - \mathbb{1}_{\eta=1} \xi \quad (7)$$

where $u(\epsilon)$ is the unemployment rate of workers of ability level ϵ , which will be determined by δ , λ_u , and also the ability distribution and firms' ESHI policies. Note that the probability a worker is currently employed at a given productivity level firm will depend on that worker's ability level since non-monotonicity of productivity in job transitions means ϵ and ρ distributions are not independent given current employment.¹⁰ This holds true only in

¹⁰Without health insurance, as in Postel-Vinay and Robin (2002), workers will always move from ρ to ρ' given that $\rho' > \rho$. The results is that, though the expected productivity level of a matched firm increases with employment duration, this is equally true of all ϵ workers, and so $\mathbb{E}(\epsilon|\rho) = \mathbb{E}(\epsilon)$, and vice-versa. However, with ESHI, if for example the same ρ does not provide health insurance but ρ' does, some workers of low ability may not be productive enough for ρ' to attract them away from ρ , given the cost of ESHI, and so for some particular ability level $\tilde{\epsilon}$, $P(\epsilon = \tilde{\epsilon}|\rho') = 0$ and so $\mathbb{E}(\epsilon|\rho') \neq \mathbb{E}(\epsilon)$.

period zero, as once a firm has employed a given ability level ϵ worker, the probability of that worker meeting a firm of productivity level ρ' is independent of the worker's ability.¹¹

Here I also introduce ξ , a stochastic cost shock introduced primarily to aid in equilibrium convergence. This can be justified as a randomized "health plan participation fee" that is a sunk cost after period zero, and has the effect of randomizing firms' ESHI decision such that $\eta(\rho) \in [0, 1]$ becomes a continuous function. I discuss this further in section 5.

3.4 Description of equilibrium

Due to the complexity of the model given multi-worker firms with committed health insurance provision decisions, an analytical solution is infeasible. As such, I solve the model computationally using the algorithm specified in section 5.3. However, here I provide some qualitative features of the equilibrium given the model's design and the calibrated parameter estimates.

Equilibrium is defined in this model as the fixed point of the mapping between wage function $\phi_0(\epsilon, \rho, \eta(\rho))$ and $\phi(\epsilon, \rho, \rho', \eta(\rho), \eta'(\rho'))$, and ESHI provision probability function $\eta(\rho)$, which is equal to the probability that the solution to equation (7) is equal to one, given the distribution of ξ . Since the indifference wages are deterministic, given a fixed $\eta(\rho)$ probability function, ϕ_0 and ϕ can be directly calculated using value function iteration.

In equilibrium, this $\eta(\rho)$ is strictly increasing in ρ as higher productivity firms offer higher wages, holding ϵ constant, and the marginal rate of substitution of income for health insurance increases with wages. This is also at least partly due to the fact that higher productivity firms earn higher revenue from each type of worker and so are more likely to be able afford the cost of providing health insurance. Similarly, the provision of health insurance has the potential to create an unemployment wedge amongst low productivity workers and firms as the cost of ESHI may exceed their joint productivity. However, in the simulated data this wedge appears to be of limited importance as the difference in the unemployment rate with the removal of ESHI is insignificant.

Another important feature of the model is that in equilibrium, though firms are able to decrease wages in exchange for providing health insurance, it should be the case that, stochastically, higher wage jobs are more likely to provide health insurance, since $\frac{d\eta(\rho)}{d\rho} > 0$. Indeed, in the simulated data, with roughly 85% of jobs providing health insurance, wages at jobs with ESHI are 74% higher than wages at jobs without ESHI.

Finally, it should be the case that in the event of a change to health insurance policy, i.e. a change in P , there should be some delay between the beginning of the policy and convergence to the new steady state equilibrium wage distribution as, generally speaking¹², wage increases require meeting with a new (and competitive) firm. Initial offers out of unemployment should see the quickest increase as the value of unemployment as an outside option increases with more affordable private insurance. The exact degree of this change however will depend

¹¹The worker's offer, and therefore the match value is dependent on the worker's ability, but the worker can still "match" with any firm. This is to say that λ_u and λ_e are independent of worker or firm productivities and if gainful employment cannot be produced, both sides are out of luck until the following period.

¹²In section 7, I discuss the exception in the first period, where workers lower in the ESHI wage distribution for whom unemployment has become preferable to employment at the current firm and wage will be able to renegotiate and raise their wage back to the indifference level.

on the relative value of the insurance and the cost to firms of providing it, as well as the average duration of employment. This is because as workers remain employed and receive additional offers, their wage rises closer to the level of their joint productivity with the firm. Conversely, the percentage of this surplus allocated to workers decreases as the job destruction rate increases.

4 Data

To estimate this model, I primarily utilize the 2008 panel of the Survey of Income and Program Participation (SIPP). The SIPP contains monthly data on individual labor market outcomes and health insurance coverage, as well as demographic and household composition variables. These data are collected during interviews conducted every four months, up to twelve times, so individuals may be tracked for up to four years. I use the core module of this data¹³ containing detailed labor market variables including earnings, employment status, number of weeks worked, industry of the job worked, as well as whether the individual changed jobs during each month of the four month survey period. Health insurance status is recorded as well, including whether insurance was provided by the employer or privately purchased, or came from Medicare or Medicaid, and whether the individual was covered under their own insurance plan or someone else's.

For estimation, I restrict the sample considerably to match the labor market environment detailed in the model. I retain only individuals participating fully in the labor force for the full 48 months, and who only receive health insurance from an employer or private purchase. For this, I drop individuals below 18 years of age and above 64 years of age, any individual enrolled in school, as well as any individual who receives Medicaid, Medicare, or obtain health insurance through a family member. I also drop from the sample any individual who is not observed to work at any point during the four year sample period. I further remove individuals working in the public sector, or self-employed, as these types of jobs do not fit the ESHI incentive structure in the model.¹⁴ I also restrict the sample to individuals not receiving social security income and not involved with the military. To remove some likely outliers in income or program participation, while continuing to maintain a good deal of worker heterogeneity in productivity, I retain only individuals who have completed at least a 10th grade education. Finally, I restrict the sample to those individuals who are in the data for the full four years, given all of the above restrictions, so as to obtain a balanced panel. This leaves an estimation sample consisting of 5,014 individuals tracked over 48 months for a total of 240,672 observations.

In addition, since the SIPP does not collect information on employers beyond the industry and a longitudinal identifier, I use data from the Kaiser Family Foundation Employer Health Benefit Survey (KFF) from 2009-2012, the years corresponding to those present in the SIPP

¹³The SIPP contains a core module, with the primary variables on labor market outcomes and health insurance coverage, as well as other public program utilization (i.e. food stamps, etc.), in addition to a topical module, with various other topics, including some expenditures such as health insurance premiums or tuition, collected annually. I do not use the topical module as the model abstracts from worker premiums for ESHI, but this data could be used for such an extension to the model in the future.

¹⁴Public sector jobs are overwhelmingly likely to provide health insurance, and do so for reasons related to public policy rather than profit maximization.

panel, to obtain estimates for employer sponsored health insurance premium costs. The KFF contains annual data, from firms in the private and public sector, on firm characteristics such as size, industry, and employee demographics, as well as, for my purposes, health insurance premiums paid. I restrict the sample to firms in the private sector, and utilize this data only to estimate the employer monthly premium parameter ψ .

4.1 Summary statistics

Summary statistics for the 2008 SIPP panel are presented in Table 1 below. As this sample is restricted to individuals that appear in all 48 months, the average months employed is quite high at 44.84, despite 30.8% that experience at least one month of unemployment. Just under half of the individuals are ever observed switching jobs, defined as starting a new job either out of unemployment or from a different employer, data for which is recorded using distinct employer IDs. Similarly, individuals are observed to transition between jobs at a rate of only 1.4%.

From the insurance coverage data, 17.9% of the sample are uninsured, which matches closely with the Kaiser Family Foundation's reported national averages, once Medicaid, Medicare, and military personnel are excluded. 80.1% of the sample receive insurance through their employer (and 85% of employed individuals), and less than 2% purchase insurance privately. It can also be seen that the average monthly wage at jobs providing health insurance is almost twice that of jobs without ESHI. This is a larger discrepancy than is reported in much of the literature, and I have little to offer as an explanation as to why this is the case here. This is also the data trend that has proved the most difficult to closely replicate in model simulations.

Table 1: SIPP 2008 Summary Statistics

Variable	Selected Sample	
	Mean	S.D.
	Employment	
Months employed	44.84	7.85
Ever unemployed	.308	.462
Ever switch jobs	.459	.498
New job in the month	.022	.146
Job-to-job transition	.0140	.117
Exited unemployment	.142	.349
Became unemployed	.0103	.101
	Insurance Coverage	
Uninsured	.168	.374
ESHI	.813	.390
Privately Insured	.019	.137
	Earnings	
UI benefits	1,171.11	590.87
Wage	4,104.98	2368.02
Wage, ESHI	4,427.39	2,351.31
Wage, uninsured	2,274.08	1,453.84
N = 5,014		

5 Estimation

In this section, I detail the estimation procedure used to obtain parameter estimates. I estimate the model using a two step process. In the first step, exogenous model primitives and several structural parameters are estimated from observed sample means. In the second step, the remaining structural parameters are estimated using the iterative algorithm detailed in section 5.3. To begin with however, I will detail the identification argument and the assumptions necessary to overcome data limitations.

5.1 Identification and assumptions

I begin with some simplifying assumptions not essential for identification but whose implications either aid in comprehension or substantially reduce computation time (or both). First, though the results should extend to a generalized concave utility function, the precise shape

of the utility of income does not qualitatively alter the functioning of the model,¹⁵ and so I assume log utility of income to eliminate an estimation parameter.¹⁶ Quantitatively, results may change slightly in the presence of very high concavity, but log utility remains firmly within the bounds of the literature’s estimation of reasonable risk aversion.¹⁷ Similarly, I normalize the unemployment benefit b to 1, since the level of this value is inconsequential, and normalizing it so allows me to identify the distribution of ϵ directly from the sample distribution of unemployment “wages.” I also parameterize the productivity distributions for both the worker and firm, $F \sim \log N(\mu_\epsilon, \sigma_\epsilon)$, $G \sim \log N(\mu_\rho, \sigma_\rho)$ as these distributions must be positive and together characterize wage distributions, of which lognormal distributions typically provide a good approximation. Note, however, that since all observed values of ρ must be above b , I assume instead $\rho - b \sim \log N(\mu_\rho, \sigma_\rho)$. With these assumptions, my parameters of interest reduce to $\{\alpha, \mu_\epsilon, \sigma_\epsilon, \mu_\rho, \sigma_\rho\}$ along with exogenous firm health insurance costs ψ , the similarly parameterized distribution of cost shock $\xi \sim \log N(\mu_\xi, \sigma_\xi)$ where I assume $\mu_\xi = 0$, and the transition parameters δ , λ_u and λ_e . Following the work of many before me, and deriving from Flinn and Heckman (1982), I do not attempt to separately identify the discount rate β , and instead fix it at 0.995, which gives an annual interest rate of 6%.¹⁸

I estimate these parameters using a two-step process whereby exogenous parameters $\{\psi, \delta, \lambda_u\}$ are identified from sample means in the data. The assumption that any employer will be able to hire any worker out of unemployment, and that workers will accept any offer that makes them at least indifferent, allows me to estimate λ_u from the mean monthly probability of transitioning out of unemployment in the SIPP data.¹⁹ The job destruction rate δ is also set equal to the mean monthly probability of becoming unemployed observed in the data. The firm cost of providing health insurance is estimated using the KFF data, whereby monthly premium ψ is set equal to the weighted sample mean monthly premium cost per employee, while the cost shock parameter σ_ξ is estimated in the second step.

The identification of the distribution of firm productivity, the utility of health insurance, and the variance of the unobserved cost shock require more justification. The parameterization of the productivity distribution to a log normal coupled with fixed job transition rates and fixed parameterized distribution of ϵ makes the employed wage distribution sam-

¹⁵This is true of incorporating an income tax function into the model as well. Including taxes changes the interpretation of α from pre-tax to post-tax percent marginal income increase. But with or without taxes in the model, the pre-tax wage value of health insurance remains unchanged.

¹⁶This also speeds up computation slightly as Julia’s built in log function evaluates more quickly than, for example, a CRRA functional form with risk aversion parameter $\gamma = 1.1$. In tests this reduced the time it takes to solve the model by about 40% and did not significantly alter the results.

¹⁷See Chetty (2006) for discussion of reasonable risk aversion parameters in labor supply models. He finds a mean estimate across studies of $\gamma = 0.71$, and estimates bounds between 0.15 and 1.78. As such, $\gamma = 1$ is selected for this paper for its computational expediency and the interpretability it lends to the estimated value of α .

¹⁸This is higher than the rate of wage growth observed in my sample and the BLS reported interest rate during this period, however, with monthly data, increasing β above 0.995 begins to substantially slow down the convergence of the value function. Increasing the interest rate while remaining at a reasonable value observed in much of the literature served to decrease computation time of the full model equilibrium by roughly 50% and did not impact the goodness of fit of the estimated model.

¹⁹I previously discussed a potential unemployment wedge that would cause this assumption to fail, but given that I do not find such an inefficiency from the estimated parameters, this assumption seems justified.

ple moments sufficient to identify μ_ρ and σ_ρ . This leaves α to be estimated as the primary determinant of wage depression resulting from health insurance, and σ_ξ as a partial determinant of the ESHI coverage rate. An increase in α increases the ESHI provision rate as providing health insurance becomes more profitable with larger compensating wage decreases, but decreases the difference between ESHI and non-ESHI mean wages as the former wages are increasingly depressed in exchange for health insurance, while the latter wages must be increased to offset the lack of ESHI. An increase in σ_ξ holding other parameters constant decreases the insurance provision rate since providing it becomes more costly on average, but also decreases the insurance-non-insurance wage gap, as higher variance in expected costs increases the probability that lower productivity firms provide health insurance and decreases the probability that higher productivity firms provide it. These comparative statics allow me to separately identify these parameters using moments concerning the difference in the wage distribution for jobs with and without ESHI, and the overall health insurance provision rate from the SIPP data.

5.2 First step

In the first step I estimate a number of model parameters from observed sample means. As stated above, I am able to estimate λ_u directly from the observed transition rates out of unemployment. The percentage of unemployed individuals in each month that switch to employed in the following month is averaged over the 48 month sample and assumed to be constant over the sample period. An analogous strategy is used to estimate the job destruction rate. Similarly, ψ is fixed at the average monthly cost per employee of providing health insurance reported in the KFF data, weighted by the SIPP sample percentages of workers on individual vs. family insurance plans. μ_ϵ and σ_ϵ are then estimated by fitting a lognormal distribution to the observed unemployment income in the SIPP selected sample.

5.3 Second step estimation algorithm

Fixing the first step parameters at their point estimates, I proceed to estimate the remaining parameters using the simulated method of moments (SMM). Broadly, this consists of solving the model for a given set of parameters, simulating a dataset from the model solution, and using this dataset to compute model moment equivalents. These moments take the place of model moments that are an analytical function of the parameters in the generalized method of moments (GMM). I use cross-sectional moments on wages, insurance provision, employment, and the interaction of these variables to compute the SMM objective function, and solve for the model parameters that minimize this objective using the Nelder-Mead algorithm. If \mathbf{m} is the vector of targeted moments, and θ is the vector of model parameters let

$$\mathbf{M}(\theta) = [\mathbf{m} - \mathbb{E}[\mathbf{m}|\theta]] \quad (8)$$

and the objective function is constructed as

$$\min_{\theta} \mathbf{M}(\theta)' \mathbf{W} \mathbf{M}(\theta) \quad (9)$$

where \mathbf{W} is a weighting matrix consisting of the inverse of the variance of the moments, adjusted individually to weight the moments more essential to obtaining reliable counterfactuals. To generate $\mathbf{M}(\theta)$, the model is solved using the following algorithm:

1. Given distribution parameters, discretize the worker and firm productivity distributions.
2. Compute unemployment value function $U(\epsilon)$, since this does not depend on the wage distribution at equilibrium.
3. Initialize the employment indifference wage function $\phi(\epsilon, \rho, \rho', \eta, \eta')$ and ESHI function $\eta(\rho)$. The initial guess for the wage function sets the wage necessary to switch equal to the maximum wage the losing firm can afford to pay the worker. The initial guess of $\eta(\rho)$ doesn't matter so much, as this adjusts quickly, and is set to 0.5 for all firms.
4. Iterate on the indifference wage function and $\eta(\rho)$ to find the fixed point
 - (a) Given productivity distributions and initial guesses, solve workers' value functions.
 - (b) Given value functions solve for indifference wages, with and without ESHI, for unemployed, and for employed given current productivity employer, $\hat{\phi}_0(\epsilon, \rho, \eta)$ and $\hat{\phi}(\epsilon, \rho, \rho', \eta, \eta')$.
 - (c) Given indifference wages, solve for ESHI probability $\hat{\eta}(\rho)$ at each ρ of the discretized distribution.
 - (d) Evaluate $d(\phi_0, \hat{\phi}_0)$, $d(\phi, \hat{\phi})$ and $d(\eta(\rho), \hat{\eta}(\rho))$, where $d()$ is a distance metric. If these are less than the specified tolerance level, the equilibrium has converged. If not, update each function according to $\phi_0 = \omega\phi_0 + (1 - \omega)\hat{\phi}_0$, $\phi = \omega\phi + (1 - \omega)\hat{\phi}$, and $\eta(\rho) = \omega\eta(\rho) + (1 - \omega)\hat{\eta}(\rho)$, for $\omega \in (0, 1)$, and return to step 4(a). Repeat until convergence.
5. From the computed model equilibrium, simulate job histories for N individuals over T months to reasonably match the horizon of the true data.²⁰
6. Evaluate the moment function and update parameter guesses. Repeat until convergence criterion are met.

This method is used to obtain point estimates of the model parameters. The standard errors of these parameters are estimated in the manner of GMM, where if $\mathbb{M}(\theta) = \mathbb{E}[\frac{\partial \mathbf{M}(\theta)}{\partial \theta'}]$ is the gradient matrix of the moment conditions evaluated at $\hat{\theta}$, and if $\mathbf{\Omega} = \mathbb{E}[\mathbf{M}(\theta)\mathbf{M}(\theta)']$ is the variance-covariance matrix of the moment conditions, then the asymptotic variance of $\sqrt{n}(\hat{\theta} - \theta)$ is

$$[\mathbb{M}(\theta)' \mathbf{W} \mathbb{M}(\theta)]^{-1} \mathbb{M}(\theta)' \mathbf{W} \mathbf{\Omega} \mathbf{W} \mathbb{M}(\theta) [\mathbb{M}(\theta)' \mathbf{W} \mathbb{M}(\theta)]^{-1}$$

²⁰In this case, N is set to 5000 and T is set to 96, of which the first half are thrown out and the final 48 months are used to calculate simulated moments. This is done to overcome the initial conditions problem that arises in simulating steady state models such as this one. I have experimented with alternate pre-estimation time horizons and after any number past 48 months (i.e. 96 total) the results were unchanged.

and the standard errors are the square root of the diagonal elements of this matrix divided by n . The estimator of the asymptotic variance cannot be simplified since \mathbf{W} cannot be said to be a consistent estimator of the efficient waiting matrix. Note also that, as in Petrin (2002) and Aizawa and Fang (2020), I assume that $\mathbf{\Omega}$ is a block diagonal matrix since different moments are derived from different sampling and sub-sampling processes.

6 Estimation Results

Here I present the results of the first and second step estimation and examine the resultant model fit. Using the above specified method, I obtain structural parameter estimates that are reported in Table 2 along with the estimates from the first stage. Few of these estimates are individually informative, but the model solution derived from them is illustrative. On its own, this estimate of α suggests that health insurance is worth the utility equivalent of an 11.7% post-tax wage increase. At the average insured wage this corresponds to an increase of \$471.65 post-tax or \$561.64 pre-tax dollars per month. At the true average uninsured wage, these numbers are only \$254.44 and \$288.59 respectively.

Table 2: Parameter estimates

First Step		Second Step	
Parameter	Estimate	Parameter	Estimate
λ_u	.1417 (.0028)	α	.117 (.00581)
δ	.0103 (.00022)	μ_ρ	1.138 (.00052)
ψ	639.79 (4.010)	σ_ρ	0.415 (.0189)
μ_ϵ	7.089 (.0065)	σ_ξ	1.727 (.0201)
σ_ϵ	.500 (.0046)	λ_e	0.0634 (.00056)

Standard errors in parentheses.

However, the flow utility is not the entire picture of the value of health insurance to the worker. Using the calibrated estimates for ϕ_0 and ϕ , I calculate the average difference in the indifference wage to entice a matched worker is \$1,023 lower at a job with health insurance than without. This is likely because of the cumulative effect of acquiring health insurance on future wages, as new firms offering health insurance will need to pay higher wages if you are already provided ESHI. This is increasingly likely as workers transition to higher productivity jobs as the probability of being offered health insurance increases rapidly with ρ .

6.1 Goodness of fit

The model does a good job of matching targeted moments, and these comparisons are presented in table 3. The average unemployment benefit is somewhat higher than the true data, but the employed mean wage fits very closely, suggesting perhaps that the ρ distribution is fit better than ϵ . Similarly the unemployment rate and mean insured wage fit reasonably well. Importantly, the job-to-job transition rate and ESHI provision rate match the true data very closely, as the parameters that govern these moments most directly impact the counterfactual predictions. The largest shortcoming of the model seems to be underestimating slightly the difference between the insured and uninsured wage, as the mean uninsured wage is a bit higher than in the true data. This likely stems from constraining the variance of the ρ distribution to match the employed wage variance more closely, as the insured-uninsured wage gap increases with σ_ρ .

Table 3: Targeted moment comparison

	Simulated	True
Mean wage	4,045.75	4,104.98
Mean unemployment benefit	1360.58	1,171.11
Unemployment rate	6.93%	6.58%
Job-to-job transition rate	1.42%	1.40%
ESHI, employed	85.53%	85.03%
Mean wage, ESHI	4,310.55	4,427.39
Mean wage, uninsured	2480.21	2,274.08
Difference	1,830.34	2,153.31

Goodness of fit based on targeted moments. Overall wages and insurance rates are fairly close fits. Underestimates the the average insured-uninsured wage difference slightly.

In addition to the comparison of means, I also present fitted distributions of the true and simulated wage data in figure 1. These distributions appear to fit closely, although the use of lognormal approximations for both depictions could be responsible for their similar shapes. The simulated data appears to have a higher variance and so more mass at the right tail and a corresponding increase in density at a lower point on the distribution, giving a closely matched average wage. Figure 2 shows the comparison of the simulated distribution of wages with and without health insurance, as well as a comparison with the same distributions in the true data. From figure 2a one can see that these distributions differ greatly, as wages with ESHI have a substantially higher variance, while non-ESHI wages are more closely grouped at the low end of the distribution. Figure 2b shows that these distributions match up reasonably well with their true data counterparts, both individually and in relation to each other. Overall, these results suggest that the model gives a relatively accurate approximation of the key features of the observed SIPP data. This allows me to use the fitted model to

conduct a counterfactual experiment on the availability of public insurance, which I present in the next section.

Figure 1

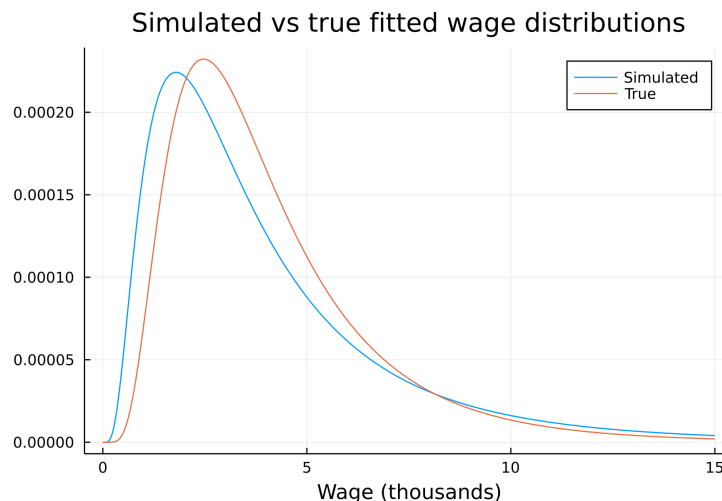



Figure 2: Wage distribution with and without ESHI



7 Counterfactual with Universal Healthcare

In this section I present the design and results of a simple policy counterfactual that simulates the introduction of a rather ambitious implementation of universal healthcare. Using the parameter estimates presented above, I conduct a counterfactual analysis in which all individuals are automatically enrolled in free, government provided health insurance. Though perhaps approaching the edge of realism in terms of government policy, this analysis merely seeks to examine the potential impact of completely severing the tie between health insurance and the labor market. Given the extreme nature of this policy counterfactual, this paper does not consider any tax scheme to pay for such a plan as any successful implementation

would be highly complex. However, a useful extension to this work would be to consider an optimal tax function such that providing universal healthcare is revenue neutral, though such an endeavor may need to incorporate many externalities that arise from changes in healthcare policy that go beyond the scope of this labor market analysis.

7.1 Counterfactual experiment design

Mechanically, this counterfactual is very simple. The automatic enrollment is achieved in the simulation program by setting the price, P , at which individuals can purchase private insurance, to zero. Given the model assumption that health insurance is a normal good, the result is that all individuals choose to purchase it.²¹ I then solve the model for a new equilibrium holding all other parameters fixed at their calibrated values, and simulate an identical set of individuals (using the same random seed) over the same time horizon as in the baseline simulation to predict any change in the wage distribution. This simple counterfactual actually works quite nicely, as the model quickly converges to a new equilibrium where no firms of any productivity level offer health insurance. This initial counterfactual simulation gives a picture of an alternate labor market where job histories progressed exactly as in the baseline simulation, but wages are determined according to this new equilibrium wage function.

Perhaps of more interest, however, is the subsequent simulation wherein I simulate the transition, over time, from the baseline wage distribution to the new equilibrium. To do this, I use the last year of the baseline model simulation as the initial conditions, and simulate ten additional years for these simulated individuals under both the baseline and the counterfactual equilibria and track the difference in the evolution of wages.

One important functional change in this counterfactual transition simulation is that with automatic enrollment in free universal healthcare, the value of unemployment increases for all workers. The result is that in the first period following the policy change, workers for whom unemployment has become preferable to employment at the initial wage and firm productivity level will renegotiate their salary to return to at least indifference between employment and unemployment, even without matching with a new firm. This will immediately increase the wages of any worker who is still employed by their first match out of unemployment and did not match with any other firms, as such a worker will have an initial wage that makes them exactly indifferent to unemployment under the old regime. Other workers who remain close to their initial wages may also see immediate gains. However, the remainder of employed individuals will be unaffected by the policy change until they receive a competitive outside offer.


7.2 Counterfactual results

The resulting counterfactual wage distribution from the first simulation is presented in figure 3 next to the baseline wage simulation. With the introduction of universal healthcare, firms can no longer depress wages in exchange for providing health insurance. However, firms that

²¹Though this is presented as a private health insurance decision, the extent to which “choosing” to purchase a free public good is in fact any choice at all is purely a semantic debate. The functional result of this modeling decision is that everyone is enrolled in health insurance.

were providing health insurance no longer pay the respective premiums, and are no longer committed to a potentially individual match-inefficient ESHI provision decision. Similarly, firms that were not providing health insurance become more competitive with relatively lower wages. The result appears to be a rather modest wage adjustment, although the difference in the average wages, represented by the vertical lines, amounts to a monthly wage increase of \$469.36, or \$5,632.32 annually. This represents an 11.6% increase in the average (pre-tax) wage from the baseline model. Some additional comparisons between the baseline and counterfactual simulations over a corresponding 4-year period are presented in table 4. In addition to the moderate average wage increase, removing ESHI from the labor market decreased wage inequality slightly, as measured by a 2.1 percentage point drop in the Gini coefficient, due to the distributional effects mentioned below. Interestingly, the unemployment rate decreased only very slightly, indicating the employment wedge caused by commitment to ESHI provision is exceedingly small. Similarly, the job to job transition rate and joint productivity was virtually the same across samples, indicating that labor market-tied health insurance does not appear to produce significant employment, job mobility, or productivity inefficiencies.

Figure 3



This \$469.36 average wage increase accounts for 73.4% of the \$639.79 per month employers are estimated to have been paying in health insurance costs, indicating roughly three quarters of the savings are passed on to workers. This figure is in line with, though slightly below estimates reported in much of the reduced form literature that more than 80% of the increase in cost of health insurance premiums are passed onto workers in the form of wage decreases or stagnation. This suggests that firms with substantial bargaining power may not be properly incentivized by market forces alone to pass efficiency gains onto their workers, even as they pass on the vast majority of the costs increases.

Looking at the transition from the baseline to the counterfactual equilibrium paints a slightly different picture. Figure 4 shows the average wage transition over a ten year period. Figure 4a shows how the counterfactual mean wage evolves from the baseline mean wage, and after around 72 months seems to settle slightly below the line representing the baseline

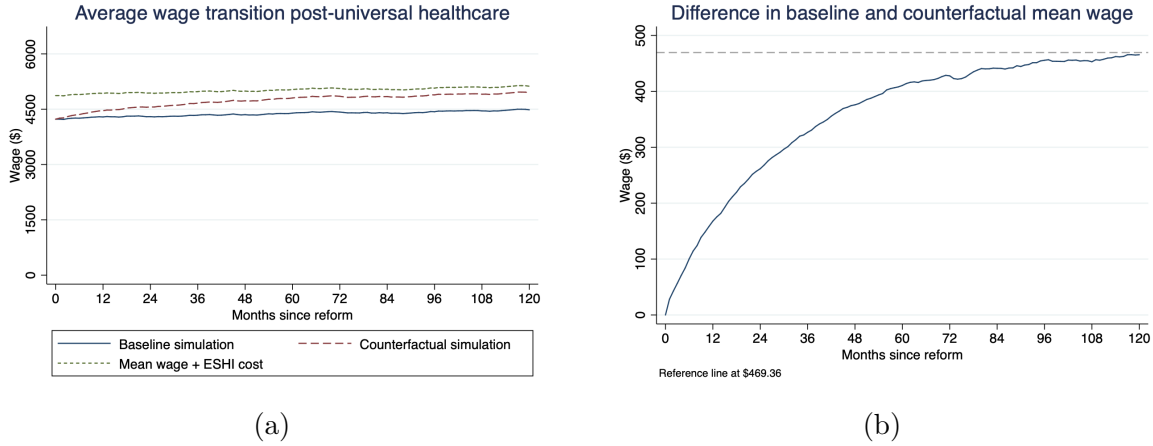
Table 4: Baseline and counterfactual wage comparison

	Simulated	Counterfactual
Mean wage	4,045.75	4,515.11
Wage inequality (Gini)	.372	.351
Unemployment rate	6.93%	6.81%
Job-to-job transition	1.42%	1.44%
Mean joint productivity	6,506.9	6,509.98

Note: Wages in pre-tax dollars per month. Joint productivity is from simulated $\epsilon \times \rho$.

wage plus the cost of ESHI. However, in figure 4b we can more clearly see the evolution of the difference in means. Here wages rise sharply initially as some workers renegotiate given the increased value of unemployment, and the rest become more likely to meet a firm capable of offering a wage that is at least above the ESHI-shaded wage from the baseline equilibrium. However, the wage growth rate quickly diminishes, and after two years the average difference is only \$261.61 per month. By the five year mark, this has increased to 87% of the total expected wage gain, at \$410.56, but the remaining wage gains take almost another five years to accrue, as the wage difference first closely approaches the full value at 114 months. As a whole, this transition simulation indicates that the pass through of cost savings to workers is fairly slow moving, as at two years only 41% of the savings have been transferred. Within five years this number is still only 64%, and the full 73% increase takes almost ten years.

Figure 4




It is also of interest to examine the distribution of these wage increases across groups. Figure 5 shows the transition of the difference in simulated and baseline wage quartiles. Look at the absolute change in figure 5a we can see that wages increase much more sharply for the highest quartile of earners, as the wage difference exceeds \$400 well within three years. The median wage appears to slightly underperform comparatively, although still exceeds

the mean wage difference in both growth rate and total increase. More significantly, we see that wage increases are substantially lower for the bottom quartile, as this difference only approaches the \$400 line only at the end of the full ten year period. This analysis suggests that the pass through of insurance savings costs benefits the highest earners the most, in terms of total wage gains. This is logical as high earners were the workers most likely to have been provided health insurance in the baseline model, and as such were most likely to have their wages reduced in return. However, from the percent change in wages presented in figure 5b we can see almost exactly the opposite result. These estimates indicate that though the highest earners see the largest wage gains, the transition to universal healthcare actually disproportionately benefits the lowest income earners. This is because increased wages at firms that would have been providing health insurance raise the wages all other firms must offer to compete, and so the wage gains are dispersed throughout the population.

Another factor we might suspect would impact the degree and rate of wage gains is the employment and health insurance status of workers at the start of the policy change. However, using these worker groupings we see virtually no difference in initial or overall wage gains. We do see a slightly quicker increase over the two to five year transition period in the unemployed and uninsured wages, but these eventually settle back towards the overall mean. These results are presented in figure 6.


Figure 5



Finally, I examine to the most basic degree the public financial implications of this policy. As the transition simulation indicates, designing a tax system that perfectly extracts the surplus generated from universal healthcare is a complex endeavor, and one I do not attempt. However, assuming such a tax were possible, I am able to obtain very crude estimates of the net increase in government spending. Making the assumption that the government would be able to provide a centralized health insurance plan at a cost that at most equals the firms' estimated insurance premium costs²² I can use the baseline employment and health insurance rates to calculate the cost of the policy. Employers that were providing health insurance receive the entirety of the ESHI cost in savings each month, and correspond to 85.5% of the

²²In reality, we might expect a nationwide health insurance system to have even lower costs due to greater risk pooling.

Figure 6



employed individuals. However, they (eventually) pass 73% of these savings onto employed individuals, who constitute 93% of the simulated labor market. This means, per simulated labor market participant, firms receive roughly \$509 in cost savings and individuals receive roughly \$437 in increased wages. This leaves roughly \$130 per labor market participant in public funding required to cover the additional cost of insuring the unemployed, and those who were previously employed but uninsured. Based on August 2021 estimates of the United States civilian labor force,²³ this amounts to an estimated \$2.1 billion per month net cost of this universal healthcare program. Of course, evaluating the true cost and benefit of universal healthcare is well beyond the scope of this paper.

8 Conclusion

I present and empirically estimate a labor search model with employer sponsored health insurance where workers of different ability levels search for jobs and are matched with firms with heterogeneous productivities. Firms hire a continuum of workers, and so must pre-commit to providing health insurance to all or none of their employees. The wage distribution and health insurance provision rates are endogenously determined in the model equilibrium, given exogenous job transition rates and health insurance premiums. Random on-the-job matching in combination with take-it-or-leave-it offers in the absence of outside options results in substantial wage dispersion that acts to somewhat temper the responsiveness of the wage distribution to changes in public health insurance policy. This model is estimated using data from the Survey of Income and Program Participation 2008 panel for worker-sided moments and the Kaiser Family Foundation Employer Health Benefits Survey from 2009-12 for firm health insurance premiums, using a simulated method of moments procedure. The estimated model performs well at replicating the wage distributions of workers with and

²³U.S. Bureau of Labor Statistics, Civilian Labor Force Level [CLF16OV], retrieved from FRED, Federal Reserve Bank of St. Louis; <https://fred.stlouisfed.org/series/CLF16OV>, September 28, 2021.

without ESHI, as well as overall employment, job-to-job transition, and health insurance provision rates observed in these datasets.

I use this estimated model to examine the effect of ESHI on wages, and the potential impact of untying health insurance from the labor market. My estimate of the utility parameter of health insurance suggests that workers value ESHI at a roughly 12 percent increase in their post tax wages. In a counterfactual experiment where workers are automatically enrolled in free universal health coverage, I find that severing the tie between healthcare and employment increases the equilibrium average pre-tax wage by about \$470 per month, equivalent to a post-tax increase of 10.5%. This average increase in monthly wage is equal to 73.3% of the estimated employer insurance premium, suggesting firms are able to keep a substantial portion of the cost savings. These findings are relatively similar to, but slightly below estimates from the reduced-form literature that upwards of 80% of historical increases in insurance costs are passed through to workers.

Counterfactual estimates also suggest that employer sponsored health insurance has little or no impact on employment or job mobility rates, or on matching efficiency in the form of average joint productivity. Overall, these findings lie somewhere between empirical literature that finds virtually no earnings or employment impacts from substantial (though incomplete) expansion of public health insurance, and structural work that suggests workers must be paid the entirety of the cost savings in order to maintain the current total level of compensation.

I also use the estimated model to simulate the transition from the baseline equilibrium to the new equilibrium without ESHI. I find that it takes almost ten years for the average wage increase to fully accrue to workers, although a little over half of the total increase occurs within 2 years as some workers are able to renegotiate to higher wages immediately following the policy change. Importantly, though the largest wage increases in levels accrue to the highest income group, the percentage increases in wage are significantly higher for the lowest earners. The result is an estimate of approximately a 2 percentage point decrease in wage inequality as measured by Gini coefficients.

The results of this paper should of course be tempered by the simplicity of the underlying model used and the many assumptions made to obtain them. This model assumes virtually complete homogeneity in both preferences for health insurance and employer costs of providing it. Furthermore the model I present is abstracted from the many underlying mechanisms of health insurance itself, from differences in quality and individual premiums, to risk pooling and the health expenditure shock distributions that are the micro foundations of its value. I also restrict my analysis to a sample of individuals who do not interact with Medicaid or Medicare, or spousal insurance, all of whose effects may certainly impact the labor market equilibrium.

Nevertheless, this paper provides an important insight into some of the mechanisms underlying the interaction between wages and health insurance. The same frictions in the labor market that produce wage dispersion even among workers with identical skills at identical jobs have the effect of introducing similar dispersion in the responsiveness of wages to healthcare policy reform. However, the substantial, though less than complete wage increase estimates presented here suggest there may be fundamentally different effects from changes to the extensive margin of health insurance provision (eliminating ESHI completely) versus the intensive margin effects presented in most previous work. Finally, the long duration of the simulated wage increases indicates that any policy proposal hoping to remain revenue

or welfare neutral must take into account not only the division of the benefits of universal healthcare but also their timing.

There are many directions for further research that arise from this paper. First, introducing additional forms of insurance provision and subsidy, or additional heterogeneity in health insurance preferences may be instrumental in fully assessing the impact of health insurance on the labor market. More directly founding health insurance benefits with a structure of healthcare costs and savings decisions is also an important avenue for the expansion of these types of search models. Finally, and perhaps most importantly, more robust and realistic policy counterfactuals and tax revenue schemes will be crucial in informing actual policy decisions as the United States seeks to expand access to healthcare for more and more of its citizens.

References

- Aizawa, N. (2019), ‘Labor market sorting and health insurance system design’, *Quantitative Economics* **10**, 1401–1451.
- Aizawa, N. and Fang, H. (2020), ‘Equilibrium labor market search and health insurance reform’, *Journal of Political Economy* **128**(11).
- Aizawa, N. and Fu, C. (2021), ‘Interaction of the labor market and the health insurance system: Employer-sponsored, individual, and public insurance’, *NBER Working Paper Series* (26713).
- Anand, P. (2017), ‘Health insurance costs and employee compensation: Evidence from the national compensation survey’, *Health Economics* **26**, 1601–1616.
- Baicker, K., Finkelstein, A., Song, J. and 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), 322–328.
- Berger, D. W., Herkenhoff, K. F. and Mongey, S. (2019), ‘Labor market power’, *NBER Working Paper Series* (25719).
- Burdett, K. and Mortensen, D. T. (1998), ‘Wage differentials, employer size, and unemployment’, *International Economic Review* **39**(2), 257–273.
- Capatina, E., Keane, M. and Maruyama, S. (2020), ‘Health shocks and the evolution of earnings over the life-cycle’, *Working Paper*.
- Chetty, R. (2006), ‘A bound on risk aversion using labor supply elasticities’, *NBER Working Paper Series* (12067).
- Currie, J. and Madrian, B. C. (1999), *Health, Health Insurance, and the Labor Market*, Vol. 3, Elsevier, chapter 50, pp. 3309–3416.

- Dey, M. S. and Flinn, C. J. (2005), ‘An equilibrium model of health insurance provision and wage determination’, *Econometrica* **73**(2), 571–627.
- Fang, H. and Shephard, A. (2019), ‘Household labor search, spousal insurance, and health care reform’, *NBER Working Paper Series* (26350).
- Flinn, C. and Heckman, J. (1982), ‘New methods for analyzing structural models of labor force dynamics’, *Journal of Econometrics* **18**, 115–168.
- Gruber, J. and Madrian, B. C. (2004), *Health Insurance, Labor Supply, and Job Mobility: A Critical Review of the Literature*, Urban Institute Press, pp. 97–178.
- Hackmann, M. B., Kolstad, J. T. and Kowalski, A. E. (2012), ‘Health reform, health insurance, and selection: Estimating selection into health insurance using massachusetts health reform’, *American Economic Review: Papers and Proceedings* **102**(3), 498–501.
- Hackmann, M. B., Kolstad, J. T. and Kowalski, A. E. (2015), ‘Adverse selection and an individual mandate: When theory meets practice’, *American Economic Review* **105**(3), 1030–1066.
- Heathcote, J., Storesletten, K. and Violante, G. (2017), ‘Optimal tax progressivity: An analytical framework’, *Quarterly Journal of Economics* **132**, 1693–1754.
- Hegenauer, C. L. (2016), ‘Are we covered? health insurance disparities in the affordable care act era’, *Michigan Sociological Review* **30**, 93–108.
- Hwang, H.-S., Mortensen, D. and Reed, R. (1998), ‘Hedonic wages and labor market search’, *Journal of Labor Economics* **16**(4), 815–847.
- Kaplan, G. (2012), ‘Inequality and the life cycle’, *Quantitative Economics* **3**(3), 471–525.
- Kolstad, J. and Kowalski, A. (2016), ‘Mandate-based health reform and the labor market: Evidence from the massachusetts reform’, *Journal of Health Economics* **47**, 81–106.
- Kolstad, J. T. and Kowalski, A. E. (2012), ‘The impact of health care reform on hospital and preventive care: Evidence from massachusetts’, *Journal of Public Economics* **96**, 909–929.
- Monheit, A., Hagan, M., Berk, M. and Farley, P. (1985), ‘The employed uninsured and the role of public policy’, *Inquiry* **22**(4), 348–364.
- Petrin, A. (2002), ‘Quantifying the benefits of new products: The case of the minivan’, *Journal of Political Economy* **110**(4), 705–729.
- Postel-Vinay, F. and Robin, J.-M. (2002), ‘Equilibrium wage dispersion with worker and employer heterogeneity’, *Econometrica* **70**(6), 2295–2350.
- See, K. (2019), ‘Medicare for all: The macroeconomic and welfare consequences of expanding public health insurance’, *Working Paper*.
- Simon, K. I. and Kaestner, R. (2004), ‘Do minimum wages affect non-wage job attributes? evidence on fringe benefits’, *ILR Review* **58**(1), 52–70.

Appendix A Tax function

I estimate the tax function following the methods outlined in Heathcote et al. (2017) and using NBER's TAXSIM program to generate post-tax wages for my sample, based on income, marital status, and dependents status. Specifically, this tax function takes the form:

$$T(w) = \tau_0 w^{1-\tau_1}$$

where $T(w)$ is after tax wages as a function of the wage, τ_1 determines tax progressivity and τ_0 determines the level of the tax. This function is differentiable and generally fits well with observed pre- and post-tax wages. It is also easy to estimate, as taking logs I obtain

$$\log T(w) = \log \tau_0 + (1 - \tau_1) \log w$$

A simple linear regression of log after-tax (monthly) wages on log post-tax (monthly) wages returns parameter estimates of $\tau_0 = 1.693$ and $\tau_1 = 0.0739$. From figure 7 shown below, this function fits the TAXSIM generated data quite closely for low income individuals and the bulk of individuals without second earners. The roughly parallel lower band of post-tax wages are from the minority of individuals who's tax burden is significantly increased by higher familial earnings. The fitted tax function is much more heavily weighted toward the former group as they vastly outnumber the latter, particularly at higher income levels. Regardless, counterfactual experiments using a model estimated with or without taxes did not meaningfully impact the results.

Figure 7

