Interaction of the Labor Market and the Health Insurance System: Employer-Sponsored, Individual, and Public Insurance

Naoki Aizawa, *University of Wisconsin–Madison and National Bureau of Economic Research (NBER)*

Chao Fu, *University of Wisconsin–Madison and NBER*

We develop and estimate an equilibrium model with heterogeneous local markets, households, and firms, highlighting the interrelationship between various components of the health insurance system—employer-sponsored health insurance (ESHI), individual health insurance (HIX), and Medicaid—and their relationship with the labor market. We estimate the model exploiting variation across states and before and after the Affordable Care Act. We consider counterfactual policies that cross subsidize between ESHI and HIX, including pure ESHI-HIX risk pooling as a special case. We find that such policies would increase household welfare and output and decrease government expenditure and would be more effective with Medicaid expansion.

I. Introduction

For most working-age households in the United States, health insurance is attainable via three channels: employer-sponsored health insurance

We thank Samuel Bailey and Shoya Ishimaru for excellent research assistance. We thank Mariacristina De Nardi, Eric French, Martin Hackmann, Tom Holmes,

Submitted June 17, 2022; Accepted February 21, 2024; Electronically published August 18, 2025.

Journal of Labor Economics, volume 43, number 4, October 2025.
 https://doi.org/10.1086/730164

(ESHI), the individual insurance market (currently known as the health insurance exchange [HIX]), and public insurance, mainly via Medicaid. Two features of this health insurance system stand out. The first feature is its tie to the labor market. Directly, ESHI covers more than two-thirds of the working-age population.[1] Indirectly, income largely determines a household's eligibility for Medicaid and the amount of HIX subsidies or credits it may receive.

The second standout feature is the segregation of risk pools across insurance channels. Current regulations (e.g., medical loss ratio regulation) require that the premium on an ESHI or HIX market should closely reflect the risk among those insured on that market. As a result, on choosing a job with ESHI or earning an income low enough to qualify for Medicaid, a household is largely segregated from the risk pool on HIX. In the current equilibrium, ESHI-covered households are healthier, or less adversely selected, than HIX enrollees. As Akerlof (1970) suggests, adverse selection may limit the affordability and even availability of HIX insurance. This leads to concerns not only about the lack of health insurance for many households but also about the potential distortion of worker-firm sorting due to ESHI jobs providing access to affordable health insurance.

Ami Ko, Jeremy Lise, Lee Lockwood, Matthew Notowidigdo, Amil Petrin, Chris Taber, and seminar and conference participants at the Allied Social Science Associations meeting, Barcelona Graduate School of Economics Summer Forum, Conference on Econometrics and Mathematical Economics (CEME) at Chicago, Chinese University of Hong Kong, Econometric Society Meetings, Federal Reserve Bank (FRB) of New York, FRB Saint Louis, Georgetown Center for Economic Research/Cemmap Conference, Hong Kong University of Science and Technology, Institute for Fiscal Studies, Purdue, Society for Economic Dynamics, Stanford Institute of Theoretical Economics, University of British Columbia, University of Minnesota, University of North Carolina, University of Southern California, University of Wisconsin–Madison, University of Wisconsin–Milwaukee, University of Western Ontario, and Yale for helpful comments. We thank Rachelle Hill at the Minnesota Census Research Data Center (RDC), Robert Thomas at the Wisconsin Census RDC, and Ray Kuntz at the Agency for Healthcare Research and Quality (AHRQ) for help with the restricted-access Medical Expenditure Panel Survey data. A part of the analysis in this paper was conducted at the Minnesota Census RDC and Wisconsin Census RDC, and the support of AHRQ is acknowledged. All results have been reviewed to ensure that no confidential information is disclosed. Any opinions and conclusions expressed herein are those of the author(s) and do not necessarily represent the views of the US Census Bureau, AHRQ, or the Department of Health and Human Services. Contact the corresponding author, Naoki Aizawa, at naizawa@wisc.edu. Information concerning access to the data used in this paper is available as supplemental material online.

[1] For example, in 2015, of the US population aged 22–64, 68% were insured via own or spousal ESHI, 10% were covered by Medicaid, and 7% were covered by HIX; the remaining 15% were uninsured. Statistics are calculated from the American Community Survey (ACS).

One natural question is whether the current segregated risk pool structure has caused nontrivial welfare losses. If the answer is yes, then redesigning the risk pool structure becomes a central issue. However, relative to the extensive discussion about policy intervention in each insurance channel (e.g., ESHI tax deductibility, HIX subsidization, and Medicaid eligibility), the issue about a proper risk pool structure has received much less attention, and our knowledge about the implication of counterfactual risk pool structures is quite limited. This paper aims to take a step forward in understanding this issue.

To achieve this goal, it is crucial to recognize that risk pool structures affect welfare mainly through affecting equilibrium sorting: the sorting of households into different insurance statuses and the sorting on the labor market. For example, we observe that highly educated workers (who tend to be healthy) are more likely to work on jobs with ESHI and that larger firms and firms with a higher fraction of skilled employees tend to offer ESHI. The extent to which such equilibrium sorting and welfare varies with risk pool structures depends fundamentally on the distribution of household heterogeneity and the distribution of firms' production technologies. Moreover, equilibrium sorting is also affected by other existing regulations. Because policy environment varies substantially across US states (e.g., Medicaid eligibility rules are state specific), the impact of adjusting the risk pool structure is likely to be different across states.

To incorporate these considerations in a coherent framework, we develop an equilibrium model where each state is a market consisting of a labor market and two insurance markets (HIX and ESHI). Markets are subject to various regulations that may vary across states and policy eras. Each state consists of a distribution of households and firms. Households differ in their demographics, health, skill levels, and tastes. Skills and tastes are unobservable to the researcher and may be distributed differently across states. A household chooses, for each adult member, between full-time jobs with and without ESHI, part-time jobs with and without ESHI, and nonemployment. It also makes decisions about Medicaid enrollment (if eligible) and individual health insurance purchases. Firms differ in their overall productivity (total factor productivity [TFP]) and the degree to which their technologies are skill biased. Each firm chooses whether to offer ESHI and the number of workers in each category (skill, full/part time), which are imperfect substitutes for one another. Equilibrium wages and premiums for HIX and ESHI clear the corresponding markets.

To estimate this model, which allows for unobserved heterogeneity across households, firms, and states, we exploit rich policy variation from the Affordable Care Act (ACA), including (1) policy variation before and after the ACA came into effect in 2014, (2) the targeted nature of certain components of the ACA that created variation in policy doses received by different firms and/or households in the same market, and (3) the differential implementation of the ACA Medicaid expansion and hence substantial differences in households' choice sets across states. We estimate the model via indirect inference, exploiting the aforementioned variation under the assumption that the state-specific distribution of unobservables is invariant to the ACA (conditional on observables).

We use both pre- and post-ACA data from the ACS, the Current Population Survey (CPS), the Medical Expenditure Panel Survey (MEPS), and the Kaiser Family Employer Health Insurance Benefit Survey. The first three datasets provide information on household characteristics, labor supply and health insurance choices, earnings, and medical expenditure; the fourth provides information on firm size, ESHI provision, and employee composition. For model validation, we leave the post-ACA data for a nonrandom sample of states out of the estimation. The estimated model matches patterns in both the estimation sample and the holdout sample.

Our estimates suggest a positive correlation between worker skill and preferences for health insurance. In the equilibrium, high-skill workers are more likely to sort into firms offering ESHI, which tend to have higher TFP and more skill-biased technologies. Households that choose to be nonemployed or earn wages low enough to be eligible for Medicaid are more likely to have low skills. The risk pool on HIX consists largely of the remaining households in the skill distribution. That is, households are ex post segregated into different risk pools largely by their skill levels. Unlike ESHI, HIX insurance is not bundled with one's job and hence is more susceptible to adverse selection, resulting in HIX premiums being higher than ESHI premiums (conditional on quality). This premium differential is regressive in nature (the less skilled face higher premiums) and may distort labor force allocation.

To explore potential improvements over the status quo, we design new schemes to break the risk pool segregation between ESHI and HIX. At face value, the phrase "risk pooling" may suggest that it requires merging ESHI and HIX markets, which would involve drastic changes to the current health insurance system. However, this is not necessary: our schemes—cross subsidization between ESHI and HIX—can implement risk pool desegregation with little change to the health insurance system. ESHI-HIX premium differentials are regulated via taxing ESHI insurers and using the tax revenue to subsidize HIX insurers. The nature of the new equilibrium is governed by this tax rate (the degree of cross subsidization), of which pure risk pooling between ESHI and HIX is a special case.

Our policy simulations suggest that the current risk pool segregation has led to losses in total output and household welfare as well as wasteful government spending. Relative to the baseline (the ACA environment in 2015), cross-subsidization policies would increase the total output by 2% to 3%, improve average household welfare by $75 to $159 (in annual consumption equivalent variation), and decrease government per-household expenditure by $14 to $49, depending on the degree of cross subsidization. Moreover, for both household welfare and government budget, ESHI-HIX cross subsidization outperforms a policy that pools risks across states for HIX. Finally, holding all initial conditions fixed, we find that ESHI-HIX cross subsidization would lead to higher welfare gains if it were combined with Medicaid expansion.

Our findings suggest that properly designing the risk pool structure across health insurance channels is a cost-effective and easily implementable way to improve welfare and labor market efficiency. Moreover, it is relevant to the current policy debate. For example, one can consider our cross-subsidization policy as a substantial expansion of the Health Reimbursement Arrangement introduced in 2020, which incentivizes small firms to insure their employees via HIX.

Our paper contributes to the literature on the link between the health insurance system and the labor market, especially studies aimed at evaluating counterfactual policies. Of these studies, one subset focuses on individual decisions (e.g., Rust and Phelan 1997; French and Jones 2011; Pohl 2018). In particular, French, Jones, and Gaudecker (2018) study the ACA's impact on retirement, savings, and welfare. Another subset uses an equilibrium approach. Dey and Flinn (2005) estimate a search and bargaining model with endogenous ESHI. Aizawa and Fang (2020), Aizawa (2019), and Fang and Shephard (2019) estimate their models using pre-ACA data and simulate the impact of various components of the ACA.

Our paper complements these studies well. First, instead of evaluating ACA policies or their variants, we examine the implication of risk pool segregation—a key feature of the US health insurance system that has received little attention. Second, departing from search models with linear production technologies, labor market efficiency in our model relies on proper matches between firms and "bundles" of labor inputs—a key determinant of aggregate productivity (Eeckhout and Kircher 2018). While previous studies find negligible effects of health insurance reforms on aggregate productivity (e.g., Dey and Flinn 2005; Aizawa 2019), we find significant effects. Third, our framework incorporates richer heterogeneity to better assess distributional effects of counterfactual policies. Finally, we exploit ACA-induced policy variation to estimate the model and rely relatively less on the model structure to examine counterfactual policies. In this regard, our paper draws insights from the design-based literature on the ACA (e.g., Kowalski 2014; Frean, Gruber, and Sommers 2017; Kaestner et al. 2017; Leung and Mas 2018).

Given that employers provide many essential benefits, large-scale reforms should account for equilibrium labor market responses (Summers 1989). Therefore, we incorporate labor market equilibrium in our model. As a cost, we model individual health insurance markets in a relatively simple way compared with studies focusing on these markets (for the survey, see Einav, Finkelstein, and Mahoney 2021).

II. Background Information

Unlike many other developed countries, the US health insurance system for the working-age population is largely employer based, with ESHI covering more than two-thirds of this population. This is partly sustained by ESHI tax deductibility: although wages are subject to payroll and income taxes, ESHI benefits are not taxed. ESHI premiums are paid (mostly) by employers as a fringe benefit for their workers. For a single-coverage (family) plan, the ESHI premium paid by firms is about 9% (22%) of the average wage among ESHI-covered workers, based on data from the ACS and Kaiser. The second insurance channel, Medicaid, is administered by each state government as an in-kind transfer for the poor, the cost of which is shared between federal and state governments. The third channel is individual health insurance markets, which suffer from significant adverse selection (e.g., Gruber 2008). The uninsured are partially insured by implicit public insurance through uncompensated care, the cost of which is born mostly by the government or charities.

The ACA was signed into law in 2010, whose major provisions came into effect in 2014. The uninsured rate among working-age Americans declined from 22% in 2012 to 15% in 2015. We use both pre- and post-ACA (2012 and 2015) data to exploit the following ACA-induced policy variation.

Individual mandate.—Starting from 2014, individuals shall have a health insurance plan meeting minimum standards or pay a tax penalty that varied with household income and size. In 2015, the penalty was the maximum of (*a*) 2% of household income in excess of the 2015 income tax filing thresholds and (*b*) $325 per adult plus $162.5 per child, up to $975 per household. This mandate was abolished in 2019.

Employer mandate.—Starting in 2015, every employer with more than N ($N = 100$ in 2015, $N = 50$ starting from 2016) full-time-equivalent employees shall provide a health insurance plan meeting minimum standards to full-time employees (average weekly hours ≥ 30) and their dependent children or pay a tax penalty. The penalty is $2,000 (indexed for future years) for each full-time employee, excluding the first 30 employees.

HIX.—Before 2014, insurers on private individual health insurance markets were allowed to price discriminate by consumers' demographics and health status. In 2014, HIX were established with significant restrictions on price discrimination: insurance premiums are subject to modified community rating and are based only on age and smoking status as specified by the government. Federal government sets the regulation, based on which state governments can set further restrictions. A household can participate in HIX only in its state. HIX insurers have to offer the same plans to every consumer; the design of HIX plans is government regulated and categorized into four plans with different levels of generosity: bronze, silver, gold, and platinum. The most popular choice is the silver plan, which is more generous than typical health insurance plans purchased in pre-ACA individual markets (app. F; apps. B–I are available online). HIX markets are also subject to the tight medical loss ratio regulation: insurers are required to spend at least 80% of their premium revenues on reimbursements and quality improvements (i.e., the premium shall closely reflect the cost of insuring the enrollees).

Income-based subsides for plans from HIX.—HIX customers may obtain both premium and coinsurance subsidies. Individuals are eligible for subsidies if (1) they are unable to get affordable coverage through an eligible employer plan meeting the minimum generosity, (2) they are ineligible for any other government health insurance program (e.g., Medicaid), and (3) their household income is between 100% and 400% of the federal poverty line (FPL). The subsidy amount varies by HIX premiums, income, family size, and states: the maximum premium contribution by the household is 2% (9%) of its income if its household income is around 100% (400%) of the FPL.[2] In addition, individuals purchasing the silver plan can obtain an income-based tax credit.

Medicaid.—Before the ACA, each state government chose Medicaid eligibility rules. Very few adults were eligible for Medicaid other than low-income parents with young children. The ACA specified (not mandated) that Medicaid expand to cover the uninsured with household income below 133% of the FPL. By 2015, 32 states (including Washington, DC) had complied.

III. A Simple Model: The Intuition

Before introducing the full model, we use a simple model to illustrate why risk pool segregation may lead to welfare loss. Consider an economy with a competitive labor market and two competitive insurance markets—ESHI and HIX—offering an identical insurance product. Workers are equally productive but differ in their health status and disutility of work. Workers are risk averse and subject to medical expenditure shocks; unhealthier workers are more likely to incur higher medical expenditure. A worker chooses between an ESHI job at wage w_1, a non-ESHI job at wage w_0, and nonemployment; if choosing either of the latter two, one also chooses whether to enroll in HIX (there is no public insurance). Firms are homogeneous, each choosing whether to offer ESHI and how many workers to hire. Neither insurance market can price discriminate. In the equilibrium, the premium on ESHI (q) and that on HIX (r) are equal to the average

[2] If states offer Medicaid to individuals whose income is below 100%, then they are not eligible for these subsidies. For additional details, see https://www.irs.gov/Affordable-Care-Act/Individuals-and-Families/Questions-and-Answers-on-the-Premium-Tax-Credit.

medical cost among enrollees on ESHI and HIX, respectively (risk pool segregation).

We show formally in appendix B the following equilibrium implications.

1. Adverse selection arises on both insurance markets because, all else being equal, the unhealthier worker benefits more from being insured.
2. For both ESHI jobs and non-ESHI jobs to exist, the wage differential ($w_0 - w_1$) cannot be larger than the cost purchasing HIX (r); otherwise, ESHI jobs would be inferior to non-ESHI jobs for all workers. Meanwhile, for firms to be indifferent between two types of jobs, it must be that $w_0 - w_1 = q$. These imply that $r \geq q$, that is, the HIX premium is (weakly) higher than the ESHI premium. If workers with poorer health incur higher disutility of work, then it is more likely that $r > q$.
3. In equilibriums with $r > q$, workers' labor supply decisions are distorted because they will face a higher insurance premium if they do not choose ESHI jobs. Moreover, all employed workers who are insured are enrolled in ESHI and all HIX enrollees are nonemployed; this creates a regressive welfare effect in that poorer households (the nonemployed in this example) face higher premiums than richer households.

Given potential welfare losses arising from risk pool segregation, it may be worth considering policies to reduce the premium differential (e.g., via risk pooling between ESHI and HIX). However, even in this simple model, it is not clear ex ante that risk pooling will improve welfare: it depends on how wages and insurance premiums adjust and how different workers change their insurance and labor supply decisions. For example, under risk pooling, previously ESHI-insured workers will face a premium higher than the original equilibrium q; some of them may choose to be uninsured. Ambiguity grows with several important real-life complicating factors. We describe these factors in the next section and account for them in our full model.

IV. Model

A. Overview of the Model

Risk pool structures may affect welfare because they affect the sorting of households across insurance statuses and the sorting on the labor market. These impacts depend fundamentally on household-side and firm-side heterogeneities. On the labor market, the room for improvement depends on the severity of existing mismatches between workers and firms; in the health insurance system, it depends on how households differ across the currently segregated risk pools. In addition, impacts of risk pool structures can also depend on specific regulations within each health insurance channel. Our full model accounts for fundamental and institutional factors that are essential for understanding sorting and for evaluating different risk pool structures.

In our model, households differ in health status, disutility of work, skills, demographics, and risk aversion. For each of its adult members, a household chooses between nonemployment and different types of jobs (part/full time, with/without ESHI). Their choices of health insurance are subject to the fact that ESHI is bundled with ESHI-providing jobs and that the eligibility for Medicaid and the net price of private insurance both depend on household income and demographics. Firms have heterogeneous technologies; they assemble workers in different categories (skill, part/full time) to produce a homogeneous good. Each firm chooses the number of workers to hire in each category and whether to offer ESHI. An equilibrium is characterized by market-clearing wages for each category (skill, part/full time) of labor for jobs with and without ESHI and by premiums that satisfy the break-even condition separately for the ESHI market and the individual health insurance market (risk pool segregation).

Risk pool structures have complex equity-efficiency implications in this model. As in Eeckhout and Kircher (2018), labor market efficiency requires proper matches between firms and bundles of labor inputs. By law, although firms' ESHI provision can differ between full- and part-time employees, it must be nondiscriminatory with respect to worker skills. As such, any policy affecting the cost of ESHI can lead firms to choose suboptimal bundles of labor inputs. These within-firm suboptimal choices can in turn distort the sorting of workers across firms. In addition, workers sort on the basis of skills and other factors. If the high skilled are healthier on average and if they tend to sort into ESHI-providing jobs in the equilibrium, then risk pool segregation would generate a regressive welfare effect. This effect could be alleviated or outweighed by HIX subsidies while exacerbated by ESHI tax exemption policies.

B. Environment

There are M isolated markets defined by state $\times$ policy era, each consisting of a labor market, an individual health insurance market, and an ESHI market. In each market m, there is a market-specific distribution of households and a distribution of firms. Households differ in $\mathbf{X}$, including each adult member's skill level s, household demographics, health status, risk preference, and work disutility. Firms differ in technology Υ, including overall productivity and the relative productivity of high- versus low-skilled labor inputs.

Households choose labor supply and health insurance status; firms choose labor inputs and ESHI provision and use labor inputs to produce a homogeneous good. Each labor market m is competitive with a vector of wages $\{w^m_{shz}\}_{(s,h,z)}$, where s is a worker's skill level, $h \in \{P, F\}$ denotes part/full time, and z is an indicator of ESHI provision. Exchanges on the labor market are blind to factors other than (s, h, z). For ease of exposition, let $w^m_{shz} = 0$ if $h = 0$ (nonemployed).

1. Insurance Status and Out-of-Pocket Health Expense

An individual's health insurance status is defined as $\text{INS} \in \{0, 1, 2, 3, 4\}$ (no insurance, ESHI, spousal ESHI, Medicaid, individual insurance); options are assumed to be mutually exclusive. We denote $\mathbf{INS} \in \{0, 1, 2, 3, 4\}^2$ as the insurance status of a couple.

A household's out-of-pocket health expense (OOP) varies with its demographics and health status (a part of $\mathbf{X}$), $\mathbf{INS}$, and the market (m) it belongs to. In addition, a household is subject to medical expenditure shocks realized after its decisions, leading to a distribution of OOP that a household may face, given by

$$\text{OOP} \sim F_{\text{OOP}}(\mathbf{X}, \mathbf{INS}, m). \tag{1}$$

Being static, our model treats health shocks as expenditure shocks. A major role of insurance is to reduce the dispersion of OOP.

2. Household's Preference and Budget

A household's utility depends on its consumption C, work status, and insurance status, specified as

$$u(C, \mathbf{h}, \mathbf{INS}; \mathbf{X}) = \frac{(C/size_{\mathbf{X}})^{1-\gamma_{\mathbf{X}}}}{1 - \gamma_{\mathbf{X}}} - D(\mathbf{h}, \mathbf{X}) + \varpi_{\mathbf{INS}}, \tag{2}$$

where $size_{\mathbf{X}}$ is an adult-equivalence factor and $\gamma_{\mathbf{X}}$ is a risk-aversion parameter, both varying by $\mathbf{X}$. The vector $\mathbf{h} = [h, h'] \in \{0, P, F\}^2$ is the labor supply status of the couple and $D(\mathbf{h}, \mathbf{X})$ is the disutility from work. Parameters $\{\varpi_{\mathbf{INS}}\}$ capture nonpecuniary values of various insurance choices (e.g., application costs and inertia).

A household uses the couple's earnings $w^m_{shz} + w^m_{s'h'z'}$ plus a net government transfer $\tilde{\text{tr}}$ to fund its consumption C and medical OOP. To capture the uncompensated care for the uninsured, we introduce a consumption floor $\underline{c}$, such that

$$C = \max\{w^m_{shz} + w^m_{s'h'z'} + \tilde{\text{tr}} - \text{OOP}, \underline{c}\},$$
$$\text{where } \tilde{\text{tr}} = -T^{\text{tax}}(w^m_{shz} + w^m_{s'h'z'}, \mathbf{X}, m) + T^{\text{ins}}(w^m_{shz} + w^m_{s'h'z'}, \mathbf{INS}, \mathbf{X}, m). \tag{3}$$

The net transfer $\tilde{\text{tr}}$ consists of federal and state income taxes net of welfare transfers ($T^{\text{tax}}(\cdot)$) and health insurance-related transfers including HIX insurance premium subsidies and penalties on the uninsured ($T^{\text{ins}}(\cdot)$). Taxes and transfers affect household sorting because they affect the relative benefits/costs of a household's options, and differently so across households.

3. *Production Function and Costs*

The output at a firm with technology Υ is given by

$$Q_{\Upsilon} = F(\mathbf{n}; \Upsilon),$$

where $\mathbf{n} = \{n_{sh}\}_{sh}$, and n_{sh} is the number of employees with skill level s and work status $h \in \{P, F\}$.

For a firm on market m that chooses labor inputs $\mathbf{n}$ and ESHI offering $\mathbf{z} = \{z_{sh}\}_{sh}$, its total cost consists of wages, payroll taxes, and the cost of ESHI, given by

$$\mathrm{TC}(\mathbf{n}, \mathbf{z}, m).$$

As detailed in appendix A, $\mathrm{TC}(\cdot)$ has four important features. First, payroll and ESHI costs both depend on employee composition $\{n_{sh}\}_{sh}$. Workers with different demands for health insurance (e.g., whether one needs coverage for themselves, their spouse, and/or child) sort into different jobs. Therefore, an ESHI-providing firm faces different total premiums for workers across (s, h). Second, wages are subject to payroll taxes, but ESHI premiums are tax deductible, providing an incentive to offer ESHI. Third, we allow for a fixed cost of ESHI provision; if it is estimated to be greater than zero, there will be a cost advantage of providing ESHI for larger (more productive) firms. Finally, because equilibrium wages and premiums are market specific, firms' total costs vary by m.

C. Household's Problem

A household chooses its labor supply and insurance status to maximize its expected utility, which can be viewed as a two-step decision problem. First, the household chooses job status $(\mathbf{h}, \mathbf{z}) \in \{(0,0), \{P, F\} \times \{0,1\}\}^2$, where each spouse can be nonemployed and hence without ESHI $(0,0)$ or working in one job category $\{P, F\} \times \{0,1\}$ (part/full time, with/without ESHI). Second, it chooses its health insurance status $\mathbf{INS}$ given $(\mathbf{h}, \mathbf{z})$. A single household's problem is similar but simpler.

The optimal choice of a household $(\mathbf{h}^*, \mathbf{z}^*, \mathbf{INS}^*)_{\mathbf{X}}$ solves the following problem:

$$\max_{(\mathbf{h},\mathbf{z}) \in \{(0,0),\{P,F\}\times\{0,1\}\}^2} \{V(\mathbf{h}, \mathbf{z}; \mathbf{X}, m) + \epsilon_{\mathbf{h},\mathbf{z}}\}, \tag{4}$$

where

$$V(\mathbf{h}, \mathbf{z}; \mathbf{X}, m) = \max_{\mathbf{INS} \in \Omega(\mathbf{X},m,\mathbf{h},\mathbf{z})} \left\{ \int u(C, \mathbf{h}, \mathbf{INS}; \mathbf{X})\, dF_{\mathrm{OOP}}(\mathbf{X}, \mathbf{INS}, m) \right\} \tag{5}$$

$$\text{s. t. the budget constraint (3).}$$

For a given job status $(\mathbf{h}, \mathbf{z})$, a household chooses its health insurance status optimally, yielding a conditionally optimal value $V(\mathbf{h}, \mathbf{z}; \mathbf{X}, m)$ defined in equation (5); the $(\mathbf{h}, \mathbf{z})$ with the highest conditional value is a household's optimal job status. In expression (4), $\epsilon_{\mathbf{h},\mathbf{z}}$ is a taste shock associated with choice $(\mathbf{h}, \mathbf{z})$, drawn from an independent and identically distributed (i.i.d.) type I extreme value distribution with a scale parameter σ_ϵ.

A household's choice set of health insurance status $\Omega(\mathbf{X}, m, \mathbf{h}, \mathbf{z})$ depends on its labor supply $(\mathbf{h}, \mathbf{z})$, reflecting the intrinsic connection between the health insurance system and the labor market (app. A, sec. A2). ESHI is directly governed by the choice of $\mathbf{z}$; if neither spouse has ESHI, the household may be eligible for Medicaid, which (partly) depends on its earnings and hence $\mathbf{h}$; if neither ESHI nor Medicaid is available, the household can choose to purchase individual insurance or stay uninsured.

D. Firm's Problem

A firm's optimal solution $(\mathbf{z}_f^*, \mathbf{n}_f^*)_\Upsilon$ (subscript f refers to "firm") solves the following problem (see app. C for the case with ESHI mandates):

$$\pi_\Upsilon^* = \max_{\mathbf{z}\in Z,\mathbf{n}\geq 0}\{F(\mathbf{n}; \Upsilon) - \mathrm{TC}(\mathbf{n}, \mathbf{z}, m) + \eta_\mathbf{z}\}, \tag{6}$$

where $\eta_\mathbf{z}$ is an i.i.d. type I extreme-value distributed shock (with a scale parameter σ_η) for choosing $\mathbf{z}$. A firm's ESHI choice is within the space Z, such that (1) the ESHI offering must be uniform conditional on work status h (required by law) and (2) ESHI is offered to part-time workers only if it is also offered to full-time workers (consistent with the data). As such, Z contains three options: offering ESHI to all employees, offering ESHI only to all full-time employees, and not offering ESHI.

E. Insurance Premiums

We assume a single product on HIX as in Hackmann, Kolstad, and Kowalski (2015) and a single product on ESHI. We endogenize equilibrium insurance premiums such that on both ESHI and HIX, insurers are break-even: on market m and $k \in \{\mathrm{ESHI}, \mathrm{HIX}\}$, the total premium is equal to the total reimbursement multiplied by a market-specific loading factor l_k^m.[3]

HIX premiums are set according to a standard age-rating curve and are otherwise nondiscriminatory. Letting r_b^m be the base premium on market m and $\Gamma(\cdot)$ be the exogenous age-rating curve, the premium faced for someone with $\mathbf{X}$ is given by

[3] Insurers may have certain market power in reality; however, our break-even assumption for HIX and ESHI is not far-fetched because premiums are regulated to closely reflect the average medical expenditure on each market. The pre-ACA individual health insurance (IHI) premium structure was much more complex. We use the pre-ACA data only for model estimation and can take observed IHI premiums as given.

$$r^m(\mathbf{X}) = \Gamma(r_b^m, age). \tag{7}$$

The base premium r_b^m adjusts to satisfy the break-even condition on HIX market m (as in Handel, Hendel, and Whinston 2015). We model the total premium paid by an ESHI-providing firm on market m as the total health insurance demands among its covered employees times a base premium q^m; q^m adjusts to satisfy the break-even condition on ESHI market m.

F. Equilibrium

DEFINITION 1. An equilibrium on market m is a tuple $\{(\mathbf{h}^*, \mathbf{z}^*, \mathbf{INS}^*)_{\mathrm{X}}, (\mathbf{z}_f^*, \mathbf{n}_f^*)_{\Upsilon}, \{w_{shz}^m\}_{shz}, r_b^m, q^m\}$ such that

1. Given $\{w_{shz}^m\}_{shz}$ and $r^m(x)$, $(\mathbf{h}^*, \mathbf{z}^*, \mathbf{INS}^*)_{\mathrm{X}}$ solves the household problem for each $\mathbf{X}$.
2. Given $\{w_{shz}^m\}_{shz}$ and q^m, $(\mathbf{z}_f^*, \mathbf{n}_f^*)_{\Upsilon}$ solves the firm problem for each Υ.
3. Equilibrium consistency:
   a. wages $\{w_{shz}^m\}_{shz}$ equate the aggregate labor demand and supply for each (s, h, z);
   b. r_b^m and $r^m(x)$ implied by equation (7) satisfy the break-even condition on HIX; q^m satisfies the break-even condition on ESHI.

The basic feature of our labor market equilibrium is similar to that in Eeckhout and Kircher (2018), who study a competitive labor market where workers differ in skills and firms differ in TFP. They analytically establish the existence of a static equilibrium and characterize it. We incorporate richer heterogeneity and model the labor market in connection to health insurance markets, the latter of which are subject to adverse selection and risk pool segregation. Adverse selection makes the supply curve (i.e., the average cost curve) on the insurance market downward sloping. The fact that both the demand curve and the supply curve are downward sloping complicates the equilibrium analysis in our setting. In addition, risk pool segregation can lead to labor market inefficiency, a source of inefficiency absent in Eeckhout and Kircher (2018). As a cost of modeling these complications, we rely on numerical methods to solve our model; we always find a unique equilibrium in our analyses.

Discussion.—Several aspects of the model deserve further discussion. First, we study the static equilibrium on a frictionless labor market. Previous studies incorporated search friction to study how health insurance availability may distort worker mobility, where firms/jobs are ranked vertically and efficiency requires that workers move toward high-productivity jobs; the distortive effect was found to be negligible (e.g., Dey and Flinn 2005; Aizawa 2019). We abstract from job mobility and focus on static worker-firm sorting, where firms' technologies cannot be ranked vertically and efficiency relies on proper matches between firms and bundles of labor inputs. In our model, labor misallocation arises from adverse selection and distorting taxes/transfers on health insurance markets.

Second, we explicitly model how households' decisions vary with their age, health, and skill, but in a cross section instead of a life cycle framework. Because we do not model health insurance's potential positive (dynamic) health effect, we may understate the value of expanding health insurance coverage, and policy implications from our model should be interpreted as short-run effects. To capture the value of health insurance beyond its effect on households' OOP, we allow for direct insurance status–specific utility terms. In a static setting, this is a reduced-form way to capture nonfinancial net benefits of various health insurance options, for example, insurance's health effect (valued by households) and other traits (e.g., application costs, stigma against Medicaid).

Third, although we model rich state-level heterogeneity, including state-specific policies and distributions of households' (observable and unobservable) characteristics, we abstract from households' migration decisions, which is a limitation. Previous studies suggest that geographical mobility responses to health policies are likely to be negligible (e.g., Goodman 2017); we do not expect our counterfactual policies to induce significant cross-state migration responses.

G. Empirical Specifications

Now we describe empirical specifications for our model; further details are in appendix A.

1. Households' Characteristics

Household characteristics are given by $\mathbf{X} = (x, \mathbf{s}, \boldsymbol{\chi})$, where x consists of marital status, the presence of young children, each spouse's health status, gender, age, and education; $\mathbf{s} \in \{1, \dots, 5\}^2$ consists of each spouse's skill level; and $\boldsymbol{\chi} \in \{1, 2\}^2$ consists of each spouse's preference type. The vector x enters the OOP distribution (1), household utility (2), and budget (3) via government transfers. Skill level s affects one's potential earnings. Preference type χ affects work disutility and risk aversion ($\gamma_{\mathbf{X}}$ in eq. [2] depend on $\mathbf{X}$ only via χ). Firms cannot discriminate workers by any component of $\mathbf{X}$ except s, HIX companies cannot discriminate customers by any component of $\mathbf{X}$ except age following equation (7), and all firms on market m face the same ESHI base premium q^m.

The researcher observes x but neither $\mathbf{s}$ nor $\boldsymbol{\chi}$. Within a household, a couple's unobservables are correlated; across households, the distribution $\Pr((\mathbf{s}, \boldsymbol{\chi})|x, \textit{state})$ varies with x and state. First, notice that we allow for a correlation between $\mathbf{s}$ and $\boldsymbol{\chi}$ to entertain the possibility that workers of different skill levels may demand health insurance differently beyond income-related reasons. Second, conditional on x, household outcomes differ across states;

this may arise partly from differences in state policies but presumably also from state-level unobservables. To account for the latter, we include a state-specific parameter in $\Pr((\mathbf{s}, \boldsymbol{\chi})|x, \mathit{state})$. Third, via the correlation between x and unobservables, our model allows the distribution of $(\mathbf{s}, \boldsymbol{\chi})$ to vary across policy eras. However, the conditional distribution $\Pr((\mathbf{s}, \boldsymbol{\chi})|x, \mathit{state})$ is assumed to be policy invariant, which is key to identifying state-level unobservables.

2. *Individual Health Insurance and Medicaid Enrollment*

When neither spouse has ESHI ($\mathbf{z} = (0, 0)$), we introduce additional preference shocks $\varepsilon_{\mathbf{INS}}$ for a household's insurance decisions, which help explain choice variation and generate realistic demand elasticity (Handel 2013). To be specific, having made its labor supply choice, a household draws a preference shock from $N(0, \sigma^2_{\varepsilon_{\mathrm{MC}}})$ for Medicaid take-up and a shock from $N(0, \sigma^2_{\varepsilon_{\mathrm{IHI}}})$ for individual insurance enrollment. When making its labor supply decisions $(\mathbf{h}, \mathbf{z})$, the household forms expectations with respect to both its OOP shocks and $\{\varepsilon_{\mathbf{INS}}\}$; the value function (5) with $\mathbf{z} = (0, 0)$ is modified to

$$V(\mathbf{X}, m, \boldsymbol{\chi}, \mathbf{s}, \mathbf{h}, \mathbf{z} = (0, 0)) = E \max_{\mathbf{INS} \in \Omega(\mathbf{X}, y, m, \mathbf{z} = (0,0))} \left\{ \int u(C, \mathbf{h}, \mathbf{INS}; x, \boldsymbol{\chi})\, dF_{\mathrm{OOP}}(x, \mathbf{INS}, m) + \varepsilon_{\mathbf{INS}} \right\}. \tag{8}$$

3. *Production Function and Firm-Specific Technology*

We specify $F(\mathbf{n}; \boldsymbol{\Upsilon})$ as a modified CES production function to include two-dimensional firm heterogeneity $\boldsymbol{\Upsilon} = (T, A)$: T is TFP, and A measures the degree of skill biasedness. Such heterogeneity helps explain the fact that ESHI-providing firms tend to be larger and have higher fractions of workers earning high wages (table 2). All else being equal, higher-T firms would demand more labor and have a cost advantage for offering ESHI if the fixed cost of ESHI is nonnegligible; higher-A firms would have higher relative demand for skilled labor and would tend to offer ESHI if skilled workers have higher demand for health insurance. Moreover, we allow T and A to be correlated, which helps explain the correlation between a firm's size and its employee composition.

V. Data

For household information, we use the 5% sample of the ACS, the CPS Annual Social and Economic Supplement, and the MEPS. We focus on the population aged 22 to 64. For firm information, we use the Kaiser Family Employer Health Benefit Survey (Kaiser), supplemented with data from Statistics of US Businesses (SUSB). We use data from 2012 (pre-ACA era) and 2015 (ACA era).

Given the inconsistent insurance information in the CPS (Pascale 2016), we use the ACS for most household information (health insurance, labor market status, demographics, and residential states) and the CPS for health status information.[4] Denoting $\tilde{x}$ as the vector of household characteristics excluding its health status ($\tilde{x}$ is observed in both the CPS and the ACS), we use the CPS to estimate a logistic function $\Psi(healthy|\tilde{x}, state)$. The distribution $\Pr(x|state)$ is the product of the empirical distribution $\Pr(\tilde{x}|state)$ from the ACS and $\hat{\Psi}(healthy|\tilde{x}, state)$.

MEPS is a set of large-scale surveys of US households, their medical providers, and employers. We use its Household Component (geocoded), a panel survey covering two full calendar years. Key to our analyses, MEPS collects detailed information on each household member's demographics, health, income, employment, medical service usage, charges and source of payments, and health insurance coverage. MEPS identifies 30 states, accounting for 89% of US households, with the remaining states encrypted. We exclude Massachusetts and Hawaii, which implemented (nearly) statewide universal coverage before the ACA. We focus our analysis on the remaining 28 identified states, 15 of which expanded Medicaid.

Kaiser is a cross-sectional survey of firms representative of US firms with at least three workers. It contains information on firm size, ESHI provision, and employee composition in terms of wage levels and full-/part-time status. Our sample consists of all private employers with nonmissing information on ESHI provision. The distribution of other firm-level variables in this sample is similar to that in the entire private-firm sample; we assume that the ESHI offering information is missing at random. Firm locations are known up to the census region (Northeast, Midwest, South, and West), which allows us to estimate firm-side parameters separately for each region.[5]

A. Summary Statistics

Panel A of table 1 provides individual-level statistics before and after the ACA and separately for Medicaid expansion and nonexpansion states. Panel A1 shows that demographic distributions are largely stable between the two years and that expansion states have a higher-educated population than nonexpansion states. Panel A2 shows that the uninsured rate declined from 19.2% to 10.8% in expansion states (a 44% decrease) and from 26.3% to 19.4% in nonexpansion states (a 26% decrease). The biggest increase in

[4] Health status is self-reported as excellent, very good, good, fair, or poor. We define the first three categories as being healthy. This variable is highly correlated with gross medical costs (app. F, sec. F1).

[5] Ideally, we would focus on the same 28 states for both firm and household sides of the data, but Kaiser does not provide firms' state ID. Using SUSB, we find that within each census region, the distribution of firm characteristics in all states is very similar to that in states included in the household sample (fig. A1, available online).

Table 1
Individual Level Summary Statistics (%)

	A. By Residential States and Year			
	Medicaid Expansion States		Nonexpansion States	
	2012	2015	2012	2015
A1. Demographics:				
Education low (below high school)	12.13	11.49	13.48	13.10
Education high (above high school)	33.35	34.89	28.63	29.77
Single	42.09	42.82	42.80	44.24
Childless	61.34	62.21	61.62	62.31
A2. Insurance status:				
Uninsured	19.20	10.76	26.28	19.42
ESHI	68.08	69.29	63.73	66.81
Medicaid	7.71	12.97	4.56	5.54
Individual insurance	5.01	6.98	5.43	8.23
A3. Work status:				
Nonemployed	22.16	20.00	22.59	21.18
Full time	70.81	73.12	71.47	72.67
Number of Individuals (ACS)	27,140	27,465	18,927	19,734
	B. Sorting			
	Status Distribution by Education (ACS)			
	High (Above High School)	Middle (High School)	Low (Below High School)	% Unhealthy by Status (CPS)
B1. Insurance status:				
Uninsured	5.98	19.79	42.54	10.83
ESHI	84.26	64.85	34.38	5.55
Medicaid	1.90	9.29	19.27	18.74
Individual insurance	7.85	6.08	3.81	7.43
B2. Work status:				
Nonemployed	13.43	22.30	37.99	16.17
Part time	5.96	6.93	6.58	8.60
Full time	80.61	70.77	55.43	5.20
Number of Individuals	31,164	51,379	10,723	97,121

NOTE.—This table reports individual-level summary statistics from the ACS, except for those in the last column in panel B, which are from the CPS.

insurance enrollment was in Medicaid for expansion states and in individual insurance for nonexpansion states. Note that the Medicaid coverage rate was already higher in expansion states in 2012. Panel A3 shows that from 2012 to 2015, employment increased by 2.2 (1.4) percentage points in expansion (nonexpansion) states. Consistent with previous studies (e.g., Leung and Mas 2018), changes in work statuses are smaller than changes in insurance statuses.

Table 2
Summary Statistics: Firms

	A. Kaiser										
			Size		Size ≥ 500 (%)		% Full-Time Workers		% High-Wage Workers		
Year	Obs	ESHI (%)	All	ESHI	All	ESHI	All	ESHI	All	ESHI	
2012	1,981	56.1	21.9	32.9	.71	1.25	74.5	84.6	23.6	32.3	
		. . .	(55.0)	(70.6)	. . .	. . .	(29.6)	(19.6)	(26.6)	(26.2)	
2015	1,852	51.4	22.1	34.5	.74	1.42	72.8	79.3	26.9	33.7	
		. . .	(56.3)	(75.5)	. . .	. . .	(30.5)	(26.5)	(28.7)	(29.0)	

	B. SUSB			
	Size ≤ 4	4 < Size < 100	100 ≤ Size < 500	Size ≥ 500
2012 (%)	59.07	35.66	2.49	2.78
2015 (%)	59.93	36.08	2.20	1.80

Panel B shows the sorting of households into different insurance and work statuses. These patterns are qualitatively the same between 2012 and 2015 and between Medicaid expansion and nonexpansion states; therefore, we present statistics pooling all observations. The difference is clear across education groups. Most individuals in the high-education group are insured via ESHI (84.2%) and work full time (80.6%); in contrast, the majority of individuals in the low-education group are either uninsured (42.5%) or insured via Medicaid (19.3%), and 38% of them are nonemployed.

The last column of Panel B shows the percentage of unhealthy individuals for a given status. Consistent with the prediction in section III, the percentage of unhealthy individuals is lower among ESHI enrollees (5.6%) than among individual insurance enrollees (7.4%). This percentage is the highest among Medicaid enrollees (18.7%). Health status is also highly correlated with one's work status: more than 16% of the nonemployed are unhealthy, while only 5% full-time workers are unhealthy. These differences have substantial cost implications: for any given insurance status, the average medical expenditure among the unhealthy is more than three times as large as that among the healthy (app. F).

Table 2 summarizes firm-level data. Panel A summarizes the Kaiser sample (cross-firm standard deviations are in parentheses). For firm size and worker compositions, we present the statistics for all firms and for firms with ESHI. We also present the fraction of firms sized over 500 (firm sizes in Kaiser are top coded at 500). Compared with average firms, firms with ESHI are larger and have more full-time and more high-wage workers.[6]

[6] Kaiser specifies only three crude divisions of wage levels: \$24,000 (\$23,000) is the upper bound for low earnings and \$55,000 (\$58,000) is the lower bound for high earnings in 2012 (2015) in real dollar terms.

Data from SUSB (panel B) show that the firm size distribution is relatively stable between the two years.[7]

VI. Estimation

A. Parameters Estimated Outside the Model

We estimate, outside the model, the OOP health expenditure distribution $F_{\mathrm{OOP}}(\cdot)$, government health care–related policies, and net transfer functions $T^{\mathrm{ins}}(\cdot)$ and $T^{\mathrm{tax}}(\cdot)$. We provide a brief description here and details in appendix F.

OOP health expenditure consists of the health insurance premium $r^m(x)$ and out-of-pocket medical costs, estimated from MEPS. For $r^m(x)$ in the estimation sample, we use the observed average premium paid by households x on market m. A household's out-of-pocket medical cost is the sum of its members' gross medical costs minus the total reimbursement based on the most common health insurance plan. We estimate each household member's gross medical cost as a stochastic function of one's own characteristics, household characteristics, and insurance status, where the distribution of the random component is market specific.

ESHI premium q^m is set at the average ESHI premium reported by firms in Kaiser on market m.

Health care–related government policies are parameterized following rules specific to each market (state × era). Using Kaiser, we specify the Medicaid eligibility and coverage rule $\mathrm{MC}(\mathbf{X}, m, \mathbf{h}, \mathbf{z})$ as a market-specific function of household characteristics and income. In modeling $\mathrm{MC}(\cdot)$, we account for cross-state variation in Medicaid rules before the ACA in all states and after the ACA in nonexpansion states. We model Medicaid eligibility rules as defined by the federal government in Medicaid expansion states after the ACA. We abstract from asset testing for Medicaid, which would require detailed asset data and nontrivial complication in our setting.

Government net transfer function $T^{\mathrm{tax}}(\cdot)$ consists of income tax, welfare benefits (TANF), and food stamps (SNAP); $T^{\mathrm{ins}}(\cdot)$ consists of consumer subsidies on HIX and tax penalties for the uninsured (ACA individual mandates). We parameterize each of these components; we follow Chan (2013) in specifying TANF and SNAP functions.

B. Structural Estimation: Overview

1. Estimation Sample and Validation

We divide the household sample into an estimation sample and a validation sample. The estimation sample, from which our auxiliary models are calculated, includes the pre-ACA data from all 28 states in our sample

[7] SUSB firm size is categorical and size ≤4 is the first category, while Kaiser data only contains firms with size ≥3 and firm sizes are top coded. The top coding of firm sizes in Kaiser data is taken into account in our estimation (app. E, sec. E4).

and the post-ACA data from all but the seven states with the lowest poverty rates. The post-ACA data for these seven states are held out for model validation.

We use the data this way because, first, information from a state in at least one policy era is necessary to identify state-specific parameters; information from multiple states in both policy eras gives us the variation to identify policy-invariant household preference parameters without having to rely entirely on the model structure. Second, several major ACA components were targeted at low-income households, leading to potentially different impacts in states with different poverty rates. It will increase the credibility of our model and its counterfactual policy implications if the model is able to fit the post-ACA patterns in this nonrandom holdout sample. The holdout sample (lowest-poverty) states are indeed quite different from the other states (panel A of table 3): they are more likely to have expanded Medicaid (5/7 vs. 10/21) and have more educated residents.

2. *Equilibrium Prices in the Estimation*

Assuming that data are generated from an equilibrium, we only need to solve each household's and each firm's decision problem during the estimation. Among equilibrium prices, we observe health insurance premiums but not $\{w^m_{shz}\}$ because skill s is unobservable (although we observe the joint distribution of x, h, z, and wages). However, since the realized equilibrium wages $\{w^m_{shz}\}$ are taken as given by households and firms, they can be treated

Table 3
Sample Split and Two-Stage Estimation

	A. State Characteristics (Sample Split)					
			Education (%)			
	Number of States	Number of Medicaid Expansion States	High	Low	Singles (%)	Childless (%)
State groups:						
Lowest-poverty states	7	5	38.5	8.0	40.7	62.0
Other states	21	10	30.9	13.3	43.3	61.8

	B. Two-Stage Estimation	
	Structural	Nonstructural
Parameters to be estimated:		
Stage 1	Distribution of household preferences, types, and skills	Realized equilibrium wages
Stage 2	Distribution of firm production technology, ESHI fixed cost	. . .

Note.—Panel A provides state-level statistics for the holdout sample (lowest-poverty states) and the estimation sample (other states). Panel B lists parameters to be estimated in each stage of the two-stage estimation procedure.

as parameters to be estimated together with structural parameters. To keep the estimation tractable, we parameterize the distribution of $\{w^m_{shz}\}$. These wage parameters and functional-form assumptions are nonstructural and used only in the estimation. In counterfactual policy simulations, wages and insurance premiums are all treated nonparametrically and obtained by solving a fixed-point problem following definition 1.

3. *Two-Stage Estimation*

We estimate the model in two stages, also shown in panel B of table 3. In stage 1, we estimate household-side parameters (Θ^H) and the (nonstructural) wage parameters by matching model-predicted household decisions to the data. In stage 2, given the parameter estimates in stage 1, we estimate firm-side parameters (Θ^F) by matching firms' decisions to the data.

In both stages, the estimation is via indirect inference. Letting $\bar{\beta}$ be our chosen set of auxiliary model parameters computed from data and $\hat{\beta}(\Theta)$ be the corresponding auxiliary model parameters obtained from simulating a large dataset from the model (parameterized by a particular vector Θ) and computing the same estimators, the structural parameter estimator is the solution

$$\hat{\Theta} = \text{argmin}_{\Theta}\left[\hat{\beta}(\Theta) - \bar{\beta}\right]' W\left[\hat{\beta}(\Theta) - \bar{\beta}\right],$$

where W is a weighting matrix. We obtain standard errors for $\hat{\Theta}$ by the delta method numerically.

C. Structural Estimation: Identification and Auxiliary Models

The ACA induced rich variation that we exploit to identify our structural model. We summarize how the directly observed variation in prices and policies relates to our model. Notice that prices are endogenous; however, during the estimation we only need to solve individual household/firm problems, which take prices as given.

On the household side, (1) variation in individual health insurance premiums $r^m(x)$ affects households' OOP if they choose to enroll; (2) Medicaid eligibility rules MC(·), and hence the insurance choice set $\Omega(\cdot)$, differ across markets (state × policy era); and (3) HIX premium subsidies and the individual mandate both affect household budget via $T^{\text{ins}}(\cdot)$.

On the firm side, ESHI premiums q^m vary across markets. The ACA employer mandate directly changed the cost of not providing ESHI.

Besides exploiting policy variation for identification, we impose functional-form assumptions. The following points should ease the concern that our counterfactual policy results may be driven mostly by functional forms. First, the fundamental economic intuition underlying our policy implications is functional-form free (sec. III). Second, our model estimates are largely consistent with those in the literature (sec. VII.A). Third, following Todd and Wolpin (2006) and Keane and Wolpin (2007), we will conduct out-of-sample model validation in section VII.B to provide further credibility for our counterfactual analysis.

In the following, we provide an overview of the identification and summarize auxiliary models we target in the estimation. Appendix E provides more details.

1. *Stage 1 Estimation: Household-Side Parameters*

In stage 1 estimation, we design auxiliary models to identify household-side parameters governing (i) the distribution $\Pr((\mathbf{s}, \boldsymbol{\chi})|x, state)$, (ii) the disutility of work, (iii) wage offers, and (iv) preference for health insurance. All parameters are identified jointly. However, to the extent that certain aspects of the data are more informative of certain groups of parameters, we discuss the identification in parts. Section H1 of appendix H provides more evidence on the mapping between data and parameters by measuring the sensitivity of auxiliary models to model parameters (Cooper 2016; Einav, Finkelstein, and Mahoney 2018).

Preference for health insurance.—First, we discuss how we identify parameters governing households' preference for health insurance, given the rest of the model primitives. Consider the case where, given its labor market outcomes o, a household faces the choice between purchasing individual health insurance or staying uninsured ($\text{IHI} \in \{1, 0\}$). Simplifying the notation, the probability of observing IHI purchase is a weighted sum of type-specific probabilities of choosing IHI, given by

$$\Pr(\text{IHI} = 1|o) = \sum_{\chi} \Pr(\chi|o) \Pr(E[u^c|o, \chi, 1] + \varpi_{\text{IHI}} + \varepsilon_{\text{IHI}} > E[u^c|o, \chi, 0]), \tag{9}$$

where $E[u^c|o, \chi, \text{IHI}]$ is the expected type-specific utility from consumption conditional on o and IHI. The weight $\Pr(\chi|o)$ is the posterior type distribution conditional on o; for now, we take it as given and will discuss its identification next. As shown in Lewbel (2000), this type of model can be semiparametrically identified if there exists a variable, excluded from the type distribution, that has a large support and affects choice-specific payoffs. In our case, ACA-induced increase in IHI generosity serves as such an excluded variable: it made IHI more attractive by increasing $E[u^c|o, \chi, 1]$ for households of the same type. However, this variable has limited support; we therefore impose parametric assumptions. With these additional assumptions, we can use the covariation of the IHI generosity and the IHI take-up rate to identify type-specific preferences over consumption. We can then use the IHI take-up rate to identify the nonpecuniary preference ϖ_{IHI}. A similar argument can be used to identify Medicaid-specific preference parameters.

Other parameters.—We use (mainly) labor market outcomes before and after the ACA to identify the rest of household-side parameters, which are in turn used to derive the endogenous $\Pr(\chi|o)$ in the previous discussion. Our identification relies on the assumption that the distribution $\Pr((\mathbf{s}, \chi)|x, state)$ is policy invariant, which allows us to identify state-level unobservables using within-state variation in labor market outcomes before and after the ACA.

Our labor supply model is essentially a generalized Roy model (Heckman and Vytlacil 2007). As discussed in French and Taber (2011), identifying this class of model requires exclusion restrictions that affect the payoff only in the relevant sector. To supplement policy variation, we allow the distribution of types and skills (and thus wage offers) to vary by education, age, gender, and marital status but not by the presence of children or health status. By itself, either excluded variable increases the disutility of work and, via medical expenses, increases the value of ESHI jobs relative to non-ESHI jobs. More importantly, both excluded variables interact with policy changes. For example, although the ACA-induced change in equilibrium wages equally affects households of the same skill type within a state, ACA individual insurance premium subsidies, for which ESHI-covered workers are not eligible, interact with the size of the households. Moreover, the dependence of insurance premiums on health was allowed before the ACA and disallowed after the ACA. As such, the ACA directly changed the value of non-ESHI jobs, differentially so for households depending on the presence of children and/or health status; this creates policy variation within the same unobservable type of household, given our exclusion restriction.

To identify the correlation between s and χ, we exploit the fact that for the same household, we observe not only its labor market outcomes but also its insurance status. Conditional on $(x, state, year)$, the correlation between insurance take-up and income is informative of how skill and preferences are correlated. All else being equal, the CRRA utility function implies a negative correlation between income and insurance take-up. This is violated in the data, which suggests a positive correlation between skills and risk aversion.

In addition, our auxiliary models exploit other policy variation, such as the targeted nature of many ACA components, to inform us of model primitives. Note that although policies target income (endogenous in our model), we can still exploit how responses to the ACA vary by education.

Auxiliary models.—All auxiliary models in stage 1 estimation are based on the estimation sample only. Following the identification argument, our first set of auxiliary models are coefficients from regressions of the following form:

$$y_{ist} = x_{ist}\alpha_1 + d_s + I(t = 2015)x_{ist}[\mathrm{MEP}_s\alpha_2 + (1 - \mathrm{MEP}_s)\alpha_3] + \epsilon_{ist}, \qquad (10)$$

where y_{ist} is an indicator of a given insurance/work status for individual i, with characteristics x_{ist} in state s and year t; d_s is a state fixed effect; and $\mathrm{MEP}_s \in \{0, 1\}$ indicates whether state s expanded Medicaid under the

ACA. The other main stage 1 auxiliary models include (1) coefficients from a regression of (log) earnings on state dummies, the ACA era dummy, and one's work and insurance status; (2) moments reflecting the correlation between choices and health status; and (3) moments of joint outcomes between couples.

2. Stage 2 Estimation: Firm-Side Parameters

Firm-side parameters include those governing (i) the production function, (ii) the distribution of firm-specific technology $\Upsilon = (T, A)$ (TFP and skill biasedness), (iii) the fixed cost of ESHI, and (iv) the distribution of the random shocks associated with ESHI choice. We estimate firm-side parameters for each census region separately.

To see how these parameters are identified, it is useful to start with a firm's choice of labor input $\{n_{sh}\}$ for a given ESHI choice. A firm chooses $\{n_{sh}\}$ to equate the marginal productivity and the marginal cost of each labor input n_{sh}. The latter consists of wage and the expected cost of ESHI, both of which are known given estimates from stage 1: wages are stage 1 parameter estimates, the ESHI base premium is data, and the household expected demand for ESHI is derived from household preference parameters. The marginal productivity of labor is known up to parameters governing the production function and firm-specific technology (T, A). These parameters govern firms' size and labor composition for a given ESHI choice. In other words, from stage 1 estimation, we know the marginal costs of various types of labor inputs in each market (state $\times$ policy era). Relating the within- and cross-market variation in these marginal costs to the variation in employee composition provides major identifying information about parameters governing the production function.

The identification of the distribution of (T, A) needs to account for the fact that ESHI is an endogenous choice. For that, notice the following. First, the ACA employer mandate introduced a penalty for non-ESHI firms, where the penalty is a known function of the number of full-time-equivalent employees. The mandate changed the optimal choice of labor inputs for a given ESHI choice and the relative profitability of different choices of ESHI in a way that is known to us up to firm-side parameters (app. C). Second, for a given ESHI choice, the ratio of different types of labor depends on A but not T. However, T directly affects the size of a firm. As such, the correlation between labor ratio and firm size is directly informative of the correlation between T and A. Finally, given the assumption that idiosyncratic shocks to firms' ESHI decisions are independent of (T, A), the fixed cost of ESHI is identified from the relationship between firm size and ESHI offer rate (e.g., Aizawa and Fang 2020).

Auxiliary models.—Auxiliary models in stage 2 estimation includes policy era–specific moments constructed for each census region. These moments are designed to capture the joint distribution of the firm size, the employee composition (high/low wage, full/part time), and the ESHI provision, both before and after the ACA. In addition, we also target the aggregate supply of labor for each (s, h, z) category derived from stage 1 estimates to guarantee equilibrium consistency.

VII. Estimation Results

A. Parameter Estimates

We report a selected set of parameter estimates in the main text and the rest in section I3 of appendix I. The standard errors, derived numerically via the delta method, are in parentheses.

Household-side parameters.—Panel A of table 4 shows selected parameters governing household preferences. The left columns show that type 1 singles and (type 1, type 1) couples have higher relative risk aversion (γ_x) compared with their type 2 counterparts; households with mixed types of spouses have γ_x closer to type 1 households. These estimated γ's are in the range of the estimates in other studies (e.g., Cohen and Einav 2007; French and Jones 2011). The estimated annual consumption floor against health expenditure shocks ($2,600) is very close to that in De Nardi, French, and Jones (2010). The nonpecuniary value is estimated to be positive for ESHI and negative for Medicaid and individual insurance, the latter presumably reflecting the inertia or psychic cost associated with applying for Medicaid and individual insurance.

The right column of panel A show that unhealthy individuals and those with children have higher disutility of work and that type 1 individuals have lower disutility of work. We find that the disutility of working full time is lower than that of working part time. This is because work disutility in this model is a composite of various factors that affect labor supply choices beyond contemporary pecuniary benefits.

Panel B reports estimates for the distribution of types and skills. We find that individuals who are younger, lower educated, married, and female are more likely to be type 2 (the less risk averse type) and that spouses are more likely to be the same type. Moreover, type 1 (more risk averse) individuals are more likely to have higher skills. Based on these estimates, panel C reports the fraction of type 1 individuals by demographic group and by state of residence.

Relating to estimates in the literature.—The elasticity of the demand and the willingness to pay for health insurance implied by our estimates are comparable to those in the literature. In particular, as in Finkelstein, Hendren, and Shepard (2019), when we focus on the population choosing between HIX and noninsurance, we find that the HIX enrollment rate would be 49% if 75% of the premium cost is subsidized and 61% if 90% of the premium cost is subsidized; the corresponding enrollment rates in Finkelstein, Hendren, and Shepard (2019) are 49% and 79%. Similarly,

Table 4
Selected Parameter Estimates: Household

	Value		Value
A. Preferences:			
γ_x: type 1 singles or (types 1, 1) couples	4.12 (.003)	Disutility of working:	
γ_x: type 2 singles or (types 2, 2) couples	2.11 (.003)	Full-time job (unhealthy)	−3.29 (.03)
γ_x: (types 1, 2) couples	3.29 (.01)	Part-time job (unhealthy)	−3.16 (.03)
Consumption floor ($10,000)	.26 (.001)	Full-time job (type 1)	−2.10 (.004)
Nonpecuniary value ϖ_{INS}: Medicaid	−.40 (.002)	Part-time job (type 1)	−2.32 (.01)
Nonpecuniary value ϖ_{INS}: individual insurance	−.13 (.001)	Full-time job (type 2)	−3.36 (.01)
Nonpecuniary value ϖ_{INS}: ESHI	1.43 (.003)	Part-time job (type 2)	−4.23 (.01)
		Full-time job (child)	−.982 (.004)
		Part-time job (child)	−.911 (.004)
B. Type and skill distribution:			
$\Pr(\chi = 2 \mid x, state)$:		$\Pr(s \mid x, \chi)$:	
Age	−.42 (.002)	Age	.82 (.001)
Education = middle	.94 (.01)	Education = middle	−2.64 (.01)
Education = low	1.10 (.01)	Education = low	−3.88 (.01)
Married	.40 (.001)	Female	−.92 (.002)
Female	.58 (.002)	$\chi = 1$	1.51 (.003)
ϱ: type correlation between a couple	.77 (.01)		
C. Simulated type distribution in the sample, $\Pr(\chi = 1 \mid \cdot)$ (%):			
By demographics:		By state of residence:	
All	85.4	Expansion states	92.0
Singles	79.1	Nonexpansion states	75.1
Education = low	75.6	State poverty rate (lowest)	94.6
Education = high	95.3	State poverty rate (Q2)	93.8
Age >40	93.1	State poverty rate (Q3)	85.2
Childless	85.8	State poverty rate (highest)	70.6

NOTE.—See sec. A1 of app. A for functional forms and table D1 for state-specific parameters. Panels A and B report household-side parameter estimates; standard errors (in parentheses) are derived via the delta method. Panel C reports the fraction of type 1 by observables, calculated given estimates in panel B and table D1.

our estimates imply that among those covered by Medicaid, the willingness to pay for Medicaid is 21% of the cost of Medicaid, which is close to the 22%–46% range found in Finkelstein, Hendren, and Luttmer (2019).

Firm-side estimates.—Table A1 reports the estimates of main firm-side parameters. We find that higher-TFP firms are also more likely to be more skill biased and hence have higher demand for high-skill workers. As shown in table 4, individuals with higher risk aversion are more likely to have higher skill levels. As a result, higher-TFP firms are more likely to offer ESHI, and high-skill workers are more likely to sort into these firms.

Given our estimated model, for each market m and insurance type $k \in \{\text{HIX}, \text{ESHI}\}$, we obtain the loading factor l_k^m (the total premium divided by the total reimbursement) from the baseline equilibrium in the post-ACA era. We use these loading factors in our counterfactual policy experiments.

B. In-Sample Model Fit and Out-of-Sample Validation

Our model fits the data well in general (tables E1–E4; tables B1, B2, C1–C5, D1, D2, E1–E6, F1–F6, G1–G3 are available online). The model replicates reasonably well, for example, the pre- versus post-ACA difference in households' insurance and work statuses in both Medicaid expansion and nonexpansion states. It captures well the observed pattern that ESHI-providing firms tend to be larger and to hire higher fractions of full-time and high-wage workers; it also replicates how firms' hiring and ESHI-provision choices differ before and after the ACA.

For out-of-sample model validation, the model can reasonably replicate the patterns observed in the lowest-poverty-rate states in the post-ACA era (the holdout sample), both overall and in Medicaid-expansion states (table E5). Relative to the out-of-sample prediction from statistical multinomial logit models specified in section F4 of appendix F, our model outperforms. Given that the holdout sample is systematically different from the estimation sample, this validation exercise lends us some confidence of the model in conducting counterfactual policy experiments.

VIII. Counterfactual Experiments

Our estimation results suggest that high-skill workers are more likely to sort into firms offering ESHI, which tend to be more skill biased, and that households that choose to be nonemployed and/or earn wages low enough to be eligible for Medicaid are more likely to have low skills. These two groups of households are largely segregated from the HIX risk pool because regulation requires that the premium on an ESHI/HIX market should closely reflect the risk among those insured on that market. A lot of discussion has centered around how to encourage the uninsured—especially the healthy ones—to participate in HIX. A natural alternative is to look beyond the uninsured (14.4% of the working-age population) and to desegregate the risk pools between ESHI and HIX.

It should be noted that the welfare implication of risk pooling is theoretically ambiguous. Under the status quo, the ESHI pool is of lower risk than the HIX pool (table 1). Pooling risk across the two markets may decrease the HIX premium, making insurance more affordable for households not covered by ESHI or Medicaid. However, it may increase the ESHI premium and hence the cost of labor for firms were they to offer ESHI. As a result, risk pooling may have broad impacts on both the supply and the demand of labor for ESHI and non-ESHI jobs and hence equilibrium wages, calling for a framework that explicitly accounts for labor market adjustment. Moreover, the welfare implication depends on how households and firms would respond in the equilibrium, which in turn depends on the distribution of household and firm heterogeneity. Our estimated model provides us with all of these ingredients.

In the following, we first describe the ESHI-HIX risk-pooling equilibrium and show how this equilibrium can be achieved via a simple cross-subsidization policy. We then present the impact of this policy on the baseline economy and how the impact varies with Medicaid expansion policies. We conduct two sets of robustness checks in section H2 of appendix H.

A. Cross Subsidization between ESHI and HIX

1. *Risk-Pooling Equilibrium*

To pool risks across ESHI and HIX, the break-even condition should hold across these two markets (i.e., the sum of total expected costs for insurers on ESHI and HIX equals the sum of total premiums on both markets). Specifically, let $\tilde{r}_b^m$ be the new base premium on HIX, which implies age-adjusted premiums as

$$\tilde{r}^m(x) = \Gamma(\tilde{r}_b^m, age), \tag{11}$$

where $\Gamma(\cdot)$ is the same age-premium rule as in the baseline; let the premium on ESHI be

$$\tilde{q}^m = \theta \tilde{r}_b^m, \tag{12}$$

where θ is a modifiable policy parameter that governs the degree of premium adjustment. For a given θ, we solve for the new equilibrium wages and insurance premiums, such that under $\tilde{r}_b^m(\theta)$, which implies $\tilde{r}^m(x;\theta)$ and $\tilde{q}^m(\theta)$, the break-even condition holds across ESHI and HIX. Such an equilibrium effectively pools the risks on ESHI and HIX; equilibrium prices and outcomes are governed by θ.

2. *Implementation: ESHI-HIX Cross Subsidization*

The idea of pooling the risks on ESHI and HIX is of little practical value if it is very hard to implement. We design a policy tool—ESHI-HIX cross subsidization—that can easily implement the risk-pooling equilibrium associated with any given $\tilde{r}_b^m(\theta)$. Specifically, for $k \in \{\text{HIX}, \text{ESHI}\}$ and market m, let $\mu_k^m(x;\theta)$ be the measure of households with characteristics x that opt for k on m in the new equilibrium associated with $\tilde{r}_b^m(\theta)$, and let $C_k^m(x;\theta)$ be the average expected cost—that is, the baseline loading factor times the expected reimbursement—among these households for the insurer. The ESHI-HIX risk-pooling equilibrium with $\tilde{r}_b^m(\theta)$ can be implemented by imposing taxes $\tau_k^m(\theta)$ defined by

$$(1-\tau_{\text{HIX}}^m(\theta))\int \mu_{\text{HIX}}^m(x;\theta)\tilde{r}^m(x;\theta)\,dF_m(x) = \int \mu_{\text{HIX}}^m(x;\theta)C_{\text{HIX}}^m(x;\theta)\,dF_m(x),$$

$$(1-\tau_{\text{ESHI}}^m(\theta))\int \mu_{\text{ESHI}}^m(x;\theta)\tilde{q}^m(\theta)\,dF_m(x) = \int \mu_{\text{ESHI}}^m(x;\theta)C_{\text{ESHI}}^m(x;\theta)\,dF_m(x).$$

The tax $\tau_k^m(\theta)$ is positive (negative) if the total cost for the insurer on k is smaller (larger) than the total premium collected on k; insurers are break-even after tax on each market.

After imposing $\tau_k^m(\theta)$ on insurers, there is no need for further intervention: HIX and ESHI markets would still operate separately, yet the equilibrium premiums on HIX and ESHI would be the desired risk-pooling equilibrium premiums. By construction, the total subsidy allocated to insurers on the riskier market is offset by the total tax collected from insurers on the healthier market.

The policy parameter θ serves to adjust the degree of cross subsidization between ESHI and HIX: a higher θ implies a larger subsidization flowing from ESHI to HIX. As a starting point, we consider a θ just enough to offset the difference between ESHI and HIX in their actuarially fair values and quality of care. We denote this special parameter value as θ°, which is calibrated at 1.45 (app. G, sec. G2). The equilibrium achieved under θ° is one that pools the risk across ESHI and HIX without further adjustment, in that premiums on the two markets are equalized conditional on quality. Then we experiment by increasing θ up to $2\theta^\circ$. Among these experiments, we find qualitatively consistent results; quantitatively, the welfare impact increases at first, peaks around $1.5\theta^\circ$, and slightly decreases afterward. In the following, we report impacts of two cross-subsidization policies, with $\theta = \theta^\circ$ and $1.5\theta^\circ$, respectively, imposed on the baseline economy with state-specific policies as implemented in 2015; we also discuss why welfare effect does not always increase with θ.

3. Policy Impacts

Panel A of table 5 shows percentage changes in equilibrium prices, averaged across states. Premiums adjust much more for HIX than for ESHI, mainly because a one-dollar transfer from a large market (ESHI) to a small market (HIX) will have a more noticeable impact on the latter. The decrease

Table 5
Cross Subsidization between ESHI and HIX: Prices, Status, Earnings, and Productivity

	$\theta = \theta^0$			$\theta = 1.5\theta^0$		
A. ΔPrices (%):						
Premiums:						
HIX		−6.15			−33.52	
ESHI		.56			2.68	
Wages:						
Non-ESHI jobs		−.19			−.07	
ESHI jobs		.38			.36	
	All	Low Education	High Education	All	Low Education	High Education
B. ΔStatus (ppt):						
Uninsured	−.02	−.01	−.03	−.21	−.16	−.13
HIX	.12	.09	.11	.52	.45	.36
ESHI	−.09	−.08	−.07	−.30	−.25	−.23
Nonwork	−.01	−.04	.01	.03	.01	.04
Full time	.02	.05	.00	.00	.02	.05
C. ΔEarnings and ΔOutput (%):						
ΔEarnings∣employed		.58			.52	
ΔTotal output		3.06			2.16	

NOTE.—This table reports equilibrium effects of ESHI-HIX cross subsidization when $\theta = \theta^0$ and $\theta = 1.5\theta^0$ on premiums and wages (panel A), insurance and work status (panel B), and earnings and total outputs (panel C). Results for single and childless adults are reported in table F2. ppt = percentage point.

in HIX premiums reduces the value of obtaining insurance from one's employer, which should reduce the compensating wage differential. However, how wages for each type of jobs would change is less clear ex ante. We find that wages for non-ESHI jobs (w_0) are barely affected, while wages for ESHI jobs (w_1) slightly increase.

Panel B shows the percentage point (ppt) changes in households' insurance and work statuses. Under both θs, cross subsidization increases the HIX coverage, reduces the ESHI coverage, and decreases the uninsured rate. The policy has very limited impact on employment status, but it increases average earnings. Furthermore, via their effects on matches between firms and bundles of labor inputs, these policies have nontrivial 2%–3% positive impacts on total outputs.

Result 1.—ESHI-HIX cross subsidization has small positive effects on the insured rate and average earnings and moderately positive effects on outputs.

Changes in insurance status are small despite the significant decrease in HIX premiums. This is because the relevant price a household faces is the premium net of HIX subsidies. Under the counterfactual policies, higher-income households, which are ineligible for HIX premium subsidies, see a larger decrease in net HIX premiums; however, they are also much more likely to be covered by ESHI. In contrast, subsidy-eligible households see a smaller change or even no change in net premiums (for an illustration, see table F1).

The upper part of table 6 shows the change in households' ex ante welfare measured by annual consumption equivalent variation (CEV) and the fraction of winning households. The pure risk-pooling policy ($\theta = \theta^0$) benefits 70% of households and increases the average household welfare by $75 annual CEV. With $\theta = 1.5\theta^0$, the policy benefits 81% of households and increases the average household welfare by $159. ESHI-HIX cross subsidization lowers the cost of health insurance on HIX and hence smooths consumption for households across different situations they may be in (e.g., with vs. without ESHI). Type 1 households, which are more risk averse, benefit more from such consumption-smoothing effects and hence are more likely to win. Households with higher education are more likely to win because (1) they are more likely to be type 1 (more risk averse); (2) they are more likely to work on ESHI jobs, the wages of which increase slightly (table 5); and (3) they are less likely to qualify for HIX subsidies and hence more likely to benefit from the decrease in HIX premiums. Notice that although risk segregation in itself implies a regressive welfare effect (sec. III), this effect is largely alleviated and even outweighed by the progressive HIX premium subsidy policy. As HIX premiums decrease, premium subsidies become less progressive.

The last two rows of table 6 show the effect on government net spending in the health insurance system, which includes expenditures on Medicaid and HIX subsidies net of revenues from insurance mandate tax penalties (the cross subsidization between ESHI and HIX is revenue neutral). This

Table 6
Cross Subsidization between ESHI and HIX: Household Welfare and Government Spending

	$\theta = \theta^0$		$\theta = 1.5\theta^0$	
Welfare	CEV ($)	Fr(winners)	CEV ($)	Fr(winners)
Overall	75.2	.70	158.8	.81
Type 1 singles or (type 1, type 1) couples	90.3	.73	183.3	.84
Type 2 singles or (type 2, type 2) couples	9.7	.54	40.8	.59
Low-education singles or (low, low) couples	40.4	.63	93.1	.69
High-education singles or (high, high) couples	119.1	.78	207.9	.88
Savings in government expenditure per household ($)	14.3		49.0	
Savings in HIX subsidies per enrolled household ($)	204.6		763.8	

Note.—This table reports equilibrium effects of ESHI-HIX cross subsidization when $\theta = \theta^0$ and $\theta = 1.5\theta^0$ on household welfare and government expenditure. Results for single and childless adults are reported in table F2.

net spending per household decreases by \$14 under $\theta = \theta^0$ and by \$49 under $\theta = 1.5\theta^0$, mostly from savings in HIX premium subsidies. If these savings were used as transfers to low-skilled households, these households would see larger welfare gains from ESHI-HIX cross subsidization than those reported in table 6.

Result 2.—ESHI-HIX cross subsidization benefits most households, moderately increases average household welfare, and lowers government expenditure.

B. ESHI-HIX Cross Subsidization versus Alternative Policy Measures

1. *Risk Pooling across States*

Besides the segmentation between ESHI and HIX, the current system also segregates HIX markets across states. An alternative is to allow households to buy HIX health insurance in any state and hence pool the risk across states for HIX. Column 1 of table 7 shows that this policy reduces the uninsured rate by 0.37 percentage points, a larger decline than those under ESHI-HIX cross subsidization (cols. 2 and 3, which repeat impacts in tables 5 and 6). However, it also reduces employment, decreases household welfare, and increases government expenditure. While cross-state risk pooling in HIX has slightly negative welfare impacts, ESHI-HIX cross subsidization leads to welfare gains. This arises from the fundamental difference between these two types of risk-pooling policies. Cross-state risk pooling benefits households in higher-premium states at the cost of those in

Table 7
ESHI-HIX Cross Subsidization versus Cross-State Risk Pooling for HIX

		ESHI-HIX Pooling	
	Cross-State Pooling	θ^0	$1.5\theta^0$
A. ΔStatus (ppt):			
Uninsured	−.37	−.02	−.21
HIX	−.21	.12	.52
ESHI	.66	−.09	−.30
Nonwork	.17	−.01	.03
Full time	−.45	.02	.00
B. Welfare gains:			
Average CEV (\$)	−1.7	75.2	158.8
Fr(winners)	.47	.70	.81
C. Savings in government expenditure (\$):			
Savings in government expenditure per household	3.2	14.3	49.0
Savings in HIX subsidies per enrolled household	18.8	204.6	763.8

NOTE.—ppt = percentage point.

lower-premium states. In contrast, ESHI-HIX pooling smooths consumption for the same household across different situations it may be in; given risk aversion, this benefits most households.

Result 3.—Relative to a policy that pools risks on HIX markets across states, ESHI-HIX cross subsidization leads to larger household welfare gains and lower government expenditure.

2. *Risk Pooling through Cross Subsidization versus Literally Pooling into One Insurance Market*

As mentioned in section VIII.A.2, welfare does not increase monotonically with θ, which suggests that welfare-maximizing θ may be finite. Since an infinite θ would essentially shut down ESHI, our results suggest that (finite) ESHI-HIX cross subsidization may outperform a more extreme policy that eliminates ESHI and literally pools all households into the HIX market.

The fundamental reason is as follows. Relative to nonemployed workers, employed workers tend to be healthier because health is correlated negatively with work disutility and positively with skills. Under ESHI-HIX cross subsidization, households in the ESHI pool (the healthier ones) subsidize households in the HIX pool (the less healthy ones), which lowers HIX premiums and expands coverage. ESHI workers are "willing" to do so as long as compensating wage differentials ($w_0 - w_1$) are not too large. As shown in table 5, relative to the large decrease in HIX premiums, changes are small in both ESHI premiums and wages, which in turn leads to only a small decrease in ESHI enrollment—the pool that subsidizes HIX.[8] By contrast, if all households are pooled onto HIX, among the originally ESHI-insured households, only those whose willingness to pay exceeds the HIX premium—potentially less healthy ones—will buy HIX. As a result, the HIX premium will remain relatively high and the coverage relatively low. In addition, there are several nonessential but quantitatively important factors that also play a role, including households' nonpecuniary preference for ESHI and ESHI tax deductibility.

C. Interaction between ESHI-HIX Cross Subsidization and Medicaid

It is unclear whether ESHI-HIX cross subsidization would be more effective with or without Medicaid expansion. On the one hand, with Medicaid expansion, fewer people would be uninsured, leaving smaller room for improvement from a decrease in HIX premiums. On the other hand, because Medicaid absorbs a disproportionately unhealthy population, when it

[8] Changes in wages are small partly because firms' choices of ESHI provision are not very elastic with respect to the cost of ESHI: with concave production technology, firms earn positive profits and some firms strictly prefer offering ESHI.

Table 8
Effects of ESHI-HIX Cross Subsidization by Medicaid Expansion Status

	Medicaid Expansion, MEP Complier States		Medicaid Expansion, Noncomplier States	
	No (1)	**Yes** (2)	**No** (3)	Yes (4)
$\theta = \theta^0$:				
Change in uninsured (ppt)	.19	−.04	.00	−.03
CEV ($)	−42.3	128.2	2.5	26.3
Fr(winner)	.52	.71	.69	.56
Savings for government per household	20.5	14.4	14.2	14.9
$\theta = 1.5\theta^0$:				
Change in uninsured (ppt)	−.02	−.24	−.16	−.15
CEV ($)	10.3	184.1	123.9	124.3
Fr(winner)	.71	.85	.74	.77
Savings for government per household	51.4	42.8	57.6	43.7

NOTE.—Given θ, the first (second) column shows ESHI-HIX cross-subsidization effects without (with) Medicaid expansion for MEP-complying states. The third (fourth) column shows ESHI-HIX cross-subsidization effects without (with) Medicaid expansion for MEP-noncomplying states. Boldface type for yes/no indicates the observed Medicaid expansion status in 2015. ppt = percentage point.

is expanded, the risk pool on HIX is relatively healthier and cross subsidization would be less distorting for ESHI premiums. We simulate the impacts of ESHI-HIX cross subsidization separately for the 15 ACA Medicaid expansion (MEP)–complying states and 13 noncomplying states under counterfactual scenarios with and without Medicaid expansion. That is, we examine the impacts of ESHI-HIX cross subsidization on different groups of states given the same hypothetical Medicaid expansion status and on the same group of states under different hypothetical Medicaid expansion statuses. For MEP-complying states, we use their 2012 state-specific Medicaid eligibility rules in the counterfactual nonexpansion scenario.

Table 8 shows these effects under $\theta = \theta^0$ (the upper panel) and $\theta = 1.5\theta^0$ (the lower panel). Within each panel, the first two columns are for the 15 MEP-complying states, showing the effect of ESHI-HIX cross subsidization without and with Medicaid expansion, respectively.[9] The third and fourth columns show the same statistics for the 13 MEP noncomplying states. For each group of states, the boldfaced yes/no status is their observed Medicaid expansion status in 2015. In almost all scenarios, ESHI-HIX cross subsidization improves welfare and lowers the uninsured rate and government expenditures. The only exception happens in column 1 in the upper panel: were Medicaid not expanded in MPE-complying states, the pure risk-pooling policy would lead to a small increase in the uninsured rate

[9] For example, to get the results shown in col. 1, for each state we compute the equilibrium if Medicaid were not expanded and there is no ESHI-HIX cross subsidization (E0), then we compute the equilibrium if Medicaid were not expanded but ESHI-HIX cross subsidization were in place (E1); col. 1 shows the difference between E1 and E0. Table A2 shows changes in wages and insurance premiums.

and a welfare loss. A more important finding is that given θ and the same group of states, welfare gains from the cross subsidization tend to be larger when Medicaid is expanded (cols. 1 vs. 2 and 3 vs. 4).

Result 4.—ESHI-HIX cross subsidization leads to higher welfare gains when it is interacted with Medicaid expansion.

IX. Conclusion

We have developed and estimated an equilibrium model of the labor market and health insurance markets to examine the inefficiency associated with risk pool segregation across various health insurance channels in the United States. The model features rich heterogeneity across local markets, workers, and firms. We estimate the model exploiting policy variation associated with the ACA. The estimated model matches the data well, including patterns in the holdout sample.

Via counterfactual policy experiments, we find that the risk pool segregation of the US health insurance system has led to a sizable loss in total output and household welfare as well as wasteful government spending. These losses arise from the fact that the risk pool on HIX is more adversely selected than that on ESHI, which implies a higher premium on HIX than that on ESHI (conditional on quality) under risk pool segregation. This premium differential has led to misallocation on the labor market. We find that ESHI-HIX cross subsidization could break the risk pool segregation, increase total labor output, improve household welfare, and lower government expenditure. Moreover, the policy would lead to higher welfare gains when it is interacted with Medicaid expansion. These findings suggest that policy tools designed to improve the risk pool structure across health insurance channels are promising avenues that warrant further exploration because they can be welfare enhancing, cost-effective, and easily implementable (e.g., via ESHI-HIX cross subsidization). More generally, these findings have also illustrated the value of a framework like ours, which enables one to study policies that regulate different parts of the health insurance system in a complementary manner. As such, this paper has made a modest step toward the goal of answering globally optimal social insurance design questions, as pointed out by Chetty and Finkelstein (2013).

This paper has several important limitations. For example, without considering the funding regime (e.g., the tax system) underlying the health insurance system and the general equilibrium effect on health care costs (e.g., responses from the health care sector), this paper cannot properly study the effect of more drastic health insurance reforms, such as "Medicare for all." We have also left several challenging extensions for future work. One extension is to embed dynamics into our framework, including household savings and potential direct effects of health insurance on one's health and hence future productivity. A second extension is to consider sources of inefficiency that we have abstracted from, such as search friction on the labor market (e.g., Dey and Flinn 2005) and choice friction on health insurance markets. Finally, we have abstracted from the potential effect of risk pooling on the quality of insurance products and enrollees' medical spending choices (e.g., Dickstein, Ho, and Mark 2024). Extending our framework along these lines would be promising for future research.

Appendix

A. Detailed Empirical Specifications

A1. Household Unobservables

We model the distribution of household unobservable skill and preference types as

$$\Pr((\mathbf{s}, \boldsymbol{\chi})|x, state) = \Pr(\boldsymbol{\chi}|x, state)\Pr(\mathbf{s}|x, \boldsymbol{\chi}).$$

Preference type.—Denote elements in x such that the subvectors x_1 and x_2 refer to the individual characteristics of spouse 1 and spouse 2 and that x_0 refers to household-level characteristics. We assume that types of a couple $\boldsymbol{\chi} \in \{1, 2\}^2$ follow a bivariate probit distribution, with the latent variables drawn from

$$N\left(\begin{bmatrix} x_0\beta_0 + x_1\beta + \xi_{state} \\ x_0\beta_0 + x_2\beta + \xi_{state} \end{bmatrix}, \begin{bmatrix} 1, \varrho \\ \varrho, 1 \end{bmatrix}\right), \tag{A1}$$

where ξ_{state} is a state-specific parameter and ϱ allows for matching on unobservables between a couple. For singles, $\Pr(\chi = 2) = \Phi(x_0\beta_0 + x_1\beta + \xi_{state})$. We do not impose any restriction on $\{\xi_{state}\}$; ξ_{state} can be freely correlated with a state's policies.

Skill.—The probability that a worker's skill is of level s follows a discretized lognormal distribution:

$$\Pr(s|x, \chi) =$$

$$\begin{cases} \Phi(\ln(k_s) - x'\lambda - \alpha_\chi) - \Phi(\ln(k_{s-1}) - x'\lambda - \alpha_\chi) & \text{for } 1 < s < S, \\ \Phi(\ln(k_s) - x'\lambda - \alpha_\chi) & \text{for } s = 1, \\ 1 - \Phi(\ln(k_{s-1}) - x'\lambda - \alpha_\chi) & \text{for } s = S, \end{cases} \tag{A2}$$

where α_χ is a type-specific parameter that allows for correlation between s and χ, with α_2 normalized to zero. The mass points of the amounts skill (k_s) are quantiles from $\ln N(\bar{x}'\lambda, 1)$, where $\bar{x}$ is the national average of x. That is, one's rank in the skill distribution is correlated with the distance of one's x from the national average. Notice that worker's skill and firm's TFP (two unobservable levels) jointly map into one observable object, that is, wage.

Therefore, one of the two unobservables (skill, TFP) needs to be normalized. We use the quantiles of $\ln N(\bar{x}'\lambda, 1)$ as the mass points of k_s levels, which serves as a normalization. The distribution of a couple's skills is given by $\Pr(\mathbf{s}|x, \boldsymbol{\chi}) = \Pr(s|x, \chi)\Pr(s'|x, \chi')$.

A2. Household Choice Set for Health Insurance

Conditional on $(\mathbf{h}, \mathbf{z})$, a household's choice set for health insurance is given by $\Omega(\mathbf{X}, m, \mathbf{h}, \mathbf{z}) \subset \{0, 1, 2, 3, 4\}^2$. In line with most households' choices, we assume the following. First, if only one spouse works on a job with ESHI, the other spouse and children will be covered. Second, if both spouses are covered by ESHI, they are indifferent between whose employer covers their children (i.e., in expectation, the burden of child health insurance will be split evenly between the two employers). Third, conditional on choosing $\mathbf{z} = (0, 0)$, if a household is eligible for Medicaid ($\mathrm{MC}(\mathbf{X}, m, \mathbf{h}, \mathbf{z}) = 1$), it chooses between using Medicaid or staying uninsured for the entire household; if $\mathrm{MC}(\mathbf{X}, m, \mathbf{h}, \mathbf{z}) = 0$, it chooses between individual health insurance and staying uninsured for the entire household. Therefore, the choice set $\Omega(\cdot)$ is given by

$$\begin{aligned}
&\Omega(\mathbf{X}, m, \mathbf{h}, \mathbf{z} = (1, 0)) = \{(1, 2)\},\\
&\Omega(\mathbf{X}, m, \mathbf{h}, \mathbf{z} = (0, 1)) = \{(2, 1)\},\\
&\Omega(\mathbf{X}, m, \mathbf{h}, \mathbf{z} = (1, 1)) = \{(1, 1)\},\\
&\Omega(\mathbf{X}, m, \mathbf{h}, \mathbf{z} = (0, 0)) = \begin{cases} \{(3, 3), (0, 0)\} & \text{if } \mathrm{MC}(\mathbf{X}, m, \mathbf{h}, \mathbf{z}) = 1, \\ \{(4, 4), (0, 0)\} & \text{if } \mathrm{MC}(\mathbf{X}, m, \mathbf{h}, \mathbf{z}) = 0. \end{cases}
\end{aligned} \tag{A3}$$

A3. Household Utility

In utility function (2), we assume the risk-aversion parameter $\gamma_{\mathbf{X}} = \gamma_{\chi}$ (i.e., $\gamma_{\mathbf{X}}$ varies with $\mathbf{X}$ only through $\boldsymbol{\chi}$). We set $size_x = 1$ for singles without children, $size_{\mathrm{x}} = 1.3$ for singles with children, $size_x = 1.5$ for couples without children, and $size_{\mathrm{x}} = 1.8$ for couples with children. The disutility from working is given by

$$D(\mathbf{h}, \boldsymbol{\chi}, x) =$$
$$\begin{cases} \sum_{l=P,F} I(h = l)(d_{\chi l} + \phi_{1l} I(kid > 0) + \phi_{2l} I(unhealthy)) & \text{if single,} \\ \upsilon \sum_{n=1}^{2} \sum_{l=P,F} I(h_n = l)(d_{\chi l} + \phi_{1l} I(kid > 0) + \phi_{2l} I(unhealthy)) & \text{otherwise,} \end{cases}$$

where $d_{\chi l}$ is a type-specific disutility of working with status $l = P, F$; φ_{1l} (φ_{2l}) is the additional disutility from working in the presence of young children (bad health). For a coupled household, the disutility is summed over each spouse's disutility, with a scale parameter υ to be estimated. Finally,

$$\varpi_{\text{INS}} = \sum_{k=1,3,4} \varpi_k I(\mathbf{INS} = k),$$

where $\{\varpi_k\}$ is household's nonpecuniary preferences for ESHI, Medicaid, and individual insurance.

A4. Firm Technology and Costs

Production function and firm-specific technology.—We model a firm's technology factor as $\Upsilon = (T, A)$ and modify the standard CES production function (e.g., Krusell et al. 2000) as follows:

$$F(\mathbf{n}; \Upsilon) = T\left[A\sum_{s\geq s^*} B_{sF} l_{sF}^{\rho} + (1 - A)\left(\sum_{s<s^*} B_{sF} l_{sF}^{\rho} + \sum_{s=1}^{S} B_{sP} l_{sP}^{\rho}\right)\right]^{\theta/\rho}, \quad \text{(A4)}$$

$$\text{s.t. } l_{sh} = k_s n_{sh},$$

where k_s denotes the amount of skills possessed by a worker of skill level s and n_{sh} is the number of employees with skill level s and work status $h \in \{P, F\}$. Firm-specific technologies are captured by the TFP parameter T and by $A \in (0, 1)$, the latter of which measures the degree of skill biasedness toward full-time high-skill labor. Empirically, we define the top 40% in the skill distribution as $s \geq s^*$ (i.e., $s^* = 4$). The parameters θ, ρ, and B are common across firms, with $\Sigma_{s,h} B_{sh} = 1$.

We allow T and A to be correlated within a firm, but $\{(T, A)\}$ are assumed to be independent across firms. Firm's TFP T follows a Pareto distribution: $T \sim \text{Pareto}(\underline{T}, \alpha_T)$, where $\underline{T}$ is the scale parameter and α_T is the shape parameter. As a convenient way to guarantee that $A \in (0, 1)$, skill-biasedness parameter A is assumed to follow a logit normal distribution:

$$\left[\ln\left(\frac{A}{1 - A}\right)|T\right] \sim N\left(\ln\left(\frac{\mu_A}{1 - \mu_A}\right) + \nu(\ln(T) - \ln(\mu_T))), \sigma_A^2\right),$$

where μ_A is the median of A for firms with $T = \mu_T \equiv E[T]$ and ν governs the correlation between T and A.

Firms' costs.—Firms' costs are modeled as

$$\text{TC}(\mathbf{n}, \mathbf{z}, m) = \sum_{s,h} n_{sh} w_{shz}^m (1 + \tau_w^m) + I(\mathbf{z} \neq 0)(\delta + \text{TP}(\{n_{sh} z_{sh}\}_{sh}, m)). \quad \text{(A5)}$$

The first cost component is total wages plus payroll taxes (at the rate τ_w^m). If ESHI is offered ($\mathbf{z} \neq 0$), the firm pays a fixed cost δ and the total ESHI premium for covered workers $\text{TP}(\{n_{sh} z_{sh}\}_{sh}, m)$. A firm's total premium is

$$\text{TP}(\{n_{sh} z_{sh}\}_{sh}, m) = q^m \sum_{s,h} z_{sh} n_{sh} \kappa_{sh}^m,$$

where q^m is the base premium on ESHI market m and κ_{sh}^m is the expected number of household members that worker with (s, h) needs ESHI coverage for (i.e., besides oneself), whether or not coverage is needed for one's spouse and/or child. This is an endogenous outcome of worker-firm sorting in the equilibrium. In modeling how workers' differential health insurance demands may affect individual firms' costs, we restrict our attention to the number of household members to be covered. Differential risks associated with other factors (e.g., age and health) will be pooled and reflected in the base premium q^m, which adjusts to satisfy the break-even condition.

A5. Equilibrium Prices in the Estimation

We assume that wages of non-ESHI jobs $\{w_{sh0}^m\}$ can be approximated by a discretized lognormal distribution. On market m, the skill-specific log wages $\{\ln(w_{sh0}^m)\}_{s=1}^S$ on non-ESHI jobs with hour status h are quintiles from

$$N\left(\omega_h^0 + \omega_{state}^0 + \omega_{year}^0, \sigma_{wh}^2\right) \quad \text{for } h \in \{P, F\}, \tag{A6}$$

where the mean is governed by the vector ω_0, including the part-/full-time dummy (ω_h^0), state dummies (ω_{state}^0), and year dummies (ω_{year}^0). To capture the idea of compensating wage differentials, we assume that wages of ESHI jobs are proportional to their non-ESHI counterparts. The wage ratio is allowed to vary with wage levels:

$$\frac{w_{sh1}^m}{w_{sh0}^m} = \frac{1}{1 + \exp(\omega_0^1 + \omega_1^1 w_{sh0}^m)}. \tag{A7}$$

Depending on the sign of the parameter ω_1^1, compensating differentials may be higher or lower for higher-skill workers. We treat the vector $\{\omega^0, \omega^1, \sigma_w\}$ in expression (A6) and equation (A7) as parameters to be estimated, which jointly imply $\{w_{shz}^m\}$.

Table A1
Selected Firm-Side Parameter Estimates

	Region			
	Northeast	Midwest	West	South
	A. TFP Distribution $T_j \sim \text{Pareto}(\underline{T}, \alpha_T)$			
Scale $\underline{T}$ (2012)	24.53 (3.52)	25.50 (4.71)	25.57 (1.58)	25.09 (4.09)
Scale $\underline{T}$ (2015)	25.11 (1.40)	25.95 (3.15)	26.41 (3.37)	25.50 (3.09)
Shape α_T	3.49 (.26)	3.76 (.51)	3.90 (.21)	4.14 (.19)
	B. Skill Bias $\ln(A_j/(1-A_j)))\vert T_j \sim N(\ln((\mu_A/(1-\mu_A))) + \nu(\ln(T_j) - \ln(\mu_T))), \sigma_A^2)$			
μ_A	.67 (.122)	.73 (.04)	.74 (.04)	.68 (.02)
σ_A	1.41 (.198)	1.61 (.35)	2.19 (.23)	1.27 (.22)
ν	.86 (.661)	.93 (.19)	1.55 (.85)	1.15 (.04)
	C. Other Selected Parameters			
ρ (CES power parameter)	.42 (.03)	.41 (.02)	.40 (.01)	.45 (.01)
Fixed cost of ESHI ($10,000)	3.57 (.53)	3.07 (.84)	2.99 (.56)	5.09 (1.84)
σ_η (ESHI shock scale parameter)	1.92 (.43)	1.87 (1.02)	1.91 (.77)	1.92 (1.19)

NOTE.—Standard errors (in parentheses) are derived using the delta method.

Table A2
Price Effects of ESHI-HIX Cross Subsidization by Medicaid Expansion Status

	Medicaid Expansion, MEP Complier States		Medicaid Expansion, Noncomplier States	
	No (1)	**Yes** (2)	**No** (3)	Yes (4)
$\theta = \theta^0$:				
Non-ESHI wages (w_0)	−.10	.11	−.55	−.52
ESHI wages (w_1)	.48	.32	.45	.36
HIX premium (r)	−5.05	−4.94	−6.82	−7.80
ESHI premium (q)	.74	.48	.60	.69
$\theta = 1.5\theta^0$:				
Non-ESHI wages (w_0)	.51	.54	−.78	−.76
ESHI wages (w_1)	−.05	−.01	.79	.64
HIX premium (r)	−32.89	−32.69	−34.47	−33.92
ESHI premium (q)	2.64	2.54	2.83	2.84

NOTE.—Given θ, the first (second) column shows ESHI-HIX cross-subsidization effects without (with) Medicaid expansion for MEP-complying states. The third (fourth) column shows ESHI-HIX cross-subsidization effects without (with) Medicaid expansion for MEP-noncomplying states. Boldface type for yes/no indicates the observed Medicaid expansion status in 2015.

References

Aizawa, Naoki. 2019. Labor market sorting and health insurance system design. *Quantitative Economics* 10, no. 4:1401–51.

Aizawa, Naoki, and Hanming Fang. 2020. Equilibrium labor market search and health insurance reform. *Journal of Political Economy* 128, no. 1) 4258–336.

Akerlof, George A. 1970. The market for "Lemons": Quality uncertainty and the market mechanism. *Quarterly Journal of Economics* 84, no. 3:488–500.

Chan, Marc K. 2013. A dynamic model of welfare reform. *Econometrica* 81, no. 3:941–1001.

Chetty, Raj, and Amy Finkelstein. 2013. Social insurance: Connecting theory to data. In *Handbook of public economics*, vol. 5, ed. Martin Feldstein, Alan J. Auerbach, Raj Chetty, and Emmanuel Saez, 111–93. Amsterdam: Elsevier.

Cohen, Alma, and Liran Einav. 2007. Estimating risk preferences from deductible choice. *American Economic Review* 97, no. 3:745–88.

Cooper, Russell. 2016. Money or grit? Determinants of mismatch. NBER Working Paper no. 22734, National Bureau of Economic Research, Cambridge, MA.

De Nardi, Mariacristina, Eric French, and John Bailey Jones. 2010. Why do the elderly save? The role of medical expenses. *Journal of Political Economy* 118, no. 1:39–75.

Dey, Matthew S., and Christopher J. Flinn. 2005. An equilibrium model of health insurance provision and wage determination. *Econometrica* 73, no. 2:571–627.

Dickstein, Michael J., Kate Ho, and Nathaniel Mark. 2024. Market segmentation and competition in health insurance. *Journal of Political Economy* 132, no. 1:96–148.

Eeckhout, Jan, and Philipp Kircher. 2018. Assortative matching with large firms. *Econometrica* 86, no. 1:85–132.

Einav, Liran, Amy Finkelstein, and Neale Mahoney. 2018. Provider incentives and healthcare costs: Evidence from long-term care hospitals. *Econometrica* 86, no. 6:2161–219.

———. 2021. The IO of selection markets. In *Handbook of Industrial Organization*, vol.5, ed. Kate Ho, Ali Hortaçsu, and Alessandro Lizzeri, 389–426. Amsterdam: Elsevier.

Fang, Hanming, and Andrew J. Shephard. 2019. Household labor search, spousal insurance, and health care reform. NBER Working Paper no. 26350, National Bureau of Economic Research, Cambridge, MA.

Finkelstein, Amy, Nathaniel Hendren, and Erzo F. P. Luttmer. 2019. The value of Medicaid: Interpreting results from the Oregon Health Insurance Experiment. *Journal of Political Economy* 127, no. 6:2836–74.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. Subsidizing health insurance for low-income adults: Evidence from Massachusetts. *American Economic Review* 109, no. 4:1530–67.

Frean, Molly, Jonathan Gruber, and Benjamin D. Sommers. 2017. Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act. *Journal of Health Economics* 53:72–86.

French, Eric, John Jones, and Hans-Martin von Gaudecker. 2018. The effect of the Affordable Care Act on the labor supply, savings, and social security of older Americans. Working paper, University College London.

French, Eric, and John Bailey Jones. 2011. The effects of health insurance and self-insurance on retirement behavior. *Econometrica* 79, no. 3:693–732.

French, Erich, and Christopher Taber. 2011. Identification of models of the labor market. In *Handbook of labor economics*, vol. 4, pt. A, ed. Orley C. Ashenfelter and David Card, 537–617. Amsterdam: Elsevier.

Goodman, Lucas. 2017. The effect of the Affordable Care Act Medicaid expansion on migration. *Journal of Policy Analysis and Management* 36, no. 1:211–38.

Gruber, Jonathan. 2008. Covering the uninsured in the United States. *Journal of Economic Literature* 46, no. 3:571–606.

Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski. 2015. Adverse selection and an individual mandate: When theory meets practice. *American Economic Review* 105, no. 3:1030–66.

Handel, Ben, Igal Hendel, and Michael D. Whinston. 2015. Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica* 83, no. 4:1261–313.

Handel, Benjamin R. 2013. Adverse selection and inertia in health insurance markets: When nudging hurts. *American Economic Review* 103, no. 7:2643–82.

Heckman, James J., and Edward J. Vytlacil. 2007. Econometric evaluation of social programs, part I: Causal models, structural models and econometric policy evaluation. In *Handbook of econometrics*, vol. 6, ed. James J. Heckman and Edward E. Leamer, 4779–874. Amsterdam: Elsevier.

Kaestner, Robert, Bowen Garrett, Jiajia Chen, Anuj Gangopadhyaya, and Caitlyn Fleming. 2017. Effects of ACA Medicaid expansions on health insurance coverage and labor supply. *Journal of Policy Analysis and Management* 36, no. 3:608–42.

Keane, Michael P., and Kenneth I. Wolpin. 2007. Exploring the usefulness of a nonrandom holdout sample for model validation: Welfare effects on female behavior. *International Economic Review* 48, no. 4:1351–78.

Kowalski, Amanda E. 2014. The early impact of the Affordable Care Act, state by state. *Brookings Papers on Economic Activity*, Fall, 277–333.

Krusell, Per, Lee E. Ohanian, José-Victor Ríos-Rull, and Giovanni L. Violante. 2000. Capital-skill complementarity and inequality: A macroeconomic analysis. *Econometrica* 68, no. 5:1029–53.

Leung, Pauline, and Alexandre Mas. 2018. Employment effects of the Affordable Care Act Medicaid expansions. *Industrial Relations: A Journal of Economy and Society* 57, no. 2:206–34.
Lewbel, Arthur. 2000. Semiparametric qualitative response model estimation with unknown heteroscedasticity or instrumental variables. *Journal of Econometrics* 97, no. 1:145–77.
Pascale, Joanne. 2016. Modernizing a major federal government survey: A review of the redesign of the Current Population Survey health insurance questions. *Journal of Official Statistics* 32, no. 2:461–86.
Pohl, R. Vincent. 2018. Medicaid and the labor supply of single mothers: Implications for Health Care Reform. *International Economic Review* 59, no. 3:1283–313.
Rust, John, and Christopher Phelan. 1997. How Social Security and Medicare affect retirement behavior in a world of incomplete markets. *Econometrica* 79, no. 3:781–831.
Summers, Lawrence H. 1989. Some simple economics of mandated benefits. *American Economic Review* 79, no. 2:177–83.
Todd, Petra E., and Kenneth I. Wolpin. 2006. Assessing the impact of a school subsidy program in Mexico: Using a social experiment to validate a dynamic behavioral model of child schooling and fertility. *American Economic Review* 96, no. 5:1384–417.