NBER WORKING PAPER SERIES

GROUP HEALTH INSURANCE VERSUS ACA MARKETPLACES: SELECTION, SUBSIDIES, AND WELFARE

Sebastián Fleitas
Caitlyn Fleming
Gautam Gowrisankaran
Anthony T. Lo Sasso

Working Paper 33938
http://www.nber.org/papers/w33938

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2025

We thank Leemore Dafny, Carol Propper, Nicholas Tilipman, and Frank Verboven for helpful comments, and Patricio Hernandez and Arundhati Tillu for excellent research assistance. Caitlyn Fleming coauthored this paper before joining Cornerstone Research. The views expressed herein are solely those of the author who is responsible for the content and do not necessarily represent the views of Cornerstone Research. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

At least one co-author has disclosed additional relationships of potential relevance for this research. Further information is available online at http://www.nber.org/papers/w33938

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Sebastián Fleitas, Caitlyn Fleming, Gautam Gowrisankaran, and Anthony T. Lo Sasso. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Group Health Insurance Versus ACA Marketplaces: Selection, Subsidies, and Welfare
Sebastián Fleitas, Caitlyn Fleming, Gautam Gowrisankaran, and Anthony T. Lo Sasso
NBER Working Paper No. 33938
June 2025
JEL No. H30, I13, L11

ABSTRACT

The Affordable Care Act (ACA) changed incentives for individual and employer-sponsored insurance. Using unique small group market (SGM) data with detailed claims and enrollment information, we analyze the welfare gains across households in the SGM from alternative formulations of ACA health insurance exchanges (HIXs). Accounting for observable heterogeneity increases estimated demand elasticities. Shutting down the SGM would reduce equilibrium adverse selection in HIX-type markets, increasing welfare for households moved to HIXs by $842 annually. The majority of households would obtain higher equilibrium welfare from HIXs, stemming from having more choices. Adding the 2014 ACA penalties only modestly decreases adverse selection.

Sebastián Fleitas
Universidad Catolica, Chile
sebastian.fleitas@uc.cl

Caitlyn Fleming
Cornerstone Research
caitlyn.fleming@cornerstone.com

Gautam Gowrisankaran
Columbia University
Department of Economics
and HEC Montreal and CEPR
and also NBER
gautamg2@gmail.com

Anthony T. Lo Sasso
DePaul University
Economics Department
alosasso@depaul.edu

1 Introduction

The 2010 Affordable Care Act (ACA) ushered in some of the largest changes to U.S. healthcare policy over the past few decades. One of the main mechanisms of the ACA was to expand the sale of individual health insurance. This was done through a set of interrelated policies: subsidies to encourage moderate-income households to purchase insurance, penalties to incentivize the healthy to purchase, and the establishment of health insurance exchanges (HIXs), which are insurance marketplaces that facilitate the sale of individual insurance.

During the debate before its enactment in 2009-10, many observers feared that the ACA would result in an important disruption of the healthcare sector. To assuage these fears, President Obama famously said, "If you like your health insurance, you can keep it," a statement that was later widely criticized. Not coincidentally, the final version of the ACA contained provisions that disadvantaged exchange plans relative to employer-sponsored health insurance. Most importantly, though the ACA subsidized exchange plans for moderate-income individuals, it did not make them tax deductible, while preserving the tax deductibility of employer-sponsored insurance (ESI).[1] It also established relatively modest penalties for not having health insurance that were supposed to grow over time, making it financially advantageous for many healthy individuals to pay the penalties instead of purchasing health insurance.[2]

Together, the ACA's fiscal changes regarding subsidies, tax deductibility, and penalties provided heterogeneous incentives across individuals for exchange plan take-up relative to ESI. These policies continue to be altered over time as different administrations impose their visions for health insurance reform. For instance, a 2019 Trump administration rule change allowed individuals to purchase exchange plans on a pre-tax basis through the use of health savings accounts starting in 2020, essentially making a part of premiums for

[1] The ACA did establish a "Cadillac Tax" on expensive ESI but the implementation of this tax was delayed twice, and then the tax was repealed.

[2] These penalties were ultimately repealed in the 2017 Tax Cut and Jobs Act.

employed individuals tax deductible (Federal Register, 2019). This change, together with many other policy proposals, underscores the relevance of understanding how fiscal and other incentives affect who chooses health insurance coverage.

The purpose of this paper is to understand the impact of policy design on health insurance take-up and welfare. We focus on households who are offered ESI through the small group market. We estimate their preferences for insurance take-up as a function of household characteristics, plan characteristics, and premiums.[3] We then evaluate how counterfactual policy design would affect take-up and welfare from exchange plans relative to ESI. Our counterfactuals focus on changes in tax deductibility, penalties, and subsidies and allow premiums to adjust in equilibrium based on enrollee selection.

The small group market (SGM) provides insurance coverage to employers with 2-50 covered lives. We analyze substitution between ESI from the small group market and exchange plans because the small group market is the sector where the ACA was most expected to disrupt existing health insurance. Specifically, while large employers provide natural risk pools, the small group market had the potential for substantial reclassification risk, which would limit the welfare of its participants (Fleitas et al., 2025). Moreover, while the ACA established penalties for large employers that did not provide sufficient employee health insurance coverage, it did not implement any penalties for non-offer or non-coverage in the small group market.

We use the Medical Expenditure Panel Survey (MEPS) data to understand more about the dynamics of this market. MEPS is a nationally representative survey regarding individual insurance decisions and healthcare expenditures. Figure 1 shows the percentage of employees of small group employers who were offered insurance and who took up insurance given an offer. The figure shows that slightly more than 50 percent of small group employers offered health insurance to their employees, with the rate remaining roughly constant over the 2013-2019 period. By contrast, nearly all employers with 100

[3]For employees who elect household coverage, we estimate their choice decision as a function of both their own and their households' characteristics.

or more workers offer health insurance (MEPS, 2019). Recognizing this, the 2019 rule change that made exchange plans partially tax deductible specifically targeted the small group market (Keith, 2019).

Our estimation primarily uses unique proprietary data from a large insurance company that offers plans on the small group market, over the period 2012–15. We refer to the insurer by a pseudonym, United States Insurance Company (USIC). USIC has provided us with claims, enrollment, and premium data for its small group market plans in 10 states. We estimate choice models for individuals offered plans in the small group market, using data on whether they selected ESI and, for those offered multiple plans, which specific plan they chose.[4] Our choice model specifies that household utility for each option is a function of plan characteristics—such as the deductible, copay, and after-tax premium—as well as household characteristics, including expected healthcare costs, age, gender, and disposable income. We use after-tax prices and disposable income to compute counterfactuals where we alter the tax deductibility of premiums.[5] We calculate the expected healthcare costs for individuals or households—which we call their risk score—using USIC claims data.[6]

Our data are unique in that they link claims, take-up decisions, plan attributes, and plan premiums to individuals, and connect individuals to households and to small group employers. The claims data allow us to compute detailed models of health plan choice that include the individual's health risk as a predictor of take-up. Moreover, by observing health risk, we can evaluate the costs of providing insurance under alternate scenarios, which we use to compute premiums in our equilibrium counterfactuals.

A limitation of our data is that, although we observe the number of individuals who did not take up coverage for each employer, we do not observe any characteristics for households who were offered ESI but declined coverage. We need these characteristics to estimate an accurate model of insurance take-up. We address this limitation with MEPS

[4] Almost all employers on the small group market contract with only one insurer.
[5] We calculate tax rates using the NBER *TAXSIM* program.
[6] We use the ACG risk score methodology developed by Johns Hopkins University.

data. Specifically, we use MEPS data to simulate the characteristics of households who were offered but declined coverage from USIC. We match individuals to the MEPS data using employer size, sector, and the ESI offer/take-up decision. We then simulate family status for every individual in the household and, for each individual, simulate health conditions (which allows us to estimate expected health risk), age, and income. The MEPS data also allow us to simulate income for households who take up ESI from USIC.

Using our combined USIC and MEPS data, we estimate nested logit choice models with multiple dimensions of heterogeneity. We then use our estimated models to simulate the welfare impact of a number of counterfactual policies. First, we simulate moving all households in the USIC small group market to the exchange market, using 2014 exchange plan characteristics. In counterfactual environments, selection into health insurance will differ from the baseline, leading to changes in premiums and further selection. Our counterfactuals impose the ACA medical-loss ratio standards for premiums and calculate the resulting equilibrium changes in premiums.[7] We next consider the impact of altering the tax deductibility, penalties, and subsidies for USIC plans in the SGM market. Finally, we evaluate counterfactual policies that move all USIC SGM enrollees to the exchange market, but also vary the above three fiscal incentives.

Summary of results. We first investigate the importance of accounting for individual household heterogeneity in estimating the demand for health insurance in the SGM. We find that controlling for household health risk and income yields a much higher estimate of the elasticity of health insurance demand. We also find substantial heterogeneity in premium elasticities across households. Households in the lowest quartile of health risk and the lowest quartile of income are the most price elastic on average, though there is substantial heterogeneity even within quartiles of health risk and income.

Turning to counterfactuals, we first consider the effect of moving households from the

[7]The ACA caps price-cost margins for insurance plans, which it denotes medical-loss ratio standards. For each counterfactual, we calculate a plan's premium imposing these caps using expected costs determined from the mean risk scores of all current USIC households choosing the plan. We then use a fixed-point process to solve for a new equilibrium. We also present counterfactuals that leave premiums the same as in the baseline.

small group market to health insurance exchanges. We find that moving SGM consumers to HIX plans reduces adverse selection and hence lowers equilibrium premiums for these plans. Nonetheless, when holding the number of offered HIX plans constant at the SGM level, our equilibrium counterfactuals show that the majority of households are better off in the SGM. The value of the HIX is in offering consumers more choice than the SGM. When we offer households the best mean HIX plan of each of the four actuarial levels—called metal tiers—we find that the majority of households are better off and the mean consumer welfare is $406 higher annually than under the SGM.

We next compare what would happen if households had to pay SGM premiums with post-tax dollars, instead of pre-tax dollars as they did during our sample. We find that removing tax deductibility makes households worse off, on average, by $551 annually. However, mean household surplus drops by less than the increase in government revenue. It also reduces SGM insurance take-up, but only by two percentage points. Adding subsidies similar to the HIX or penalties similar to the those in 2014 would have very little impact on insurance take-up.

Finally, we consider counterfactuals that both move households to HIXs and alter the fiscal environment of the HIX. Adding 2014-level penalties to encourage health insurance purchase reduces non-equilibrium welfare by a mean of $124. Penalties cause 1.2 percent more people to buy health insurance, but only modestly reduce adverse selection, causing annual equilibrium premiums conditional on purchase to drop by $3. Making HIX plans tax deductible increases take-up by 1.8 percentage points, but increases consumer surplus less than the cost to the government.

Relation to literature. This paper builds on three literatures. First, we add to a literature that considers individuals' choices of health plans, often with the goal of understanding health insurance policy reform (Dafny et al., 2010, 2013; Curto et al., 2021; Kim and Lyons, 2019). A number of papers in this literature specifically consider the design of the ACA and its impact on selection (Ericson and Starc, 2015; Saltzman, 2019; Tebaldi, 2025; Hackmann et al., 2015; Frean et al., 2017; Einav et al., 2024; Panhans, 2019; Finkelstein

et al., 2019). Our innovation is in incorporating individual and plan heterogeneity and in considering the impact of plan design. While this literature has generally not focused on the interaction between the small group and individual markets, closest to our paper is Dickstein et al. (2024), which also considers this interaction, focusing on HIX plan design and moral hazard. Our paper differs from this work in our data—and hence focus—on small group enrollees, which allows us to understand the impact of small group plan characteristics on enrollee choice. Second, we build on a literature that examines employers' decisions of health plan offerings (Gruber and Poterba, 1996; Gruber and Lettau, 2004; Lurie et al., 2021). Finally, we build on a literature that considers the design of taxes in differentiated products markets (Dubois et al., 2020; Allcott et al., 2019; Griffith et al., 2018, 2019).

2 Model and Estimation

We consider individuals who worked for an employer that offered health insurance and that had 50 or fewer employees. Before 2014, such an individual could purchase insurance for herself and her household through the small group market (SGM), as well as through options other than employer-sponsored insurance (ESI). Starting in 2014, the ACA created health insurance exchanges (HIXs) and these individuals could also purchase insurance through this channel. In 2014, over 70% of individuals working in small employers who were offered ESI chose this option, in plans that were largely risk related (Fleitas et al., 2022). In contrast, a very small fraction (2%) of these individuals took up insurance in the HIX.[8]

The lack of purchase from the individual market was due, in large part, to two factors. First, ESI was paid with pre-tax dollars while individual insurance was not. Second, employers often subsidize at least some part of ESI (Currie and Madrian, 1999; Olson, 2002; DeVaro and Maxwell, 2014; Lubotsky and Olson, 2015). In contrast, individuals

[8]Calculations from authors' tabulation of the Medical Expenditure Panel Survey (MEPS).

earning 400% or less of the federal poverty level qualified for premium subsidies when purchasing a plan from the HIX, providing a reason to purchase insurance in this market for relatively low-income employed individuals.

The goal of our study is to evaluate the distribution of consumer welfare, government spending, and insurance take-up if we expose households who were in the small group market to the exchange market, under its contemporaneous design, and under counterfactual designs that vary in tax deductibility, subsidies, and penalties. Accordingly, we develop a plan choice framework that explicitly models these three features. We then use the plan choice framework to simulate premiums, take-up, and welfare under alternative policy environments. We now detail these components.

Estimation sample. Our sample consists of employees who work for a small employer that offers health insurance from USIC. In our model, each employee decides whether to buy insurance and if so, which product to buy. We assume that if an employee decides to buy insurance, she does it for herself and her household and that she purchases the same option for all members of her household. The insurance purchase decision is a function of household income, health status, and other household characteristics, as well as plan characteristics. Modeling heterogeneity in choice decisions across households is important in our context, given that the utility from health insurance and fiscal incentives both vary considerably across household characteristics.

Consider a marketplace with a set of households, indexed $i = 1, \ldots, I$. Each household has access to a set of insurance plans, indexed $j = 1, \ldots, J_i$. Let $j = 0$ denote the "outside option" of not purchasing insurance. Denote the price of each plan for covering all members of the individual's household—gross of subsidies or tax deductibility—as p_{ij}. We normalize $p_{i0} = 0$.

Income and income tax. We define different measures of income and income tax. First, let Y_i^G be gross income of the household, which we treat as exogenous. Since health insurance may be paid with pre-tax dollars, taxable income depends on the insurance

choice: $Y_{ij}^T = Y_i^G - d \times p_{ij}$, for $j = 0, \ldots, J_i$, where d is an indicator for whether insurance is paid with pre-tax dollars. We specify that the income tax payable is a function of taxable income, $\tau\left(Y^T\right)$. Let s be an indicator for whether there are ACA subsidies available and $\phi\left(Y_i^G, p_{ij}\right)$ be the amount of the subsidy. Also, let f be an indicator for whether there is a penalty for not buying health insurance and $\pi(Y_i^G)$ be the 2014 ACA penalty from non-purchase of insurance. Finally, we define disposable income, Y_{ij}^D to be income net of taxes and of the purchase and penalty cost of health insurance. Formally,

$$Y_{ij}^D = Y_i^G - \tau\left(Y_i^G - d \times p_{ij}\right) - p_{ij} + s \times \phi(Y_i^G, p_{ij}) - \mathbb{1}\{j = 0\} \times f \times \pi(Y_i^G). \tag{1}$$

In (1), the first term is gross income, the second term is taxes, the third term is premiums, the fourth term is subsidies, and the final term is the penalty for non-purchase of insurance.

Equation (1) allows us to exploit the different regimes of tax deductibility, subsidies, and penalties present in the SGM and HIX. We discuss some particular cases. The first case is when $d = 1$, $f = 0$, and $s = 0$. Under these assumptions, (1) reduces to $Y_{ij}^D = Y_i^G - \tau(Y_{ij}^G - p_{ij}) - p_{ij}$. This is the situation of the SGM prior to the ACA regulations, where premiums (if health insurance is purchased) are paid with pre-tax dollars, there are no penalties for non-purchase, and no subsidies available to consumers. A second case is when $d = 0$, $f = 0$, and $s = 1$. In this case, $Y_{ij}^D = Y_i^G - \tau(Y_i^G) - p_{ij} + \phi(Y_i^G, p_{ij})$. This is the situation of the HIX in 2019, where the premiums are paid with post-tax dollars, there are no penalties for non-purchase, and subsidies are available. Finally, Equation (1) allows us to investigate the impact of counterfactual combinations of taxes and subsidies. For example, it allows for the possibility of having the set of plans from the SGM, premiums paid with pre-tax dollars, and subsidies as in the HIX, by setting $d = s = 1$ and calculating outcomes with the set of plans that is present in the small group market.

Choice model. We model health insurance demand with a discrete choice model that includes multiple dimensions of heterogeneity. First, as above, households are heteroge-

neous in their disposable income. Second, households differ in their risk scores, which indicate their expected healthcare costs. A household's risk score, r_i, is the mean of the individual risk scores for all (enrolled) household members. An individual's risk score, in turn, is a function of her age, gender, and previous health conditions and healthcare utilization. We predict individual claims risk based on detailed health records and take the mean across household members to obtain r_i. Third, plans are heterogeneous, with different coinsurance levels, copays, deductibles, out-of-pocket maxima, and premiums.

Households make a choice between purchasing one of the J_i available insurance products or purchasing the outside option of no USIC insurance, which we denote $j = 0$.[9] We specify a nested logit demand model in the tradition of McFadden (1974) and Berry (1994). We model two groups, with group 0 containing the outside option and group 1 containing the J_i insurance products. Let $g_i(j)$ be the function that assigns option j to group g for household i and let $\mathcal{G}_{ij}$ denote the subset of household i's options that are in the same group as option j.

The utility of each option, $j = 0, \ldots, J_i$, can be written as:

$$u_{ij} = \delta_{ij} + \nu_{ig_i(j)} + (1 - \sigma)\varepsilon_{ij},$$

where δ_{ij} is the mean utility of plan j for household i (which we discuss further below), $0 \leq \sigma < 1$ is the nesting parameter, ν_{ig} is common to all products in group g and is distributed $C(1-\sigma)$; and ε_{ij} is distributed type 1 extreme value $i.i.d.$ across products. Cardell (1997) shows that, given our distributional assumptions on ν_{ig} and ε_{ij}, $\nu_{ig} + (1 - \sigma_j)\varepsilon_{ij}$ is also an extreme value random variable. Note that as σ approaches one, the within-group correlation of utility levels goes to one. For $\sigma = 0$, the within group correlation is zero and we return to the standard logit model.

We estimate a nested logit model rather than a richer random coefficients model because our data include many dimensions of observable heterogeneity, which enter into δ_{ij}.

[9]This outside option includes forgoing health insurance or, for example, obtaining coverage from a spouse. We hold its utility constant in counterfactuals.

In our main specification, we define δ_{ij} to include disposable income of the household, in-network out-of-pocket costs, coinsurance rate, and deductible, and out-of-network deductible and out-of-pocket maximum, all interacted with the risk score of the household (r_i) and household size (h_i), as a way to incorporate the different effects of these plan characteristics on different types of households. For $j = 1, \dots J_i$, we write:

$$\begin{aligned}
\delta_{ij} &= \alpha_1 \times \log(Y_{ij}^D) + \alpha_2 \times \log(Y_{ij}^D) \times r_i + \alpha_3 \times \log(Y_{ij}^D) \times h_i \\
&+ \beta_1 \times oopm_j + \beta_2 \times coins_j + \beta_3 \times ded_j + \beta_4 \times copay_j + \beta_5 \times OONded_j + \beta_6 \times OONoopm_j \\
&+ \theta_1 \times r_i \times oopm_j + \theta_2 \times r_i \times coins_j + \theta_3 \times r_i \times ded_j \\
&+ \theta_4 \times r_i \times copay_j + \theta_5 \times r_i \times OONded_j + \theta_6 \times r_i \times OONoopm_j \\
&+ \omega_1 \times POS_j + \omega_2 \times PPO_j + \omega_3 \times EPO_j + \omega_4 \times HMO_j + \omega_5 \times h_i + \omega_{t(i)},
\end{aligned} \tag{2}$$

where $oopm_j$ is the in-network out-of-pocket maximum, $coins_j$ is the in-network coinsurance, ded_j is the in-network deductible, $copay_j$ is the in-network copay, $OONded_j$ is the out-of-network deductible, $OONoopm_j$ is the out-of-network maximum, HMO, POS, PPO, and EPO are plan types (with the category other plan types excluded), and the $\omega_{t(i)}$ are year fixed effects. The α, β, θ, and ω terms are parameters to be estimated. Corresponding to Equation (2), for $j = 0$, we can write:

$$\delta_{i0} = \alpha_1 \times \log(Y_{i0}^D) + \alpha_2 \times \log(Y_{i0}^D) \times r_i + \alpha_3 \times \log(Y_{i0}^D) \times h_i.$$

Given the distributional assumption for the ε terms, we can define a closed form for the probability of plan choice. First, for group $g \in \{0, 1\}$, define the inclusive value of the group as:

$$\overline{\delta}_{ig} = (1 - \sigma) \log \left(\sum_{j \in \mathcal{J}_g} \exp \left(\frac{\delta_{ij}}{1 - \sigma} \right) \right). \tag{3}$$

Then the choice probability that household i chooses plan j is:

$$\Pr_{ij} = \frac{\exp\left(\overline{\delta}_{ig}\right)\exp\left(\frac{\delta_{ij}}{1-\sigma}\right)}{\left(\sum_{g'\in\{0,1\}}\overline{\delta}_{ig'}\right)\left(\sum_{k\in\mathcal{J}_g}\exp\left(\frac{\delta_{ik}}{1-\sigma}\right)\right)}.$$

Estimation. If we observed δ_{ij}, estimation via maximum likelihood would be standard. In that case, the log likelihood for household i could be defined as:

$$\log\mathcal{L}_i(\alpha,\beta,\theta,\omega,\sigma) = \log\left(\sum_{j=0}^{J_i}\mathbb{1}\{D_i=j\}Pr_{ij}\right),$$

where D_i indexes the chosen option.

However, our main data have one central limitation: we do not perfectly observe household characteristics. While we know the number of employees who are eligible for insurance, we do not observe the health risk (r_i), household size (h_i), or income (Y_i^G) for households who choose $j = 0$. For households who choose $j \geq 1$, we observe claims data, age, and gender—which allows us to compute health risk—and we assume that the household composition is given by the household members that take up insurance. Even for these households, we do not observe income.

To address this issue, we integrate over these characteristics using an auxiliary source of data, the MEPS data. With the auxiliary data, we simulate the distribution of household characteristics conditional on employer characteristics and on the take-up decision. We then maximize a simulated likelihood function using these distributions.

Formally, define c_i to be the observable employer characteristics of household i that we can match to the auxiliary data. These include the industrial sector, the number of employees, and the fact that the employer offered health insurance to its employees. Let $F^A(Y_i^G, r_i, h_i|c_i, \mathbb{1}\{j=0\})$ be the distribution function of Y_i^G, h_i, and r_i in the auxiliary data conditional on c_i and not taking up ESI (when offered). Similarly, let $F^B(Y_i^G|c_i, \mathbb{1}\{j>0\})$ be the distribution function of Y_i^G conditional on c_i and choosing ESI. Then, we can

express the log-likelihood by integrating over r_i, h_i, and Y_i^G:

$$\log \mathcal{L}_i(\alpha, \beta, \theta, \omega, \sigma) = \mathbb{1}\{D_i = 0\} \log \left(\int Pr_{i0|Y_i^G, r_i, h_i} dF^A(Y_i^G, r_i, h_i | c_i, \mathbb{1}\{j = 0\}) \right) + \sum_{j=1}^{J_i} \mathbb{1}\{D_i = j\} \log \left(\int Pr_{ij|Y_i^G} dF^B(Y_i^G | c_i, \mathbb{1}\{j > 0\}) \right). \quad (4)$$

Equation (4) displays the log likelihood for each employee who is eligible for insurance, first for households that do not take up insurance and then for ones that do. If the employee chooses $D_i = 0$, we simulate her r_i, h_i, and Y_i^G from the matched auxiliary data. If she chooses $D_i \neq 0$, we simulate her income only from the matched auxiliary data. We estimate our structural parameters with a simulated maximum likelihood estimation using equation (4).[10]

One final complication is that the auxiliary data do not directly include the risk score, r_i. However, they include information on age, gender, health conditions, and healthcare utilization that can be used to predict r_i, similar to what we do for households in the USIC data. Thus, we estimate r_i for households who chose the outside option in three steps. First, we perform a flexible regression with the USIC data of r_i at the individual level on conditions reported in both datasets, age, and gender. Second, we use the regression coefficients to predict the MEPS household's score. Third, for households that chose the outside option, we match them to an individual in the MEPS data (as above) and then to that individual's household risk score.

Counterfactuals. We use these consumer preferences to analyze counterfactuals in which we change the policy and health plan choice environment. In particular, we analyze how much consumer welfare and insurance take-up would change if we moved all SGM consumers to HIXs, under different fiscal incentives of subsidies, taxes, and penalties. These counterfactuals allow us to analyze the role of plan characteristics offered in

[10]Our approach is conceptually similar to Petrin (2002) who uses micro moments with an auxiliary data source to estimate random coefficients in a GMM setting, but differs in that our main data are at the individual level and that we use a maximum likelihood estimator.

the different markets and of tax, penalty, and subsidy incentives. Since our approach allows for heterogeneity, we compute welfare gains at the household level and report statistics of these gains. Our approach can be used to evaluate the heterogeneity in welfare changes based on income and health status.

In all counterfactuals, we express per capita expected utility for household i in the SGM, EU_i^{SGM}, using the standard formula:

$$EU_i^{SGM} = \mathbb{E}[\max_{j \in \mathcal{J}_i} u_{ij}] = \log \left(\sum_{g \in \{0,1\}} \exp\left(\overline{\delta}_{ig}\right) \right) + \gamma, \tag{5}$$

where $\overline{\delta}_{ig}$ is calculated using equation (3) and γ is Euler's constant.

To calculate the expected utility from the HIX market, EU_i^{HIX}, we recompute δ_{ij} with the characteristics of HIX plans and recompute utility with these plans using equation (5). We then define a measure of the change in consumer surplus in dollars from the change in available plans as:

$$\Delta CS_i = \frac{EU_i^{HIX} - EU_i^{SGM}}{\frac{\partial EU_i^{SGM}}{\partial Y_i^G}}, \tag{6}$$

where ΔCS_i in equation (6) is the change in utility divided by the marginal utility of income, which is $\frac{\partial EU_i^{SGM}}{\partial Y_i^G} = \sum_{j=0}^{J_i} Pr_{ij} \times \frac{1}{Y_{ij}^D} \times (\alpha_1 + \alpha_2 r_i + \alpha_3 h_i) \times \left(1 - \frac{\partial \tau}{\partial Y_i^G}\right)$. As Petrin (2002) noted, this is an approximation because the marginal utility of income is not constant but varies with disposable income, which varies across counterfactuals.

In addition to welfare, we report insurance take-up and government expenditures under counterfactual policies. We calculate take-up using the choice probabilities, Pr_{ij}, computed from our model. We compute government revenues from household i as penalties plus taxes minus subsidies:

$$GR_i = Pr_{i0} \times f \times \pi(Y_i^G) + \sum_{j=0}^{J_i} Pr_{ij} \left[\tau \left(Y_i^G - d \times p_{ij} \right) - s \times \phi(Y_i^G, p_{ij}) \right].$$

We compute both non-equilibrium and equilibrium counterfactuals. In the non-equilibrium counterfactuals, we take the choice set of plans, the characteristics, and the premiums from the HIX as given, and we compute the welfare of households conditional on those characteristics. In our equilibrium counterfactuals, we simulate a market similar to the HIX in which individuals can buy insurance for themselves and their households. In this market, we take the offerings and the characteristics of the plans as given but we allow the premiums for HIX plans to vary based on the expected health costs to the insurer.[11] We assume that the markup of plans follows the Medical Loss Ratio rule that was implemented under the ACA. This rule forced health plans in the individual market to spend at least 80% of their premiums on claims costs, or rebate these premiums (Cicala et al., 2019). Given this, our equilibrium counterfactuals adopt the simplifying assumption that plan premiums would be exactly at the MLR margin.[12]

In order to compute premiums for our equilibrium counterfactuals we solve for a fixed point where the premium for each plan is equal to the inverse of the Medical Loss Ratio times the expected mean medical and pharmaceutical costs for households in our sample that choose the plan given the premiums in the market. Thus, in equilibrium, we are replacing the current potential HIX customers with households who were offered USIC SGM plans. Formally,

$$p_j = \frac{1}{MLR} \times \frac{\sum_i H_i \times Pr_{ij}(p_1, \ldots, p_{J_i})}{\sum_i Pr_{ij}(p_1, \ldots, p_{J_i})}, \tag{7}$$

where H_i are the total insured costs of the household given the household's risk score, household composition, and plan's metal tier, i.e., Bronze, Silver, Gold, Platinum. Equation (7) explicitly shows that Pr_{ij} is a function of all plan premiums $(p_1, ..., p_{J_i})$ offered in the choice set that each household i faces.

From equation (7), to solve for equilibrium counterfactuals, we need to predict health

[11] As we note in Appendix A, we keep the choice set for each consumer the same as in the corresponding non-equilibrium counterfactual.

[12] An alternative assumption, employed by Dickstein et al. (2024), is to endogenize the choice of HIX plans offered by assuming that the market reflects an Azevedo and Gottlieb (2017) competitive equilibrium.

costs borne by the insurance company for each household across different health plans. To do this, we use data from households who opt into USIC SGM plans to estimate *total* household healthcare costs (medical and pharmaceutical) in a hedonic model with flexible interactions of plan characteristics and the household characteristics r_i and h_i. We then use these total expenditures to compute counterfactual HIX premiums for a given plan as the mean of total expenditures for households that choose that plan multiplied by the coverage ratio based on the plan's metal tier. Thus, our equilibrium assumes that moral hazard is similar across markets, plans, and counterfactuals.[13]

3 Data

Our primary data are from a large health insurance provider, which we refer to as the United States Insurance Company, or USIC. USIC provides health insurance coverage to employers in the small group market. The data cover the years of 2013 through 2015 for ten different states: Arkansas, Delaware, Illinois, Missouri, Oklahoma, Pennsylvania, Tennessee, Texas, Wisconsin, and Wyoming. We observe data at the employer-year and enrollee-year levels as well as linked enrollee claims data.

The employer-year level data include the number of employees eligible for health insurance coverage, the number of employees enrolled, the number of plans offered by the employer, characteristics of each plan (including coinsurance, physician copay, in-network deductible, in-network out-of-pocket maximum, out-of-network deductible, and out-of-network out-of-pocket maximum), and the total premium paid to the insurer for each plan. Table 1 shows the characteristics of employers in our data. Employers in our sample offer a mean of 1.23 health insurance plans in which 67% of employees enroll. The average employer has 17.51 eligible employees. There are 70,724 unique employer-year observations in our data.

The enrollee-year data include age, gender, and the choice of health plan from the

[13] Appendix A provides more details on the computation of equilibrium counterfactuals.

employer's offerings. There are about 265,000 employees covered each year. We calculate a per-enrollee premium by dividing the total premium paid to USIC in a year for a plan by the number of enrollees (employees and dependents) at that employer and plan during that year. We calculate the premium paid by a household by multiplying the enrollee premium amount by the number of enrollees in each household.

Our claims data include the medical and pharmaceutical records for each enrollee, and provide diagnoses, procedures, and dates of service. We use our claims data to calculate the health risk, r_i: we use the "ACG risk score" as our measure of individual health risk, and then aggregate to the household level by taking the mean. Developed at the Johns Hopkins Bloomberg School of Public Health, the ACG system software uses claims-level data to calculate this score for each enrollee in a year. The scores are based on age, gender, medical, and pharmaceutical expenditures, diagnoses, and prescription drug consumption in the previous year. The score indicates the predicted relative healthcare cost for an individual over the year, and has a mean of one in a reference group chosen by the ACG software.

As noted in Section 2, our data have two primary limitations with respect to our estimation strategy. First, we do not have characteristics (i.e., age, gender, ACG scores, income) for individuals who are offered USIC insurance through their employers but decline this coverage. Second, we do not have income for individuals enrolled in a USIC plan through their employer. Thus, we supplement the USIC data with the Medical Expenditure Panel Survey (MEPS) Household Component. The MEPS is a nationally representative sample of households and contains information on demographic characteristics, health conditions, health insurance coverage, income, and employment. Our MEPS sample contains individuals who worked for a small group employer and were offered health insurance from that employer.

We first match individuals in our USIC data who declined health insurance to the MEPS data with an iterative matching process on a subset of this MEPS sample that includes individuals who declined offered employer-sponsored insurance. Due to the

limited number of individuals who satisfy the criteria of (a) being employed at small employers, (b) being offered health insurance by their employer, and (c) declining that offer, we sort MEPS individuals into "buckets" based on their employer size and industry. For each employer size-industry cell in our primary dataset, we create a bucket of at least twenty individuals in the MEPS. We then randomly select twenty MEPS individuals from the bucket so that each individual who selects the outside option in our primary dataset is represented by twenty possible draws of individuals from the MEPS. For each MEPS individual employee draw, we use information from all the individuals in the employee's household.

The MEPS data do not contain ACG risk scores, but include a number of diagnoses, health status measures, and demographics. We thus construct a predicted risk score for each MEPS observation by regressing ACG scores for our USIC data on MEPS measures and then using the regression prediction for each MEPS observation. Our risk score regression uses a flexible specification that interacts diagnoses reported by MEPS, industry fixed effects, employer size, age, and gender. We include the following diagnoses in our regression: attention deficit disorder and attention deficit hyperactivity disorder, angina, arthritis, asthma, cancer, coronary heart disease, high cholesterol, diabetes, emphysema, high blood pressure, heart attack, and stroke. We simulate the individual risk score as the predicted score plus a random shock drawn from a normal distribution with the estimated mean squared error, and take the mean across individuals in the household to get a draw of the household risk score. We perform this process for each of the twenty simulation draws corresponding to each individual in our sample who chose the outside option $j = 0$.

We also match individuals in our USIC data who took up health insurance to MEPS, in this case to construct income. We assume that these USIC individuals are married if and only if their spouse is covered by USIC insurance and—for use in our income tax calculation described below—that they have the number of children that are covered by their USIC insurance. For each of these individuals, we condition on employer size,

industry, and marital status, and create draws of income from the five observations with the closest age from the MEPS subsample of individuals who accept their ESI.

Our mean utility specification uses disposable income, Y_{ij}^D, as described in equation (1). We use *TAXSIM* to compute tax rates, τ_{ij}, under both the scenario in which premiums are paid pre-tax and in counterfactuals where premiums are paid with after-tax dollars (Feenberg and Coutts, 1993). We assume that individuals file taxes jointly if they are married and that each child is claimed as a dependent. We let Y_i^G be the total wage income for the individual and their spouse and let total premiums, p_{ij}, be proportional to the number of household members enrolled in the plan. Table 2 compares characteristics for households who select into the USIC data and those who choose the outside option. Employees who choose an inside option are on average older, more likely to be male, healthier, less likely to be married or have children, and have higher incomes than employees who choose the outside option.

Beginning January 1, 2014, the Affordable Care Act (ACA) provided subsidies to households in the form of premium tax credits. These subsidies are available to households who purchase coverage in the HIX and have income between 100% and 400% of the federal poverty level (FPL), which is determined by household characteristics, including gross income, state, household size, marital status, and number of children.

We compute subsidies, ϕ_{ij}, assuming the same household structure we used for the *TAXSIM* calculations. Based on household income and state, we first determine a household's income relative to the FPL. This then allows us to evaluate subsidy amounts since households contribution to HIX insurance is capped based on their FPL. Specifically, households at 100–133% of FPL have a cap of 2.0% of income. This rises in income up to 9.5% for individuals at 300–400% FPL.[14] The maximum subsidy an individual can receive is equal to the premium of the second lowest cost Silver plan available in the individual's state HIX less the annual individual contribution amount.[15]

[14]The individual contribution schedule is described in Table 2 the IRS Instructions for Premium Tax Credit 8962 (https://www.irs.gov/pub/irs-prior/i8962–2014.pdf).

[15]The ACA also specified cost-sharing reductions whereby insurers offering HIX plans must provide alterations of each standard Silver plan that increase their actuarial values for lower income enrollees.

For use in our counterfactuals, we also compute the 2014 penalties, $\pi(Y_i^G)$, that households would face if they did not fulfill the ACA individual mandate to hold health insurance. In 2014, individuals were fined the higher of $95 or 1% of household income for not having insurance. The penalty increased to the higher of $325 or 2% of household income in 2015.

Many of our counterfactuals move households from the plan choices available to them in the SGM to the HIX. We use the Health Insurance Exchange Public Use Files made available by the Center for Consumer Information and Insurance Oversight (CCIIO) at the Centers for Medicare and Medicaid Services (CMS) to construct HIX plans available in our ten sample states. To match our choice model in equation (2), we construct HIX plan characteristics on coinsurance, physician copay, in-network deductible, in-network out-of-pocket maximum, out-of-network deductible, out-of-network out-of-pocket maximum, and individual annual premium.

Table 3 compares the plan characteristics for plans offered by employers in the SGM to those on the HIX. On average, the SGM plans have lower mean annual premiums and in-network out-of-pocket maximums ($3,467 for SGM and $3,656 for HIX). However, HIX plans have lower deductibles, copays and out-of-network out-of-pocket maximums, though higher in-network out-of-pocket maximums. HIX plans have higher average levels of coinsurance, which we report as percent paid by the insurer.

4 Results

This section presents the results of our model of health insurance demand. We display results from multiple specifications, which differ in how much heterogeneity we allow. This section also compares our price elasticities with estimates from the literature.

Specifically, while the standard Silver plan has an actuarial value of 70%, this increases to 73% for household income between 200% and 250% of the Federal Poverty Line (FPL); to 87% for income between 150% and 200% of the FPL; and to 94% for income between 100% and 150% of the FPL. See `https://www.healthinsurance.org/obamacare/the-acas-cost-sharing-subsidies/`. We adjust the characteristics of Silver plans by income level to reflect these altered characteristics, as reported in our HIX data.

Table 4 summarizes the estimates of our demand model. In all specifications, we control for different values of the outside option for different years using year fixed effects. We also control for plan type—PPO, HMO, POS, EPO—allowing for household mean utility to vary by type. Our estimates are based on a sample of about 33 million simulated individual household-choice observations. Each individual household in our estimation sample includes five simulated draws for households in our data who chose a USIC plan and twenty draws for households who chose the outside option.

Column I presents a specification without heterogeneity in observables across households. We include this specification as a starting point to then evaluate the importance of individual household heterogeneity. Because we do not allow for any observable heterogeneity, we estimate this specification using only USIC data. Our main regressor of interest is the log premium that households pay for insurance.

This specification, and all our others, control for *unobserved* heterogeneity by using the nested logit model. The nested logit model is equivalent to a random coefficient model where there is heterogeneity on an indicator for inside option plans (Berry, 1994). In our case, allowing for correlation among the utility shocks of the inside option plans is a way to account for heterogeneity in consumers' preferences for inside option plans. The nesting parameter, σ, is identified by the extent of substitution away from the outside option as the number of inside options increase.

For this specification, we find a coefficient of -0.02 on log price. We estimate a value of 0.16 for $1-\sigma$, which implies relatively high correlations in preferences for inside option plans. By not interacting the premium with any household characteristic, this model restricts the price coefficient to be the same across households. Using these estimates in combination with the sample data, we find an estimated weighted mean own-price elasticity of -0.06. We weight the elasticity by households' predicted choices to allow for plans chosen more often to have more influence in the elasticity calculation.

Column II introduces observed individual heterogeneity by adding household income to our specification. In order to do this, we use both our USIC data and the auxiliary

dataset from MEPS. We compute disposable income (net of health insurance premiums) as discussed in Section 3, and include the log of the disposable income in the specification of the mean utility to allow for income effects. We find a coefficient of 0.78 on log disposable income; the positive coefficient here is equivalent to a negative coefficient on price. The weighted average own-price elasticity is -0.11, larger than in the previous specification. Our estimate of $1-\sigma$ is 0.29. This again implies a high correlation in the unobservables between the inside option plans in this specification, which may reflect the fact that this specification still has limited observable household heterogeneity.

Column III adds more observable heterogeneity by including the household risk score, r_i. Since the household risk score is a function of age, gender, and past consumption of health services for individuals in the household, this specification effectively introduces observable heterogeneity in many demographic dimensions. The log disposable income coefficient is 14.17 but we also interact the log disposable income with household size and risk score. The estimates show that higher risk households and larger households are relatively less elastic. The weighted average own price elasticity is -0.38. The estimate of $1-\sigma$ is 0.72, which implies a positive, but much smaller, correlation in the unobservables between the plans in the inside option than for Columns I and II.

Finally, Column IV introduces even further heterogeneity by using the specification in Column III but additionally interacting household risk score, r_i, with all plan characteristics. The estimates are very similar to those in Column III, with the log disposable income coefficient being 13.22 and the estimates suggesting a similar pattern of the change in elasticity by household risk and size. The weighted average own price elasticity is -0.36 and the estimate of $1-\sigma$ is 0.70, very similar to Column III. We view Column IV—which has the maximum level of heterogeneity—as our preferred specification.

The Table 4 estimates highlight three important facts. First, mean price elasticities are significantly larger when we incorporate individual household characteristics using the MEPS data and simulation methods. Specifications without individual household characteristics cannot incorporate income effects and tax effects. Omission of income makes

households appear much more inelastic to premium changes than in specifications where we include these variables. The lower price elasticity may result from measurement error in the price coefficient. It may also be due to utility being concave in income, implying that the average household has a higher income—and hence a lower price elasticity—than the weighted sum of price elasticities across the population.

Second, the estimates show that including heterogeneity via random coefficients may partially reflect correlations that can be better modeled with individual household characteristics. Specifically, when we incorporate observable heterogeneity by using these characteristics, the nesting parameter σ becomes smaller and the role of unobservable heterogeneity is smaller.

Third, the specifications with a high level of observed individual household heterogeneity—Columns III and IV—yield similar mean price coefficients and weighted average own price elasticities. Still, we find differences between these specifications in the type and extent of heterogeneity across households. This heterogeneity is important, because it may significantly affect the impact of the counterfactual fiscal incentives that we consider.

The elasticity estimates from our preferred specification are in the range of previous estimates in the literature for other employer-sponsored insurance settings. For example, Dafny et al. (2013) use data from a large employer and estimate an average elasticity of -0.28, ranging between -0.08 and -0.45. Our estimates are also similar to those of two other studies which identify plan elasticities using the plan choices of university employees following implementation of a fixed-contribution arrangement, under which employees bear the incremental cost of their plan choices. Royalty and Solomon (1999) report estimates between -0.28 and -0.62 for Stanford employees and Cutler and Reber (1998) provide estimates for Harvard employees of -0.30 to -0.60. Bundorf et al. (2012) study plan choices of employees in eleven small and midsize employers during 2004-05 and find elasticities of approximately -0.57 (based on reported semi-elasticities).

Figure 4 presents the smoothed histogram of the estimates of household own premium elasticities by plan for Column IV. The figure shows a large heterogeneity around the

weighted mean elasticity of -0.36. Figure 5a presents the smoothed distribution of the average elasticities by household by different quartiles of household risk score r_i. The figure presents two interesting facts of the setting. First, households with higher risk are in general more inelastic, since the distributions move to the right for the quartiles with higher risk. Second, these distributions show that there is a significant amount of heterogeneity even within the same risk quartile. In this sense, using the quartile of risk as a proxy for preferences would considerably reduce the estimated heterogeneity.

Figure 5b presents the distribution of the average elasticities by household by different quartiles of gross income (Y_i^G). Again, the distributions by quartile provide evidence of differences across quartiles. Higher income households are more inelastic, although there is also much heterogeneity within income quartiles. These different patterns of substitution are dependent on rich household characteristics and play a fundamental role in our counterfactuals, where households face different fiscal incentives based on their income and make different choices based on their income and risk score.

Not included in Table 4, for our preferred specification, Column IV, we find that potential enrollees prefer a smaller out-of-pocket maximum: a one percent increase in the out-of-pocket maximum leads to a drop in the average probability of choosing a plan of 0.8 percentage points. Additionally, households dislike higher deductibles and also value a lower out-of-network deductible, which is consistent with them considering the possibility of having to visit an out-of-network physician or hospital. Finally, we find that HMOs and PPOs are popular plan types across households, compared to the other options available.

5 Counterfactuals

This section presents several counterfactuals that evaluate the impact of alternative policies on the distribution of welfare for our sample, which includes households with small group market options. We consider three potential sets of counterfactuals: moving house-

holds from the small group market into HIX plans, changing the tax-subsidy-penalty environment to which health insurance is subject, and the interaction of both. Our counterfactuals all use the estimated coefficients from our preferred specification of the SGM demand model, Column IV of Table 4.

To evaluate moving households into the HIX, we use plan offerings during 2014 for all the HIX markets in our USIC SGM sample. For each counterfactual, we present descriptive statistics of the welfare changes and the changes in net government revenues. We evaluate three potential fiscal changes, and the combination of multiple changes. First, we consider changing the tax deductibility of premiums. Second, we evaluate the impact of imposing penalties for non-enrollment. Third, we investigate the impact of HIX subsidies. As described in Section 3, we compute each household's counterfactual taxes, penalties, and subsidies using our dataset from USIC supplemented with MEPS data.

Our counterfactuals address two complications of our setting. First, households generally have many more plans available in the HIX than in the SGM. Thus, when we move households from the SGM to the HIX, this expands their choice sets considerably. Since the consumer utility function includes a logit unobservable, expanding the offered choice set then overestimates the incremental value to consumers of this change (Ackerberg and Rysman, 2005). In order to avoid this problem, we follow a strategy similar to Dafny et al. (2013) and investigate three counterfactual scenarios that together bound the welfare gains. In the first scenario, which we call *Swap to mean best plan*, we keep the same number of plans that households have in the choice set in the SGM, but we choose the plans that provide the highest average utility to all households in the market. Under this scenario, households are offered the same number of plans in the data as in the counterfactual. The second scenario, called *Choice of best mean plan in each metal tier*, increases the number of plans in the choice set to four plus the outside option, where we let the plans with the highest average utility in each metal tier of the HIX be the inside option plans. This scenario is a compromise between giving households more options than under the SGM while also keeping a limited set of plans whose characteristics are constant across

households in a market. Finally, in the third scenario, called *Choice of all HIX plans*, we allow consumers to choose among all the plans offered in their geographical market in the HIX. These counterfactuals may make HIXs look valuable simply because they have a large number of similar options.

A second complication is that plan premiums may adjust in equilibrium based on households moving. For instance, if low health-risk households moved to HIX plans, this may cause premiums to drop. Our counterfactuals that simulate moving households to the HIX thus evaluate two types of counterfactuals, that vary based on whether we allow premiums to change. First, we conduct counterfactuals without adjustments in premiums, which we call *non-equilibrium counterfactuals*. These counterfactuals are the correct approach for the goal of evaluating the incentives of one household faced with an alternative set of plans, since one household switching will not change the plan premiums in HIXs. Second, we allow for premiums to adjust in equilibrium, which we denote *equilibrium counterfactuals*. As discussed in Section 2, our equilibrium counterfactuals set premiums equal to the ACA medical-loss ratio standard—where the households that can choose these plans are those in our SGM sample—and then calculate premium changes with a fixed-point approach.

We now investigate our three sets of counterfactuals in turn.

5.1 Value of Switching to HIX Plans

We first analyze the relative value of the Small Group Market compared to the HIX, if the fiscal policy in terms of the tax-subsidy-penalty environment stays constant with its 2019 environment. Thus, in our first counterfactual, everyone pays the premiums with pre-tax dollars in the SGM and with post-tax dollars in HIXs, eligible households receive subsidies in HIXs, and households do not have to pay a penalty if they choose the outside option and are uninsured.

Table 5 presents the results of these counterfactuals. There are two main takeaways from this exercise. First, the main value of HIXs is in creating choice for consumers. Mov-

ing households to HIXs but restricting the number of plans to the number that consumers face in the SGM creates welfare losses for consumers. Column I of Table 5 presents both non-equilibrium (Panel A) and equilibrium (Panel B) counterfactuals, and show mean and median negative consumer welfare changes for most households. I.e., only 15% of consumers are better off in the non-equilibrium case. Prices drop when we account for equilibrium since the HIX market is adversely selected, incorporating higher health risk than the SGM. However, even with the lower equilibrium prices, we find that only 20% of consumers are better off when the number of plans remain the same.

We next turn to the *Choice of best mean plan in each metal tier*, which is more realistic in the sense that it offers households more choices with heterogeneity in characteristics but limits the number of extra plans. Even when households do not like the offerings of HIXs they do like the additional choices. The share of households that are better off in this scenario is 48% in the non-equilibrium counterfactuals and 60% in the equilibrium counterfactuals (Panels A and B of Column II of Table 5, respectively). The equilibrium mean consumer welfare gain, with every consumer moving and the prices adjusting, is better for households than the non-equilibrium result ($\$-436$ vs. \$406 annually relative to the baseline). This \$842 welfare gain is consistent with annual premiums conditional on buying falling from \$3,802 to \$2,719. The effect is due to the equilibrium movement of SGM enrollees mitigating the adverse selection in HIXs. However, achieving the equilibrium counterfactual requires coordination, since individual household incentives to move away from the SGM are smaller than the equilibrium effects.

Turning to Column III, if we allow households to have access to all HIX plans, then 92% of them are better off in the non-equilibrium counterfactuals while 96% of them are better off under the equilibrium counterfactuals. However, we view this counterfactual as overstating the value of HIXs, by giving households access to many $i.i.d.$ unobservable shocks even though these options likely have correlated utilities. On balance, we believe that Column II provides the most accurate view of the impact of HIXs, as it gives households access to plans with different risk-sharing arrangements but also limits the number

of plans.

As discussed in the introduction and presented in Figure 1, SGM enrollment was very stable around the introduction of the ACA. Focusing on Column II and Panel A, our results may explain why we do not observe a decline in the SGM after the creation of HIXs: the majority of households who are offered SGM plans are worse off if we move them to HIXs.

On the government side, in all three columns——and with and without equilibrium effects——we find that moving households to HIXs increases taxes but increases subsidies even more. In fact, this creates more government involvement in the market and is worse than budget neutral for the government. In equilibrium counterfactuals, the government loss is lower than in non-equilibrium counterfactuals because it generates more advantageous selection: the increased probability that lower-risk households purchase insurance makes prices go down, which in turn lowers subsidies, reducing government expenditures.

5.2 Impact of Changing Tax-Subsidy-Penalty Incentives in the SGM

Table 6 presents counterfactuals that alter the tax-subsidy-penalty structure of the SGM. For these results, we do not change consumer choice sets but instead, for each household, we use their choice set of SGM plans from the data. Thus, unlike in Section 5.1, we investigate a single plan choice set, rather than three different choice sets. Also unlike in Section 5.1, we keep premiums constant in these counterfactuals, rather than computing the impact of equilibrium premium setting. We do this because our data contain information from only one SGM insurer, USIC, and would not allow us to credibly understand counterfactual premium-setting policies.

We compare outcomes given the current SGM to outcomes under three counterfactual policies. In all three counterfactuals, households pay premiums with post-tax dollars, instead of the baseline pre-tax dollars. The counterfactuals vary in policies besides removing tax deductibility. Column I compares a scenario where we only remove tax

deductibility to the baseline. Column II evaluates removing tax deductibility and adding ACA subsidies to the baseline. Finally, Column III considers removing tax deductibility and adding the ACA penalties.

Our takeaway from all three counterfactuals is that removing tax deductibility makes potential SGM participants worse off. However, across counterfactuals, the increase in government revenues is always larger than consumer losses. In this sense, having SGM health insurance be tax deductible does not increase welfare.

Focusing on Column I, taking away the tax deductibility creates a mean welfare loss of $551 and a median welfare loss of $482, which are both somewhat lower than the increase in government revenues per person, $629. Removing tax deductibility modestly reduces the fraction of households that choose an SGM plan, by 2 percentage points. However, this comes with important distributional effects. Removing tax deductibility disproportionately hurts households with higher income. This can be seen from the fact that the mean welfare loss from removing tax deductibility is more than the median welfare loss, implying a right-skewed distribution of welfare losses.

Column II differs from Column I in that it adds subsidies (also removing tax deductibility as in Column I). Many households are slightly better off now compared to Column I, since they would receive the subsidy. The mean welfare loss goes down from $551 to $548. However, the median welfare loss does not change at all, implying that the median household in the SGM would not receive a subsidy in the HIX.

Finally, Column III examines what would occur if ACA penalties were imposed for households that did not purchase SGM insurance when offered, still removing tax deductibility. We find that mean household welfare would go down by $670, or $119 more than in Column I. However, government revenue would increase by $198 more.

5.3 Moving to HIX and Changing Tax-Subsidy-Penalty Incentives

Our final set of counterfactuals combines the movement to HIX plans from Section 5.1 with changes in the tax-subsidy-penalty environment, as in Section 5.2. This allows us

to understand how changes in tax-subsidy-penalty policies would impact the welfare changes from moving people to the HIX. We use the *Choice of best mean plan in each metal tier* set of plans, since we believe this presents the most accurate view of HIXs, as discussed in Section 5.1. We again consider non-equilibrium and equilibrium counterfactuals.

Table 7 presents the results of these counterfactuals. Column I repeats the results of Table 5, Column II, which shows the baseline 2019 environment of $d = 0$, $s = 1$, and $f = 0$ for our preferred set of HIX plans.

Column II considers the introduction of penalties. Without equilibrium, adding penalties lowers mean household welfare by \$124, changing it from -\$436 to -\$560 relative to the SGM. The equilibrium drop in mean welfare is somewhat smaller, at \$108 (moving welfare from \$406 to \$298). On the one hand, adding penalties yields almost no change in the non-equilibrium premiums paid conditional on purchasing insurance since these premiums are not directly linked to health risk among households that choose insurance. On the other hand, in equilibrium, penalties are linked to health risk. Imposing penalties causes slightly more low-risk households to purchase insurance, which yields a modest \$3 mean drop (from \$2,719 to \$2,716) in equilibrium premiums paid. As with our results in Section 5.2, penalties increase government revenues significantly. In the equilibrium counterfactuals, they increase government revenues by \$175 per household. Most of this effect is coming from the increase in penalties but there is also a small increase in subsidies.

Column III presents the results of adding tax deductibility to our baseline. Tax deductibility is relevant for SGM households because their income levels imply that many have reasonably high marginal tax rates. These households would benefit from tax deductibility significantly: equilibrium mean and median consumer welfare increase by \$551 and \$395, respectively. However, implementing tax deductibility only increases take-up by about 2 percentage points. Finally, Column IV shows that the marginal effect of adding penalties is similar with and without tax deductibility.

We can use these results to understand the extent to which the Tax Cuts and Jobs Act of 2017 caused adverse selection in the HIX market, by eliminating ACA penalties. Specifically, we can do this by comparing the equilibrium to non-equilibrium counterfactual premiums conditional on purchase in Column II to those in Column I (and analogously for Columns IV and III with tax deductibility).[16] As noted above, the difference in annual premiums paid conditional on purchasing health insurance is only $3 in equilibrium and roughly $0 without equilibrium. With tax deductibility, it is even smaller at $2. Thus, ACA penalties helped with adverse selection in individual health insurance, but the effect of the repeal for our sample is very modest.

6 Conclusions

This paper analyzes the interplay between employer-sponsored insurance for the small group market and the ACA health insurance exchange marketplaces. Specifically, we evaluate who benefits from being able to purchase insurance in the SGM relative to having this market eliminated and being faced with HIX plans. We model the demand for health insurance plans with a flexible discrete choice specification. We use household disposable income (net of health insurance premiums) as our main regressor to capture price elasticity, which allows us to evaluate the heterogeneous implications of health insurance taxability and other fiscal incentives.

To estimate the model, we use a unique dataset from a large U.S. health insurer, "United States Insurance Company" (USIC), with premium information on over 30,000 employers and more than 700,000 enrollees in the SGM. This dataset contains detailed data on plan offerings, characteristics, and premiums, linked to medical and pharmaceutical claims for individuals that purchase insurance through USIC. Since this dataset does not include enrollee income information or household characteristics for those choosing the outside option, we complement these data with information on health characteris-

[16] Premiums reflect average health risk by our equilibrium pricing assumption.

tics and income from the Medical Expenditure Panel Survey (MEPS). We estimate the model via simulated maximum likelihood, where we integrate over the distributions of household characteristics in order to obtain choice probabilities. This allows us to recover individual household heterogeneity in price coefficients and elasticities of substitution in terms of health conditions, income, age and gender. We find that this heterogeneity is important in obtaining credible demand elasticities and substitution patterns.

We simulate counterfactuals to evaluate the questions of how the tax-subsidy-penalty environment affects the value of the small group market and what would occur if this market were simply eliminated. We obtain four main conclusions.

First, when holding the number of plan options fixed, most households tend to prefer the SGM, based on plan characteristics and equilibrium premium setting in this market. This is true particularly for households with individuals who are healthier and have higher income.

Second, a feature of the HIX is that it offers households more choices. Even in the case where households could choose from the best mean HIX plan per metal tier, fewer than half of SGM enrollees would be better off from moving to HIX plans in the absence of equilibrium effects. An equilibrium where the SGM was shut down—and all enrollees in our sample were forced to move to HIX plans to obtain insurance—would increase welfare for 60% of enrollees in our sample. However, reaching such a new equilibrium may be difficult because it would require coordination.

Third, alternative designs of the SGM that remove the tax deductibility but provide subsidies can improve consumer welfare. Such policies would imply important distributional effects, favoring lower income households over relatively high income ones.

Finally, imposing penalties as in the ACA only slightly addresses the adverse selection problems of individual health insurance. Specifically, adding 2014 ACA penalties decreases the difference in equilibrium premiums relative to non-equilibrium premiums by only $3 annually (reflecting a similar drop in health risk) relative to without the penalties. We are not aware of any model-based evaluation similar to ours of the equilibrium

effects of penalty elimination, which were implemented by the Tax Cuts and Jobs Act of 2017.

References

Ackerberg, D. A. and Rysman, M. (2005). Unobserved product differentiation in discrete-choice models: Estimating price elasticities and welfare effects. *RAND Journal of Economics*, 36:771–788.

Allcott, H., Lockwood, B. B., and Taubinsky, D. (2019). Regressive sin taxes, with an application to the optimal soda tax. *The Quarterly Journal of Economics*, 134(3):1557–1626.

Azevedo, E. M. and Gottlieb, D. (2017). Perfect competition in markets with adverse selection. *Econometrica*, 85(1):67–105.

Berry, S. T. (1994). Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25:242–262.

Bundorf, M. K., Levin, J., and Mahoney, N. (2012). Pricing and welfare in health plan choice. *American Economic Review*, 102(7):3214–48.

Cardell, N. S. (1997). Variance components structures for the extreme-value and logistic distributions with application to models of heterogeneity. *Econometric Theory*, 13(2):185–213.

Cicala, S., Lieber, E. M., and Marone, V. (2019). Regulating markups in US health insurance. *American Economic Journal: Applied Economics*, 11(4):71–104.

Currie, J. and Madrian, B. C. (1999). Health, health insurance and the labor market. *Handbook of labor economics*, 3:3309–3416.

Curto, V., Einav, L., Levin, J., and Bhattacharya, J. (2021). Can health insurance competition work? Evidence from Medicare Advantage. *Journal of Political Economy*, 129(2):570–606.

Cutler, D. M. and Reber, S. J. (1998). Paying for health insurance: The trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, 113(2):433–466.

Dafny, L., Ho, K., and Varela, M. (2010). An individual health plan exchange: Which employees would benefit and why? *American Economic Review*, 100(2):485–89.

Dafny, L., Ho, K., and Varela, M. (2013). Let them have choice: Gains from shifting away from employer-sponsored health insurance and toward an individual exchange. *American Economic Journal: Economic Policy*, 5(1):32–58.

DeVaro, J. and Maxwell, N. L. (2014). The elusive wage-benefit trade-off: The case of employer-provided health insurance. *International Journal of Industrial Organization*, 37:23–37.

Dickstein, M. J., Ho, K., and Mark, N. (2024). Market segmentation and competition in health insurance. *Journal of Political Economy*, 132(1):96–148.

Dubois, P., Griffith, R., and O'Connell, M. (2020). How well targeted are soda taxes? *American Economic Review*, 110(11):3661–3704.

Einav, L., Finkelstein, A., and Tebaldi, P. (2024). Market design in regulated health insurance markets: Risk adjustment vs. subsidies. Technical report, National Bureau of Economic Research.

Ericson, K. M. and Starc, A. (2015). Pricing regulation and imperfect competition on the Massachusetts Health Insurance Exchange. *Review of Economics and Statistics*, 97(3):667–682.

Federal Register (2019). Health reimbursement arrangements and other account-based group health plans. 84 Fed. Reg. No. 119 (June 20, 2019), Rules and Regulations (45 C.F.R. pts. 144, 146, 147, and 155).

Feenberg, D. and Coutts, E. (1993). An introduction to the TAXSIM model. *Journal of Policy Analysis and Management*, 12(1):189–194.

Finkelstein, A., Hendren, N., and Shepard, M. (2019). Subsidizing health insurance for low-income adults: Evidence from Massachusetts. *American Economic Review*, 109(4):1530–1567.

Fleitas, S., Gowrisankaran, G., and Lo Sasso, A. (2022). Incumbent regulation and adverse selection: You can keep your health plan, but at what cost? *Journal of Public Economics*, 205:104556.

Fleitas, S., Gowrisankaran, G., and Lo Sasso, A. (2025). Reclassification risk in the small group health insurance market. *Review of Economics and Statistics*, page forthcoming.

Frean, M., Gruber, J., and Sommers, B. D. (2017). Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act. *Journal of Health Economics*, 53:72–86.

Griffith, R., Nesheim, L., and O'Connell, M. (2018). Income effects and the welfare consequences of tax in differentiated product oligopoly. *Quantitative Economics*, 9(1):305–341.

Griffith, R., O'Connell, M., and Smith, K. (2019). Tax design in the alcohol market. *Journal of Public Economics*, 172:20–35.

Gruber, J. and Lettau, M. (2004). How elastic is the firm's demand for health insurance? *Journal of Public Economics*, 88(7-8):1273–1293.

Gruber, J. and Poterba, J. M. (1996). Tax subsidies to employer-provided health insurance. In *Empirical foundations of household taxation*, pages 135–168. University of Chicago Press.

Hackmann, M. B., Kolstad, J. T., and Kowalski, A. E. (2015). Adverse selection and an individual mandate: When theory meets practice. *American Economic Review*, 105(3):1030–66.

Keith, K. (2019). Final rule on health reimbursement arrangements could shake up markets. Health Affairs Blog, June 14, 2019. DOI: 10.1377/hblog20190614.388950.

Kim, G. and Lyons, S. (2019). HISIM2—CBO's new health insurance simulation model.

Lubotsky, D. and Olson, C. A. (2015). Premium copayments and the trade-off between wages and employer-provided health insurance. *Journal of Health Economics*, 44:63–79.

Lurie, I. Z., Sacks, D. W., and Heim, B. (2021). Does the individual mandate affect insurance coverage? Evidence from tax returns. *American Economic Journal: Economic Policy*, 13(2):378–407.

McFadden, D. (1974). Conditional logit analysis of qualitative choice behavior. In Zarembka, P., editor, *In Frontiers in Econometrics*. Academic Press, New York.

MEPS (2019). MEPS insurance component presentation 2019; accessed: 06.16.2024. `https://www.meps.ahrq.gov/data_files/publications/cb24/cb24a.shtml`.

Olson, C. A. (2002). Do workers accept lower wages in exchange for health benefits? *Journal of Labor Economics*, 20(S2):S91–S114.

Panhans, M. (2019). Adverse selection in ACA exchange markets: Evidence from Colorado. *American Economic Journal: Applied Economics*, 11(2):1–36.

Petrin, A. (2002). Quantifying the benefits of new products: The case of the minivan. *Journal of Political Economy*, 110(4):705–729.

Royalty, A. B. and Solomon, N. (1999). Health plan choice: price elasticities in a managed competition setting. *Journal of Human Resources*, 34:1–41.

Saltzman, E. (2019). Demand for health insurance: Evidence from the California and Washington ACA exchanges. *Journal of Health Economics*, 63:197–222.

Tebaldi, P. (2025). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the ACA. *Review of Economic Studies*, 92(1):586–620.

Tables and Figures

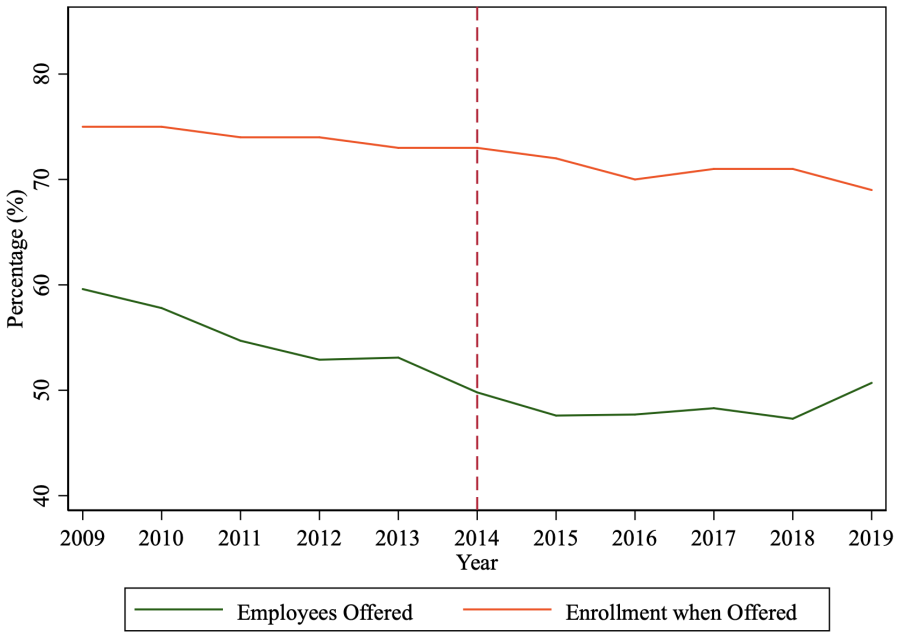

Figure 1: Offering and Enrollment in the Small Group Market

Note: graph constructed by authors using Medical Expenditure Panel Survey, available at `https://www.meps.ahrq.gov/data_files/publications/cb24/cb24a.shtml`.

Figure 2: States in our estimation sample

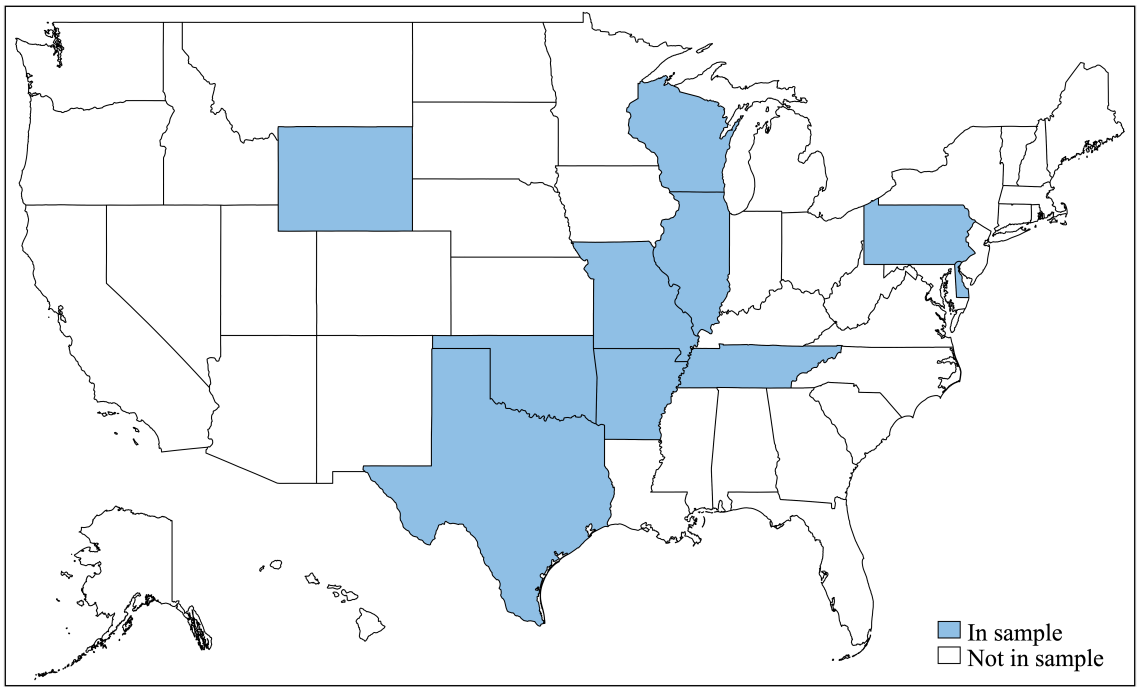

Figure 3: Distribution of plans offered by employers in SGM

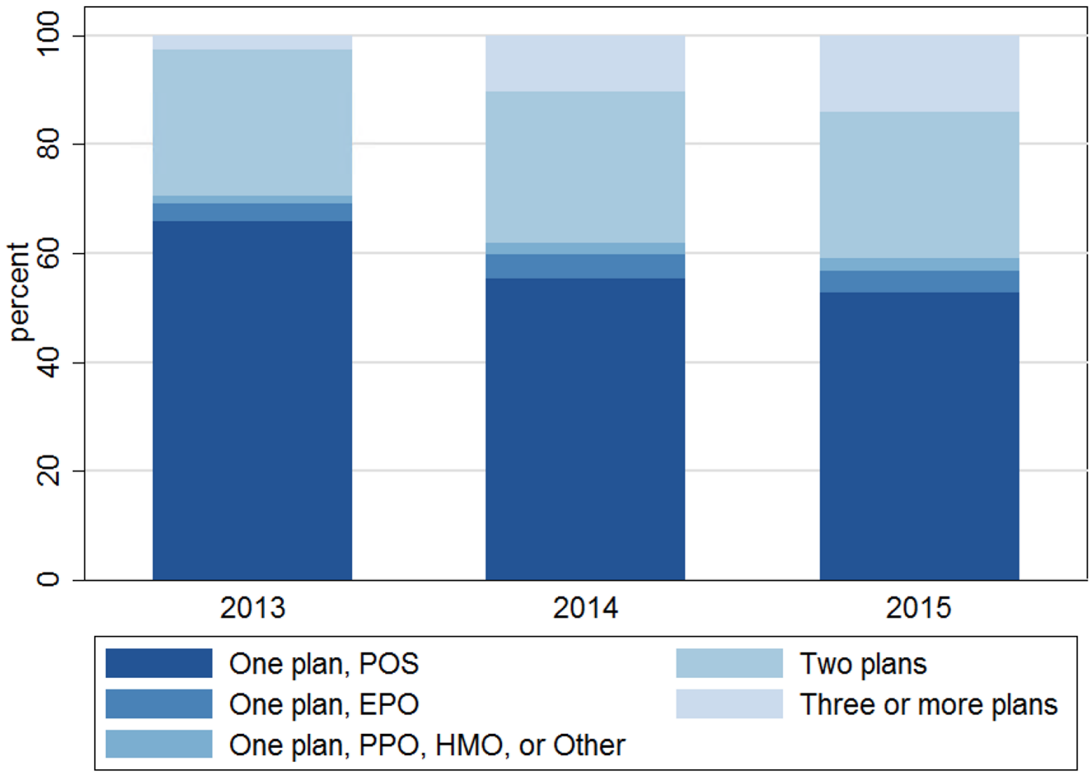

Note: graph constructed by authors using USIC data and estimation sample.

Figure 4: Weighted Average Plan Own-Price Elasticity

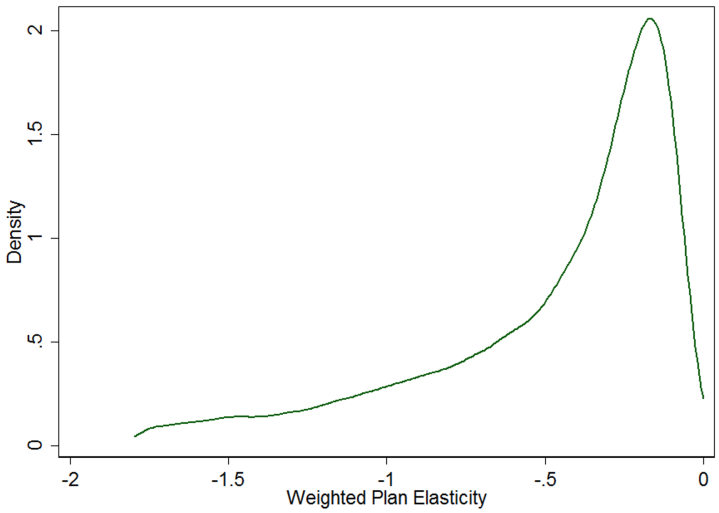

Note: histogram constructed based on preferred specification estimates from Column IV of Table 4.

Figure 5: Heterogeneity in Elasticities

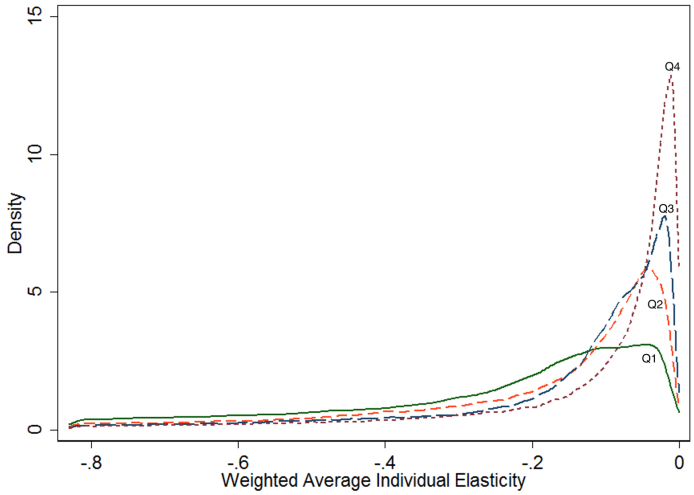

(a) Heterogeneity by r_i quartile

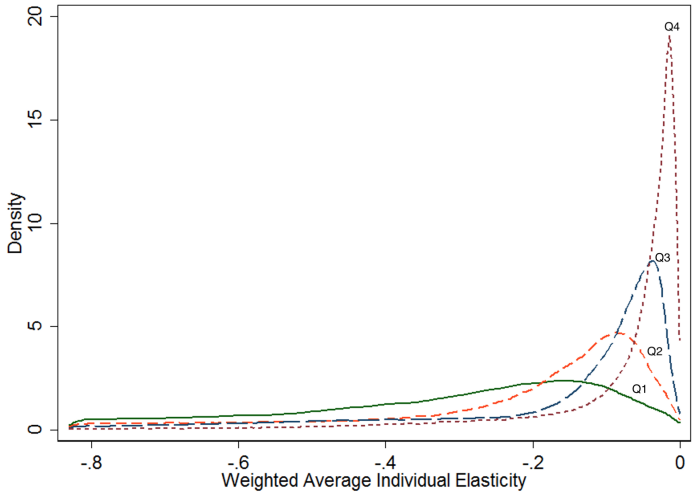

(b) Heterogeneity by Y_i^D quartile

Note: histogram constructed based on preferred specification estimates from Column IV of Table 4.

Table 1: Employer Characteristics

Firm size	17.51
	(18.74)
Employee take-up	0.67
	(0.29)
Number of plans offered	1.23
	(0.50)
Industries	
Natural Resources, Mining, Construction, and Manufacturing	0.25
Transportation, Utilities, and Communications	0.04
Wholesale and Retail Trade	0.19
Finance, Insurance, Real Estate, and Professional and Business Services	0.08
Services	0.42
Public Administration and Military	0.01
Observations (employer-years)	70,724

Note: Each observation is one employer during one year, 2013-2015. Table reports mean values with standard deviations in parentheses.

Table 2: Individual Household Characteristics

	Overall	Insiders	Outsiders
Household size	2.09	1.83	2.54
	(1.11)	(1.32)	(0.27)
Household health risk, r_i	0.80	0.91	0.63
	(1.10)	(1.38)	(0.9)
Household income, Y_i^G	61,872	58,929	66,934
	(24,522)	(29,596)	(9,481)
Employee age	43	45	39
	(10)	(12)	(3)
Employee female	0.35	0.36	0.32
	(0.40)	(0.48)	(0.18)
Employee health risk	0.99	0.90	1.15
	(0.74)	(0.89)	(0.30)
Observations (individual-year)	1,274,584	805,209	469,375
Individuals			
2013	432,429	266,896	165,533
2014	451,869	279,305	172,564
2015	390,286	259,008	131,278

Note: Each observation is one household during one year, 2013-2015. Table reports mean values with standard deviations in parentheses. "Insiders" are households who chose a USIC plan and "outsiders" are households who chose the outside option.

Table 3: Plan characteristics

	SGM	HIX
Annual premium	3,373	4,596
	(1,710)	(1,113)
In-network deductible.	2,165	1,071
	(1,435)	(1,842)
Copay	24	0
	(11)	(0)
Coinsurance (%)	87	92
	(17)	(14)
In-network out-of-pocket maximum	3,467	3,656
	(2,041)	(2,220)
Out-of-network deductible.	4,116	1,914
	(2,872)	(3,635)
Out-of-network out-of-pocket maximum	8,194	6,659
	(4,468)	(7,490)
Observations (plan-year)	86,766	1,801

Note: Each observation is one plan during one year, 2013-2015 for the case of the SGM and 2014 for the case of the HIX. Coinsurance is reported as percent of claims paid by insurer. For silver plans, we report averages across all cost-sharing reduction variants of the same plan. Table reports means with standard deviations in parentheses.

Table 4: Demand estimation results

	(I)	(II)	(III)	(IV)
log(p_{ij})	-0.02			
	(0.01)			
log(Y_{ij}^D)		0.78	14.17	13.22
		(0.02)	(0.07)	(0.07)
log(Y_{ij}^D) $\times$ r_i			-6.57	-5.50
			(0.04)	(0.04)
log(Y_{ij}^D) $\times$ $Family\ Size$			-2.37	-2.30
			(0.02)	(0.16)
$1-\sigma$	0.16	0.29	0.72	0.70
	(0.004)	(0.004)	(0.004)	(0.004)
Wgt. avg. own-price elasticity	-0.06	-0.11	-0.38	-0.36
Main Variable	log(p_{ij})	log(Y_{ij}^D)	log(Y_{ij}^D)	log(Y_{ij}^D)
Main Variable $\times$ r_i	No	No	Yes	Yes
Plan Characteristics	Yes	Yes	Yes	Yes
Plan Characteristics $\times$ r_i	No	No	No	Yes
Plan Type FE	Yes	Yes	Yes	Yes
Year FE	Yes	Yes	Yes	Yes
Observations	32,830,705	32,830,705	32,830,705	32,830,705

Note: All specifications include plan characteristics (coinsurance, deductible, out of pocket maximum, copay, out of network deductible, out of network out of pocket maximum and out of network deductible), plan type indicators and fixed effects by year. Standard errors are in parentheses.

Table 5: Counterfactuals on Health Insurance Exchanges Versus Small Group Market

$\Delta WS = HIX - SGM$	(I)	(II)	(III)
Counterfactual HIX:			
Tax Deductibility (d)	0	0	0
ACA Subsidies (s)	1	1	1
Penalties (f)	0	0	0
Scenario:	Swap to mean best plan	Choice of best mean plan in each metal tier	Choice of all HIX plans
Panel A: Non-Equilibrium Counterfactuals			
Mean welfare ($)	-3,389	-436	9,523
Median welfare ($)	-1,992	-95	5,096
% Positive or equal	15%	48%	92%
Percentage point change in take-up	-6.0	-3.3	15.1
Mean premium paid ($)	1,994	2,498	3,434
Mean premium paid ($) cond. on purchase	3,175	3,802	4,043
Net change in gov't. revenue ($)	-264	-231	-495
Net change in taxes ($)	219	219	219
Average subsidy amount ($)	483	450	714
% Receiving subsidy	40%	37%	40%
Average subsidy ($) cond. on receipt	1,215	1,210	1,796
Panel B: Equilibrium Counterfactuals			
Mean welfare ($)	-3,127	406	10,892
Median welfare ($)	-1,717	534	6,449
% Positive or equal	20%	60%	96%
Percentage point change in take-up	-3.7	0.2	18.1
Mean premium paid ($)	1,756	1,896	2,437
Mean premium paid ($) cond. on purchase	2,678	2,719	2,780
Net change in gov't. revenue ($)	-182	-96	-409
Net change in taxes ($)	218	218	218
Average subsidy amount ($)	400	314	627
% Receiving subsidy	33%	30%	35%
Average subsidy ($) cond. on receipt	1,200	1,061	1,804

Note: Change in take-up is the average change in the probability of take-up when moving from the SGM to the HIX. In each panel, we present the results dropping the first and last percentile of the welfare change based on the *Choice of best mean plan in each metal tier* scenario in Column (II).

Table 6: Counterfactuals on Alternative Formulations of the Small Group Market

$\Delta\, WS = SGM^{Counter} - SGM$	(I)	(II)	(III)
$Counterfactual\ SGM:$			
Tax Deductibility (d)	0	0	0
ACA Subsidies (s)	0	1	0
Penalties (f)	0	0	1
Mean welfare ($)	-551	-548	-670
Median welfare ($)	-482	-482	-593
% Positive or equal	2%	3%	0%
Percentage point change in take-up	-2.2	-2.1	-1.1
Mean premium paid ($)	1,947	1,949	1,980
Mean premium paid ($) cond. on purchase	2,946	2,946	2,946
Net change in gov't. revenue ($)	629	623	827
Net change in taxes ($)	629	629	629
Average subsidy ($)	0	6	0
Average subsidy ($) cond. on receipt	0	405	0
Average penalty ($)	0	0	198

Note: Change in take-up is the average change in the probability of take-up when comparing the three scenarios to the baseline. We present the results dropping the first and last percentile of the welfare change.

Table 7: Counterfactuals on Health Insurance Exchanges with Different Fiscal Incentives

	(I)	(II)	(III)	(IV)
$\Delta WS = HIX - SGM$				
Scenario: Choice of best mean plan in each metal tier				
Counterfactual HIX:				
Tax Deductibility (d)	0	0	1	1
ACA Subsidies (s)	1	1	1	1
Penalties (f)	0	1	0	1
Panel A: Non-Equilibrium Counterfactuals				
Mean welfare ($)	-436	-560	243	129
Median welfare ($)	-95	-177	404	334
% Positive or equal	48%	47%	57%	56%
Percentage point change in take-up	-3.3	-2.2	-0.7	0.3
Mean premium paid ($)	2,498	2,538	2,645	2,683
Mean premium paid ($) cond. on purchase	3,802	3,802	3,858	3,858
Net change in gov't. revenue ($)	-231	-43	-1,034	-870
Net change in taxes ($)	219	219	-561	-572
Average subsidy ($)	450	458	473	482
% Receiving subsidy	37%	37%	37%	37%
Average subsidy ($) cond. on receipt	1,210	1,232	1,273	1,295
Average penalty ($)	0	196	0	184
Average penalty ($) cond. on purchase	0	207	0	193
Panel B: Equilibrium Counterfactuals				
Mean welfare ($)	406	298	957	855
Median welfare ($)	534	468	929	870
% Positive or equal	60%	58%	66%	65%
Percentage point change in take-up	0.2	1.2	2.0	2.9
Mean premium paid ($)	1,896	1,920	1,940	1,962
Mean premium paid ($) cond. on buying	2,719	2,716	2,711	2,709
Net change in gov't. revenue ($)	-96	79	-729	-569
Net change in taxes ($)	218	218	-406	-413
Average subsidy ($)	314	319	323	327
% Receiving subsidy	30%	30%	29%	29%
Average subsidy ($) cond. on receipt	1,061	1,083	1,106	1,122
Average penalty ($)	0	180	0	171
Average penalty ($) cond. on paying	0	189	0	180

Note: Change in take-up is the average change in the probability of take-up when moving from the SGM to the HIX. In each panel, we present the results dropping the first and last percentile of the welfare change based on the *Choice of best mean plan in each metal tier* scenario in Column (II).

A Appendix: Equilibrium Counterfactuals

In the equilibrium counterfactuals, we allow for the premium in the HIX plans to vary based on health costs to the insurer and the probability that households choose a plan.

We let the markup of plan premiums over total medical expenditures follows the Medical Loss Ratio (MLR) rule implemented under the ACA. The ACA requires insurers to spend at least 80% of premiums on medical care. We assume that plans' premiums would be at the MLR margin, and so premiums are set to be no more than 125% of expenditures.

We begin by predicting total health expenditures for each health plan. Using data from households who opt in to their USIC ESI, we predict annual expenditure by performing a flexible regression of expenditure on plan characteristics and their interaction with individual household characteristics of r_i, and h_i. We use the regression coefficients to predict insurer expenditure for each of the health plans available in the HIX.

We compute the total expenditure by plan in the HIX as the total expenditure by household multiplied by the coverage of the plan's metal tier. The choice set is determined by the plans each consumer had to choose from in the non-equilibrium scenario. Therefore, our equilibrium and non-equilibrium results show the welfare change from the same set of HIX plans, but the equilibrium results allow for premiums to adjust.

For our main analysis, we limit the plans available to the *Choice of best mean plan in each metal tier* plans. Based on welfare from non-equilibrium counterfactuals, we select the plans that provide the highest average utility plan across individuals in each metal tier offered in each geographical market.

We then allow for premiums to update based on the individuals who select into a plan and their annual expenditures (weighted by the probability of choosing a plan), times the markup. As we explained in equation 7, we solve for a fixed point where the premium for each plan is equal to the inverse MLR times total expenditures for the integral of the share of that plan over households given the current premiums in the market. We scale individual costs by the actuarial value of the metal tier, where the insurer pays 60%

of medical costs for a Bronze plan, 70% of costs for a Silver plan, 80% for a Gold plan, and 90% for a Platinum health plan. Our fixed-point process sets a tolerance of $0.1 and we iterate until the difference between new premiums and old premiums is below the tolerance level.