Firm Learning in a Selection Market

Claudio Lucarelli[1] and Evan Saltzman[*2]

[1]University of Pennsylvania
[2]Emory University

November 9, 2021

Abstract

A prevalent mechanism for implementing large social programs is creating new markets. Assuming firms have full information upon market inception is commonplace in the literature. We develop an adaptive learning model with selection to study how firms' knowledge of demand and cost affects conclusions about the market equilibrium. Assuming firms have full information leads to a more favorable conclusion about the equilibrium in the California ACA exchange; annual per-capita welfare is \$249 higher and total welfare is \$593 million higher. Taxpayers and disadvantaged subpopulations disproportionately bear the social cost of firm misinformation. Firm information has substantial implications for policy design.

Keywords: Adaptive learning, adverse selection, health insurance.

JEL Codes: I11, I13, L51, L88, H51

[*]We thank Ulrich Doraszelski, Ben Handel, Neale Mahoney, Esfandiar Maasoumi, Ian McCarthy, Mark Pauly, Dan Polsky, Juan Rubio-Ramirez, Ashley Swanson, Pietro Tebaldi, and Bob Town for advice and helpful feedback. Yisroel Cahn and Sam Wunderly provided excellent research assistance.

1 Introduction

Large-scale social programs in the United States are increasingly being delivered through the private sector. Prominent examples occur in education, child care, health care services, and especially health insurance. An increasing share of Americans obtain health insurance through publicly-supported private health insurance markets. As of 2020, approximately 53.9 million people participate in the Medicaid Managed Care program for the poor, 24.1 million people are enrolled in the Medicare Advantage program for the elderly, 46.5 million people have a Medicare Part D prescription drug plan, and 11.4 million get health coverage through the Affordable Care Act (ACA) exchanges (Kaiser Family Foundation, 2020). The current trend of delivering social programs through private markets is expected to continue for the foreseeable future (Gruber, 2017).

Implementation of these social programs requires establishing new markets. A significant challenge for participating firms is that they initially have little knowledge of the relevant market characteristics such as consumer preferences and competitors' behavior for making optimal decisions. In markets with selection, firms face the additional challenges of forecasting cost and understanding how cost is correlated with demand. Suboptimal decisions while firms are learning may have significant welfare implications. Therefore, potential sources of firm uncertainty should be considered when designing and evaluating social programs. However, the empirical industrial organization (IO) literature and policy-makers usually assume firms have full information and the market is in equilibrium upon market inception (Doraszelski et al., 2018). This assumption is a significant shortcoming of the literature given the increasing prevalence of social programs delivered through new private markets.

In this paper, we study how firms' knowledge about their demand and cost affects conclusions about the market equilibrium and the potential welfare implications for market design. Similar to Doraszelski et al. (2018), we estimate an adaptive learning model that allows firms to progressively learn about demand preferences in a new market. We extend Doraszelski et al. (2018)'s adaptive learning model by allowing firms to learn about cost and the correlation between demand and cost. We apply our framework to the state-based health insurance exchanges created in 2014 under the ACA, where eligible consumers can receive government subsidies for purchasing insurance from private firms. The ACA setting has two important features that make it particularly appealing for studying firm learning. First, the ACA setting provides rich data on consumer plan choices, firm costs, and firms' own predictions about cost from the establishment of the exchanges in 2014. We obtain consumer-level administrative data on consumer plan choices from the California ACA exchange. Our California data account for approximately 13% of nationwide enrollment in the ACA exchanges (Kaiser Family Foundation, 2020) and contain nearly 10 million consumer plan choices between 2014 and 2019. We also use data on firms' predictions about their costs and actual costs from insurer rate filings. The rate review process requires firms to provide actuarial justification for their proposed premiums, including a detailed explanation of their cost forecast. The availability of credible data on firms' own cost predictions is a particularly novel feature of our setting. We document that firms' predictions about their costs converged to their realized costs over the first several years of the exchange, suggesting that firms were learning. Second, firms faced considerable uncertainty in predicting consumer preferences for health insurance and the cost of insuring their enrollees. Potential enrollees came from two very distinct sources: those with coverage in the pre-ACA individual health insurance market (i.e., the market where consumers buy insurance directly from an insurer) and those without insurance (Gruber, 2017). New ACA restrictions that prohibited firms from using health status to set premiums created additional sources of uncertainty (Pauly et al., 2015, 2020).

We make three primary contributions to the literature: (1) we extend the empirical IO literature on firm learning to a market with selection, where firms need to learn about the correlation of demand and cost; (2) we demonstrate that in our setting, assuming firms have full information upon market inception leads to a more favorable conclusion about the market equilibrium, largely because of firms' initial imprecision in estimating cost and consumer premium sensitivity; (3) we show firm learning has substantial implications for the efficacy of key program design features.

Our paper contributes to the empirical IO literature on firm learning in oligopoly markets (see Aguirregabiria and Jeon (2020) for a recent review of this literature), which mostly focuses on how consumers learn about their demand (Ackerberg, 2003; Dickstein, 2018) or how firms learn about their cost (Benkard, 2000; Zhang, 2010; Conley and Udry, 2010; Newberry, 2016). Jeon (2020) studies how firms in the container shipping industry learn about their demand. Several papers study whether the market converges to an equilibrium. Joskow et al. (1998) study the market for sulfur dioxide emissions following passage of 1990 Clean Air Act and find the market had become reasonably efficient by mid-1994. Hortaçsu and Puller (2008) analyze the bidding behavior of firms in the Texas electricity spot market from 2001 to 2003, finding that large firms made bids that were close to optimal. Hortaçsu et al. (2019) extend this work by examining the impact of large firms' superior strategic ability on market efficiency. Huang et al. (2021) study how firms learn about consumer demand in the Washington state liquor market following deregulation in 2012, finding that prices converge to levels consistent with profit maximization. Doraszelski et al. (2018) use adaptive learning and fictitious play models to study how firms learn about their demand and competitors' behavior in the U.K. electricity market following deregulation. They find that it takes several years before firms' behavior is consistent with a complete information Nash equilibrium and convergence to equilibrium is better described with learning models than with standard IO models.

We extend this literature by applying adaptive learning to a selection market where firms not only need to learn about their demand and cost, but also how demand and cost are correlated. To the best of our knowledge, our paper is the first to empirically study firm learning and the convergence to equilibrium in a selection market, where uncertainty is particularly acute. In our model, firms use only the *available* information on demand and cost to form expectations about the future.[1] Our framework is a straightforward extension of standard approaches in the empirical IO literature and

[1]The field of macroeconomics has a long history of including adaptive learning in dynamic general equilibrium models (Sargent, 1993; Evans and Honkapohja, 2001)

can be applied with data on demand, cost, and firm predictions. The model accounts for adverse selection and moral hazard and endogenizes consumer choices, premiums, plan risk, and claims. We estimate the model using our data on firms' own cost predictions and consumer-level enrollment data. Our parameter estimates indicate firms initially underestimated inertia and premium elasticities for exchange coverage, but overestimated own-plan premium elasticities.

We next use our estimated learning model to determine how much assuming firms have full information affects conclusions about the market equilibrium. The full information assumption is commonplace in the previous literature evaluating the design of government-created health insurance markets. In previous ACA exchange studies, Tebaldi (2020), Saltzman (2021), Polyakova and Ryan (2021), and Einav et al. (2019) assume full information. The full information assumption is also made in studies of Medicare Advantage (Town and Liu, 2003; Lustig, 2009; Curto et al., 2020; Miller et al., 2019), Medicare Part D (Abaluck and Gruber, 2011, 2016; Ketcham et al., 2015; Decarolis et al., 2020; Fleitas, 2017; Lucarelli et al., 2012), Medigap (Starc, 2014), and the pre-ACA Massachusetts exchange (Ericson and Starc, 2015; Geruso et al., 2019; Hackmann et al., 2015; Finkelstein et al., 2019; Jaffe and Shepard, 2020). Given the incidence of government-created health insurance markets over the last two decades, it is especially important to assess the validity of assuming full information.

Assuming full information in the California ACA exchange setting, compared to modeling firm uncertainty, leads to a more favorable conclusion about the equilibrium; premiums are lower, total enrollment is higher, and social welfare is higher. The differences between the equilibrium results using full information and using only the available information decline over time, an indication that firms are learning. Relative to the observed equilibrium, average premiums (including subsidies) are 14.4% lower in 2015 and 9.0% lower in 2018 when firms have full information. Annual per-capita social welfare is \$592 higher in 2015 and \$156 higher in 2018; annual total social welfare is \$1.38 billion higher in 2015 and \$371 million higher in 2018. Welfare gains associated with assuming full information largely accrue to the government because of a reduction in spending on premium-linked subsidies. This finding suggests that taxpayers bear a large share of the social welfare cost of firm misinformation. The welfare of disadvantaged subpopulations that are more price sensitive is disproportionately overestimated when assuming firms have full information. Average annual per-capita consumer surplus increases \$102 for Black consumers and \$78 for Hispanic consumers under full information relative to available information. This result highlights the importance of closely monitoring the welfare of disadvantaged subpopulations when creating new markets.

We then investigate mechanisms for the more favorable conclusion about the market equilibrium under full information. We find biased estimates of the cost parameters and premium parameters are the primary drivers. Premiums would have been 8.0% lower if firms had known the cost parameters and 8.1% lower if firms had known the premium parameters. Conversely, premiums would have been 6.9% higher if firms had not underestimated inertia.

We also show that firm information has substantial implications for program design. We use our learning model to simulate the impact of four program design features with and without full information, including: (1) regulation that prohibits or attenuates firms' ability to set premiums by consumer characteristics, known as community rating; (2) inertia, the persistence of health plan choices over time despite changes in premiums; (3) risk adjustment, which transfers money from plans with low-risk consumers to plans with high-risk consumers to discourage cherry-picking behavior by firms; and (4) the individual mandate, which requires most consumers to purchase insurance. Modifying community rating has a larger impact when firms have full information because firms initially underestimated premium elasticities for exchange coverage and community rating directly affects absolute premiums (i.e., relative to the outside option). Eliminating inertia also has a stronger effect under full information because firms initially underestimated inertia. Eliminating risk adjustment has a smaller impact under full information because risk adjustment directly affects relative premiums (i.e., between exchange plans) and firms initially overestimated own-premium elasticities. Eliminating the mandate has little impact on the equilibrium regardless of whether learning is modeled because the mandate penalty is small relative to premiums.

Our results have a number of important implications, both for researchers and policymakers. Firm learning may be relevant for analyses that forecast the impact of social programs, such as those conducted regularly by the Congressional Budget Office (CBO), and studies that retrospectively analyze program impact and the efficacy of certain program design features. Researchers should consider whether firm uncertainty is relevant for their specific setting and whether it is likely to affect estimates of key model parameters. Our study also suggests policymakers should adopt policies that promote information sharing between firms to reduce uncertainty. States can promote information sharing by adopting an "active purchasing" model where the exchange actively solicits information from firms or by creating an all-payer claims database. Because California has already adopted an active purchasing model, the effect of firm information might be even larger in other states.

The remainder of this paper is organized as follows. Section 2 describes the data and ACA setting. Section 3 develops a model of the ACA exchanges. Section 4 discusses estimation. Section 5 presents the model parameter estimates. Section 6 uses the model to simulate the impact of learning. Section 7 uses the model to simulate policy counterfactuals. Section 8 concludes.

2 Data and Policy Background

The Affordable Care Act (ACA) seeks to expand health care access coverage by promoting subsidized access to health insurance. A key mechanism for accomplishing this objective was the establishment of state-based health insurance exchanges in 2014. Eligible exchange consumers can receive subsidies to purchase health insurance from private insurance firms. Firms must comply with numerous regulations, including limitations on price discrimination.

To study these exchanges, we use two primary sets of data: (1) 2014-2018 plan-market-level data on firm costs and predictions about cost from insurer rate filings and (2) 2014-2019 consumer-level data on enrollee choices from the California ACA exchange. We describe these data sources in the following two subsections.

2.1 Data on Firm Costs and Predicted Cost

We obtain data on firm costs and predictions about cost from insurer rate filings. All participating California exchange insurers must submit their proposed premiums for actuarial review at the Department of Managed Health Care (DMHC). Insurers are required to include detailed supporting data justifying premium increases, including past medical claims and expected trends. DMHC does not have the authority to reject premium increases, but can find the insurer's rate filing "unreasonable" if the supporting data do not support the rate increase and the insurer refuses to adjust their rates accordingly. Insurers must notify enrollees of an unreasonable finding. As part of the rate filing, insurers must include an independent actuarial certification which confirms its actuarial methodologies were audited by an independent firm. Because rate filings are subject to extensive scrutiny by both DHMC and independent auditors, we assume insurers truthfully report their projected costs and cannot strategically misreport in order to gain a competitive advantage.

The DMHC rate review process usually begins the summer before the new plan year when the proposed premiums take effect and can last several months. Firms submit their premiums for plan year t in the summer of year $t-1$ using experience data (i.e., supporting data) from plan year $t-2$. For example, rate filings for 2016 are submitted in the summer of 2015 and report experience from 2014, the most recent complete year of experience. The new premiums for 2016 take effect on January 1, 2016. Firms cannot adjust premiums in the middle of the plan year. Similarly, consumers can only switch exchange plans once a year during a period called "open enrollment."

Insurers did not have any experience data from the exchanges in 2014 to make projections. Most insurers developed their 2014 premium rates using experience from other lines of business, including the pre-ACA individual market and the small group market. Although these were useful starting points, a substantial portion of the potential exchange population consisted of consumers who were uninsured. Insurers had to estimate both the size and health status of the uninsured population that would enroll. As part of its rate filing, Blue Shield indicated that it used the U.S. Census Bureau's Current Population Survey (CPS) to estimate the size of the uninsured population by age, income, and geography. Blue Shield estimated the uninsured population's take-up of insurance by calibrating premium sensitivity factors with its experience data for each age-income group. The firm assumed for each age group that the health status distribution of the uninsured population was the same as the health status distribution in its experience data.

The insurer rate filings provide key plan-market-level financial information, including data on enrollee medical claims and two important ACA risk mitigation programs – reinsurance and risk adjustment. Reinsurance was a temporary ACA program in effect from 2014 until 2016 that provided "insurance to insurers" for any enrollees with very high medical claims. The federal government served as the reinsurer and funded the program through a tax on all private insurance plans, including employer-sponsored plans. Risk adjustment is a permanent program where plans with lower-than-average risk make transfer payments to plans with higher-than-average risk. ACA risk adjustment transfers sum to zero, whereas the reinsurance program provides an inflow of funds to the ACA exchanges. Other risk adjustment programs, such as the one used in Medicare Advantage, may benchmark risk adjustment payments to the risk of those choosing the outside option (e.g., traditional Medicare) and also provide an inflow of funds to the market. The objective of risk adjustment is to disincentivize firms from cherry-picking the lowest-risk consumers to reduce cost (Handel et al., 2015; Layton, 2017; Mahoney and Weyl, 2017). Cherry-picking may result in the unraveling of the most generous, high-cost plans. Risk adjustment discourages strategic variation in premiums by plan generosity, but does not explicitly restrict such variation. In the next section, we discuss the calculation of ACA risk adjustment transfers.

Closely related to the reinsurance and risk adjustment programs was the ACA's implementation of medical loss ratio (MLR) requirements and a temporary risk corridor program. The MLR is the share of premiums spent on medical claims or efforts to improve quality of care (i.e., not profit distributions or plan administrative costs). ACA insurers must send rebates to their enrollees if the MLR falls below 80%. The MLR requirement does not appear to be binding in the California ACA exchange; only once did a California exchange insurer (out of 13) make MLR rebate payments across 5 years of data. We therefore omit MLR constraints in the model developed in the next section. The ACA's risk corridor program, in place between 2014 and 2016, reduced both insurer gains and losses. Insurers with substantial gains paid into the program, whereas insurers with substantial losses drew from the program. Profit and loss reduction were symmetric such that the risk corridor program had no impact on expected profit. Because the model developed in the next section assumes insurers are risk-neutral profit-maximizers and entry decisions are exogenous, risk corridors have no impact in our model.

A unique feature of the rate filing data is the ability to compare firms' predictions about their costs with their realized costs. We refer to the difference between the firm's predicted and realized average costs as the *cost prediction error*, where cost is the sum of claims, risk adjustment, and reinsurance. The cost prediction error for year t is the difference between the predicted average cost reported in the year t rate filing and the realized average cost reported in the year $t + 2$ rate filing. For example, the 2016 cost prediction error uses predicted cost data from the 2016 rate filing and the realized cost data for 2016 as reported two years later in the 2018 rate filing.

Figure 1 reports the firms' cost prediction error. In the first year of the ACA exchanges, Blue Shield and Kaiser over-predicted average monthly costs by $96 and $115, respectively, whereas Health Net under-predicted its average monthly cost by $118. The firms' prediction error narrowed considerably over the first five years of the exchanges. During this period, the direction of the prediction error reversed for all four firms, most strikingly for Kaiser. This reversal suggests that the firms were not strategically misleading regulators with their predictions. By 2018, Anthem and Blue Shield were able to predict their average costs to within $5 of their actual costs. Kaiser also had its smallest cost prediction error in 2018. Health Net reduced its prediction error by more then half from 2014 to 2018. We interpret this convergence of predicted and actual costs as evidence of firm learning. Morrisey et al. (2017) also find anecdotal evidence of substantial initial uncertainty and firm learning in interviews with insurance firm representatives from 5 states, including California.

Figure 1: Cost Prediction Error By Year

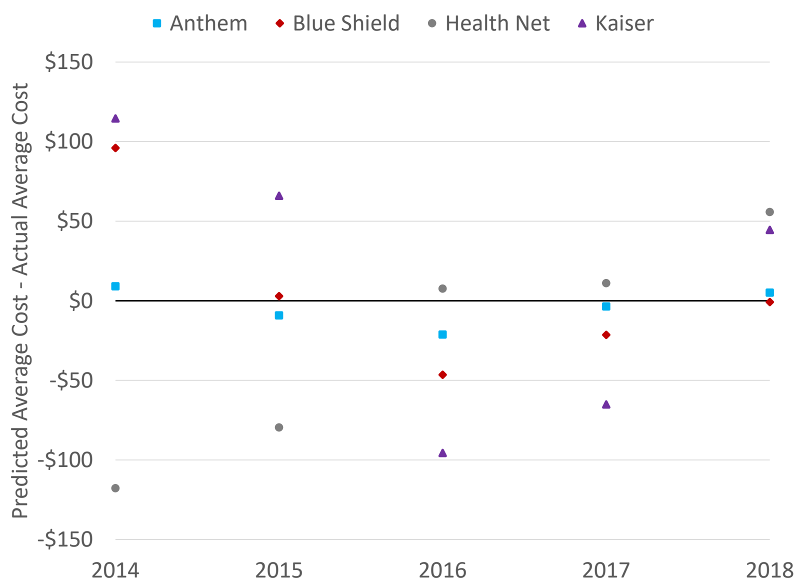

Notes: Figure shows the evolution of the average cost prediction error for the four large firms. Average cost equals average claims minus the average risk adjustment transfer received and average reinsurance received.

2.2 Data on Enrollee Choices

We obtain consumer-level enrollment data from the California ACA exchange. There are approximately 10 million records in our enrollment data between 2014 and 2019. Our enrollment data include every enrollee's chosen plan and key enrollee characteristics, but not enrollee utilization. The data provide sufficient information define every household's complete choice set and the household-specific premium paid for each plan in its choice set.

Appendix Table A1 summarizes enrollee characteristics by plan year. About 90% of exchange enrollees are eligible for premium subsidies. Premium subsidies are available to consumers who (1) have income between 100% and 400% of the federal poverty line (FPL); (2) are citizens or legal residents; (3) are ineligible for public insurance such as Medicare or Medicaid; and (4) lack access to an "affordable plan offer" through employer-sponsored insurance. Most households in California with income below 138% of FPL are eligible for Medicaid and therefore ineligible for premium subsidies. A plan is defined as "affordable" if the employee's contribution to the employer's single coverage plan is less than 9.5% of the employee's household income in the 2014 plan year. This percentage increases slightly each year. The next section discusses the complex ACA formula used to calculate premium subsidies.

Exchange consumers have access to a diverse set of plans that varies by geographic market and age. Figure 2a shows that 4 firms – Anthem, Blue Shield, Health Net, and Kaiser – dominate the California exchange. There are also 9 regional firms that offer exchange plans.[2] Anthem's market share declined substantially in 2018 when it exited most of the state. Consumers can select a plan from one of the four actuarial value (AV) or "metal" tiers: bronze (with 60% AV), silver (with 70% AV), gold (with 80% AV), and platinum (with 90% AV). Individuals under age 30 can buy a basic catastrophic plan, but premium subsidies cannot be used to purchase catastrophic plans. Consequently, Figure 2b indicates that only 1% of consumers select a catastrophic plan. In contrast, about 60% of consumers choose a plan from the silver tier because eligible consumers must choose silver to receive cost sharing reductions (CSRs) that reduce that reduce deductibles, copays, etc. CSRs increase the AV of the silver plan from 70% to (1) 94% for consumers with income below 150% of the federal poverty level (FPL); (2) 87% for consumers with income between 150% and 200% of FPL; and (3) 73% for consumers with income between 200% and 250% of FPL. Consumers with income above 250% of FPL are ineligible for CSRs. Approximately two-thirds of California consumers are eligible for CSRs.

Consumers also have an outside option to forgo insurance. Although any citizen or legal resident can purchase an exchange plan, only individuals without access to public or employer-sponsored insurance purchase exchange plans in practice because of premium subsidy eligibility rules and the

[2]These firms include Chinese Community Health Plan, Contra Costa, L.A. Care Health Plan, Molina Healthcare, Oscar, Sharp Health Plan, United Healthcare, Valley Health Plan, and Western Health Advantage.

Figure 2: Market Share By Year

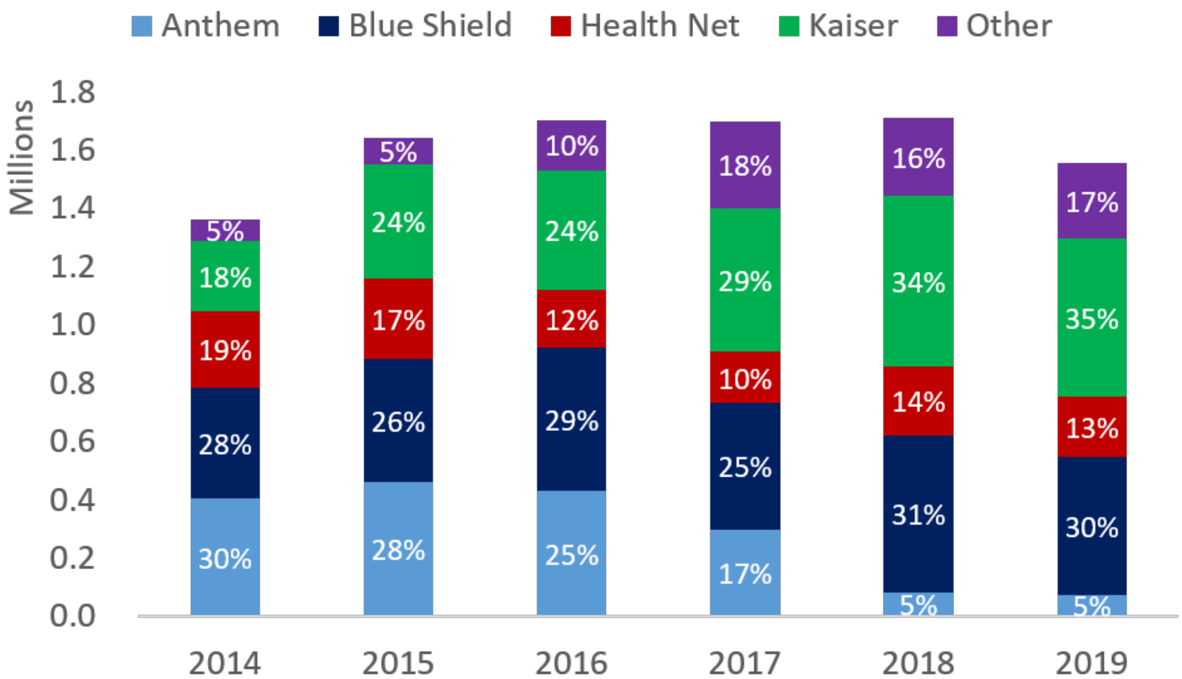

(a) By Insurer

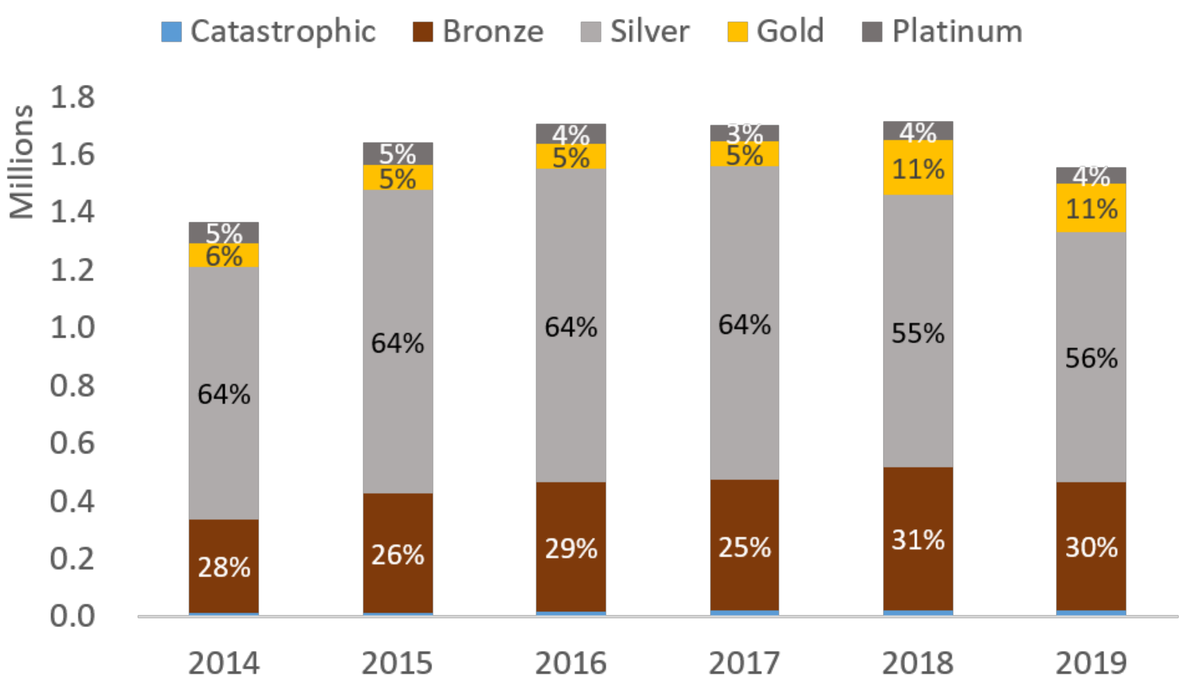

(b) By Metal Tier

prohibitive cost of exchange plans relative to other insurance. We follow the procedure originally developed in Saltzman et al. (2021) to construct the exchange-exchange eligible population using panel data from the Survey of Income and Program Participation (U.S. Census Bureau, 2019). Appendix C provides a summary of this procedure.

Consumers without insurance may be subject to a penalty under the ACA's individual mandate. The individual mandate penalty was phased in between 2014 and 2016. The penalty for a single person was the greater of \$95 and 1% of income (exceeding the tax filing threshold) in 2014 and the greater of \$695 and 2.5% of income in 2016. After passage of the Tax Cuts and Jobs Act of 2017, the penalty was set to 0 starting in 2019. Exemptions from the ACA's individual mandate are made for certain groups, including (1) those with income below the tax filing threshold and (2) individuals who lack access to a health insurance plan that is less than 8% of their income in 2014 (this percentage changes slightly each year).

Although our focus is firm learning, a natural concern is whether consumers learn and adjust their plan choices accordingly. In related work, Saltzman et al. (2021) identify two significant features of this market that mitigate the concern of consumer learning: (1) annual churn into and out of the market is substantial and (2) switching between plans is minimal despite highly volatile premiums during the study timeframe. High levels of churn suggest limited opportunities for consumers to learn and low levels of switching indicate consumers are not adjusting their plan choices over time. Hence, we do not model consumer learning.

3 Model

Saltzman (2021) specifies a structural model of the California ACA exchange and estimates it assuming firms have full information about their demand and cost. We follow this underlying model, but relax the assumption of full information in estimation as discussed in the next section. Using this model allows us to isolate the impact of the full information assumption and quantify potential specification errors from using a standard IO approach.

Consider a two-stage game where in each period t (1) insurers set premiums simultaneously and (2) consumers choose plans. Below we detail how we model each of these stages, beginning with consumer plan choice.

3.1 Consumer Plan Choice

Households select the plan that maximizes their (indirect) utility function

$$U_{ijt}(p;\boldsymbol{\beta}) \quad \equiv \quad \beta_i^p p_{ijt}(p) + \beta_i^y y_{ij(t-1)} + x_{ij}'\beta^x + \xi_j + \epsilon_{ijt}^d \tag{1}$$

where $p = \mathbf{p}_t$ is the vector of plan base premiums set by all insurers in each market in year t, $p_{ijt}(p)$ is household i's premium for plan j in year t, $y_{ij(t-1)}$ indicates whether household i chose plan j in the previous year, x_{ij} is a vector of observed product characteristics including the plan AV, ξ_j is a vector of unobserved product characteristics, and ϵ_{ijt}^d is an error term. We allow the household's premium parameter $\beta_i^p = \beta^p + w_{it}'\phi$ to vary with household characteristics w_{it} and the household's inertia parameter $\beta_i^y = \beta^y + x_{ij}'\kappa + w_{it}'\nu$ to vary with household and product characteristics. Premium subsidies reduce the household's premium $p_{ijt}(p)$ as discussed below. CSRs increase the AV of silver plans in equation (1). The utility of the outside option $U_{i0t} = \beta_i^p \rho_{it} + \epsilon_{i0t}$, where ρ_{it} is the household's penalty for not purchasing insurance in year t.

3.1.1 Calculating Household Premiums

The household's premium $p_{ijt}(p)$ is calculated as

$$p_{ijt}(p) \quad = \quad \max\left\{ \underbrace{\sigma_{it}p_{jmt}}_{\substack{\text{full}\\\text{premium}}} - \underbrace{\max\{\sigma_{it}p_{bmt} - \zeta_{it}, 0\}}_{\text{premium subsidy}}, 0 \right\} \tag{2}$$

where σ_{it} is the household's rating factor, p_{jmt} is the base premium of plan j in market m and year t, p_{bmt} is the base premium of the benchmark plan, and ζ_{it} is the household's income contribution cap. The product of the rating factor and the plan's base premium equals the household's full or unsubsidized premium.

Household rating factors are subject to the ACA's "modified community rating" regulations. California insurers cannot use health status to rate plan premiums and are only permitted to use age and geographic residence of the household's members.[3] Figure 3a compares the age rating curve in effect between 2014-2017[4] with average cost differences by age and gender (Yamamoto, 2013). Insurers are able to charge a 64-year-old up to 3 times as much as a 21-year-old (i.e., the age rating factor in Figure 3a is 3 for a 64-year-old and 1 for a 21-year-old). However, Figure 3a indicates that 64-year-old females cost insurers an average of 4 times as much as 21-year-old females and 64-year-old males cost insurers an average of 6 times as much as 21-year-old males. Insurers therefore undercharge older adults (particularly females) and must overcharge younger adults (particularly males) relative to their expected cost, creating the potential for adverse selection. Females tend to have higher medical costs during their child-bearing years, whereas males have higher medical costs over age 60. Figure 3b shows the partition of California's 58 counties into 19 rating areas. An insurer's premium must be the same for all consumers of the same age within a rating area.

Premium subsidies are calculated as the difference the household's unsubsidized premium for the benchmark plan ($\sigma_{it} p_{bmt}$) and the household's income contribution cap ζ_{it} as specified by the ACA. The ACA's premium subsidy is endogenous because it depends on the benchmark plan premium. The ACA defines the benchmark plan as the second-cheapest silver plan available to the household. The benchmark plan varies across households because of heterogeneous firm entry across markets. The income contribution cap ranged from 2% of annual income for consumers earning 100% of the federal poverty level (FPL) and 9.5% of annual income for consumers earning 400% of FPL in 2014. The contribution caps were set initially by the ACA and are updated annually by the Internal Revenue Service (IRS). Because the ACA's subsidy formula uses the second-

[3]The ACA also permits rating by tobacco usage, but California prohibits tobacco rating.

[4]The age rating curve for 2018 used slightly higher age rating factors for children under 21

Figure 3: Modified Community Rating in the California Exchange

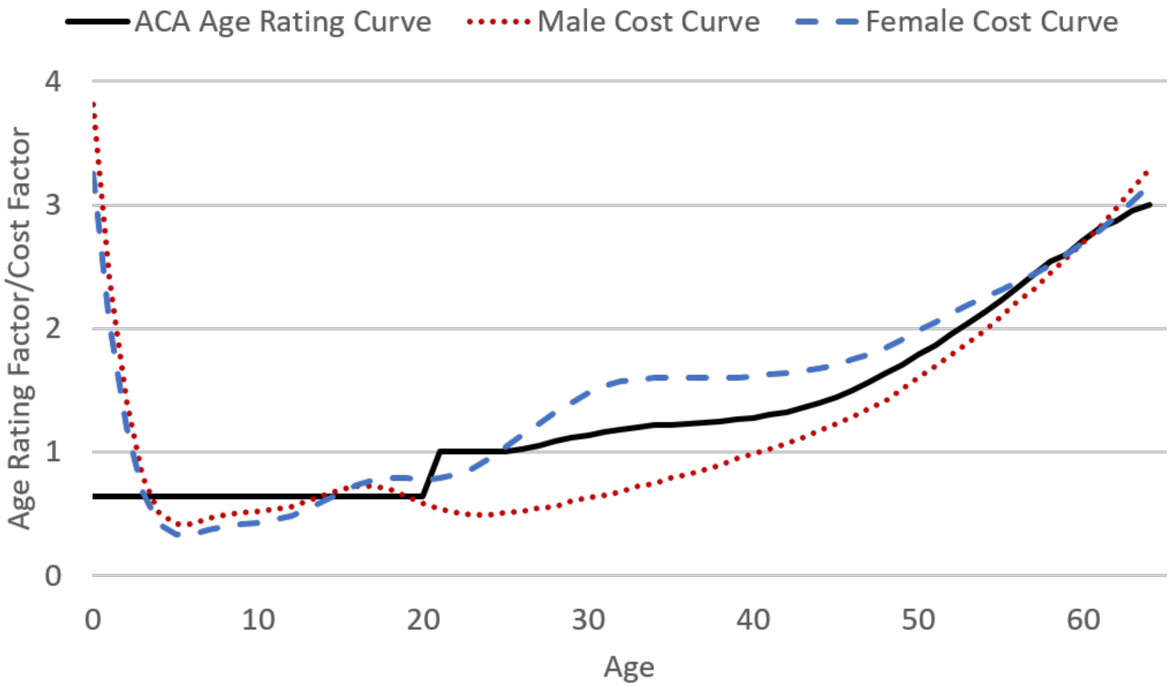

(a) ACA Age Rating Curve vs. Observed Age Cost Curves By Gender

(b) Premium Rating Regions in California

Notes: Panel (a) compares the ACA's age rating curve with the observed age cost curves by age and gender (Yamamoto, 2013). By design, a 21-year-old is assigned a rating factor of 1 and a 64-year-old is assigned a rating factor of 3. A 64-year-old can therefore be charged 3 times as much as a 21-year-old. Panel (b) shows the partition of California's 58 counties into 19 rating areas (Department of Managed Health Care, 2016).

cheapest silver plan premium as the benchmark, the premium subsidy may exceed the full premium of some bronze plans; the subsidy is reduced in these cases to ensure the premium is nonnegative. As discussed in the next section, this nonlinearity in the ACA's subsidy formula creates exogenous variation in relative premiums that we use to identify the premium parameter in equation 1.

3.1.2 Calculating Demand

We assume that the vector of error terms ϵ_i has the generalized extreme value distribution so that equation (1) is a nested logit model with two nests. The first nest contains all exchange plans and the second nest contains the outside option. This nest structure captures the primary substitution channel between silver plans (which must be selected to receive CSRs) and the outside option. Under the assumption that ϵ_i has the generalized extreme value distribution, the household choice probabilities are

$$q_{ijt}(p;\boldsymbol{\beta}) \quad = \quad \frac{e^{V_{ijt}(p;\boldsymbol{\beta})/\lambda}\left(\sum_j e^{V_{ijt}(p;\boldsymbol{\beta})/\lambda}\right)^{\lambda-1}}{1+\left(\sum_j e^{V_{ijt}(p;\boldsymbol{\beta})/\lambda}\right)^{\lambda}} \tag{3}$$

where $V_{ijt}(p;\boldsymbol{\beta}) \equiv \beta_i^p p_{ijt}(p)+\beta_i^y y_{ij(t-1)}+x_{ij}'\beta^x+\xi_j$ and λ is the nesting parameter. The household choice probabilities in equation (3) converge to the standard logit choice probabilities when $\lambda \to 1$.

The sensitivity of a subsidized consumer's demand $q_{ijt}(p)$ to a premium change is

$$\frac{\partial q_{ikt}(p)}{\partial p_{jmt}} \quad = \quad \sum_{l\in J_{mt}} \frac{\partial q_{ikt}(p)}{\partial p_{ilt}(p)}\frac{\partial p_{ilt}(p)}{\partial p_{jmt}}$$

for all plans j, k, where J_{mt} is the set of available plans. If the subsidy does not exceed the full premium, then the sensitivity of the consumer's premium to the insurer's base premium is

$$\frac{\partial p_{ilt}(p)}{\partial p_{jmt}} = \begin{cases} 0 & l=j, j=b \\ \sigma_{it} & l=j, j\neq b \\ -\sigma_{it} & l\neq j, j=b \\ 0 & l\neq j, j\neq b \end{cases} \tag{4}$$

An increase in a plan's base premium results in consumers paying more for that plan, unless it is the benchmark plan. A small increase in the benchmark plan base premium increases the subsidy by the same amount. Hence, the consumer's contribution to the benchmark plan premium remains constant, but the larger subsidy reduces what consumers pay for all other plans. Modeling this endogenous subsidy design poses substantial computational issues for estimation and simulation. We model the ACA's endogenous subsidy because of the key role premium subsidies play in determining which economic agents assume the cost of learning.

3.2 Firm Premium-Setting

Assume firms are risk-neutral and maximize expected profit. A firm sets the vector of base premiums for the plans that it sells to maximize

$$\begin{aligned} \pi_{ft}(p;\boldsymbol{\theta}) &= R_{ft}(p;\boldsymbol{\beta}) - C_{ft}(p;\boldsymbol{\theta}) + RA_{ft}(p;\boldsymbol{\theta}) + RI_{ft}(p;\boldsymbol{\theta}) - V_{ft}(p;\boldsymbol{\beta}) - FC_{ft} \\ &= R_{ft}(p;\boldsymbol{\beta}) - (1-\iota_{ft})C_{ft}(p;\boldsymbol{\theta}) + RA_{ft}(p;\boldsymbol{\theta}) - V_{ft}(p;\boldsymbol{\beta}) - FC_{ft} \end{aligned} \tag{5}$$

where $R_{ft}(\cdot)$ is total premium revenue, $C_{ft}(\cdot)$ is total claims, $RA_{ft}(\cdot)$ is risk adjustment received, $RI_{ft}(\cdot)$ is reinsurance received, $V_{ft}(\cdot)$ is variable administrative cost (e.g., commissions or fees), FC_{ft} is fixed cost, and ι_{ft} indicates the AV of the reinsurance contract (i.e., the expected percentage of claims paid by the reinsurer). The model parameters $\boldsymbol{\theta} \equiv (\boldsymbol{\beta}, \boldsymbol{\gamma}, \boldsymbol{\mu}, \boldsymbol{\eta})$, where $\boldsymbol{\beta}$ are the demand parameters (as defined above), $\boldsymbol{\gamma}$ are the risk score parameters, $\boldsymbol{\mu}$ are the average claims parameters, and $\boldsymbol{\eta}$ are the predicted cost parameters. As discussed above, the risk corridor program makes a positive monotonic transformation of firm profit and hence does not affect the optimal solution, assuming firms are risk-neutral and maximize expected profit. We also ignore MLR constraints because the empirical evidence suggests that they are not binding. The next two subsections discuss calculation of the risk adjustment transfer and the model equilibrium.

3.2.1 Calculating Risk Adjustment Transfers

We now discuss the calculation of risk adjustment transfers under the ACA. Under the ACA's single risk pool provisions, risk adjustment occurs at the state level for all metal plans (i.e., platinum, gold, silver, and bronze) in the individual market. Catastrophic plans have a separate risk adjustment pool.

Pope et al. (2014) derive the ACA risk adjustment transfer formula.[5] In our notation, the average risk adjustment transfer $ra_{jmt}(p)$ is

$$ra_{jmt}(p) = \left(\frac{r_{jmt}(p) \sum_{m \in M, l \in J_{mt}} q_{lmt}(p)}{\sum_{m \in M, l \in J_{mt}} r_{jmt}(p) q_{lmt}(p)} - \frac{h_j \sum_{m \in M, l \in J_{mt}} q_{lmt}(p)}{\sum_{m \in M, l \in J_{mt}} h_l q_{lmt}(p)} \right) \overline{p}$$

where $r_{jmt}(p)$ is the plan risk score, h_j is an exogenous expected utilization factor set by regulation that accounts for the plan AV and any associated moral hazard, and $\overline{p}$ is the weighted average premium in the market. The plan's total risk adjustment transfer $RA_{jmt}(p)$ equals

$$RA_{jmt}(p; \boldsymbol{\theta}) \;\; = \;\; ra_{jmt}(p) q_{jmt}(p) \;\; = \;\; \left[rs_{jmt}(p) - us_{jmt}(p) \right] R_t(p; \boldsymbol{\beta}) \tag{6}$$

where $rs_{jmt}(p)$ is the plan's "risk share" of total claims, $us_{jmt}(p)$ is the plan's "utilization share" of total claims, and $R_t(p) = \sum_f R_{ft}(p; \boldsymbol{\beta})$ is total premium revenue across all plans. The total risk adjustment transfer for the firm equals the sum of the risk adjustment transfers for the plans that it sells (i.e., $RA_{ft}(p) = \sum_{m \in M, j \in J_{fmt}} RA_{jmt}(p)$, where J_{fmt} is the set of all plans offered by firm f in market m and year t).

The plan's risk share includes the combined effects of adverse selection, moral hazard, and plan AV, whereas the plan's utilization share only includes the effects of moral hazard and plan AV. Thus, the difference between the risk share and utilization share captures the plan's relative risk due to adverse selection only. The risk share equals

$$rs_{jmt}(p) = \frac{r_{jmt}(p) q_{jmt}(p)}{\sum_{m \in M, l \in J_{mt}} r_{lmt}(p) q_{lmt}(p)}$$

where J_{mt} is the set of all plans offered in market m and year t, and the plan risk score $r_{jmt}(p)$ is a

[5] We start with Pope et al. (2014)'s transfer formula as derived in their first appendix.

function of enrollee characteristics and the plan AV. We do not directly observe plan risk scores in the insurer rate filing data. However, we observe all other variables in formula (6) including each plan's risk adjustment transfer and can therefore solve formula (6) for the plan risk scores. The plan's utilization share $us_{jmt}(p)$ equals

$$us_{jmt}(p) = \frac{h_j q_{jmt}(p)}{\sum_{m \in M, l \in J_{mt}} h_l q_{lmt}(p)}$$

If the risk share exceeds the utilization share, then the plan has high risk relative to its expected utilization and receives a risk adjustment transfer. If the risk share is less then the utilization share, then the plan has low risk relative to its expected utilization and pays a risk adjustment transfer.

3.2.2 Equilibrium

Now we find the first-order necessary conditions for a Nash equilibrium. Differentiating equation (5) yields the first-order conditions

$$MR_{jmt}(p;\boldsymbol{\beta}) + MRA_{jmt}(p;\boldsymbol{\theta}) \quad = \quad (1 - \iota_{ft})MC_{jmt}(p;\boldsymbol{\theta}) + MV_{jmt}(p;\boldsymbol{\beta}) \tag{7}$$

for all markets m in which plan j is offered by the firm in year t. We define marginal revenue $MR_{jmt}(p;\boldsymbol{\beta}) \equiv \frac{\partial R_{ft}(p;\boldsymbol{\beta})}{\partial q_{jmt}(p;\boldsymbol{\beta})}$, marginal claims $MC_{jmt}(p;\boldsymbol{\theta}) \equiv \frac{\partial C_{ft}(p;\boldsymbol{\theta})}{\partial q_{jmt}(p;\boldsymbol{\beta})}$, marginal risk adjustment transfer $MRA_{jmt}(p;\boldsymbol{\theta}) \equiv \frac{\partial RA_{ft}(p;\boldsymbol{\theta})}{\partial q_{jmt}(p;\boldsymbol{\beta})}$, and marginal variable administrative cost $MV_{jmt}(p;\boldsymbol{\beta}) \equiv \frac{\partial V_{ft}(p;\boldsymbol{\beta})}{\partial q_{jmt}(p;\boldsymbol{\beta})}$.

Appendix A shows that every variable in equations (5) and (7) can be written in terms of three estimable variables: (1) household choice probabilities $q_{ijt}(p;\boldsymbol{\beta})$; (2) plan risk scores $r_{jmt}(p;\boldsymbol{\theta})$; and (3) average claims $c_{jmt}(p;\boldsymbol{\theta})$. Household choice probabilities are computed using equation (3). We calculate plan risk scores as a function of observable enrollee characteristics and the plan generosity using the estimating equation

$$\ln r_{jmt}(p;\boldsymbol{\theta}) = \sum_{d\in D}\gamma^d s_{djmt}(p;\boldsymbol{\beta}) + MT_j'\gamma^{MT} + \epsilon_{jmt}^r \tag{8}$$

The predicted demographic share $s_{djmt}(\cdot)$ is the share of plan j's enrollment in market m and year t with demographic characteristic d, MT_j is a vector metal tier fixed effects, ϵ_{jmt}^r is an error term, and the vector of risk score parameters $\boldsymbol{\gamma} = (\gamma^d, \gamma^{MT}, \gamma^n)$. The demographic shares are computed by aggregating the household choice probabilities. We calculate plan average claims as a function of the plan risk score using the estimating equation

$$\ln c_{jmt}(p;\boldsymbol{\theta}) = \mu^r \ln r_{jmt}(p;\boldsymbol{\theta}) + x_j'\mu^x + \mu^l l_t + n_m'\mu^n + \epsilon_{jmt}^c \tag{9}$$

where $r_{jmt}(\cdot)$ is the predicted risk score computed using equation (8), x_j are product characteristics (not including plan AV), l_t is a linear trend, n_m' are market fixed effects, ϵ_{jmt}^c is an error term, and $\boldsymbol{\mu} = (\mu^r, \mu^x, \mu^l, \mu^n)$ are the claims parameters. Equation 16 in Appendix A provides a formula for how average claims respond to a change in premiums. If a plan is adversely selected, then $\frac{\partial c_{jmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}} > 0$.

4 Estimation

In this section, we explain how we use the generalized method of moments (GMM) to estimate the parameter vector $\boldsymbol{\theta}$. We first review a standard approach in the IO literature used by Saltzman (2021) that assumes firms have full information. We then relax the full information assumption to accommodate firm learning.

4.1 Standard Approach

In the standard approach, the econometrician pools data from all years to estimate the demand parameters $\boldsymbol{\beta}$, the risk score parameters $\boldsymbol{\gamma}$, and the average claims parameters $\boldsymbol{\mu}$. The predicted cost parameters $\boldsymbol{\eta}$ that determine how the firm forecasts cost from its past experience are not estimated. This is because the econometrician assumes firms have full knowledge and do not need to forecast future costs. Saltzman (2021) estimates $\boldsymbol{\beta}$, $\boldsymbol{\gamma}$, and $\boldsymbol{\mu}$ by creating four sets of moment conditions: (1) demand moments that match observed choices and predicted household choice probabilities; (2) risk score moments that match observed and predicted risk scores; (3) average claims moments that match observed and predicted average claims; and (4) the first-order conditions for profit maximization in equation (7). Denote N^{IJT} as the number of plans available to all households in all years, N^{JMT} as the number of plans available in all markets and years, and N_{jt}^{M} as the number of markets where plan j is offered in year t. Let χ_{ijt} be an indicator of whether household i chose plan j at time t, r_{jmt} be the observed plan risk score, and c_{jmt} be the observed plan average claims. Define the risk score covariates $\mathbf{z}^{r}_{jmt}(p;\boldsymbol{\beta}) \equiv (s_{djmt}(p,\boldsymbol{\beta}), MT_j)$ and the average claims covariates $\mathbf{z}^{c}_{jmt}(p;\boldsymbol{\theta}) \equiv (\ln r_{jmt}(p;\boldsymbol{\theta}), x_j, u_t, n_m)$. The moment conditions are

$$
\begin{aligned}
\frac{1}{N^{IJT}} \sum_{i\in I, j\in J, t\in T} \frac{\chi_{ijt}\partial \ln q_{ijt}(p;\boldsymbol{\beta})}{\partial \boldsymbol{\beta}} &= \mathbf{0} \\
\frac{1}{N^{JMT}} \sum_{j\in J, m\in M, t\in T} \mathbf{z}^{r}_{jmt}(p;\boldsymbol{\theta}) \left(\ln r_{jmt} - \boldsymbol{\gamma}'\mathbf{z}^{r}_{jmt}(p;\boldsymbol{\theta})\right) &= \mathbf{0} \\
\frac{1}{N^{JMT}} \sum_{j\in J, m\in M, t\in T} \mathbf{z}^{c}_{jmt}(p;\boldsymbol{\theta}) \left(\ln c_{jmt} - \boldsymbol{\mu}'\mathbf{z}^{c}_{jmt}(p;\boldsymbol{\theta})\right) &= \mathbf{0} \\
\frac{1}{N^{M}_{jt}} \sum_{m\in M} g_{jmt}(p;\boldsymbol{\theta}) &= 0, \qquad \forall j \in J_t, t\in T
\end{aligned} \tag{10}
$$

where the first-order condition values

$$
g_{jmt}(p;\boldsymbol{\theta}) \equiv MR_{jmt}(p;\boldsymbol{\beta}) - (1-\iota_{ft})MC_{jmt}(p;\boldsymbol{\theta}) + MRA_{jmt}(p;\boldsymbol{\theta}) - MV_{ft}(p;\boldsymbol{\beta})
$$

Because model (10) over-identifies the model parameters, we use two-step feasible GMM to find the values of $\boldsymbol{\theta}$ that minimize the GMM objective $[\mathbf{m}(\boldsymbol{\theta})]'\mathbf{W}^{-1}[\mathbf{m}(\boldsymbol{\theta})]$, where $\mathbf{m}(\boldsymbol{\theta})$ is the vector of moment values in model (10) and the optimal weight matrix $\mathbf{W}$ is a consistent estimate of the variance-covariance matrix of the moment values.

The primary estimation challenge is to identify the effect of premiums on household choices (i.e., the parameter β_i^p). Two primary sources of exogenous variation are used to identify the premium parameter β_i^p, including: (1) exogenous variation in absolute premiums (i.e., relative to the outside option) that results from the phasing-in of the mandate penalty between 2014 and 2016 and elimination of the penalty in 2019; (2) exogenous variation in relative premiums (i.e., between plans) that results from kinks in the household premium formula (2). As discussed above, some bronze plans may be "free" to low-income consumers if the subsidy exceeds the full premium (i.e., the second-cheapest silver plan available to the consumer may exceed the premium of some bronze plans). The set of free plans varies by market, time, and household characteristics, including age, income, and household composition. While we believe these sources of exogenous variation cleanly identify the premium parameter, we also estimate equation (1) with insurer-market fixed effects. These fixed effects control for unobservables at the insurer's discretion such as provider networks and formularies that could be correlated with premiums. Ho and Pakes (2014) and Tebaldi (2020) follow a similar approach. Our estimates are similar when including insurer-market fixed effects.

Another identification challenge is that we do not observe patient medical conditions that are used to predict plan risk scores. Estimates of the risk score parameter γ^d may be biased by omitting patient medical conditions. We address this potential source of bias by computing predicted demographic shares using the estimated consumer-level choice probabilities from equation (3) instead of the observed demographic shares, which may be endogenous. The identifying assumption is that the predicted demographic shares are based on exogenous determinants of consumer plan demand. Choice model (3) can be interpreted as the first-stage of an IV regression for computing unbiased estimates of plan risk scores. A similar empirical strategy is widely used in the hospital choice literature to compute measures of hospital market concentration (e.g., Kessler and McClellan (2000)).

A third identification challenge is to compute an unbiased estimate of the average claims parameter μ^r. We compute predicted plan risk scores using equation (8) instead of the observed plan risk scores, which may be endogenous. Enrollee characteristics should affect average claims through the plan risk score only and not directly affect average claims. This may not be the case if the ACA risk score is an imprecise measure of plan claims risk.

4.2 Adaptive Learning

Now we relax the full information assumption to model how firms learn about the parameter vector $\boldsymbol{\theta}$. Firm participation in a new insurance market involves two principal sources of uncertainty: demand and cost uncertainty. Demand uncertainty arises when firms do not know consumer preferences. Cost uncertainty arises when firms do not know the cost of insuring their enrollees. A defining feature of a selection market such as insurance is that demand and cost uncertainty are correlated. Our model accommodates this correlation in estimating equations (8) and (9).

We assume that in any plan year t, firms use data from years $2014, \ldots, t-1$ to estimate $\boldsymbol{\theta}$ and determine its predicted average cost $a_{jmt}(p; \boldsymbol{\theta})$ for year t. To capture the key features of the firm's forecast of average cost, we use the estimating equation

$$a_{jmt}(p; \boldsymbol{\theta}) \quad = \quad \left[(1 - \iota_t) \left(c_{jm(t-1)}(p; \boldsymbol{\theta}) + \mu^l\right) + RA_{jm(t-1)}(p; \boldsymbol{\theta})\right] x_j' \boldsymbol{\eta} + \epsilon^a_{jmt} \tag{11}$$

where $c_{jm(t-1)}(p; \boldsymbol{\theta})$ is predicted average claims in period $t-1$ using equation (9), μ^l is the long-run linear trend from equation (9), $RA_{jm(t-1)}(p; \boldsymbol{\theta})$ is the predicted risk adjustment transfer in period $t-1$ using equation (6), x_j are product characteristics, and ϵ^a_{jmt} is an error term. The vector of predicted cost parameters $\boldsymbol{\eta}$ represents deviations from the long-run linear trend. These deviations, which vary by firm and plan network type, may be the result of anticipated technological innovations, pharmaceutical introductions, or government policy changes.

Equation 11 captures the basic idea of how firms forecast cost by applying a trend to past experience. However, it is a simplification of the lengthy actuarial justification in the rate filings of how firms trend past experience to forecast future cost. To mitigate any error introduced by this simplification, we form moment conditions that match the observed predicted average cost a_{jmt} in the rate filings with predicted average cost $a_{jmt}(p;\boldsymbol{\theta})$ as calculated in equation 11.

Define the previous period set $T_t \equiv \{2014,\ldots,t-1\}$ and denote $\boldsymbol{\theta}_t$ as the firm's estimate of $\boldsymbol{\theta}$ at time t. Let N^{JM}_{ft} be the total number of plans across all markets sold by firm f at time t. To estimate $\boldsymbol{\theta}_t$, we create the following moment conditions:

$$
\begin{aligned}
\frac{1}{N^{IJT_t}} \sum_{i\in I, j\in J, \tau\in T_t} \frac{\chi_{ij\tau}\partial \ln q_{ij\tau}(p;\boldsymbol{\beta}_t)}{\partial \boldsymbol{\beta}_t} &= \mathbf{0} \\
\frac{1}{N^{JMT_t}} \sum_{j\in J, m\in M, \tau\in T_t} \mathbf{z}^r_{jm\tau}(p;\boldsymbol{\beta}_t)\left(\ln r_{jm\tau} - \boldsymbol{\gamma}_t'\mathbf{z}^r_{jm\tau}(p;\boldsymbol{\beta}_t)\right) &= \mathbf{0} \\
\frac{1}{N^{JMT_t}} \sum_{j\in J, m\in M, \tau\in T_t} \mathbf{z}^c_{jm\tau}(p;\boldsymbol{\theta}_t)\left(\ln c_{jm\tau} - \boldsymbol{\mu}_t'\mathbf{z}^c_{jm\tau}(p;\boldsymbol{\theta}_t)\right) &= \mathbf{0} \\
\frac{1}{N^{M}_{j\tau}} \sum_{m\in M} g_{jm\tau}(p;\boldsymbol{\theta}_t) &= 0, \qquad \forall j\in J_\tau, \tau\in T_t \\
\frac{1}{N^{M}_{jt}} \sum_{m\in M} x_j'\left(a_{jmt} - a_{jmt}(p;\boldsymbol{\theta}_t)\right) &= 0, \qquad \forall j\in J_{ft} \\
\frac{1}{N^{JM}_{ft}} \sum_{j\in J_{ft}, m\in M} x_j'\left(a_{jmt} - a_{jmt}(p;\boldsymbol{\theta}_t)\right) &= 0, \qquad \forall f\in F \\
\frac{1}{N^{M}_{jt}} \sum_{m\in M} \overline{g}_{jmt}(p;\boldsymbol{\theta}_t) &= 0, \qquad \forall j\in J_t
\end{aligned} \tag{12}
$$

The first four sets of moment conditions are similar to the moment conditions in the standard approach, except that we omit data from year t. The fifth set of moment conditions matches each plan's observed predicted average cost a_{jmt} as reported in the rate filings with the predicted average cost from the model $a_{jmt}(p;\boldsymbol{\theta}_t)$. The sixth set of moment conditions matches each firm's observed predicted average cost with the model predicted average cost. The final set of moment conditions are the first-order conditions for year t. We compute the first-order condition values $\overline{g}_{jmt}(p;\boldsymbol{\theta}_t)$ by applying the predicted cost parameters to average claims, risk adjustment, and reinsurance.

To maintain tractability, we do not model other potential sources of uncertainty that are less relevant for our setting. Our model does not allow for structural or strategic uncertainty that arises when firms have private information about their demand and cost primitives. In this market, firms have ample access to their competitors' rate filings and the regulatory rate review process occurs over several months, providing firms numerous opportunities to learn about their competitors' proposed rates. We also assume consumers are myopic and do not learn over time. As discussed in the data section, evidence of consumer learning appears to be minimal in our setting.

5 Parameter Estimates

Table I summarizes the adaptive learning parameter estimates $\widehat{\boldsymbol{\theta}}_t$ for $t \in \{2015, 2016, 2017, 2018\}$ and the standard approach estimates $\widehat{\boldsymbol{\theta}}$. Detailed parameter estimates are provided in Table A2 in Appendix D. For the most part, firms underestimated the key demand parameters in the ACA's initial years. Firms significantly underestimated the premium parameter in 2015 when only data from 2014 were available. Estimates of the premium parameter more than double in magnitude in 2016 and are close to the standard approach estimates by 2017. Figure 4 shows the mean own-premium elasticities and exchange coverage elasticities of demand implied by these parameter estimates. Firms initially overestimated the sensitivity of a plan's demand to its own premium, but underestimated the sensitivity of total exchange enrollment to a change in all exchange premiums. Firms also initially underestimated inertia (i.e., the previous choice parameter). The previous choice parameter estimate in 2016 was only two-thirds of the previous choice parameter estimate using the standard approach. The trend for plan generosity (i.e., the plan AV) is similar to the trend for the premium parameter estimate. Firms initially underestimated the effect of plan generosity, but were close to the standard approach estimates by 2017. We also find firms initially underestimated the nesting parameter, expecting less substitution between the exchange nest and outside option nest.[6]

Learning estimates of the supply-side parameters generally converge over time toward the standard approach estimates. Estimates of the silver, gold, and platinum parameters in the risk score

[6]Recall that our nested logit choice model converges to the standard logit as $\lambda \to 1$.

regression converge non-monotonically toward the standard approach estimates. As expected, platinum plans have the greatest exposure to claims risk. Estimates of the young adult (under age 55) share parameter are negative as expected, but quite volatile. The standard errors are large initially, but decline over time. Consumers of Hispanic origin have considerably less claims risk than other racial and ethnic groups, but firms substantially overestimate this effect initially. Firms initially underestimated the relationship between average claims and the plan risk score; the estimated parameter increases from 0.866 using data available in 2014 to 1.024 using all data. Estimates of the HMO parameter in the average claims regression are highly imprecise in 2015 and 2016. The estimated HMO parameter has the wrong (positive) sign in 2018. Using the data available in 2017 and all data, HMO plans are predicted to have statistically significant lower claims as expected.

The learning estimates for 2015, when only 2014 data were available, have a significant limitation. Although inertia was a key consideration for firms in determining premiums for 2015, it was not possible to estimate inertia with only a single year of data. In our 2015 specification, we omit the previous choice variable, implicitly assuming the previous choice parameter is zero. This implicit assumption may be a source omitted variable bias, resulting in a downward-biased estimate of the premium parameter. Firms may have accounted for inertia using alternative methods, such as estimating inertia using data from the pre-ACA individual market. Because of this limitation, we use the 2016 learning estimates instead of the 2015 learning estimates to conduct policy simulations below. The 2016 learning estimates are more similar to the standard approach estimates than the 2015 learning estimates, and therefore, we expect this modeling decision to provide a conservative estimate of the impact of learning.

To evaluate our estimated model, we compute predicted costs for each firm by year using our model parameter estimates and compare our estimates to the firms' predicted costs in the rate filing data. Figure 5 summarizes this comparison by plotting the firms' cost prediction errors. The observed cost prediction errors are the same as in Figure 1. We create 95% confidence intervals around our model point estimates of the cost prediction error by taking 1,000,000 draws from each

Table I: Summary of Parameter Estimates

	$\widehat{\boldsymbol{\theta}}_{2015}$	$\widehat{\boldsymbol{\theta}}_{2016}$	$\widehat{\boldsymbol{\theta}}_{2017}$	$\widehat{\boldsymbol{\theta}}_{2018}$	$\widehat{\boldsymbol{\theta}}$
Demand Parameters ($\widehat{\boldsymbol{\beta}}_t$)					
Monthly Premium ($100)	-0.084^{***}	-0.199^{***}	-0.244^{***}	-0.261^{***}	-0.245^{***}
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
Previous Choice		0.270^{***}	0.355^{***}	0.407^{***}	0.419^{***}
		(0.014)	(0.013)	(0.011)	(0.010)
AV	0.300^{***}	0.650^{***}	0.777^{***}	0.834^{***}	0.806^{***}
	(0.022)	(0.017)	(0.013)	(0.011)	(0.010)
Nesting Parameter	0.034^{***}	0.086^{***}	0.111^{***}	0.126^{***}	0.132^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Risk Score Parameters ($\widehat{\gamma}_t$)					
Silver	0.652^{***}	0.492^{***}	0.656^{***}	0.605^{***}	0.569^{***}
	(0.064)	(0.043)	(0.045)	(0.033)	(0.028)
Gold	0.734^{***}	0.655^{***}	0.842^{***}	0.794^{***}	0.785^{***}
	(0.149)	(0.093)	(0.092)	(0.065)	(0.050)
Platinum	1.033^{***}	0.978^{***}	1.179^{***}	1.136^{***}	1.125^{***}
	(0.150)	(0.084)	(0.086)	(0.062)	(0.049)
Share Ages 18 to 54	-1.197	-0.698	-1.292	-1.525^{***}	-1.136^{***}
	(1.212)	(0.837)	(0.838)	(0.561)	(0.412)
Share Hispanic	-2.000^{***}	-1.251^{***}	-0.903^{**}	-0.938^{***}	-1.042^{***}
	(0.651)	(0.413)	(0.406)	(0.269)	(0.213)
Average Claims Parameters ($\widehat{\boldsymbol{\mu}}_t$)					
Log Risk Score	0.866^{***}	0.939^{***}	0.978^{***}	0.991^{***}	1.024^{***}
	(0.033)	(0.008)	(0.005)	(0.004)	(0.005)
HMO	-0.174	0.027	-0.114^{***}	0.092^{***}	-0.138^{***}
	(0.166)	(0.034)	(0.014)	(0.012)	(0.008)

Notes: Table summarizes the adaptive learning parameter estimates $\widehat{\boldsymbol{\theta}}_t$ for $t \in \{2015, 2016, 2017, 2018\}$ and the standard approach estimates $\widehat{\boldsymbol{\theta}}$. Robust standard errors are in parentheses (*** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level). We compute the household-specific monthly premium and previous choice parameters for each household using the demographic interaction terms and report an average across all households in this table. The raw parameter estimates are available in Table A2.

Figure 4: Estimated Premium Elasticities of Demand By Year

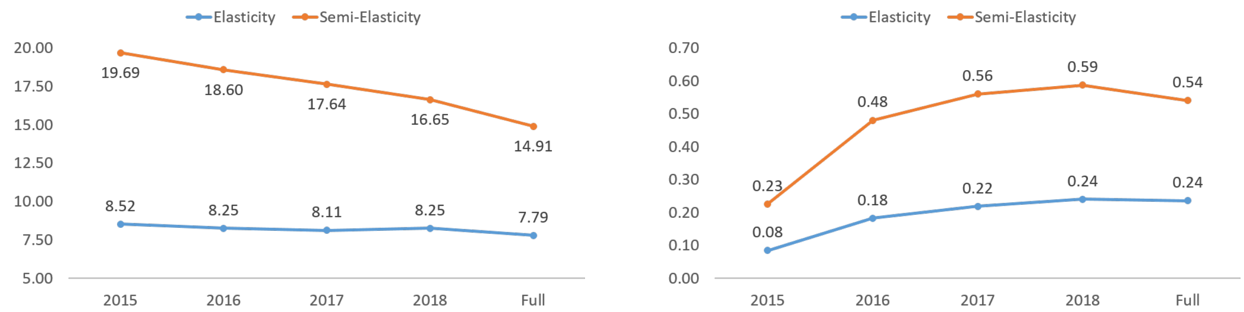

(a) Mean Own-Premium Elasticities

(b) Mean Coverage Elasticities

Notes: Figure shows the premium elasticities of demand implied by the learning parameter estimates $\widehat{\theta}_t$ for $t \in \{2015, 2016, 2017, 2018\}$ and the standard approach (or full information) estimates $\widehat{\theta}$. Panel (a) shows how a plan's demand responds to a change in its own premium. Panel (b) shows how total exchange enrollment responds to a change in all exchange premiums. Semi-elasticities are calculated for a \$100 change in annual premiums.

estimated parameter distribution[7], recomputing the cost prediction error, and finding the 2.5%- and 97.5%-quantiles. Our estimated learning model is effective in capturing firm learning behavior. The largest difference in the cost prediction error between the data and the model is \$6 or about 1.5% of the firm's average predicted cost. In most cases, the difference is less than \$2 or about 0.5% of the firm's average predicted cost. The 95% confidence intervals tighten over time, reflecting decreasing firm uncertainty as more data become available.

6 Impact of Assuming Full Information

6.1 Simulation Methodology

In this section, we simulate the impact of assuming full information on the estimated model equilibrium. We do this by (1) replacing the firms' learning estimates $\widehat{\theta}_t$ for $t \in \{2015, 2016, 2017, 2018\}$ with the 'full information parameters," which we assume are the estimates $\widehat{\theta}$ obtained using all

[7]We assume the parameter distribution is normal with mean equal to the parameter point estimate and standard deviation equal to the standard error in Table I.

Figure 5: Cost Prediction Error By Year and Insurer

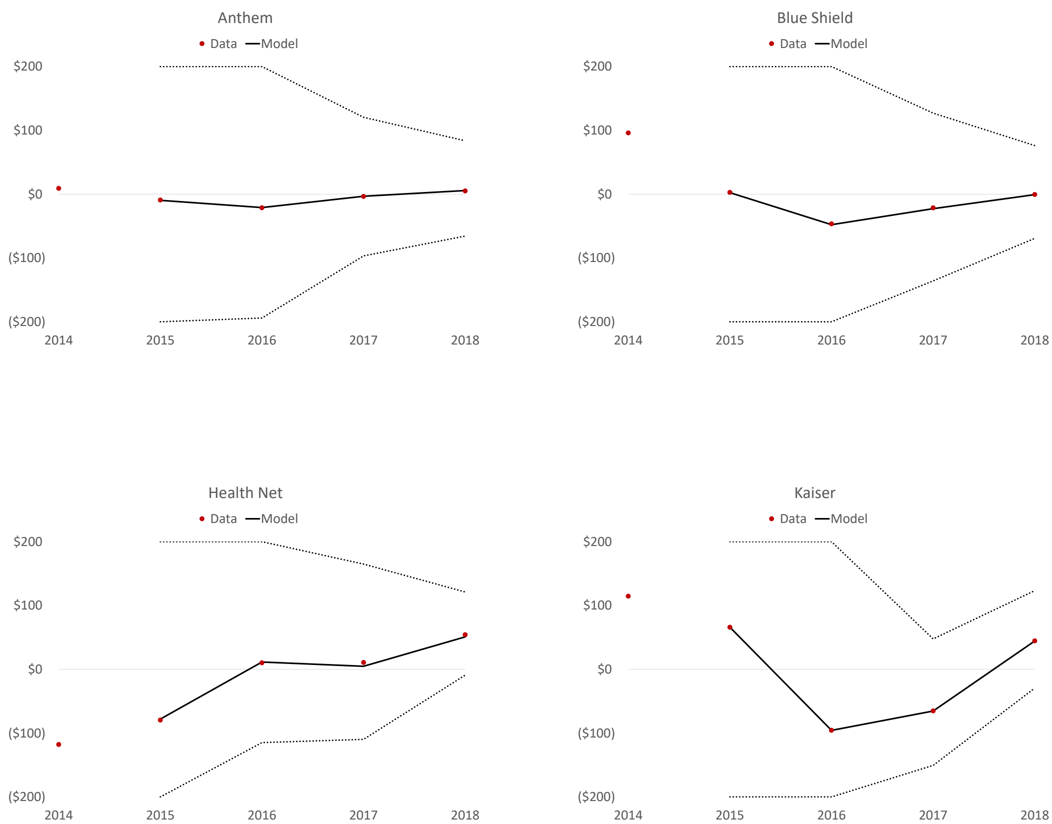

Notes: Cost prediction error for the data equals predicted average cost as reported in the rate filing data minus realized average cost as reported in the rate filing data (i.e., the same as in Figure 1). Cost prediction error for the model equals predicted average cost estimated by the model minus realized average cost as reported in the rate filing data. The dashed lines indicate the lower and upper limit of the 95% confidence interval for the cost prediction error point estimates from the model. The confidence interval limits are capped at −$200 and $200 for presentational purposes.

available data (i.e., the last column in Table I); (2) solving for the new vector premiums that satisfy the firms' first-order conditions in equation (7); and (3) computing several measures of the new equilibrium outcome, including premiums, enrollment, and social welfare.

We compute total social welfare in year t as

$$SW_t \;=\; CS_t + \pi_t - \delta GS_t$$

where CS_t is total consumer surplus, π_t is total firm profit, and GS_t is net government spending adjusted by the deadweight loss of taxation δ that results from distortions in prices and consumer behavior. The deadweight loss of taxation corresponds to the additional compensation consumers need in order to obtain their original utility levels (i.e., before government spending on the premium subsidies, CSRs, etc.) at the distorted prices (Hausman and Poterba, 1987). Following Hausman and Poterba (1987) and Decarolis et al. (2020), we multiply government spending by a factor of 1.3 to account for the deadweight loss of taxation. We compute total consumer surplus

$$CS_t \;=\; -\sum_{i\in I}\frac{1}{\beta_i^p}\ln\left(\sum_{j\in J}\exp\left(V_{ijt}(p;\boldsymbol{\beta}_t)/\lambda\right)^{\lambda}+\exp\left(\beta_{it}^p\rho_{it}\right)\right)+\sum_{j\in J}\left[q_{ijt}(p)*\frac{\beta_{ij}^y * y_{ij(t-1)}}{\beta_i^p}\right] \qquad (13)$$

where the first term of equation (13) is the standard nested logit formula for consumer surplus and the second term "corrects" the first term to remove gains in welfare that result from inertia. Total firm profit is $\pi_t = \sum_{t\in T}\pi_{ft}(p;\boldsymbol{\theta}_t)$, where $\pi_{ft}(p;\boldsymbol{\theta}_t)$ is defined in equation (5). Net government spending GS_t equals the sum of spending on premium subsidies, CSRs, and uncompensated care for the uninsured minus revenue collected from the mandate penalty. Premium subsidy spending is the sum of subsidies received by each consumer in equation (2). Spending on CSRs is computed as

$$CSR_t \;=\; \sum_{i\in I, j\in J} s_j^g q_{ijt}(p)c_{jmt}(p;\boldsymbol{\theta}_t)$$

where s_j^g is the expected share of claims paid by the government for plan j.[8] We calculate spending

[8]Ignoring moral hazard, the government's expected outlay is $94-70=24\%$ of claims for the 94% CSR plan,

on uncompensated care by multiplying the number of uninsured that we estimate in each scenario by \$2,025, the estimated annual uncompensated care cost per uninsured[9], and a factor accounting for the change in the uninsured population's risk score. Penalty revenue collected by the government equals $\sum_{i \in I} q_{i0t}\rho_{it}$, where q_{i0t} is the household's probability of choosing the outside option.

6.2 Impact of Assuming Full Information

Detailed simulation results are presented in Table A3 of Appendix E. Figure 6 summarizes the impact of assuming full information over time on average subsidized premiums (Figure 6a), average unsubsidized premiums (Figure 6b), average unsubsidized premiums by metal tier (Figure 6c), annual per-capita social welfare (Figure 6d), total exchange enrollment (Figure 6e), and plan market share by metal tier (Figure 6f).

For all years in our data, using the parameters estimated with full information ($\widehat{\theta}$) instead of the parameters estimated with only the available information ($\widehat{\theta}_t$ for $t \in \{2015, 2016, 2017, 2018\}$) results in a generally more favorable conclusion about the model equilibrium. When the full information parameters are used instead of the available information parameters, average subsidized premiums are between 3.1% and 14.4% lower and average unsubsidized premiums are between 6.0% and 39.3% lower. Premiums declines are generally larger for the less generous plans (e.g., bronze and silver) and smaller for the more generous plans (e.g., gold and platinum); platinum premiums even increase in 2016 and 2017. The impact of assuming full information on total exchange enrollment is relatively small, but positive for all years. Despite minimal differences in total exchange enrollment, there are considerable shifts in enrollment across the metal tiers. Because gold

$87 - 70 = 17\%$ of claims for the 87% CSR plan, and $73 - 70 = 3\%$ of claims for the 73% CSR plan. To account for moral hazard, we follow Pope et al. (2014) and assume there is no moral hazard for consumers in the 73% plan, while consumers in the 87% and 94% plans increase consumption by 12%. Including moral hazard, the $s_j^g = 26.88\%$ for the 94% CSR plan, $s_j^g = 19.04\%$ for the 87% CSR plan, and $s_j^g = 3\%$ for the 73% CSR plan.

[9] We multiply the per-capita amount of medical costs that are paid on behalf of the nonelderly uninsured as estimated by Coughlin et al. (2014) by an inflation factor using data from the National Health Expenditure Accounts to adjust the estimates to the timeframe of this study (Centers for Medicare and Medicaid Services, 2018).

and platinum plans become more expensive relative to bronze and silver plans under full information, gold and platinum plan market share decreases between 2.5 and 4.9 percentage points.

The more favorable equilibrium market conditions predicted under full information result in annual per-capita social welfare estimates that are between \$592 higher for 2015 and \$156 higher for 2018. Assuming full information in the model overestimates annual total social welfare (i.e., the product of the per-capita welfare estimates and market size) by between \$1.38 billion and \$0.37 billion. Lower premiums under full information directly benefit consumers and indirectly benefit the government because the ACA's subsidies are price-linked. Gains in consumer and taxpayers' welfare are partially offset by a reduction in firm profit, which falls because of the decline in premiums.

Figure 6 allows us to assess the impact of assuming full information over time. In general, the impact is largest in 2015 and declines over time. This trend is particularly evident in Figure 6d. Assuming full information leads to an estimate of annual-per capita social welfare that is \$592 higher in 2015, but only \$156 higher in 2018. The welfare differences decline monotonically over the study time period. Figures 6c and 6f also indicate that the premium and enrollment distortions between bronze and platinum plans decline over time. Assuming full information increases the estimate of bronze market share by 4.0 percentage points in 2015 and 0.5 percentage points in 2018 and decreases the estimate of platinum market share by 3.1 percentage points in 2015 and 1.1 percentage points in 2018. Although the impact of assuming full information on average subsidized and unsubsidized premiums is largest in 2015, it is slightly larger in 2018 than in 2016 or 2017.

Figure 7 indicates that in 2016, the gains in consumer surplus from assuming firms have full information are heterogeneous across the exchange population. Lower-income households gain the least consumer surplus because they pay the least in premiums and are shielded from premium changes by the ACA's endogenous subsidy. Disadvantaged subpopulations that are more price sensitive also realize substantial gains in welfare. Average annual per-capita consumer surplus increases \$102 and \$78 for Black and Hispanic consumers, respectively. The consumer surplus

Figure 6: Impact of Assuming Full Information By Year

(a) Average Subsidized Premiums

(d) Change in Annual Per-Capita Social Welfare

(b) Average Unsubsidized Premiums

(e) Total Exchange Enrollment

(c) Change in Unsubsidized Premiums By Metal

(f) Change in Market Share By Metal

Notes: Figure shows the equilibrium impact of using the full information parameters instead of the learning parameters by year. Panel (a) shows the impact on average subsidized premiums, panel (b) shows the change in annual per-capita social welfare when switching to the full information parameters, panel (c) shows the impact on average unsubsidized premiums, panel (d) shows the impact on total exchange enrollment, panel (e) shows the change in average unsubsidized premiums by metal tier, and panel (f) shows the change in market share by metal tier.

Figure 7: Change in Average Annual Per-Capita Consumer Surplus by Household Characteristics

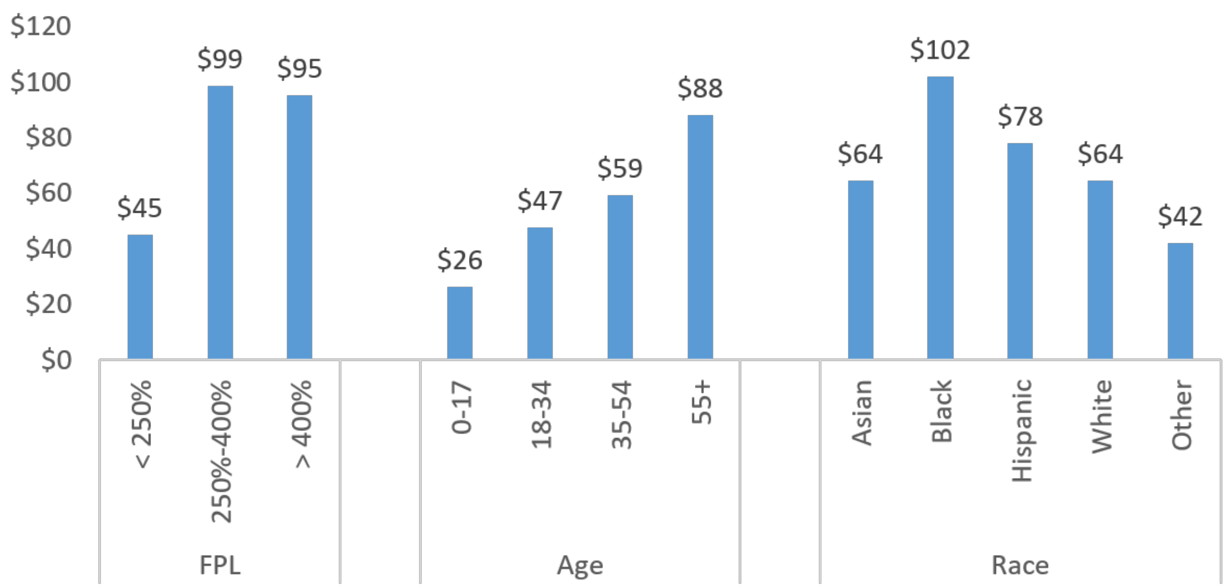

Notes: Figure shows the change in average annual per-capita consumer surplus that results from using the full information parameters instead of the 2016 learning parameters. Consumer surplus changes are shown by income measured as a percentage of the federal poverty level, age, and race/ethnicity.

gains are also increasing in age. Consumers over age 55 realize average annual per-capita consumer surplus gains of \$88, compared to only \$47 for young adults between ages 18 and 34.

6.3 Mechanisms

We now investigate whether certain parameters drive the impact of assuming full information. Figure 8 shows the progressive impact of switching from the 2016 learning parameters to the full information parameters on average subsidized premiums (Figure 8a), average unsubsidized premiums (Figure 8b), average unsubsidized premiums by metal tier (Figure 8c), annual per-capita social welfare (Figure 8d), total exchange enrollment (Figure 8e), and plan market share by metal tier (Figure 8f). Starting with the 2016 learning parameters in the far left of each panel, we first simulate replacing the 2016 cost parameters with the full information cost parameters from equations (9) and (11). Moving to the right of each panel, we then simulate switching to the full information risk score parameters from equation (8), the full information demand parameters in equation (1) (except for the premium and inertia parameters), the full information inertia parameters in equation (1), and finally the full information premium parameters in equation (1). The numbers reported on the far left and far right of each panel (labeled "2016 information" and "Full information", respectively) are the same as reported in Figure 6 for the year 2016. As discussed in the previous section, we use the year 2016 to conduct these simulations instead of the year 2015 because firms could not estimate inertia in 2015.

By almost every measure reported in Figure 8, replacing the 2016 cost parameters with the full information cost parameters has the largest impact. Average subsidized premiums decline 10.0% from \$128 to \$117 and average unsubsidized premiums decline 8.0% from \$403 to \$373. Substantial declines in bronze premiums result in bronze plan market share increasing by 3.2 percentage points. Annual per-capita social welfare increases \$261. These results indicate that firms' uncertainty of the cost parameters is an important driver of the effect of learning.

Replacing the 2016 premium parameters with the full information premium parameters is also a key mechanism for the impact of learning. Firms underestimated premium sensitivity in 2016 (Table I), resulting in excessively high markups. Switching to the full information premium parameters increases the estimate of premium sensitivity, reducing average subsidized premiums by 3.4% from \$121 to \$117 and average unsubsidized premiums by 8.1% from \$409 to \$378. Substantial declines in bronze premiums result in bronze plan market share increasing by 1.1 percentage points. Annual per-capita social welfare increases \$91.

Whereas using the full information cost and premium parameters reduces premiums, using the full information inertia parameters instead of the 2016 inertia parameters increases premiums. Higher premiums under full information are the result of firms initially underestimating inertia, a key source of firm market power (Saltzman et al., 2021). Replacing the 2016 inertia parameters with the full information inertia parameters increases average subsidized premiums by 4.9% from \$115 to \$121 and average unsubsidized premiums by 6.9% from \$381 to \$409. Significant increases in bronze premiums result in bronze plan market share decreasing by 1.2 percentage points. Annual per-capita social welfare decreases \$127, partially offsetting the welfare gains from using the full information cost and premium parameters.

Taken together, our results suggest that assuming full information results in a more favorable conclusion about the model equilibrium. Relative to estimation with only the available information, premiums are lower, total enrollment is higher, and social welfare is higher under full information. The difference between the available and full information results generally declines over time. Imprecise estimates of the cost and premium parameters are the key drivers of the more favorable conclusion about the model equilibrium under full information. The equilibrium conclusion would be even more favorable if firms had not underestimated inertia.

Figure 8: Determining Which Parameters Drive the Impact of Assuming Full Information

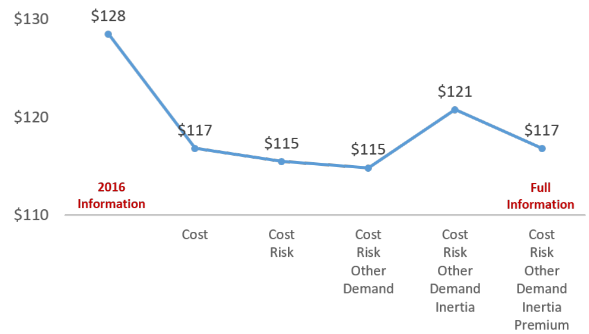

(a) Average Subsidized Premiums

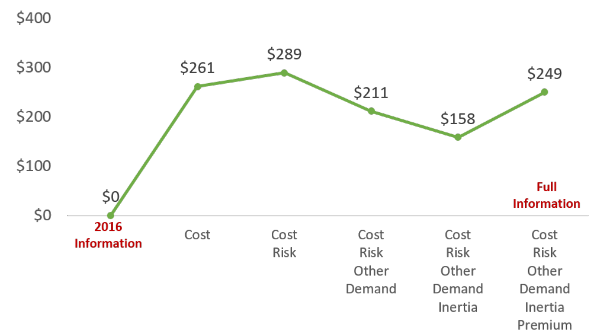

(d) Change in Annual Per-Capita Social Welfare

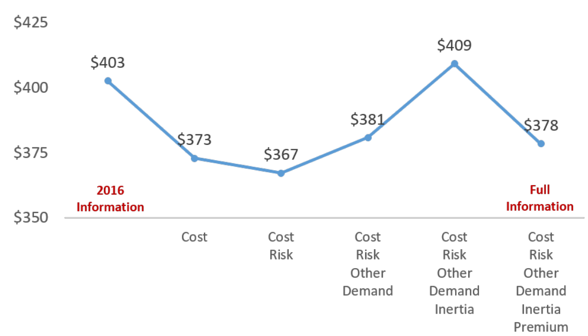

(b) Average Unsubsidized Premiums

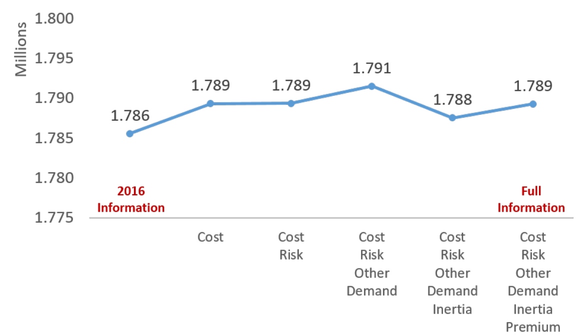

(e) Total Exchange Enrollment

(c) Change in Unsubsidized Premiums By Metal

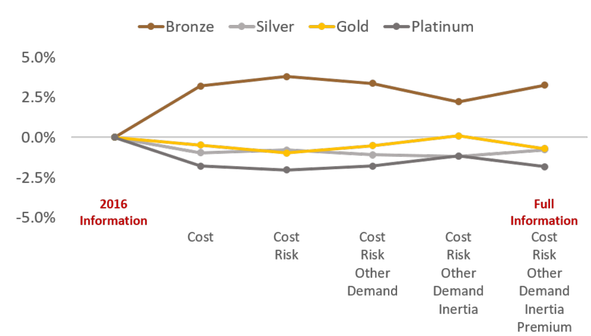

(f) Change in Market Share By Metal

Notes: Figure shows the equilibrium impact of using the full information parameters instead of the 2016 learning parameters. The x-axis of each graph indicates which parameters are known to the firm, including (1) the "cost" parameters from the average claims and predicted cost equations; (2) the "risk" parameters from the risk score equation; (3) all "other demand" parameters in the utility equation besides the inertia and premium parameters; (4) the "inertia" parameters in the utility equation; and (5) the "premium" parameters in the utility equation. The figures show the equilibrium impact on subsidized premiums (panel a), annual per-capita social welfare (panel b), unsubsidized premiums (panel c), total exchange enrollment (panel d), average unsubsidized premiums by metal tier (panel e), and market share by metal tier (panel f).

7 Policy Simulations

Conclusions about the efficacy of government regulations may depend on whether learning is modeled. We now study how firm knowledge affects the impact of key ACA regulations and market features, including community rating, inertia, risk adjustment, and the individual mandate. As before, we simulate these changes for 2016 instead of 2015 because firms could not estimate inertia in 2015. Figure 9 summarizes our results. Tables A5 and A6 in Appendix E provide detailed results.

Overall, our results indicate that policy assessments depend substantially on whether learning is modeled. Assuming full information can either lead to underestimating or overestimating the impact of an intervention, depending on which parameters are most relevant for the intervention. Modifying community rating rules has a larger impact when firms have full information because community rating directly affects absolute premiums (i.e., relative to the outside option) and firms initially underestimated exchange coverage elasticities (see Figure 4b). Eliminating inertia also has a stronger effect under full information because firms initially underestimated the inertia parameters. Conversely, eliminating risk adjustment has a smaller impact under full information because risk adjustment directly affects relative premiums (i.e., between exchange plans) and firms initially overestimated own-premium elasticities (see Figure 4a). Eliminating the individual mandate has little impact on the model equilibrium, regardless of whether learning is modeled, because of the penalty's small size relative to premiums and the ACA's robust price-linked subsidies. The following four subsections discuss in detail how learning interacts with each intervention.

7.1 Community Rating

The ACA's modified community rating rules allow firms to set premiums by age and geography. We simulate two changes to modified community rating: (1) relaxing restrictions on age and gender rating and (2) requiring firms to charge all consumers the same premium, regardless of age and geography (i.e., pure community rating). We simulate the relaxation of age-gender rating restrictions by replacing the ACA age rating factors with the male and female cost factors in Figure 3a and solving for the new vector of premiums that satisfy the firms' first-order conditions. We simulate pure community rating by setting all individual rating factors to 1 and solving for the new vector of premiums that satisfy the firms' first-order conditions.

For the most part, modifying community rating rules has a larger impact when firms have full information. The first-order effect of modifying community rules is to change the premiums of *all* plans charged to a given consumer by the same adjustment to the rating factor. Community rating rule modifications therefore affect absolute premiums paid by a given consumer, rather than relative premiums between exchange plans. Because firms initially underestimated exchange coverage elasticities, we observe a larger impact under full information.

As expected in a market with adverse selection, relaxing community rating decreases average premiums and requiring pure community rating increases average premiums. Relaxing community rating decreases average subsidized premiums by 13.8% under 2016 information and by 15.9% under full information; average unsubsidized premiums fall by 6.2% under 2016 information and by 5.2% under full information. Conversely, requiring pure community rating increases average subsidized premiums by 14.7% under 2016 information and by 15.7% under full information; average unsubsidized premiums increase by 0.3% under 2016 information and by 2.7% under full information.

The effect of modifying community rating rules on total exchange enrollment is similar under 2016 information and full information. Relaxing community rating increases enrollment by 1.6% under 2016 information and by 1.2% under full information, whereas requiring pure community rating decreases enrollment by 2.0% under 2016 information and by 1.7% under full information. Relaxing community rating has minimal effects on the enrollment distribution across metal tiers.

Welfare effects are also slightly larger under full information. Relaxing community rating decreases annual per-capita social welfare by $46 under 2016 information and by $96 under full information. Requiring community rating increases annual per-capita social welfare by $128 under

2016 information and by \$161 under full information. The effect of requiring community rating on consumers is very heterogeneous, as shown in Figure 10. Annual per-capita consumer surplus for people over age 55 increases by \$1,279 under 2016 information and \$1,005 under full information, but decreases for all other age groups. Total welfare increases because welfare gains for older consumers outweigh welfare losses for younger consumers and large government subsidies for older consumers are reduced.

7.2 Inertia

A prominent feature of health insurance markets is consumer inertia. We simulate the elimination of inertia by setting the inertia parameter vector β_i^y to zero in equation (1) and solving for the new vector of premiums that satisfy the firms' first-order conditions.

As for modifying community rating, eliminating inertia has a considerably more significant impact under full information. Firms initially underestimated the inertia parameters. Setting the inertia parameters to zero in our counterfactual therefore has a larger effect under full information than under 2016 information.

Eliminating inertia has a very substantial negative impact on premiums, particularly under full information. Averaged subsidized premiums decrease by 3.0% under 2016 information and by 17.0% under full information. Average unsubsidized premiums decrease by 12.3% under 2016 information and 15.0% under full information.[10] These results support the argument that inertia is a significant source of firm market power (Saltzman et al., 2021).

Regardless of whether learning is modeled, total exchange enrollment falls by about 2%, because a small share of consumers find it optimal to forgo insurance in the absence of inertia. Enrollment shifts towards the less generous bronze plans when inertia is eliminated, particularly under full

[10] Our "full information" results do not exactly match the results in Saltzman et al. (2021), primarily because we simulate the elimination of inertia for the year 2016, whereas Saltzman et al. (2021) simulate the elimination of inertia for the years 2015-2018 and take a simple average across the four years.

information. Bronze market share increases by 1.5 percentage points under 2016 information and by 3.7 percentage points under full information.

Gains in annual per-capita social welfare from eliminating inertia are substantial, increasing \$750 under 2016 information and \$655 under full information. Consumers benefit directly from reoptimizing without choice frictions and the government benefits from lower premiums, which reduces its spending on premium subsidies. These welfare gains are partially offset by a reduction in firm profit.

7.3 Risk Adjustment

We now evaluate the impact of risk adjustment, an important ACA program designed to mitigate firm risk selection. We simulate the elimination of risk adjustment by removing the marginal risk adjustment transfer $MRA_{jmt}(p; \boldsymbol{\theta})$ from equation (7) and solving for the vector of premiums that satisfy the resulting first-order conditions in equation (7).

In contrast to modifying community rating rules or eliminating of inertia, removing risk adjustment generally has a stronger impact under 2016 information than under full information. Eliminating risk adjustment transfers from plans that attract low-risk consumers (e.g., bronze plans) to plans that attract high-risk consumers (e.g., platinum plans) has the first-order effect of reducing the cost of offering bronze plans and increasing the cost of offering platinum plans. As a result, risk adjustment leads to significant shifts in relative premiums between plans. Because firms initially overestimated own-premium elasticities, we observe a larger impact of eliminating risk adjustment under 2016 information.

Eliminating risk adjustment has a substantial negative impact on premiums. Average subsidized premiums decrease by 23.6% under 2016 information and by 22.2% under full information. Average unsubsidized premiums decrease by 17.8% under 2016 information and by 13.5% under full information. A key driver for the decrease is intensive margin adverse selection (i.e., selection between plans). Average bronze premiums decline by 32.1% under 2016 information and by 27.7% under full information. These reduced bronze plan premiums result in bronze plan enrollment increasing 12.0 percentage points under 2016 information and 9.1 percentage points under full information. Conversely, platinum plan enrollment completely unravels under 2016 information and declines to 1.2% market share under full information. Lower bronze plan premiums attract new consumers to the exchange. Total exchange enrollment increases 0.5 percentage points under 2016 information and 0.6 percentage points under full information. We therefore find evidence that risk adjustment involves a tradeoff between underinsurance and underenrollment under both 2016 information and full information.

Annual per-capita social welfare increases when risk adjustment is eliminated by \$346 under 2016 information and by \$144 under full information. Consumers benefit directly from lower premiums and the government benefits from lower subsidy spending, particularly under 2016 information. These welfare gains are slightly offset by a reduction in firm profit.

7.4 Individual Mandate

Finally, we consider elimination of the individual mandate penalty in the year 2016, the first year that the penalty was in full effect. We do so by setting the penalty ρ_{it} to zero in equation (1) and solving for the vector of premiums that satisfy the firms' first-order conditions in equation (7).

In contrast to the other changes considered, eliminating the mandate penalty has a minimal impact, regardless of whether learning is modeled. Although the relatively limited effect of the mandate may be surprising given its prominence in policy discussions, it is consistent with the observed impact between 2018 and 2019 when the mandate penalty was set to zero. The relatively small penalty combined with robust price-linked subsidies that shield consumers from premium shocks limit the mandate's impact.

Figure 9 summaries the impact of eliminating the mandate penalty. Average subsidized premiums are largely unaffected and average unsubsidized premiums increase by 0.6%, with or without learning. Enrollment shifts between metal tiers are minimal, but there is a total enrollment drop of about 63,000 or 3.5% under 2016 information and about 53,000 or 2.9% under full information. These enrollment declines are the result of marginal consumers reoptimizing and deciding to forgo insurance when the penalty is eliminated. The mandate penalty has slightly less effect on enrollment when firms know the full information parameters because equilibrium premiums are lower prior to its elimination (i.e., comparing the "Base" scenario bars in Figure 9a). Annual per-capita social welfare declines by \$327 under 2016 information and by \$334 under full information when the mandate is eliminated. Reductions in premium subsidy spending and gains in consumer surplus from not compelling marginal consumers to purchase insurance are more than offset by a loss of revenue from penalty collections, as well as increased uncompensated care payments.

8 Conclusion

Large-scale social programs, including recent public health insurance expansions, are increasingly being implemented by creating new markets with private sector participation. In new markets, the standard IO assumptions of market equilibrium and complete information might be unrealistic (Doraszelski et al., 2018). We study the effects of relaxing these standard assumptions by estimating an adaptive learning model in a selection market using data from the California ACA exchange. Firms initially faced considerable uncertainty in predicting who would enroll and how much their enrollees would cost. Our setting is appealing because we observe the creation of a new market and can exploit data on firms' predictions about their costs, as well as their actual costs.

Our principal finding is that the full information assumption leads to a more favorable conclusion about the market equilibrium in the California ACA exchange setting. Firms' initial lack of knowledge of cost and consumer premium sensitivity are the primary drivers of the more favorable conclusion about the market equilibrium. Although this finding does not necessarily generalize to

Figure 9: Policy Impact of Omitting Learning

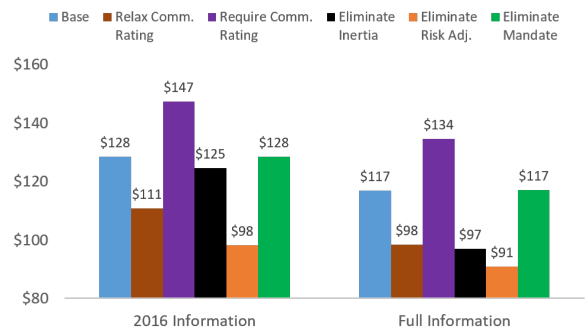

(a) Average Subsidized Premiums

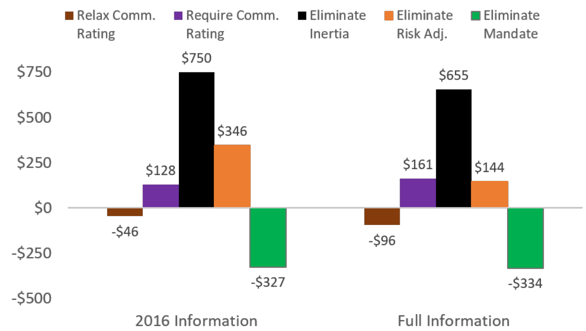

(c) Change in Annual Per-Capita Social Welfare

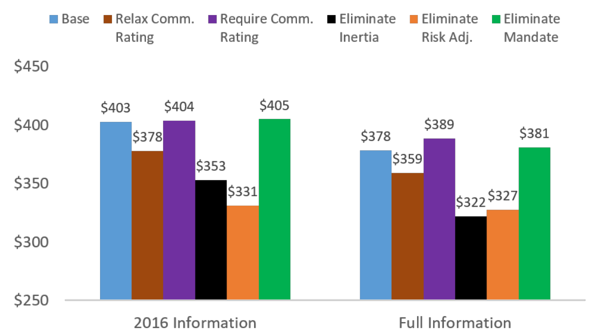

(b) Average Unsubsidized Premiums

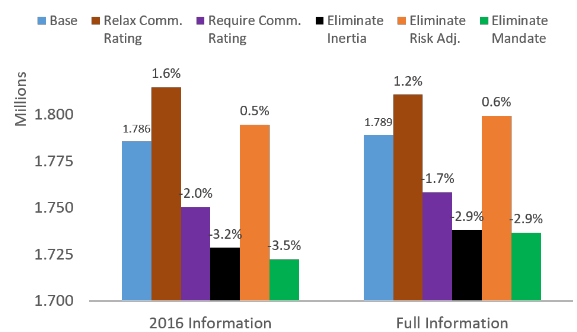

(d) Total Exchange Enrollment

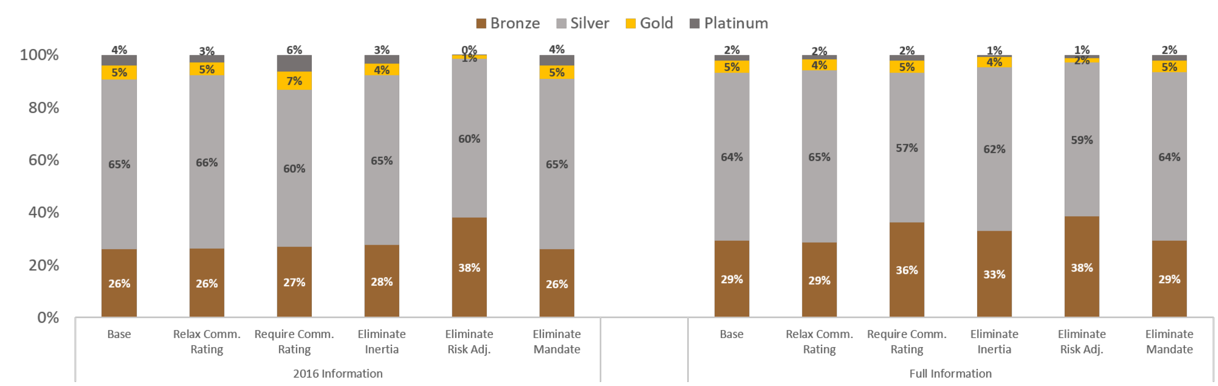

(e) Market Share By Metal

Notes: Figure shows the impact of using the full information parameters instead of the 2016 learning parameters on 5 different market feature changes, including: (1) eliminating risk adjustment; (2) eliminating the individual mandate; (3) eliminating inertia; (4) relaxing ACA modified community rating; and (5) requiring pure community rating. The figures show the equilibrium impact on subsidized premiums (panel a), annual per-capita social welfare (panel b), unsubsidized premiums (panel c), total exchange enrollment (panel d), and market share by metal tier (panel e).

Figure 10: Effect of Requiring Pure Community Rating on Consumer Surplus by Age Group

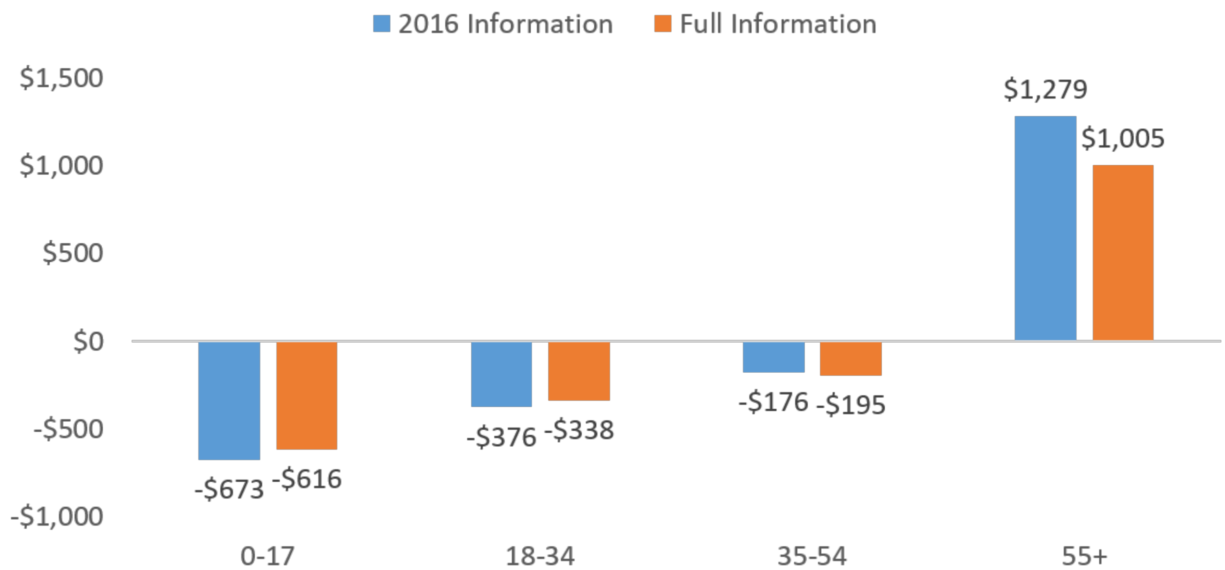

Notes: Figure shows the change in average annual per-capita consumer surplus that results from requiring pure community rating. Consumer surplus changes are shown by age group and by the information available.

other settings, it does suggest that researchers should consider whether firm uncertainty is relevant for their specific setting when analyzing a new or recently-created social program. Our study results also suggest that policymakers should adopt policies that promote information sharing between firms to reduce uncertainty. State exchanges can accomplish this by adopting an "active purchasing" model where the exchange actively solicits cost information from firms. Because California has adopted one of the most aggressive active purchasing models, the effect of firm information might be even larger in other state exchanges. The creation of all-payer claims databases may also facilitate information exchange and reduce firm uncertainty. The U.S. Supreme Court's recent decision in *Gobeille v. Liberty Mutual* to place significant limits on states' collection of health care claims data runs counter to these goals.

We also find that firm information has substantial implications for key program design features. The direction and magnitude of the effects depend on how each program feature interacts with market characteristics that are initially uncertain to firms. We find that pricing regulation and policies targeting inertia have a larger impact when firms have full information. In contrast, risk adjustment has a smaller impact when firms have full information. The individual mandate has minimal impact, regardless of whether firms have full information.

Our study has some limitations. Although the primary focus of our study is firm learning, a natural concern is whether consumers also learn and adjust their plan choices accordingly form year to year (Ketcham et al., 2012; Miravete, 2003; List, 2003, 2004, 2006; List and Millimet, 2008). A significant feature of the ACA exchanges is high consumer churn due to exogenous reasons, such as job status changes or substantial income shocks. We also find minimal evidence of consumers switching plans despite highly volatile premiums during our study period. Hence, consumers have limited opportunities and incentives to learn in our study setting. A useful extension of our study may consider both consumer and firm learning. Similar to Doraszelski et al. (2018), our model is not fully dynamic and does not allow firms to be forward-looking when setting premiums. Given the existing empirical tools, it is unclear how to estimate time-specific learning parameters and consistent equilibrium beliefs in a model of learning. Even in models without learning, modeling forward-looking behavior in health insurance markets is particularly challenging and would require significant compromises on key institutional details (Fleitas, 2017; Miller, 2019).

We expect the establishment of new insurance markets to be an increasingly important mechanism for expanding access to health insurance and reducing health care costs, especially under recent proposals to transform Medicare into a premium support or defined contribution program. The methods used in this paper might be useful for analyzing the potential impact of these markets and the impact of proposed regulation while firms are still learning.

References

Abaluck, J. and J. Gruber (2011, May). Heterogeneity in choice inconsistencies among the elderly: Evidence from prescription drug choice. *The American Economic Review 101*(3), 377–381.

Abaluck, J. and J. Gruber (2016). Evolving choice inconsistencies in choice of prescription drug insurance. *American Economic Review 106*(8), 2145–2184.

Ackerberg, D. (2003). Advertising, learning, and consumer choice in experience good markets: An empirical examination. *International Economic Review 44*(3), 1007–1040.

Aguirregabiria, V. and J. Jeon (2020). Firms' beliefs and learning: Models, identifcation, and empirical evidence. *Review of Industrial Organization 56*, 203–235.

Benkard, C. L. (2000). Learning and forgetting: The dynamics of aircraft production. *The American Economic Review 90*(4), 1034–1054.

Centers for Medicare and Medicaid Services (2018, January). *National Health Expentire Data*. https://www.cms.gov/Research-Statistics-Data-and-Systems/Statistics-Trends-and-Reports/NationalHealthExpendData/NationalHealthAccountsHistorical.html.

Conley, T. and C. Udry (2010). Learning about a new technology: Pineapple in ghana. *The American Economic Review 100*(1), 35–69.

Coughlin, T., J. Holahan, K. Caswell, and M. McGrath (2014, May). *Uncompensated Care for Uninsured in 2013: A Detailed Examination*. https://kaiserfamilyfoundation.files.wordpress.com/2014/05/8596-uncompensated-care-for-the-uninsured-in-2013.pdf.

Curto, V., L. Einav, J. Levin, and J. Bhattacharya (2020, June). Can health insurance competition work? evidence from medicare advantage. Accepted, Journal of Political Economy.

Decarolis, F., M. Polyakova, and S. Ryan (2020, May). Subsidy design in privately provided social insurance: Lessons from medicare part d. *Journal of Political Economy 128*(5).

Department of Managed Health Care (2016). *Premium Rate Review Filings*. http://wpso.dmhc.ca.gov/ratereview/.

Dickstein, M. (2018). Efficient provision of experience goods: Evidence from antidepressant choice.

Domurat, R. (2017). How do supply-side regulations in the aca impact market outcomes? evidence from california.

Doraszelski, U., G. Lewis, and A. Pakes (2018). Just starting out: Learning and equilibrium in a new market. *The American Economic Review 108*(3), 565–615.

Einav, L., A. Finkelstein, and P. Tebaldi (2019). Market design in regulated health insurance markets: Risk adjustment vs. subsidies.

Ericson, K. and A. Starc (2015). Pricing regulation and imperfect competition on the massachusetts health insurance exchange. *Review of Economics and Statistics 97*(3), 667–682.

Evans, G. and S. Honkapohja (2001). *Learning and Expectations in Macroeconomics*. Princeton University Press.

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

Fleitas, S. (2017, May). Dynamic competition and price regulation when consumers have inertia: Evidence from medicare part d. Working paper.

Geruso, M., T. Layton, G. McCormack, and M. Shepard (2019). The two margin problem in insurance markets.

Gruber, J. (2017). Delivering public health insurance through private plan choice in the united states. *Journal of Economic Perspectives 31*(4), 3–22.

Hackmann, M., J. Kolstad, and A. Kowalski (2015, March). Adverse selection and an individual mandate: When theory meets practice. *The American Economic Review 105*(3), 1030–1066.

Handel, B., I. Hendel, and M. Whinston (2015). Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica 83*(4), 1261–1313.

Hausman, J. and J. Poterba (1987). Household behavior and the tax reform act of 1986. *Journal of Economic Perspectives 1*(1), 101–119.

Ho, K. and A. Pakes (2014, December). Hospital choices, hospital prices, and financial incentives to physicians. *The American Economic Review 104*(12), 3841–3884.

Hortaçsu, A., F. Luco, S. L. Puller, and D. Zhu (2019). Does strategic ability affect efficiency? evidence from electricity markets. *American Economic Review 109*(12), 4302–42.

Hortaçsu, A. and S. L. Puller (2008). Understanding strategic bidding in multi-unit auctions: a case study of the texas electricity spot market. *The RAND Journal of Economics 39*(1), 86–114.

Huang, Y., P. Ellickson, and M. Lovett (2021). Learning to set prices. Working paper.

Jaffe, S. and M. Shepard (2020). Price-linked subsidies and imperfect competition in health insurance. *American Economic Journal: Economic Policy 12*(3), 279–311.

Jeon, J. (2020). Learning and investment under demand uncertainty in container shipping.

Joskow, P. L., R. Schmalensee, and E. M. Bailey (1998). The market for sulfur dioxide emissions. *American Economic Review*, 669–685.

Kaiser Family Foundation (2020, April). *State Health Facts: Health Reform*. `https://www.kff.org/statedata/`.

Kessler, D. and M. McClellan (2000). Is hospital competition socially wasteful? *Quarterly Journal of Economics 115*, 577–615.

Ketcham, J., C. Lucarelli, E. Miravete, and C. Roebuck (2012). Sinking, swimming, or learning to swim in medicare part d. *American Economic Review 102*(6), 2639–2673.

Ketcham, J., C. Lucarelli, and C. Power (2015). Paying attention or paying too much in medicare part d. *American Economic Review 105*(1), 204–233.

Layton, T. (2017). Imperfect risk adjustment, risk preferences, and sorting in competitive health insurance markets. *Journal of Health Economics*.

List, J. A. (2003). Does market experience eliminate market anomalies? *The Quarterly Journal of Economics 118*(1), 41–71.

List, J. A. (2004). Neoclassical theory versus prospect theory: Evidence from the marketplace. *Econometrica 72*(2), 615–625.

List, J. A. (2006). The behavioralist meets the market: Measuring social preferences and reputation effects in actual transactions. *Journal of political Economy 114*(1), 1–37.

List, J. A. and D. L. Millimet (2008). The market: Catalyst for rationality and filter of irrationality. *The BE Journal of Economic Analysis & Policy 8*(1).

Lucarelli, C., J. Prince, and K. Simon (2012). The welfare impact of reducing choice in medicare part d: A comparison of two regulation strategies. *International Economic Review 53*(4), 1155–1177.

Lustig, J. (2009). Measuring welfare losses from adverse selection and imperfect competition in privatized medicare.

Mahoney, N. and E. G. Weyl (2017). Imperfect competition in selection markets. *The Review of Economics and Statistics 99*(4), 637–651.

Miller, K. (2019, June). Estimating costs when consumers have inertia: Are private medicare insurers more efficient? Working paper.

Miller, K., A. Petrin, R. Town, and M. Chernew (2019, February). Optimal managed competition subsidies. Working paper.

Miravete, E. J. (2003, March). Choosing the wrong calling plan? ignorance and learning. *American Economic Review 93*(1), 297–310.

Morrisey, M. A., A. M. Rivlin, R. P. Nathan, and M. A. Hall (2017). Five-state study of aca marketplace competition: A summary report. *Risk Management and Insurance Review 20*(2), 153–172.

Newberry, P. (2016). An empirical study of observational learning. *RAND Journal of Economics 47*(2), 394–432.

Pauly, M., S. Harrington, and A. Leive (2015). "sticker shock" in individual insurance under health reform? *American Journal of Health Economics 1*(4), 494–514.

Pauly, M., A. Leive, and S. Harrington (2020). Losses (and gains) from health reform for non-medicaid uninsureds. *Journal of Risk and Insurance 87*(1), 41–66.

Polyakova, M. and S. Ryan (2021). Subsidy targeting with market power. *The National Bureau of Economic Research*.

Pope, G., H. Bachofer, A. Pearlman, J. Kautter, E. Hunter, D. Miller, and P. Keenan (2014). Risk transfer formula for individual and small group markets under the affordable care act. *Medicare and Medicaid Research Review 4*(3), 1–46.

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

Saltzman, E. (2021). Managing adverse selection: Underinsurance vs. underenrollment. *The RAND Journal of Economics 52*(2), 359–381.

Saltzman, E., A. Swanson, and D. Polsky (2021, July). Inertia, market power, and adverse selection in health insurance: Evidence from the aca exchanges. Working Paper 29097, National Bureau of Economic Research.

Sargent, T. (1993). *Bounded Rationality in Macroeconomics*. Oxford University Press.

Starc, A. (2014). Insurer pricing and consumer welfare: Evidence from medigap. *RAND Journal of Economics 45*(1), 198–220.

Tebaldi, P. (2020). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca.

Town, R. and S. Liu (2003). The welfare impact of medicare hmos. *The RAND Journal of Economics 34*(4), 719–736.

U.S. Census Bureau (2019, August). *Survey of Income and Program Participation*. `https://www.census.gov/sipp/`.

Yamamoto, D. (2013, June). Health care costs - from birth to death. *Society of Actuaries*. `https://www.cms.gov/CCIIO/Programs-and-Initiatives/Premium-Stabilization-Programs/Downloads/June-30-2016-RA-and-RI-Summary-Report-5CR-063016.pdf`.

Zhang, J. (2010). The sound of silence: Observational learning in the u.s. kidney market. *Marketing Science 29*(2), 315–335.

A Mathematical Formulas in ACA Exchange Model

In this appendix, we write the variables in equations (5) and (7) in terms of three variables and associated partial derivatives: (1) the household choice probabilities $q_{ijt}(p)$; (2) the risk scores $r_{jmt}(p)$; and (3) plan average claims $c_{jmt}(p)$.

Household Choice Probabilities:

The household choice probabilities as defined in equation (3) are

$$q_{ijt}(p;\boldsymbol{\beta}) = \frac{e^{V_{ijt}(p;\boldsymbol{\beta})/\lambda}\left(\sum_j e^{V_{ijt}(p;\boldsymbol{\beta})/\lambda}\right)^{\lambda-1}}{1+\left(\sum_j e^{V_{ijt}(p;\boldsymbol{\beta})/\lambda}\right)^{\lambda}}$$

The (k,j) element of the Jacobian matrix of the household choice probability is

$$\frac{\partial q_{ikt}(p;\boldsymbol{\beta})}{\partial p_{ijt}} = \begin{cases} \beta_i^p q_{ijt}(p;\boldsymbol{\beta})\left[\frac{1}{\lambda}+\frac{\lambda-1}{\lambda}q'_{ijt}(p;\boldsymbol{\beta})-q_{ijt}(p;\boldsymbol{\beta})\right] & k=j \\ \beta_i^p q_{ijt}(p;\boldsymbol{\beta})\left[\frac{\lambda-1}{\lambda}q'_{ijt}(p;\boldsymbol{\beta})-q_{ijt}(p;\boldsymbol{\beta})\right] & k\neq j \end{cases} \tag{14}$$

where $q'_{ijt}(p;\boldsymbol{\beta})$ is the probability of choosing j, conditional on choosing a plan. Household i's demand partial derivative with respect to the firm's base plan premium p_{jmt} is

$$\frac{\partial q_{ikt}(p;\boldsymbol{\beta})}{\partial p_{jmt}} = \sum_{l\in J_{mt}} \frac{\partial q_{ikt}(p;\boldsymbol{\beta})}{\partial p_{ilt}(p)}\frac{\partial p_{ilt}(p)}{\partial p_{jmt}}$$

where $\frac{\partial p_{ilt}(p)}{\partial p_{jmt}}$ is given in equation (4).

Plan Risk Scores:

We define the plan risk score in equation (8) as

$$\ln r_{jmt}(p;\boldsymbol{\theta}) = \sum_{d\in D}\gamma^d s_{djmt}(p;\boldsymbol{\beta}) + MT'_j\gamma^{MT} + \epsilon^r_{jmt}$$

The (k,j)-element of the Jacobian matrix of the plan risk score equals

$$\frac{\partial r_{kmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}} = \frac{r_{kmt}(p;\boldsymbol{\theta})}{q_{kmt}(p;\boldsymbol{\beta})}\sum_{d\in D}\gamma^d\left[\frac{\partial q_{dkmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}} - s_{dkmt}(p;\boldsymbol{\beta})\frac{\partial q_{kmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}\right] \tag{15}$$

Plan Average Claims:

We define log average claims in equation (9) as

$$\ln c_{jmt}(p;\boldsymbol{\theta}) \quad = \quad \mu^r \ln r_{jmt}(p;\boldsymbol{\theta}) + x_j'\mu^x + \mu^l l_t + n_m'\mu^n + \epsilon_{jmt}^c$$

The (k,j)-element of the Jacobian matrix of plan average claims equals

$$\frac{\partial c_{kmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}} \quad = \quad \mu^r \frac{c_{kmt}(p;\boldsymbol{\theta})}{r_{kmt}(p;\boldsymbol{\theta})} \frac{\partial r_{kmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}} \tag{16}$$

Plan and Firm Demand:

Total plan demand $q_{jmt}(p;\boldsymbol{\beta})$ and total firm demand $q_{ft}(p;\boldsymbol{\beta})$ equal

$$q_{jmt}(p;\boldsymbol{\beta}) = \sum_{i \in I} \left(\mathbb{I}_{i,m,t}\right) q_{ijt}(p;\boldsymbol{\beta})$$

$$q_{ft}(p;\boldsymbol{\beta}) = \sum_{i \in I, k \in J_f} q_{ikt}(p;\boldsymbol{\beta})$$

The plan and firm demand partial derivatives are

$$\frac{\partial q_{kmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}} = \sum_{i \in I} \left(\mathbb{I}_{i,m,t}\right) \frac{\partial q_{ikt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}$$

$$\frac{\partial q_{ft}(p;\boldsymbol{\beta})}{\partial p_{jmt}} = \sum_{i \in I, k \in J_f} \frac{\partial q_{kmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}$$

Firm Revenue:

Total premium revenue earned by the firm is

$$R_{ft}(p) = \sum_{i \in I, k \in J_{fmt}} \sigma_{it} p_{kmt} q_{ikt}(p)$$

and marginal revenue $MR_{jmt}(p;\boldsymbol{\beta}) \equiv \frac{\partial R_{ft}(p;\boldsymbol{\beta})}{\partial q_{jmt}(p;\boldsymbol{\beta})}$ is

$$MR_{jmt}(p;\boldsymbol{\beta}) \quad = \quad \left(\frac{\partial q_{jmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}\right)^{-1} \sum_{i \in I, k \in J_{fmt}} \sigma_{it} \left(q_{ijt}(p;\boldsymbol{\beta}) \; + p_{kmt} \frac{\partial q_{ikt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}\right)$$

Firm Claims:

Total claims paid by the firm are

$$C_{ft}(p) = \sum_{k \in J_{fmt}} c_{kmt}(p) q_{kmt}(p)$$

Marginal claims $MC_{jmt}(p;\boldsymbol{\theta}) \equiv \frac{\partial C_{ft}(p;\boldsymbol{\theta})}{\partial q_{jmt}(p;\boldsymbol{\beta})}$ is

$$MC_{jmt}(p;\boldsymbol{\theta}) \quad = \quad \left(\frac{\partial q_{jmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}\right)^{-1} \sum_{k \in J_{fmt}} \left[c_{kmt}(p;\boldsymbol{\theta}) \frac{\partial q_{kmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}} + q_{kmt}(p;\boldsymbol{\beta}) \frac{\partial c_{kmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}}\right] \tag{17}$$

Firm Variable Administrative Cost:

Total variable administrative cost is

$$V_{ft}(p) = v_{ft} q_{ft}(p)$$

where v_{ft} is the variable administrative cost per-member per-month. Marginal variable administrative cost $MV_{jmt}(p;\boldsymbol{\beta}) = \frac{\partial V_{ft}(p;\boldsymbol{\beta})}{\partial q_{jmt}(p;\boldsymbol{\beta})}$ is

$$MV_{jmt}(p;\boldsymbol{\beta}) \quad = \quad v_{ft} \frac{\partial q_{ft}(p;\boldsymbol{\beta})/\partial p_{jmt}}{\partial q_{jmt}(p;\boldsymbol{\beta})/\partial p_{jmt}} \tag{18}$$

Firm Risk Adjustment:

The firm's risk adjustment transfer is

$$RA_{ft}(p) = R_t(p) \sum_{m \in M, k \in J_{fmt}} \left[rs_{kmt}(p) - us_{kmt}(p)\right]$$

The marginal risk adjustment transfer $MRA_{jmt}(p;\boldsymbol{\theta}) = \frac{\partial RA_{ft}(p;\boldsymbol{\theta})}{\partial q_{jmt}(p;\boldsymbol{\beta})}$ is

$$\begin{aligned} MRA_{jmt}(p;\boldsymbol{\theta}) \quad &= \quad \left(\frac{\partial q_{jmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}\right)^{-1} \sum_{k \in J_{fmt}} \left[\frac{\partial R_t(p;\boldsymbol{\beta})}{\partial p_{jmt}} \left(rs_{kmt}(p;\boldsymbol{\theta}) - us_{kmt}(p;\boldsymbol{\beta})\right)\right. \\ &\quad + \left. R_t(p;\boldsymbol{\beta}) \left(\frac{\partial rs_{kmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}} - \frac{\partial us_{kmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}\right)\right] \end{aligned} \tag{19}$$

where

$$\frac{\partial R_t(p;\boldsymbol{\beta})}{\partial p_{jmt}} \quad = \quad \sum_{l \in J_{mt}} MR_{lmt}(p;\boldsymbol{\beta}) \frac{\partial q_{lmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}$$

$$
\begin{aligned}
\frac{\partial us_{kmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}} &= \left(\sum_{m\in M,l\in J_{mt}} h_l q_{lmt}(p;\boldsymbol{\beta})\right)^{-1}\Bigg[h_k \frac{\partial q_{kmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}} \\
&\quad - \frac{h_k q_{kmt}(p;\boldsymbol{\beta})}{\sum_{m\in M,l\in J_{mt}} h_l q_{lmt}(p;\boldsymbol{\beta})}\sum_{l\in J_{mt}} h_l \frac{\partial q_{lmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}}\Bigg]
\end{aligned}
$$

$$
\begin{aligned}
\frac{\partial rs_{kmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}} &= \left(\sum_{m\in M,l\in J_{mt}} r_{lmt}(p;\boldsymbol{\theta}) q_{lmt}(p;\boldsymbol{\beta})\right)^{-1}\Bigg[\left(r_{kmt}(p;\boldsymbol{\theta})\frac{\partial q_{kmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}} + q_{kmt}(p;\boldsymbol{\beta})\frac{\partial r_{kmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}}\right) \\
&\quad - \frac{r_{jmt}(p;\boldsymbol{\theta}) q_{jmt}(p;\boldsymbol{\beta})}{\sum_{m\in M,l\in J_{mt}} r_{lmt}(p;\boldsymbol{\theta}) q_{lmt}(p;\boldsymbol{\beta})}\sum_{l\in J_{mt}}\left[r_{lmt}(p;\boldsymbol{\theta})\frac{\partial q_{lmt}(p;\boldsymbol{\beta})}{\partial p_{jmt}} + q_{lmt}(p;\boldsymbol{\beta})\frac{\partial r_{lmt}(p;\boldsymbol{\theta})}{\partial p_{jmt}}\right]\Bigg]
\end{aligned}
$$

B Summary Statistics

C Construction of the Uninsured Population

This appendix describes the procedure originally developed in Saltzman et al. (2021) to construct the exchange-eligible population. We model switching both between plans and into and out of the exchange market using five years of longitudinal data on exchange customers. Previous studies of the exchanges have treated demand as static, merging administrative data on exchange enrollees with survey data such as the American Community Survey (ACS) on the uninsured to form the universe of potential exchange consumers (Tebaldi, 2020; Domurat, 2017; Saltzman, 2019). The sample of uninsured in the ACS is limited and ACS geographic identifiers are difficult to match with those in our administrative data. In contrast, we form the uninsured population using data on consumers in the California data for years in which they did not have exchange coverage. For example, consumers with exchange coverage in 2016 and 2017 are considered uninsured in 2014, 2015, and 2018 if they remained eligible for the exchange market. Consumers might lose exchange market eligibility if they gain access to employer-sponsored insurance or public insurance (e.g.,

Table A1: Demographic Distribution By Year

	2014	2015	2016	2017	2018	2019	Overall
Market Size	2,197,669	2,420,764	2,461,389	2,444,685	2,429,209	2,272,457	14,226,173
Total Enrollment	1,362,316	1,639,923	1,702,160	1,697,074	1,710,469	1,553,374	9,665,316
Income							
138% FPL or less	4.7%	3.5%	3.3%	4.0%	4.0%	3.5%	3.8%
138% FPL to 150% FPL	14.1%	14.3%	14.6%	14.7%	14.4%	14.0%	14.4%
150% FPL to 200% FPL	32.8%	32.8%	31.9%	30.3%	28.8%	28.4%	30.8%
200% FPL to 250% FPL	16.8%	16.7%	16.3%	16.3%	16.7%	16.7%	16.6%
250% FPL to 400% FPL	22.4%	23.4%	23.6%	23.6%	25.8%	27.4%	24.4%
400% FPL or greater	9.3%	9.3%	10.3%	11.0%	10.3%	9.9%	10.0%
Subsidy Status							
Subsidized	89.6%	88.8%	87.5%	86.5%	87.3%	87.7%	87.8%
Unsubsidized	10.4%	11.2%	12.5%	13.5%	12.7%	12.3%	12.2%
Age							
0-17	5.7%	6.0%	6.2%	6.7%	7.3%	7.3%	6.5%
18-25	11.1%	11.3%	11.1%	10.7%	10.5%	10.0%	10.8%
26-34	16.3%	16.9%	17.4%	17.6%	17.7%	17.3%	17.2%
35-44	16.6%	15.9%	15.3%	15.1%	15.2%	15.1%	15.5%
45-54	24.4%	23.5%	22.8%	22.2%	21.4%	21.0%	22.5%
55+	25.8%	26.3%	27.2%	27.8%	27.9%	29.3%	27.4%
Gender							
Female	52.6%	52.2%	51.9%	52.2%	52.5%	52.5%	52.3%
Male	47.4%	47.8%	48.1%	47.8%	47.5%	47.5%	47.7%
Race							
Asian	22.8%	21.8%	22.0%	22.6%	23.0%	23.4%	22.6%
Black/African American	2.7%	2.5%	2.4%	2.4%	2.4%	2.4%	2.5%
Hispanic	27.5%	28.2%	28.0%	28.3%	28.4%	27.8%	28.0%
Non-Hispanic White	39.4%	39.5%	39.6%	38.5%	37.1%	36.8%	38.5%
Other Race	7.7%	7.9%	7.9%	8.2%	9.1%	9.6%	8.4%

Medicaid or Medicare).

Because our data do not indicate when enrollees lose exchange market eligibility, we use data from the Survey of Income and Program Participation (SIPP) to impute consumer eligibility. The SIPP is well-suited for the imputation because (1) it asks the insurance status of respondents for every month over a three-year period (2013-2015) and (2) it includes detailed information on the chief reasons for consumers' coverage status, such as whether the respondent obtained or lost an offer of employer-sponsored insurance, moved in or out of California, or became eligible or ineligible for Medicare or Medicaid. For SIPP respondents who newly obtained or gave up individual market coverage, we construct a *transitioned* variable that indicates whether the respondent gained or lost eligibility for the individual market. The *transitioned* variable takes value 1 if the respondent (1) belongs to a household that obtained or lost an offer of employer-sponsored insurance; (2) moved out of or into California; (3) the respondent turned 65 and qualified for Medicare; and (4) the respondent became eligible for Medicaid following a drop in income. We estimate a logit model regression of the *transitioned* variable on the demographic variables available in both the SIPP and the California data, including age, income, gender, race, and household size. We use the estimated logit to predict whether California consumers observed for only some years of the study timeframe transitioned into or out of the exchange market. Consumers transitioning into or out of the exchange market are removed from the study population during years when they are not enrolled in an exchange plan.

D Model Parameter Estimates

Table A2: Estimated Parameters

Demand Parameters ($\widehat{\beta}_t$)

	$\widehat{\boldsymbol{\theta}}_{2015}$	$\widehat{\boldsymbol{\theta}}_{2016}$	$\widehat{\boldsymbol{\theta}}_{2017}$	$\widehat{\boldsymbol{\theta}}_{2018}$	$\widehat{\boldsymbol{\theta}}$
Monthly Premium ($100)	-0.041^{***}	-0.100^{***}	-0.128^{***}	-0.140^{***}	-0.134^{***}
	(0.003)	(0.003)	(0.003)	(0.002)	(0.002)
250% to 400% of FPL	0.008^{***}	0.025^{***}	0.035^{***}	0.041^{***}	0.041^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
> 400% of FPL	0.018^{***}	0.053^{***}	0.076^{***}	0.089^{***}	0.083^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Ages 0 to 17	-0.065^{***}	-0.142^{***}	-0.161^{***}	-0.168^{***}	-0.148^{***}
	(0.005)	(0.005)	(0.005)	(0.004)	(0.003)
Ages 18 to 34	-0.064^{***}	-0.148^{***}	-0.182^{***}	-0.192^{***}	-0.179^{***}
	(0.005)	(0.004)	(0.003)	(0.003)	(0.002)
Ages 35 to 54	-0.028^{***}	-0.068^{***}	-0.083^{***}	-0.090^{***}	-0.085^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Male	-0.007^{***}	-0.020^{***}	-0.024^{***}	-0.023^{***}	-0.020^{***}
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Family	0.002^{***}	0.004^{***}	0.005^{***}	0.002^{**}	-0.001
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
Asian	-0.020^{***}	-0.047^{***}	-0.054^{***}	-0.059^{***}	-0.056^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)
Black	-0.010^{***}	-0.022^{***}	-0.026^{***}	-0.029^{***}	-0.022^{***}
	(0.002)	(0.003)	(0.004)	(0.003)	(0.003)
Hispanic	-0.030^{***}	-0.066^{***}	-0.077^{***}	-0.082^{***}	-0.068^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Other race	-0.002^{**}	-0.006^{***}	-0.006^{**}	-0.003	-0.005^{***}
	(0.001)	(0.002)	(0.002)	(0.002)	(0.002)
AV	0.300^{***}	0.650^{***}	0.777^{***}	0.834^{***}	0.806^{***}
	(0.022)	(0.017)	(0.013)	(0.011)	(0.010)
Silver	0.032^{***}	0.082^{***}	0.107^{***}	0.128^{***}	0.131^{***}
	(0.002)	(0.003)	(0.002)	(0.002)	(0.002)
HMO	0.009^{***}	0.048^{***}	0.075^{***}	-0.016^{***}	-0.033^{***}
	(0.001)	(0.003)	(0.003)	(0.002)	(0.001)
Anthem	0.061^{***}	0.178^{***}	0.218^{***}	0.115^{***}	0.089^{***}
	(0.005)	(0.006)	(0.005)	(0.003)	(0.002)
Blue Shield	0.062^{***}	0.175^{***}	0.223^{***}	0.122^{***}	0.100^{***}
	(0.005)	(0.006)	(0.005)	(0.003)	(0.002)
Kaiser	0.059^{***}	0.142^{***}	0.158^{***}	0.150^{***}	0.135^{***}
	(0.004)	(0.004)	(0.003)	(0.002)	(0.002)
Health Net	0.029^{***}	0.074^{***}	0.068^{***}	0.033^{***}	0.021^{***}
	(0.002)	(0.003)	(0.002)	(0.001)	(0.001)
Anthem x HMO	-0.062^{***}	-0.183^{***}	-0.254^{***}	-0.189^{***}	-0.180^{***}
	(0.005)	(0.006)	(0.006)	(0.004)	(0.004)
Nesting Parameter	0.034^{***}	0.086^{***}	0.111^{***}	0.126^{***}	0.132^{***}
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)

	$\widehat{\boldsymbol{\theta}}_{2015}$	$\widehat{\boldsymbol{\theta}}_{2016}$	$\widehat{\boldsymbol{\theta}}_{2017}$	$\widehat{\boldsymbol{\theta}}_{2018}$	$\widehat{\boldsymbol{\theta}}$
Previous Choice		0.212^{***}	0.257^{***}	0.233^{***}	0.222^{***}
		(0.014)	(0.013)	(0.011)	(0.010)
250% to 400% of FPL		0.016^{***}	0.024^{***}	0.030^{***}	0.016^{***}
		(0.004)	(0.003)	(0.003)	(0.003)
> 400% of FPL		0.019^{***}	0.037^{***}	0.046^{***}	0.031^{***}
		(0.005)	(0.005)	(0.004)	(0.004)
Ages 0 to 17		0.035^{***}	0.026^{***}	0.022^{***}	0.014^{**}
		(0.011)	(0.009)	(0.008)	(0.007)
Ages 18 to 34		-0.018^{***}	-0.021^{***}	-0.015^{***}	-0.014^{***}
		(0.004)	(0.004)	(0.003)	(0.003)
Ages 35 to 54		-0.020^{***}	-0.023^{***}	-0.018^{***}	-0.016^{***}
		(0.004)	(0.003)	(0.003)	(0.003)
Male		0.022^{***}	0.027^{***}	0.032^{***}	0.034^{***}
		(0.004)	(0.004)	(0.003)	(0.003)
Family		-0.025^{***}	-0.035^{***}	-0.046^{***}	-0.046^{***}
		(0.003)	(0.003)	(0.002)	(0.002)
Asian		-0.017^{***}	-0.039^{***}	-0.052^{***}	-0.052^{***}
		(0.004)	(0.003)	(0.003)	(0.003)
Black		-0.035^{***}	-0.039^{***}	-0.048^{***}	-0.040^{***}
		(0.010)	(0.009)	(0.008)	(0.007)
Hispanic		-0.013^{***}	-0.029^{***}	-0.033^{***}	-0.031^{***}
		(0.004)	(0.003)	(0.003)	(0.003)
Other race		-0.009	-0.018^{***}	-0.025^{***}	-0.025^{***}
		(0.006)	(0.006)	(0.005)	(0.004)
Anthem		-0.052^{***}	-0.044^{***}	0.024^{***}	0.115^{***}
		(0.009)	(0.008)	(0.007)	(0.006)
Blue Shield		0.000	0.020^{**}	0.089^{***}	0.206^{***}
		(0.010)	(0.009)	(0.008)	(0.007)
Kaiser		-0.046^{***}	-0.042^{***}	-0.002	0.044^{***}
		(0.007)	(0.006)	(0.005)	(0.004)
Health Net		-0.110^{***}	-0.148^{***}	-0.089^{***}	0.021^{***}
		(0.008)	(0.007)	(0.005)	(0.004)
HMO		0.080^{***}	0.087^{***}	0.123^{***}	0.143^{***}
		(0.007)	(0.006)	(0.006)	(0.006)
AV		0.327^{***}	0.434^{***}	0.476^{***}	0.398^{***}
		(0.017)	(0.015)	(0.013)	(0.011)
Silver		-0.108^{***}	-0.137^{***}	-0.165^{***}	-0.182^{***}
		(0.004)	(0.004)	(0.004)	(0.003)

Risk Score Parameters ($\widehat{\gamma}_t$)

	$\widehat{\boldsymbol{\theta}}_{2015}$	$\widehat{\boldsymbol{\theta}}_{2016}$	$\widehat{\boldsymbol{\theta}}_{2017}$	$\widehat{\boldsymbol{\theta}}_{2018}$	$\widehat{\boldsymbol{\theta}}$
Silver	0.652^{***}	0.492^{***}	0.656^{***}	0.605^{***}	0.569^{***}
	(0.064)	(0.043)	(0.045)	(0.033)	(0.028)
Gold	0.734^{***}	0.655^{***}	0.842^{***}	0.794^{***}	0.785^{***}
	(0.149)	(0.093)	(0.092)	(0.065)	(0.050)
Platinum	1.033^{***}	0.978^{***}	1.179^{***}	1.136^{***}	1.125^{***}
	(0.150)	(0.084)	(0.086)	(0.062)	(0.049)
Share Ages 18 to 54	-1.197	-0.698	-1.292	-1.525^{***}	-1.136^{***}
	(1.212)	(0.837)	(0.838)	(0.561)	(0.412)
Share Hispanic	-2.000^{***}	-1.251^{***}	-0.903^{**}	-0.938^{***}	-1.042^{***}
	(0.651)	(0.413)	(0.406)	(0.269)	(0.213)

Trend Parameters ($\widehat{\boldsymbol{\eta}}_t$)

	$\widehat{\boldsymbol{\theta}}_{2015}$	$\widehat{\boldsymbol{\theta}}_{2016}$	$\widehat{\boldsymbol{\theta}}_{2017}$	$\widehat{\boldsymbol{\theta}}_{2018}$	$\widehat{\boldsymbol{\theta}}$
Intercept	1.331^{***}	1.365^{***}	1.241^{***}	1.819^{***}	
	(0.125)	(0.066)	(0.040)	(0.055)	
Anthem	-0.056	-0.095^{*}	-0.104^{***}	-0.738^{***}	
	(0.123)	(0.057)	(0.040)	(0.058)	
Blue Shield	0.087	-0.120^{*}	-0.016	-0.391^{***}	
	(0.131)	(0.065)	(0.038)	(0.055)	
Health Net	-0.154	-0.223^{***}	0.382^{***}	0.088^{*}	
	(0.125)	(0.055)	(0.041)	(0.050)	
Kaiser	0.363^{**}	-0.198^{***}	0.008	-0.411^{***}	
	(0.172)	(0.061)	(0.044)	(0.047)	
HMO	-0.484^{***}	-0.302^{***}	-0.146^{***}	-0.253^{***}	
	(0.109)	(0.041)	(0.016)	(0.040)	

Average Claims Parameters ($\widehat{\boldsymbol{\mu}}_t$)

	$\widehat{\boldsymbol{\theta}}_{2015}$	$\widehat{\boldsymbol{\theta}}_{2016}$	$\widehat{\boldsymbol{\theta}}_{2017}$	$\widehat{\boldsymbol{\theta}}_{2018}$	$\widehat{\boldsymbol{\theta}}$
Intercept	5.484^{*}	5.819^{***}	5.933^{***}	6.804^{***}	5.980^{***}
	(3.235)	(0.343)	(0.123)	(0.031)	(0.083)
HMO	-0.174	0.027	-0.114^{***}	0.092^{***}	-0.138^{***}
	(0.166)	(0.034)	(0.014)	(0.012)	(0.008)
Log risk score	0.866^{***}	0.939^{***}	0.978^{***}	0.991^{***}	1.024^{***}
	(0.033)	(0.008)	(0.005)	(0.004)	(0.005)
Trend		-0.054^{***}	-0.026^{***}	-0.034^{***}	0.023^{***}
		(0.007)	(0.002)	(0.002)	(0.002)
Anthem	0.071	0.211^{***}	-0.010	0.018	0.150^{***}
	(0.248)	(0.034)	(0.020)	(0.021)	(0.014)
Blue Shield	-0.036	0.119^{***}	-0.068^{***}	-0.071^{***}	0.100^{***}
	(0.287)	(0.043)	(0.022)	(0.022)	(0.023)
Health Net	-0.254	-0.009	-0.094^{***}	-0.047^{***}	0.079^{***}
	(0.229)	(0.038)	(0.022)	(0.017)	(0.018)
Kaiser	0.036	0.113^{***}	0.102^{***}	-0.147^{***}	0.203^{***}
	(0.218)	(0.038)	(0.022)	(0.023)	(0.024)

Notes: Robust standard errors are in parentheses (*** indicates statistical significance at the 1% level, ** at the 5% level, and * at the 10% level). Parameter estimates for the market fixed effects in equations (1) and (9) are omitted.

E Detailed Simulation Results

Table A3: Impact of Omitting Learning By Year

	2015		2016		2017		2018	
	Avail.	Full	Avail.	Full	Avail.	Full	Avail.	Full
Monthly Premiums								
Bronze	$327	$165	$332	$298	$366	$336	$394	$339
Silver	$407	$258	$420	$405	$469	$435	$583	$508
Gold	$442	$326	$454	$443	$488	$469	$566	$520
Platinum	$488	$442	$507	$551	$540	$565	$631	$591
Anthem	$402	$237	$411	$383	$466	$450	$650	$650
Blue Shield	$414	$218	$419	$368	$484	$441	$573	$484
Health Net	$354	$330	$371	$412	$408	$351	$484	$373
Kaiser	$401	$276	$417	$457	$448	$429	$515	$497
Other Insurer	$358	$281	$353	$307	$370	$338	$467	$387
HMO	$376	$281	$387	$387	$418	$385	$504	$447
PPO	$411	$226	$416	$374	$483	$445	$581	$480
Average	$395	$239	$403	$378	$444	$411	$526	$458
Subs. Average	$122	$105	$128	$117	$138	$134	$123	$112
Enrollment								
Total Coverage	1,727,136	1,736,076	1,785,543	1,789,239	1,764,992	1,766,390	1,766,046	1,779,395
% Enrolled	73.8%	74.2%	75.1%	75.2%	74.6%	74.6%	74.3%	74.9%
Bronze	23.7%	27.7%	26.1%	29.4%	27.3%	28.9%	30.5%	31.0%
Silver	64.8%	65.7%	64.7%	63.9%	63.8%	64.7%	57.3%	59.6%
Gold	6.3%	4.5%	5.3%	4.6%	5.5%	4.5%	8.3%	6.5%
Platinum	5.1%	2.1%	3.9%	2.1%	3.4%	1.9%	4.0%	2.9%
Change in Annual Per-Capita Social Welfare								
Cons. Surplus		$86		$62		$4		$46
Profit		($1054)		$43		($178)		($504)
Government Spending								
Prem. Subsidies		($1212)		($107)		($260)		($479)
CSRs		$14		($4)		$2		$12
Penalties		($4)		($1)		($1)		($6)
Uncomp. Care		($6)		($2)		($1)		($11)
Social Welfare		$592		$249		$162		$156

Notes: Table shows the equilibrium impact of using the parameters estimated with “Full” information instead of the learning parameters estimated with “Available” information by year. The first panel shows the effect on enrollee-weighted average premiums by metal tier, insurer, and plan network type. The second panel reports the impact on insurance coverage. The third panel shows the change in annual per-capita social welfare for each year when switching from the parameters estimated with “Full” information to the learning parameters estimated with “Available” information.

Table A4: Determining Which Parameters Drive the Impact of Omitting Learning

	Base	(1)	(2)	(3)	(4)	(5)	(6)
Full Inf. Parameters							
Cost		✓	✓	✓	✓	✓	✓
Risk			✓	✓	✓	✓	✓
Other Demand				✓	✓	✓	✓
Inertia						✓	✓
Premium					✓		✓
Monthly Premiums							
Bronze	$332	$292	$285	$296	$272	$331	$298
Silver	$420	$399	$395	$409	$386	$433	$405
Gold	$454	$437	$438	$446	$429	$464	$443
Platinum	$507	$549	$549	$553	$557	$555	$551
Anthem	$411	$374	$365	$378	$351	$417	$383
Blue Shield	$419	$363	$358	$373	$349	$397	$368
Health Net	$371	$408	$403	$412	$384	$446	$412
Kaiser	$417	$450	$446	$453	$422	$494	$457
Other Insurer	$353	$304	$302	$315	$291	$333	$307
HMO	$387	$383	$379	$391	$363	$417	$387
PPO	$416	$367	$361	$375	$350	$405	$374
Average	$403	$373	$367	$381	$355	$409	$378
Subs. Average	$128	$117	$115	$115	$112	$121	$117
Enrollment							
Total Enrollment	1,785,543	1,789,295	1,789,310	1,791,469	1,792,189	1,787,515	1,789,239
% Enrolled	75.1%	75.2%	75.2%	75.3%	75.3%	75.1%	75.2%
Bronze	26.1%	29.3%	29.9%	29.5%	30.4%	28.3%	29.4%
Silver	64.7%	63.7%	63.9%	63.6%	64.3%	63.5%	63.9%
Gold	5.3%	4.8%	4.3%	4.8%	4.1%	5.4%	4.6%
Platinum	3.9%	2.1%	1.9%	2.1%	1.2%	2.8%	2.1%
Change in Annual Per-Capita Social Welfare							
Cons. Surplus		$64	$68	$72	$85	$46	$62
Profit		($13)	($40)	$49	($132)	$277	$43
Government Spending							
Prem. Subsidies		($157)	($196)	($64)	($269)	$132	($107)
CSRs		($4)	($3)	($3)	($0)	($6)	($4)
Penalties		($2)	($2)	($3)	($3)	($1)	($1)
Uncomp. Care		($2)	($2)	($4)	($4)	($1)	($2)
Social Welfare		$261	$289	$211	$306	$158	$249

Notes: Table shows the impact of replacing the 2016 learning parameter estimates with the full information parameters of the model. The first panel indicates which parameters are known to the firm, including (1) the parameters from the average claims and predicted cost equations; (2) the parameters from the risk score equation; (3) all parameters in the utility equation except the inertia and premium parameters; (4) the inertia parameters in the utility equation; and (5) the premium parameters in the utility equation. The second panel shows the effect on enrollee-weighted average premiums by metal tier, insurer, and plan network type. The third panel reports the impact on insurance coverage. The fourth panel shows the change in annual per-capita social welfare relative to the base scenario.

Table A5: Policy Simulations w/ 2016 Learning Parameters

	Base	(7)	(8)	(9)	(10)	(11)
Scenario Definitions						
Risk Adjustment	✓		✓	✓	✓	✓
Individual Mandate	✓	✓		✓	✓	✓
Inertia	✓	✓	✓		✓	✓
Modified Comm. Rating	✓	✓	✓	✓		
Pure Comm. Rating						✓
Monthly Premiums						
Bronze	$332	$225	$334	$292	$313	$316
Silver	$420	$394	$423	$369	$401	$405
Gold	$454	$445	$457	$419	$425	$455
Platinum	$507	$703	$510	$465	$476	$512
Anthem	$411	$332	$414	$360	$390	$398
Blue Shield	$419	$377	$422	$378	$401	$417
Health Net	$371	$299	$373	$318	$349	$355
Kaiser	$417	$353	$419	$368	$393	$400
Other Insurer	$353	$262	$356	$314	$332	$346
HMO	$387	$316	$389	$335	$365	$372
PPO	$416	$345	$419	$370	$395	$406
Average	$403	$331	$405	$353	$381	$389
Subs. Average	$128	$98	$128	$125	$123	$135
Enrollment						
Total Coverage	1,785,543	1,794,564	1,722,191	1,728,436	1,796,472	1,782,624
% Enrolled	75.1%	75.4%	72.4%	72.7%	75.5%	74.9%
Bronze	26.1%	38.1%	26.0%	27.6%	27.1%	28.3%
Silver	64.7%	60.4%	64.8%	64.7%	64.7%	60.1%
Gold	5.3%	1.4%	5.3%	4.4%	5.0%	6.4%
Platinum	3.9%	0.1%	3.9%	3.3%	3.2%	5.1%
Change in Annual Per-Capita Social Welfare						
Cons. Surplus		$99	$32	$491	$16	$27
Profit		($245)	($19)	($349)	($118)	($165)
Government Spending						
Prem. Subsidies		($361)	($66)	($478)	($133)	($191)
CSRs		($16)	($8)	($14)	$0	($20)
Penalties		($4)	($289)	$28	($5)	$2
Uncomp. Care		($6)	$46	$52	($2)	($2)
Social Welfare		$346	($327)	$750	$66	$142

Notes: Table shows the impact of 5 policy changes using the 2016 learning parameter estimates. The policy changes include eliminating risk adjustment (scenario 7), eliminating the individual mandate (scenario 8), eliminating inertia (scenario 9), relaxing community rating (scenario 10), and requiring community rating (scenario 11). The second panel shows the effect on enrollee-weighted average premiums by metal tier, insurer, and plan network type. The third panel reports the impact on insurance coverage. The fourth panel shows the change in annual per-capita social welfare relative to the base scenario.

Table A6: Policy Simulations w/ Full Information Parameters

	Base	(12)	(13)	(14)	(15)	(16)
Scenario Definitions						
Risk Adjustment	✓		✓	✓	✓	✓
Individual Mandate	✓	✓		✓	✓	✓
Inertia	✓	✓	✓		✓	✓
Modified Comm. Rating	✓	✓	✓	✓		
Pure Comm. Rating						✓
Monthly Premiums						
Bronze	$298	$215	$300	$248	$285	$289
Silver	$405	$393	$407	$354	$389	$406
Gold	$443	$478	$447	$406	$413	$485
Platinum	$551	$494	$555	$479	$508	$608
Anthem	$383	$327	$385	$308	$363	$356
Blue Shield	$368	$342	$371	$331	$357	$379
Health Net	$412	$336	$414	$360	$388	$423
Kaiser	$457	$357	$459	$379	$427	$447
Other Insurer	$307	$255	$308	$288	$295	$318
HMO	$387	$322	$389	$317	$366	$395
PPO	$374	$331	$376	$324	$358	$365
Average	$378	$327	$381	$322	$361	$376
Subs. Average	$117	$91	$117	$97	$113	$124
Enrollment						
Total Coverage	1,789,239	1,799,366	1,736,538	1,738,001	1,795,498	1,785,686
% Enrolled	75.2%	75.6%	73.0%	73.1%	75.5%	75.1
Bronze	29.4%	38.5%	29.4%	33.1%	29.7%	32.9
Silver	63.9%	58.8%	64.0%	62.3%	63.8%	60.4
Gold	4.6%	1.6%	4.6%	3.9%	4.5%	4.6
Platinum	2.1%	1.2%	2.1%	0.7%	1.9%	2.2
Change in Annual Per-Capita Social Welfare						
Cons. Surplus		$102	$29	$463	$15	$1
Profit		($263)	($15)	($304)	($123)	($37)
Government Spending						
Prem. Subsidies		($214)	($53)	($389)	($116)	($97)
CSRs		($18)	($6)	($14)	($1)	($11)
Penalties		($5)	($288)	$25	($3)	$2
Uncomp. Care		($7)	$38	$47	($0)	$0
Social Welfare		$144	($334)	$655	$40	$106

Notes: Table shows the impact of 5 policy changes using the full information parameter estimates. The policy changes include eliminating risk adjustment (scenario 12), eliminating the individual mandate (scenario 13), eliminating inertia (scenario 14), relaxing community rating (scenario 15), and requiring community rating (scenario 16). The second panel shows the effect on enrollee-weighted average premiums by metal tier, insurer, and plan network type. The third panel reports the impact on insurance coverage. The fourth panel shows the change in annual per-capita social welfare relative to the base scenario.