BECKER FRIEDMAN INSTITUTE
FOR RESEARCH IN ECONOMICS
THE UNIVERSITY OF CHICAGO

Health Economics Series
No. 2017-05

Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA

Pietro Tebaldi

August 2017

Keywords: subsidies, health insurance, health reform, ACA, health exchanges
JEL Codes: I11, I13, I18, L51, H51, L88

Becker Friedman Institute
for Research in Economics

Contact:
773.702.5599
bfi@uchicago.edu
bfi.uchicago.edu

Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA*

Pietro Tebaldi+

August 11, 2017

Abstract. To design premium subsidies in a health insurance market it is necessary to estimate consumer demand, cost, and study how different subsidy schemes affect insurers' incentives. I combine data on household-level enrollment and plan-level claims from the Californian Affordable Care Act insurance exchange with a model of insurance demand and insurers' competition to assess equilibrium outcomes under alternative subsidy designs. I estimate that younger households are significantly more price sensitive and cheaper to cover. Consequently, counterfactuals show that providing more generous subsidies to this group leads to equilibria where all buyers are better off and per-person public spending is lower.

Keywords: subsidies, health insurance, health reform, ACA, health exchanges

JEL Classification Codes: I11, I13, I18, L51, H51, L88

*Financial support from the Kapnick Foundation through a grant to the Stanford Institute for Economics Policy Research and from the Becker-Friedman Institue Healthcare Initiative is gratefully acknowledged. Special thanks go to Liran Einav, Tim Bresnahan, Matt Jackson, and Mark Duggan for their support and research advise, and to four anonymous referees for their constructive feedback. Sofronis Clarides, Mark Shepard, Jonathan Williams, Martin Hackmann, Kate Ho, and Ben Handel provided helpful discussions. I benefited from research assistance from Hanbin Yang and Honglin Li at the University of Chicago, and I received constructive comments from Matt Backus, Lanier Benkard, Stephane Bonhomme, Francesco Decarolis, Michael Dickstein, Michael Dinerstein, Matthew Gentzkow, Josh Gottlieb, Michael Greenstone, Jon Gruber, Han Hong, James Heckman, Caroline Hoxby, Jon Levin, Davide Malacrino, Joe Orsini, Ali Hortacsu, Ariel Pakes, Luigi Pistaferri, Maria Polyakova, Paulo Somaini, Ben Sommers, Bob Town, Alonso Villacorta, Glen Weyl, Michael Whinston, and participants in numerous seminars and conferences. All remaining mistakes are mine.

+University of Chicago, Department of Economics. Email: ptebaldi@uchicago.edu.

1 Introduction

Welfare losses from adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976; Stiglitz, 1987; Einav, Finkelstein, and Cullen, 2010), consumption externalities (Pauly, 1970; Summers, 1989; Mahoney, 2015), and affordability considerations (Wagstaff and van Doorslaer, 2000; Bundorf and Pauly, 2006) justify the growing role of governments in regulating and supporting premium payments in private health insurance markets (Colombo and Tapay, 2004). However, despite the wide adoption of means-tested subsidies in what is typically referred to as "government-sponsored health insurance" (see also Einav and Levin, 2015; Decarolis, Polyakova, and Ryan, 2015; Jaffe and Shepard, 2017; Finkelstein, Hendren, and Shepard, 2017), the response of market outcomes to alternative subsidy schemes is still largely unexplored. A recent and large-scale example of such market design decision is found in the low-income subsidy (tax credit) introduced by the 2010 US health care reform (Patient Protection and Affordable Care Act; ACA). Between 2014-2016, under this program the federal government spent approximately $40 billion each year to provide discounts on health insurance premiums to more than 10 million US citizens (Anthony et al., 2015). Knowledge of the relationship between subsidy design and outcomes such as coverage levels and public spending is critical to evaluate the success of the ACA, and for the design of similar programs in the future.

In this paper, I study the dependence of equilibrium outcomes on how subsidies interact with three important features of private health insurance markets: demand from subsidized households, insurers' price competition, and adverse selection generated by the correlation between willingness-to-pay and expected health cost. Characteristics of demand determine the extent to which subsidies increase insurance enrollment. Pricing incentives and market power of imperfectly competitive insurers react to these changes in demand, but also to corresponding changes in expected cost driven by differences in the composition of enrollment pools.

To account for these effects, and compare different subsidy designs, I use data on individual-level enrollment and plan-level claims from the first year of the Californian ACA marketplace, in which over 90% of the 1.3 million buyers received federal subsidies. I discuss identification and estimation of demand and cost, exploiting details of the regulatory environment and variation in the composition of buyers across different contracts. I then use these estimates as inputs in a model of insurers' competition customized to ACA rating regulations. Within this framework, I study equilibrium pricing under different subsidy designs, comparing prices, enrollment, markups, and public spending. My results imply that the ACA subsidy scheme leaves room for improvements that are quantitatively significant and consistent with theoretical predictions. The alternatives I explore here emphasize the importance of increasing incentives for young and healthy buyers to participate, hence lowering average cost, prices, and public spending.

The paper makes three main contributions. First, I use individual-level administrative data from a large ACA exchange to estimate demand for coverage among the low-income uninsured. For identification, I leverage the ACA rating regulations, by which households in the same geographic market face identical choice sets but different annual premiums for (exogenous) regulatory reasons such as differences in age, income, or household size; this is similar to the approach of Ho and Pakes (2014); Geruso (2016). Existing research that estimates demand in the ACA context primarily exploits cross-sectional variation in aggregate prices and enrollment across geographic markets (Dafny, Gruber, and Ody, 2015; Abraham, Drake, Sacks, and Simon, 2017), or changes in uninsurance across different waves of representative surveys (Frean, Gruber, and Sommers, 2017; Sacks, 2017).[1] Adding to this work, my empirical strategy, together with the size and granularity of the California data, allows me to estimate a discrete-choice model with rich observable and unobservable heterogeneity in preferences.

Across logit, nested logit, and mixed logit specifications (Berry, 1994), I find that relatively younger (under-50) households are, on average, twice as price sensitive as their older (50-64) counterparts: the mean drop in demand if all premiums increase by $100/year is estimated to be between 2-4% for the former group, and approximately 1.5% for the latter. These average figures aggregate a large dispersion within each group. In particular, I find that the difference between under and over-50 is largely driven by a left tail in the distribution of coverage drop induced by a $100 lower subsidy. Documenting this heterogeneity across and within demographics is important, since in the insurance context differences in preferences can be related to differences in expected cost.

The second contribution of the paper is to estimate insurers' cost, combining the enrollment data and the demand model with information on realized plan-level average claims.[2] Upon observing individual-level enrollment, and using estimates of demand heterogeneity across households, cost differences across buyers are identified by projecting average claims on the composition of enrollment pools in terms of demographics and estimated willingness-to-pay for insurance. Emphasizing the importance of adverse selection in this market, I estimate a strong relationship between expected cost, household characteristics, and willingness-to-pay for coverage.

My estimates imply that a 10-year increase in the age of a buyer leads to a 20% increase in expected cost (with a 4:1 ratio from 21-64). Even conditioning on age, going from the bottom 20% to the top 20% of the distribution of willingness-to-pay corresponds to a 25% increase in expected cost. The resulting joint distribution of observables, preferences, and cost describes selection. Without controlling for household demographics, the estimated correlation between expected cost and the probability of purchase at a premium equal to $3,000/year (median cost in the population) is 0.44. After controlling for age, income, household size, and geographic region, the correlation

[1]More recent work uses individual-level data from small exchanges, or adopts different identification strategies (Panhans, 2017; Saltzman, 2017), without estimating a flexible demand model in which households with the same observable characteristics can differ in their willingness-to-pay for coverage.

[2]The source is 2016 regulatory rate-review filings from the Center of Medicaid and Medicare Services (CMS).

between probability of purchase and cost is as high as 0.81.

Despite the obvious advantage of estimating cost from observed claims, I also consider a different strategy, identifying cost parameters combining demand estimates with —instead of claims data— assumptions on how insurers set premiums to form a Nash equilibrium (c.f. Nash-in-prices; see also Lustig, 2010; Bundorf, Levin, and Mahoney, 2012; Starc, 2014; Ericson and Starc, 2015; Decarolis, Polyakova, and Ryan, 2015; Jaffe and Shepard, 2017).[3] The use of this strategy is twofold: on the one hand it can be used to study selection markets when cost data is not available,[4] on the other hand, as is relevant for this paper, it provides a heuristic criterion to support the supply assumptions for counterfactual analysis. Indeed, although a (static) Nash-in-prices assumption abstracts away from important institutional details of the ACA, and from behavioral and dynamic considerations on insurers' strategies, I find that cost estimates obtained under this assumption are similar to those obtained from realized claims (without relying on any assumptions on insurers' behavior). Although this falls very short from providing a formal test, it supports the extent to which analyzing the impact of different subsidy designs under Nash-in-prices can capture first-order incentives, and be informative from a normative perspective.

Following this, the third contribution amounts to comparing premiums, coverage, and spending under different subsidy designs. I use demand and cost estimates obtained from enrollment and claims data within a Nash-in-prices supply model adapted to the ACA regulations. My main result shows that a policy change a policy change that reduces subsidies to relatively older buyers (high-demand and high-cost) and increases subsidies to "young invincibles" (low-demand and low-cost) can make all buyers better off, increase profits, and lower government spending. Intuitively, shifting subsidy generosity from the high-cost, high-demand group to the low-cost, low-demand group changes the relative composition of enrollment pools, lowering average cost and increasing aggregate elasticity. This puts downward pressure on equilibrium prices, and increases quantity purchased for all groups while also reducing public spending. Since the group receiving a lower subsidy can also be made better off, the benefits of heterogeneous subsidization can be achieved while avoiding redistributive concerns.

Quantitatively, within my empirical framework I find that lowering subsidies for the over-50 households by \$25/month and increasing subsidies for the under-50 by \$50/month increases total enrollment by 13% (+149,000), and consumer surplus by 15% (+\$492 million), while average cost decreases by 9% (-\$315/year), and per-enrollee public spending by 10% (-\$307/year). Most importantly, in the new equilibrium with this change in monthly subsidies, all buyers are weakly

[3]The main intuition is that one needs to observe variation in the composition of marginal buyers across otherwise similar contracts. Variation in marginal revenues (obtained from the demand system) corresponding to variation in the composition of marginal buyers (also obtained from the demand system) identifies how buyers imply different expected costs for the insurer. Appendix A formalizes this argument.

[4]This was the case for an earlier version of this paper.

better off since they face lower net-of-subsidy premium. Over a wide range of variations in monthly subsidies, I find this mechanism to be robust: increasing participation of the young invincibles with higher discounts lowers premiums sufficiently to allow for lower discounts for older buyers without making this group worse off.

In my last exercise adopting a more standard utilitarian approach, I find the changes to ACA discounts (up to $50/month) that maximize consumer surplus, while keeping total spending lower than under the ACA. The solution to this problem implies again lower subsidies for over-50 and relatively high-income (250-400% of the FPL) households and higher subsidies for under-50 low-income (130-250% of the FPL). Compared to the equilibrium outcomes in my model of the ACA, this optimal combination of subsidies leads to a 2% decrease in total spending (-$58 million) and an 11% increase in consumer surplus ($356 million), while average cost drops by 7% ($243/year).

To carry on these counterfactual exercise, I consider a situation in which buyers are provided a fixed voucher varying with household size, age, and income. This is different from the "price-linked" subsidy used under the ACA, where discounts are computed as a function of premiums to guarantee that a specific (benchmark) plan[5] is affordable for a fixed fraction of income. Vouchers simplify my analysis since equilibrium computations are significantly more tractable. Moreover, from a theoretical perspective, fixed vouchers are less distortionary, since price-linked subsidies create extra incentives for insurers to charge higher markups. This distortion is the focus of Jaffe and Shepard (2017), where using estimates from the pre-ACA Massachusetts exchange they find that markups under price-linked subsidies are between 5-10% higher than under fixed vouchers. Corroborating their finding, in my context I find that going from the ACA price-linked subsidy to a fixed voucher (of equal amount) would lower average markups by 11%. Importantly, here I abstract away from implementation issues, whereas Jaffe and Shepard (2017) emphasize how price-linked subsidies require the government to know less about the market: regulators only need to choose the right "affordable amount", with no need to predict equilibrium under different voucher schemes. Despite this important consideration, they find that, even for a substantial degree of regulator's uncertainty on market primitives, and thus imperfectly calibrated vouchers, welfare is still higher than under a price-linked subsidy. I therefore use vouchers as my benchmark case, and focus instead on how the generosity of subsidies across different age groups directly impacts enrollment and spending.

My work also relates to several other studies that use different data and institutional contexts to study the interaction between regulations and welfare in private health insurance. Many use data from the Massachusetts' health exchange, a setting similar to ACA marketplaces, to study special incentives for the enrollment of young adults (Long, Yemane, and Stockley, 2010), the impact of product standardization and narrow provider networks on consumer choice (Ericson and

[5] The second-cheapest Silver in the region, where plans are grouped into Bronze, Silver, Gold, and Platinum.

Starc, 2012b,a, 2013, 2014), and adverse selection driven by differential preferences for high-quality medical providers (Shepard, 2017). Hackmann, Kolstad, and Kowalski (2015) measure the welfare effect of an insurance mandate (which in a simple model is akin to premium subsidies), and Ericson and Starc (2015) estimate a demand-supply model to study several ACA-like regulations, with the main focus on the effect of age-rating regulations in the market for unsubsidized, high-income buyers. Other studies combine theoretical results with simulations based on estimates from non-individual health insurance: the impact of insurance mandates and minimum coverage provisions is studied by Azevedo and Gottlieb (2017); the relationship between risk-adjustment and insurers' competition by Mahoney and Weyl (2014); the long-run welfare impact of community rating rules by Handel, Hendel, and Whinston (2015); the interaction between exchange design and labor markets by Aizawa (2015); the interaction between market structure and negotiated prices between insurers and physicians by Ho and Lee (2017); the scope for welfare gains from long-term insurance contract by Handel, Hendel, and Whinston (2017). Outside the ACA setting, the cases of Medicare Advantage, Medicare Part D, Medigap, and Medicaid are the main focus of, among many others, Duggan and Hayford (2013), Curto et al. (2014), Duggan, Starc, and Vabson (2014), Starc (2014), Clemens (2015). Notably, Decarolis (2015) shows distortions of insurers' decisions due to the design of subsidies in Medicare Part D, and Decarolis, Polyakova, and Ryan (2015) focus on welfare-optimal market (including subsidy) design in the same market. They consider the comprehensive question of how different types of subsidization and other government interventions impact social benefits and costs, finding that current Part D subsidies are overly generous toward buyers who, on the margin, place a lower value on the services than the social cost of their premium discounts. Despite the differences across institutional contexts, this fact is consistent with my findings here and with the recent work of Finkelstein, Hendren, and Shepard (2017), as we also find that the vast majority of low-income households are willing to pay for insurance less than their own expected cost.

2 ACA marketplaces

2.1 Institutional context and federal regulations

As of 2013, 17 percent of US citizens younger than 65 did not have health insurance coverage (Smith et al., 2014). In 2014, the ACA instituted health insurance marketplaces in each of the fifty states. A marketplace is a market in which private insurers offer a variety of coverage options, and the federal government provides subsidies for low-income participants. After the first three years of their operation, approximately 85 percent of the 18 million buyers in the marketplaces receive premium subsidies (Layton, Montz, and Shepard, 2017), with annual government disbursements of approximately $40 billion (Anthony et al., 2015).

ACA marketplaces operate in each state separately, but they all follow similar institutions and regulations. Each state is divided into geographic rating regions —groups of counties or zip codes— defining the level at which decisions by buyers and insurers take place (Dickstein et al., 2015). Every spring, insurers announce their interest in offering plans in each region in the subsequent calendar year. Entrants undergo a certification process, after which they offer different coverage options, classified into five coverage levels: Minimum Coverage, Bronze, Silver, Gold, and Platinum. Minimum Coverage indicates plans with very high deductible, which cannot be purchased by subsidized buyers, nor by buyers older than 35. The four metal tiers represent increasing generosity of insurance, measured (and advertised) as an estimate of the actuarial value of the plan: 60% for Bronze, 70% for Silver, 80% for Gold, and 90% or more for Platinum. Products and prices are set and made public at the end of every summer, and individuals can then compare and purchase plans in their region during the "open enrollment" period in the late months of each year. Coverage then lasts for the subsequent calendar year.

Pricing regulations. One important provision of the ACA is that insurers are not allowed to arbitrarily vary prices depending on buyers' observable characteristics. Characteristics that can affect annual premiums are the buyer's age (see also Orsini and Tebaldi, 2017) and, in some states, tobacco use, but even these adjustments are done in a pre-specified way.

For this, each plan j offered in region r is associated with a single base price b_{jr}. This is translated to age-specific (pre-subsidy) premiums using given age adjustment factors A^τ, equal for all products: when covering a buyer of age τ the insurer receives

$$P_{jr}^\tau = A^\tau \cdot b_{jr}. \tag{1}$$

Age adjustments start from 0.635 up to 20-year-olds, equal 1 for 21-year-old buyers, increase smoothly to 1.4 at age 45, and finally reach 3 at age 64. Details for all ages are shown in Figure A in the Online Appendix.

Premium subsidies. Although P_{jr}^τ is the premium received by the seller when a τ-year-old buyer enrolls in the plan, subsidies are provided for all households with annual income below four times the federal poverty level (FPL; approximately \$47,000 for a single individual). For this, the law establishes a cap on the premium amount the household should pay for the second-cheapest Silver plan (benchmark plan) in each region. This cap is a function of the household income (see Table 1), ranging —for single buyers— from \$684 per-year for the lowest income group to \$4,368 for the highest income group. Importantly, given income this cap amount does not vary with age.

Formally, this subsidy scheme defines a premium discount $V(\theta, b_r)$ available for each household with characteristics $\theta = (\tau, y, h)$, where τ is the average age of household members, y is income as percentage of the FPL, and h is household size. As a function of base prices b_r the discount is:

$$V(\theta, b_r) = \max\left\{h \cdot A^{\tau} \cdot \widehat{b}_r^{2s} - \overline{P}^{yh}, 0\right\}, \tag{2}$$

where $\overline{P}^{yh}$ is the premium cap for households of size h with income y, and $\widehat{b}_r^{2s}$ is the base price of the benchmark plan in the region. Since the law does not allow for negative prices, the price of plan j for a household with characteristics θ is equal to

$$P_{jr}^{\theta} = \max\left\{\underbrace{h \cdot A^{\tau} \cdot b_{jr}}_{\text{Insurer revenue}} - \underbrace{V(\theta, b_r)}_{\text{Federal subsidy}}, 0\right\}. \tag{3}$$

There are two properties of this subsidy scheme that are important for my empirical strategy and counterfactual analysis. First, for a given income level, households of the same size can find one Silver plan for exactly the same premium, independently of age. Second, the difference in premium, across both insurers and levels of coverage, increases in the age of household members, while it does not vary with income.

Cost-sharing subsidies. Another important regulation in ACA marketplaces is the provision of cost-sharing subsidies, available for households purchasing a Silver plan if their income is lower than 250% of the FPL. For them, the federal government covers part of deductible and out-of-pocket expenses, increasing the actuarial value of Silver plans from 70% to 95% for income levels between 100-150% of the FPL, 88% for income levels between 150-200% of the FPL, and 74% for income levels between 200-250% of the FPL.

Cost-sharing subsidies do not directly affect prices, yet make Silver plans more attractive the lower the household income (see also DeLeire et al., 2016).[6] A second consequence is that, although the insurer covers approximately 70% of the health expenses, buyers' utilization when enrolled in Silver plans will be as if the plans provided higher coverage, and therefore likely to be higher (c.f. "moral hazard" in health insurance, see e.g. Manning et al., 1987; Einav et al., 2013). The extent to which utilization is distorted and how insurers respond to this through premium adjustments pose an interesting question for future research. Since here I will estimate cost directly from average claims realized under cost-sharing subsidies, I will hold this regulation unaltered throughout my counterfactuals.

Risk adjustment, reinsurance, and risk corridors. The ACA introduced three programs to mitigate insurers' incentives to cream skim healthy patients, and to facilitate the stabilization of

[6] I emphasize this in Figure B in the Online Appendix, showing how the share of households purchasing Silver steadily declines as income increases, from over 0.9 to 0.4, with discontinuous jumps at the income cutoffs determining eligibility for cost-sharing subsidies.

the new markets. The programs are called risk-adjustment (permanent), re-insurance (2014-2016 only), and risk-corridors (2014-2016 only), often referred to as “the three R’s”.[7]

Risk-adjustment under the ACA determines monetary transfers from insurers with ex-ante relatively less risky enrollees to those with ex-ante relatively more risky enrollees. The main role of the program is then to mitigate insurers’ incentives to select the healthiest within market participants. Importantly, this is a budget neutral program. The government calculates these transfers through a risk-adjustment formula developed by the Department of Health and Human Services (see Kautter et al., 2014). If risk profiles (in terms of preexisting conditions, age, gender, tobacco use) do not differ across insurers, risk-adjustment implies no transfers. This differs from risk-adjustment in other government-sponsored markets such as Medicare Advantage and Medicare Part D. In these systems, when a high risk consumer picks a plan, the insurer receives a payment from the government regardless of the overall risk composition of the market. Such risk-adjustment compensates insurers for changes in risk at the extensive margin: if all enrollees are risky (e.g. relatively old), all insurers receive payments and have incentives to reduce premiums as a result. Under the ACA, instead, if enrollees distribute uniformly (in terms of risk) across competitors, there are no risk-adjustment payments and pricing incentives are not affected.

The omission of risk-adjustment should not have a major impact on my counterfactuals for the following reasons. First, I focus on policies that change the overall risk composition of the market rather than the distribution of risk across insurers. Second, in states like California, where most product characteristics are standardized and all insurers must offer all metal tiers, insurers have limited ability to differentially skim healthy buyers away from competitors. Indeed, as shown by Bindman et al. (2016), in the first year of the Californian exchange that I analyze here, the four largest insurers (who covered 91% of enrollees) had average risk scores varying from 0.98 to 1.03, and risk-adjustment transfers represented a very small fraction of total revenues. Nevertheless, it is still important to underline how the relevance of risk-adjustment under the ACA is likely to be very different in other states. Where insurers can freely choose which plans to offer as well as the characteristics of these plans such as deductible and out-of-pocket payments, the risk distribution across competitors can show a higher dispersion, increasing the importance of accounting for risk-adjustment when studying pricing incentives.

Re-insurance and risk-corridors are temporary programs facilitating market stabilization in the early years, reimbursing insurers for the ex-post realized riskiness of their pools independently from the one of their competitors. Re-insurance collects a fixed amount for every health insurance policy sold by any issuer in any market in the US,[8] and compensates every insurer for individual claims exceeding an attachment point ($45,000 in 2014-15, and $90,000 in 2016) until a cap of $250,000.

[7] See e.g. http://kff.org/health-reform/issue-brief/explaining-health-care-reform-risk-adjustment-reinsurance-and-risk-corridors/.

[8] State high-risk pools are excluded from the program. The annual per-contract amounts were set at $63 in 2014, $44 in 2015, and $27 in 2016.

The coinsurance rates are 100% in 2014, and 50% in 2015-16. By covering part of the right tail of risk, this program limits insurers' incentives to set high-premiums in fear of incurring losses due to the riskiness of the newly insured.

While re-insurance reimburses the cost of covering high-cost patients, risk-corridors are intended to facilitate the target of a 20% (variable) profit margin. Every insurer who (across all markets served in the state) does not spend at least 77% of premiums in claims and administrative costs must pay into the program. The payment is proportional to the difference between 80% of premiums and the amount spent. Symmetrically, every insurer who spends more than 83% is eligible for reimbursement, with amounts being again proportional to the difference between spending and the 80% target. Importantly, this program is not guaranteed to pay out, since it is possible that the payments due to less profitable insurers are larger than the dues of the more profitable ones.[9]

The role of the exchanges: active purchasers and clearinghouses. In terms of governance of the marketplaces, state regulators can decide whether to set up a state-based organization or to have local insurers offer plans through a federal-run market platform. There are two main models of governance: the exchange as an active purchaser (CA, CT, KY, MD, MA, NV, NY, OR, RI, VT), and the exchange as a clearinghouse (all other states). As illustrated in Krinn, Karaca-Mandic, and Blewett (2015), clearinghouses accept all health plans that meet published criteria, while active purchasers negotiate conditions for entry, premiums, provider networks, number of plans, and benefits. Overall, evidence on how the chosen model of governance affects outcomes is mixed, with the usual complications arising from cross-states comparisons; see Krinn, Karaca-Mandic, and Blewett (2015) and Scheffler et al. (2016).

As a leading example of active purchaser, the California exchange needs to approve entry of insurers in any given rating region,[10] and insurers who want to offer individual coverage must do so through the exchange. Additionally, the exchange imposes strict limits on the number of contracts that each insurer must offer, and opts for fully standardized combinations of deductibles and co-pays within each metal tier (details in Table 2). In this situation, within a rating region, insurers are differentiated only in their brand name, the structure of their provider networks, and the associated premiums. This will be important for my analysis, because deductibles and co-pays are exogenous plan characteristics not determined by insurers.

The process through which base prices (and thus premiums) are set represents a major difference between the two models of governance. In clearinghouses, the exchange has no role: rates are set by insurers as posted prices as long as they comply with existing (federal and local) regulations of

[9] For example, in 2014 insurers were due a total of $2.8 billion while only owing $362 million. Therefore the program paid only 12.5% of what was due to insurers who realized lower-than-expected variable margins.

[10] See for instance http://www.latimes.com/business/la-fi-obamacare-unitedhealth-20150116-story.html on the decision on UnitedHealth entry in year three of the market.

health insurance markets. On the other hand, when exchanges act as active purchasers, they have an important role in the process. Using the words of the administrators of the Californian exchange,[11] the exchange "jawbones down premiums to the extent it can, leveraging its private information on risk mix, competitor rates, and the price elasticity of demand". In my conversations with exchange staff, I learned that exchange officials provide participating insurers with estimates of how changes to base prices would affect the number and risk composition of their enrollees. Valuing market stability across years, the exchange does not want insurers to be surprised, hence announces when competitors are planning to raise or decrease rates, and provides carriers with market analyses to facilitate the pricing process.

Although I do not model the negotiations and information sharing between insurers and the exchange (which remains a very interesting topic for future work), the extent to which exchange officials directly helped insurers in forming beliefs about demand, cost, and competitors' strategies, is consistent with my finding that Nash-in-prices approximates pricing patterns in California fairly well. Indeed, over the first three years of the ACA, the Californian exchange has been remarkably stable in terms of risk composition, premiums, and insurers participation.[12] This is not the case when looking at clearinghouses, in which the early years of the market showed substantial instability, with major events of insurers entry, exit, and large price adjustments from one year to the other. Consequently, the analysis in this article might be more suitable, in terms of external validity, for stable exchanges with an active role along dimensions not constrained under the federal law.

2.2 Data from the 2014 Californian exchange

The California exchange (Covered California) is among the three largest ACA marketplaces, along with Florida and Texas. With over 1.5 million enrollees it accounted for over 13% of national enrollment in 2015.[13] Thanks to its size and standardization of plan characteristics (Table 2), this exchange provides a useful setup to estimate estimate the demand-cost model to quantify the effect of alternative subsidy designs.

Enrollment. The state is divided into 19 rating regions (map in Figure C in the Online Appendix), and in 2014-2015 there were between 3 and 6 active insurers in each region out of a total of 11 participants. The main data source for my analysis is an extract of the official records of Covered California, obtained via Public Records Act (CA Gov §6250). This contains non-identifiable individual-level enrollment information for every purchase in the exchange during the first open

[11] See http://healthaffairs.org/blog/2015/10/02/whither-health-insurance-exchanges-under-the-affordable-care-act-active-purchasing-versus-passive-marketplaces/.

[12] See http://healthaffairs.org/blog/2015/10/26/health-risk-continues-to-improve-in-covered-california/. In the Online Appendix Table A I show that demand estimates are indeed robust to using different years of enrollment and/or different definitions of potential buyers.

[13] See http://kff.org/health-reform/issue-brief/data-note-how-has-the-individual-insurance-market-grown-under-the-affordable-care-act/ .

enrollment period (October 2013 - April 2014). Each record shows unique individual and household identifiers, age, rating region of residence, bins for annual household income as % of the FPL, unique identifier for the selected plan (insurer, region, network type, metal tier), total premium paid by the household, and information on coverage termination (if occurred). To use these data, I first restrict my analysis to households with less than 6 members (99.5% of the buyers are in this group). Second, I compute a continuous measure of household income by inverting formula (3) for the subsidy calculation (I observe the discount and all elements of the formula to derive the price ceiling, which can then be mapped to the annual income of the household). This returns a complete individual-level dataset of 1,291,214 enrollment records (877,365 households),[14] summarized in panel (a) of Table 3.

Potential buyers. My second data source is an extract of the 2013 American Community Survey (ACS) accessed via IPUMS (Ruggles et al., 2015). For the analysis I use only the California sample. Applying household weights, I construct a dataset containing individual and household identifiers, health insurance coverage information, age, gender, household annual income, and geographic area that can be mapped directly to rating regions. I define potential buyers for Covered California as those who did not have any coverage or had individually purchased coverage at the end of 2013. This returns a dataset of 6,122,167 potential buyers (3,392,942 households),[15] summarized in panel (b) of Table 3. The table shows the difference between potential buyers and enrollees in the exchange: the latter are, on average, younger, live in smaller households, and have lower income when eligible for subsidies. Those with income higher than four times the FPL represent 35% of potential buyers in the ACS, but amount to less than 6% of exchange enrollees. For this reason this group plays a minor role in my empirical findings.

Average claims. The last data source consists of the rate-review filings collected by the Center for Medicare and Medicaid Serices (CMS).[16] Every year, insurers must justify their updates to base prices with "previous experience" in the market. In particular, pricing decisions for plans

[14] The dataset originally contains 1,450,477 person-purchase records. I discard all records with missing age or age over 64 (11,461), records with negative premium or subsidized premium higher than non-subsidized premium (935), records with multiple purchases per-household or different plan selections within a household (61,235), records with plan selections outside the region of residence (30,942), households with more than 5 members (7,295), and records with missing variables. As explained later, an additional 0.69% of households from areas with very low population are excluded because they cannot be matched to potential buyers in the ACS.

[15] This construction is consistent with external sources, see e.g. the Small Area Health Insurance Estimates (https://www.census.gov/did/www/sahie/), or estimates from the California Health Care Foundation (http://www.chcf.org/publications/2016/12/californias-uninsured). I omit combinations of demographic and micro-area in which exchange enrollment is larger than the number of potential buyers under my definition; this happens in small areas for which the measurement in the ACS is very imprecise, losing representativeness. This excludes from my analysis an additional 0.69% of households.

[16] See: https://www.cms.gov/CCIIO/Resources/Data-Resources/ratereview.html.

covering 2016 —taken during 2015— must be accompanied by average plan-level claims data from coverage during 2014. Therefore, for every plan still offered in 2016, the data summarizes 2014 enrollment and 2014 average incurred claims for which the insurer was responsible. Using these filings, I am able to match average claims for 476 out of 490 plans offered in 2014 Covered California, covering 98.9% of enrollment. The two missing insurers (each active in one region, for a total of 10 plans) covered only 15,313 individuals. Summary statistics for plan-level enrollment and claims are reported in panels (c) and (d) of Table 3, respectively.

Prices and choices by age and income My empirical strategy relies on variation in choices corresponding to exogenous variation in premiums induced by regulations, within a rating region, across households who face otherwise identical choice sets. The raw premium variation for each plan in Covered California is represented in Figure 1. For each region, the x-axis represents specific plans (e.g. Kaiser, HMO, Bronze), and each point in the scatter plot corresponds to the annual premium (y-axis) for a specific demographic group. To see more closely how premiums vary by age, income, and metal tier, premiums that are relevant for specific groups are summarized by panels (a) through (d) of Table 4. Panels (c) and (d) show that premiums for non-subsidized high-income older than 50 are approximately 3 times larger than those for their counterparts younger than 30, because of age-rating adjustment. These are equal to the amounts received by insurers. Because of the subsidy formula, however, this monotonicity does not hold for households eligible for subsidies.[17] For them, the ACA subsidy design implies that Silver plans are available for approximately the same amount for all ages; the second-cheapest Silver is available for *exactly* the same amount for a given income level and household composition. For lower coverage, premiums decrease in age, while the opposite is true for Gold and Platinum plans, a pattern that is mechanically implied by the formula of the subsidy scheme in equations (2) and (3).

Following these differences in premium, the data shows a large heterogeneity in choices, and the degree by which older individuals are more willing to pay for insurance than younger ones is key for my results. Evidence for this can be by comparing the age distribution of the potential buyers from the ACS data to the age distribution of Covered California enrollees; see Figure 2. This alone suggests the degree by which older households are more likely to purchase coverage in the exchange.

Figures 3 and 4 summarize further the relationship between premium, choice, and household characteristics. The key takeaway is that the data rejects a model in which preferences are invariant across demographics. If this was the case, conditioning on income, age would not affect participation decisions, since the level of prices for subsidized buyers is constant in age. At the same time, upon participating, older buyers would buy less generous coverage since the price difference between

[17] Here I focus on income between 200-220% of FPL, but recall that income only changes the level of prices, not the differences across tiers.

high and low-deductible plans increases in age. Both of these predictions fail in the data, providing evidence that older households are less price-sensitive and more willing to pay for insurance.[18]

Market shares and regional differences across carriers. To summarize market structure, Table 5 shows the share of regional enrollment specific to each carrier. The number of insurers in each region goes from three to six, for a total of eleven different participants. Four are large players —Anthem Blue Cross, Blue Shield, HealthNet, and Kaiser—, operating almost everywhere in the state and together covering 91% of Covered California enrollees. The remaining seven are smaller local insurers offering coverage only in a small number of regions.

Insurers are differentially attractive in different regions, following differences in provider networks, brand, and the set of competitors. Among the four largest carriers, each captures on average between 15-36% of subsidized buyers, yet these shares range from being negligible to a maximum of 35% for HealthNet, 50% for Kaiser and Blue Shield, and over 90% for Anthem. Part of these differences in the success of large insurers can be explained by the role played by local competitors. For instance, Chinese Community Health Plan captures up to 30% of buyers in the San Francisco Bay Area, while Sharp captures 10% of enrollment in San Diego.

3 Econometric model

3.1 Primitives

There are R geographic markets (or rating regions), indexed by $r = 1, ..., R$. In each region, J health insurance plans are offered to a population of households by N insurers, indexed by $n = 1, ..., N$.[19] For each n, $J_n \subset J$ is the set of products offered by n, and with a slight abuse of notation $n(j)$ denotes the insurer offering plan j. Within J_n, plans differ vertically in generosity of coverage as determined by deductibles, out-of-pocket maxima, and co-pays. Other insurer-region specific characteristics such as brand and provider networks are constant within a region for all plans offered by the same carrier.

[18] Specifically, for the extensive margin decision on whether or not to purchase coverage in the exchange, Figure 3 plots the share of potential buyers choosing to buy coverage (x-axis) against the average premium paid (y-axis). Each curve represents a different type of households, by age and size, and along each curve different points correpond to different income bins (the labels indicate the lower bound of income as a percentage of FPL). This highlights how, even conditioning on income, older households are significantly more likely to purchase a plan in Covered California despite paying the same premium. Figure 4 considers instead the intensive margin decision on whether to spend more to obtain higher coverage. (I distinguish between high-deductible Bronze plans and plans with lower deductibles, recalling that for most buyers cost-sharing subsidies imply that Silver is the most generous plan.) Per ACA regulations, older buyers face a higher premium difference between levels of coverage, but nevertheless the figure shows that they are more likely to pay this difference to lower the annual deductible.

[19] Having an equal number of insurers and products across regions is a simplification to keep notation uncluttered. This does not affect the empirical application.

A potential buyer (or household) i in region r is defined by a pair (θ^i, ν^i); superscripts are used throughout to index buyers, while subscripts index regions, insurers, and products. The type θ^i collects three household observables: $\theta^i = (\tau^i, y^i, h^i)$. τ^i is the average age of the individuals in the household, y^i is the annual household income (as % of the FPL), and h^i is the composition of the household (e.g. single, couple without kids, couple with kids).

Preferences of a buyer are described by the vector $\nu^i = (\nu^i_1, ..., \nu^i_J) \in \mathbb{R}^J$. This collects i's willingness-to-pay for each of the J products relative to the outside option $j = 0$. For buyers of type θ, $P^\theta_r = (P^\theta_{1r}, ..., P^\theta_{Jr})$ denotes the vector of differences between the price of each j and the price of the outside option (e.g. tax penalty for the lack of insurance). Household i chooses j when $\nu^i \in \mathcal{D}_j(P^\theta_r)$, where

$$\mathcal{D}_j(P^\theta_r) = \left\{ \nu \in \mathbb{R}^J : j \in \underset{k \in J}{\operatorname{argmax}} \left\{ \nu_k - P^\theta_{kr} \right\} \text{ , and } \nu_j \geq P^\theta_{jr} \right\}. \tag{4}$$

If a buyer (θ, ν) purchases plan j in region r, the insurer cost is $c^i \geq 0$ (total claims after risk-adjustment and re-insurance payments, if any). This is drawn i.i.d. across buyers from the distribution $\eta_{jr}(c|\theta, \nu)$. Differences in η across j and r may reflect different underlying contracts with health providers, differences in administrative and medical costs, differences in the generosity of coverage, or differences in expected utilization of health services. The dependence of η on (θ, ν) captures selection, allowing the distribution of cost to vary with buyers' characteristics and preferences.

3.2 Observables

In every region, the demographic composition of potential buyers is observed, with $G_r(\theta) \geq 0$ denoting the fraction of households with characteristics θ in region r. For each household the econometrician observes the corresponding choice, consisting of either a plan in the exchange or the outside option $j = 0$. It is then possible to construct, for every θ, the share of buyers choosing j in r among those of type θ, denoted as S^θ_{jr}.

All products are observed with a price P^θ_{jr} and a set of characteristics z^θ_{jr}, where these include the generosity of the plan (deductible, out-of-pocket maxima, and co-pays, possibly summarized in the actuarial value of the product; c.f. Table 2), as well as region and insurer identifiers. Within a region, prices of the same plan vary with θ because of subsidies, age-adjustments, and differences in household size. Plan characteristics may vary with θ because of cost-sharing subsidies, since households with lower income receive more generous insurance when purchasing Silver plans.

The last observable available to the econometrician consists of ex-post realized plan-level average claims. Specifically, given a plan j in region r, one observes C_{jr}, equal to average per-person medical claims paid by the insurer $n(j)$ during the coverage period.

3.3 Assumptions

I mantain the following assumptions on the data generating process:

Random utility model. For every household i in region r, the vector ν^i of willingness-to-pay for different plans is such that

$$\nu_j^i = \beta^i z_{jr}^\theta + \sigma^i \epsilon_{jr}^i, \tag{5}$$

where $\delta^i = \left(\beta^i, \sigma^i\right)$ is drawn i.i.d. from the distribution $F_r\left(\delta \middle| \theta^i\right)$, and ϵ_{jr}^i is an idiosyncratic preference shock drawn i.i.d from a Type I extreme value distribution. I assume the support of $F_r\left(\delta|\theta\right)$ is finite (Berry, Carnall, and Spiller, 1996; Train, 2008), and $\sigma^i > 0$ with probability one.

It follows that, conditional on $\theta^i = \theta$, the probability that household i chooses j in r is

$$\mathcal{S}_{jr}^\theta\left(P_r^\theta\right) = \sum_\delta F_r\left(\delta|\theta\right) \times \mathcal{S}_{jr}^{\theta,\delta}\left(P_r^\theta\right), \tag{6}$$

where

$$\mathcal{S}_{jr}^{\theta,\delta}\left(P_r^\theta\right) = \left(\frac{\exp\left(-\frac{1}{\sigma}P_{jr}^\theta + \frac{\beta}{\sigma}z_{jr}^\theta\right)}{1 + \sum_{k \in J}\exp\left(-\frac{1}{\sigma}P_{kr}^\theta + \frac{\beta}{\sigma}z_{kr}^\theta\right)}\right), \tag{7}$$

and the total market share of plan j in region r is $\sum_\theta G_r(\theta)\mathcal{S}_{jr}^\theta\left(P_r^\theta\right)$.

Different assumptions on the distribution $F_r|\theta$ correspond to several standard cases, including simple logit and nested logit. The standard logit with θ-specific coefficients corresponds to the case in which F_r is constant across regions and degenerate ($|\text{Supp}F_r|\theta| = 1$). The nested logit with two nests, one of which containing the outside option $j = 0$, corresponds to the case in which only the constant term in z_{jr}^θ has a random coefficient, and all other product characteristics have deterministic coefficients (see Berry, 1994).

Cost model and independence from logit errors. I allow cost to depend on observable demographics θ and on the preference parameters δ, but I assume that it is independent from the idiosyncratic preference shocks collected in ϵ. Formally,

$$c \perp \epsilon, \text{ implying that } c \sim \eta_{jr}\left(c|\theta,\nu\right) = \eta_{jr}\left(c|\theta,\delta\right). \tag{8}$$

I also assume a parametric form for the way in which expected cost may vary with buyer, product, and market characteristics. For a given function ϕ with unknown parameter κ:

$$\int c\, \mathrm{d}\eta_{jr}\left(c|\theta,\delta\right) = \phi(\theta,\delta,z_{jr};\kappa). \tag{9}$$

One can then rewrite the cost for the insurer when a household of type (θ, δ) chooses j in r as $c^i = \phi(\theta, \delta, z_{jr}; \kappa) + \widetilde{\omega}^i_{jr}$, where $\widetilde{\omega}^i_{jr}$ is a mean-zero error term independent from household characteristics.

With this notation, the observed average cost C_{jr} can be written as:

$$C_{jr} = \sum_{\theta,\delta} \left(\frac{G_r(\theta) F_r\left(\delta|\theta\right) \mathcal{S}^{\theta,\delta}_{jr}\left(P^{\theta}_r\right)}{\sum_{\widehat{\theta}} G_r\left(\widehat{\theta}\right) \mathcal{S}^{\widehat{\theta}}_{jr}\left(P^{\widehat{\theta}}_r\right)} \right) \times \phi(\theta, \delta, z_{jr}; \kappa) + \omega_{jr}, \tag{10}$$

where ω_{jr} is i.i.d. mean-zero. In words, up to an error term, the observed average cost of plan j in region r is a convex combination of expected cost conditional on household type (including both demographics and preference parameters) with weights equal to the relative share of different household types choosing the plan. The extent to which ϕ varies with θ and δ captures selection, allowing buyers with different characteristics and preferences for insurance to imply different expected cost when purchasing the same plan.

4 Identification

The primitives to be identified (and estimated) are the collection of $F_r|\theta$ for all r and θ —determining the demand system—, and the cost parameters collected in κ.

4.1 ACA regulations and within-region demand identification

The institutional details of Covered California provide an argument for identification of demand. Within a region, all buyers face identical choice sets, including unobserved features of available options. These might include advertisement as well as provider networks associated with each plan. At the same time, as discussed in Section 2, ACA subsidies and rating regulations generate granular variation in prices across households with different age, size, and income. Similarly, cost-sharing subsidies imply that the generosity of Silver plans —summarized by the actuarial value— varies discontinuously when the household's income crosses the thresholds of 150, 200, and 250% of the FPL.

Taken together, these regulations generate exogenous variation in prices and product characteristics —not corresponding to variation in unobservables— *within* the region-plan level, at which non-price supply decisions such as entry, advertising, or determination of provider networks take place. This institutional feature is very convenient for identifying demand parameters, since it is not necessary to rely on variation in prices and choice sets across different markets. Simply put, one can think of estimating demand with fixed-effects at the plan-region level, capturing the unobservables that are likely to affect pricing decisions. A similar approach is adopted in Ho and Pakes (2014); Geruso (2016), and it is originally outlined in Chamberlain (1980).

With this intuition, I consider alternative assumptions on $F_r|\theta$, and corresponding sufficient conditions for identification of demand. All assumptions serve the purpose of imposing restrictions on the way in which preferences may vary with the household observable characteristics θ. With these restrictions, variation in choice probabilities S_r^θ corresponding to the (exogenous) variation in P_r^θ and z_r^θ across θ's identifies demand. In choosing between alternative assumptions, the trade-off is between allowing preferences to vary flexibly with θ on the one hand, and preserving identifying variation in (P_r^θ, z_r^θ) on the other.

One possibility is to assume that, within a given group of observables (e.g. all single buyers, aged 20-25, earning between 15,000-17,500 USD, and living in Los Angeles), buyers have the same (distribution of) preferences. Formally, for a partition of observables in disjoint sets Θ^t, with $t = 1, ..., T$, the assumption is that, if $\theta, \theta' \in \Theta^t$, $F_r|\theta = F_r|\theta' = F_r^t$. The condition for identification is then intuitive, and it follows directly from the arguments in Thompson (1989); Berry and Haile (2014): for each t, F_r^t is identified if the support of $\left(P_r^\theta, z_r^\theta \middle| \theta \in \Theta^t\right)$ is sufficiently large; specifically, larger than the support of F_r^t.

I adopt this assumption for my logit and nested logit specifications. For these models, I let the preference parameters δ vary across households of different (average) age in bins of 10 years, household size, and subsidy eligibility. Moreover, for the nested logit specification, I assume that the characteristics in θ do not affect the distribution of the random coefficient on the constant term in z_r^θ.

A second alternative is to assume a parametric form for how $F_r|\theta$ varies with θ. If for a given parametric family $\psi(\theta;\mu)$, known up to a finite-dimensional parameter μ, $F_r|\theta = \psi(\theta;\mu_r)$, the entire variation in $\left(S_r^\theta, P_r^\theta, z_r^\theta\right)$ across θ's and j (holding r fixed) can be exploited to identify μ_r. In this case μ_r is the only unknown of the demand system: if $F_r(\delta|\theta) = \psi^\delta(\theta;\mu_r)$,

$$S_{jr}^\theta = \sum_\delta \psi^\delta(\theta;\mu_r)\mathcal{S}_{jr}^{\theta,\delta}\left(P_r^\theta\right), \tag{11}$$

where variation in (P_r^θ, z_r^θ) across θ's maps directly into corresponding variation in $\mathcal{S}_{jr}^{\theta,\delta}\left(P_r^\theta\right)$. This is also the framework considered in Fox, il Kim, and Yang (2016). Identification of μ_r relies on standard invertibility conditions for identification of a nonlinear parametric model (Newey and McFadden, 1994). If, moreover, the function ψ is continuously differentiable in θ, discontinuities in the mapping $\theta \to (P_r^\theta, z_r^\theta)$ (or in its derivatives) provide an additional source of identification of the parameter vector μ_r. Such discontinuities are present in the ACA regulations, and are exploited by DeLeire et al. (2016) to study demand responses to changes in insurance generosity. A similar argument have also been adopted by Ericson and Starc (2015), and more recently by Finkelstein, Hendren, and Shepard (2017), in the context of the pre-ACA Massachusetts health insurance exchange, where the regulation implied discontinuities in premium across specific age and income thresholds.

I consider this alternative assumption in my richest specification, a finite-type mixed logit (Berry, Carnall, and Spiller, 1996; Train, 2008) in which preference parameters δ can belong to different latent classes. As in Bhat (1997), I assume that the probability of being in each class varies with demographics θ following a multinomial logit model with unknown parameter μ_r:

$$F_r(\delta|\theta) = \psi^{\delta}(\theta;\mu_r) = \frac{\exp\left(\mu_r^{\delta}\theta\right)}{\sum_{\delta'}\exp\left(\mu_r^{\delta'}\theta\right)}. \tag{12}$$

Heterogeneity in demand parameters within similar demographics is identified off violations of the IIA assumption imposed by the standard logit and nested logit specifications.[20] When comparing θ's with similar prices and characteristics —ideally identical, as it is the case for 20-22 year old buyers—, differences in choice probabilities are informative about how the distribution of class membership varies with θ.

4.2 Cost identification from average claims

A challenge for identification of the model's primitives is that insurers' costs may not vary only across products and markets; this would be standard in the Bresnahan (1981); Berry, Levinsohn, and Pakes (1995) framework. Instead, to model selection it is important to allow the (expected) cost of the insurer to vary also across different buyers, depending on characteristics such as observable demographics and preferences.

For identification of cross-buyer cost heterogeneity, I leverage the plan-level cost information, the detailed composition of plans' enrollment in terms of demographics, and estimates of the composition of plans' enrollment in terms of preferences. This is similar to Bundorf, Levin, and Mahoney (2012). Intuitively, to identify how different buyers (e.g. over 50 and under 50) imply a different cost for the same insurer-plan pair (e.g. Silver plan offered by Anthem), one needs to observe plans with similar characteristics (e.g. two Silver plans offered by Anthem in two similar —ideally identical— markets) but different composition of enrollment (e.g. only over 50 buyers in one market and only under 50 buyers in the other). Differences in realized average costs across these plans is then informative about differences in cost across different buyers.

More formally, the left-hand side of equation (10) is observed in the data. When demand is identified, the right-hand side is also fully known up to the parameter κ. Then, given the assumed independence between c and ϵ, identification of κ follows standard rank conditions. In particular, to identify the difference in expected cost between two groups of buyers, say (θ, δ) and (θ', δ'), it is necessary to observe at least two products jr and $j'r'$ with the same product and

[20] This is also discussed in Berry, Carnall, and Spiller (1996): IIA requires that the values of δ estimated using the relationship between prices and choices considering only Bronze and Silver plans, for instance, are equal to the δ estimated considering Bronze and Platinum plans. Violations of this equivalence is informative about how buyers who purchase Platinum have different preferences from those purchasing lower coverage.

market characteristics ($z_{jr} = z_{j'r'}$), and the following matrix (representing the number of buyers by products and groups) must be full column rank:

$$\begin{bmatrix} \frac{G_r(\theta)F_r(\delta|\theta)\mathcal{S}_{jr}^{\theta,\delta}\left(P_r^{\theta}\right)}{\sum_{\widehat{\theta}} G_r\left(\widehat{\theta}\right)\mathcal{S}_{jr}^{\widehat{\theta}}\left(P_r^{\widehat{\theta}}\right)} & \frac{G_r(\theta')F_r(\delta'|\theta')\mathcal{S}_{jr}^{\theta',\delta'}\left(P_r^{\theta'}\right)}{\sum_{\widehat{\theta}} G_r\left(\widehat{\theta}\right)\mathcal{S}_{jr}^{\widehat{\theta}}\left(P_r^{\widehat{\theta}}\right)} \\ \\ \frac{G_{r'}(\theta)F_{r'}(\delta|\theta)\mathcal{S}_{j'r'}^{\theta,\delta}\left(P_{r'}^{\theta}\right)}{\sum_{\widehat{\theta}} G_{r'}\left(\widehat{\theta}\right)\mathcal{S}_{j'r'}^{\widehat{\theta}}\left(P_{r'}^{\widehat{\theta}}\right)} & \frac{G_{r'}(\theta')F_{r'}(\delta'|\theta')\mathcal{S}_{j'r'}^{\theta',\delta'}\left(P_{r'}^{\theta'}\right)}{\sum_{\widehat{\theta}} G_{r'}\left(\widehat{\theta}\right)\mathcal{S}_{j'r'}^{\widehat{\theta}}\left(P_{r'}^{\widehat{\theta}}\right)} \end{bmatrix}. \tag{13}$$

This requirement extends naturally to higher dimensions of heterogeneity, where one can think of a cost type K as a set of θ's and δ's for which it is assumed that $\phi(\theta, \delta, z_{jr}; \kappa) \equiv \phi^K(z_{jr}; \kappa)$. Differences in cost types are identifiable if and only if one observes sufficient variation in the composition (in terms of K's) of enrollment across similar (ideally identical) products.

Ultimately, one needs then to observe, within a class of plans with otherwise equal cost functions, shifters of buyers' composition that are excluded from cost functions. Examples of these shifters include variation in the set of competing plans (a version of "BLP instruments", c.f. Berry, Levinsohn, and Pakes, 1995), or variation in the composition of potential buyers in terms of demographics affecting demand (a version of "Waldfogel instruments", c.f. Waldfogel, 2003).

4.3 Cost identification from equilibrium assumptions

One important advantage of observing both, choices and claims, is that one can estimate the demand-cost system allowing for cost heterogeneity across buyers without making assumptions on insurers' behavior (Einav, Finkelstein, and Cullen, 2010). However, supply side assumptions on, for example, how insurers set prices, are still needed to analyze (and quantify) market responses to counterfactual policies that are not directly observed in the data. A different consideration is that data on premiums, plan characteristics, and enrollment are often easier to access than claims data. Often privacy concerns pose severe limits to the diffusion of individual-level claims, and negotiated prices between insurers and providers are a key pieces of private information that regulators "leave" to the supply side of a health insurance market.[21]

There is then a twofold interest in discussing identification of expected costs in a health insurance market using, instead of claims data, assumptions on the way in which insurers set prices. First, as I will do in this article, comparing cost estimates obtained from claims data to those obtained under supply side assumptions can shed light on the validity (or lack thereof) of these assumptions for

[21] Access to comprehensive claims data is particularly challenging when considering larger markets in which several private insurers compete. Consequently, many prominent papers in the literature are based on proprietary claims datasets from a single employer (e.g. Einav, Finkelstein, and Cullen, 2010; Handel, Hendel, and Whinston, 2015). Important exceptions in the recent literature include Curto et al. (2014) in the context of Medicare Advantage, Ho and Lee (2017) looking at insurance options offered to public employees in California, and Shepard (2017) and Finkelstein, Hendren, and Shepard (2017) in the context of the pre-ACA Massachusetts health insurance exchange.

counterfactual analyses. Second, when claims data are not available, the same approach enables one to follow the tradition of the empirical industrial organization literature (c.f. Rosse, 1970; Bresnahan, 1981; Berry, Levinsohn, and Pakes, 1995), and carry on cost estimation as well as counterfactual analyses under maintained supply side assumptions.[22]

A tractable, widely used assumption on insurers' pricing behavior is the so-called Nash-in-prices assumption, by which every insurer sets premiums to maximize expected profits during the coverage period taking the competitors' decisions as known and fixed. This is adopted in Bundorf, Levin, and Mahoney (2012), Starc (2014), Curto et al. (2014), Ericson and Starc (2015), Decarolis, Polyakova, and Ryan (2015), Jaffe and Shepard (2017), and later in this article.

In its simplest version, with single-plan insurers and uniform pricing across θ, Nash-in-prices implies that in every region r, and for every product j,

$$\frac{\partial \Pi_{jr}(P_r)}{\partial P_{jr}} = 0\,, \quad \text{where} \tag{14}$$

$$\Pi_{jr}(P_r) = \sum_{\theta,\delta} G_r(\theta) F_r\left(\delta|\theta\right) \mathcal{S}^{\theta,\delta}_{jr}\left(P_r\right)\left(P_{jr} - \phi(\theta,\delta,z_{jr};\kappa)\right). \tag{15}$$

The first-order condition in (14) can be rewritten as an equality between (expected) marginal revenues and (expected) marginal cost: $MR_{jr}(P_r) = MC_{jr}(P_r;\kappa)$, where

$$MR_{jr}(P_r) = \sum_{\theta} G_r(\theta)\left(S^{\theta}_{jr}(P_r) + P_{jr}\frac{\partial S^{\theta}_{jr}(P_r)}{\partial P_{jr}}\right)\,, \text{ and} \tag{16}$$

$$MC_{jr}(P_r;\kappa) = \sum_{\theta,\delta} \underbrace{\left(G_r(\theta) F_r\left(\delta|\theta\right) \frac{\partial S^{\theta,\delta}_{jr}(P_r)}{\partial P_{jr}}\right)}_{\text{Marginal buyers of type } (\theta,\delta)} \phi(\theta,\delta,z_{jr};\kappa). \tag{17}$$

Importantly, with risk-neutral insurers, the marginal cost of a plan is a linear combination of cost across different types of buyers, with weights equal to the (measure of) marginal buyers of each given type. With this observation in hand, the argument for identification of κ when demand is identified is intuitive.

While one needs variation in the composition of enrollment pools across products with the same cost functions when using observed average claims, when imposing Nash-in-prices one needs variation in the composition of marginal buyers. The intuition is as follows: suppose there are two products assumed to have the same cost functions (or equivalently the same z_{jr}), and the premium of both products increase by \$1, if all buyers reacting to the price change of one product are "old", and all buyers reacting to the price change of the other product are "young", the difference in

[22]In the health insurance context this approach was adopted by Lustig (2010), and in an earlier version of this article. Several other papers circumvented the need for claims data using representative health expenditure surveys (see e.g. Starc, 2014).

marginal revenues (obtained from demand estimates) between the two products is equal to the difference in cost between "old" and "young".

More formally (a complete, self-contained argument can be found in Appendix A), assume again that the space of possible (θ, δ) is partitioned in $\mathcal{K}$ subsets (or cost types), each denoted by K. When $(\theta, \delta), (\theta', \delta') \in K$, by assumption $\phi(\theta, \delta, z_{jr}; \kappa) = \phi(\theta', \delta', z_{jr}; \kappa) = \phi^K(z_{jr}; \kappa)$. Then κ is identified if, for any given z_{jr}, there are at least $\mathcal{K}$ products $j'r'$ with $z_{j'r'} = z_{jr}$, and for which the following matrix describing the composition of marginal buyers across cost types and products is full column rank:

$$\begin{bmatrix} \sum_{(\theta,\delta)\in K^1} G_r(\theta) F_r(\delta|\theta) \frac{\partial S^{\theta,\delta}_{jr}(P_r)}{\partial P_{jr}} & \dots & \sum_{(\theta,\delta)\in K^{\mathcal{K}}} G_r(\theta) F_r(\delta|\theta) \frac{\partial S^{\theta,\delta}_{jr}(P_r)}{\partial P_{jr}} \\ & \dots & \\ \sum_{(\theta,\delta)\in K^1} G_{r'}(\theta) F_{r'}(\delta|\theta) \frac{\partial S^{\theta,\delta}_{j'r'}(P_{r'})}{\partial P_{j'r'}} & \dots & \sum_{(\theta,\delta)\in K^{\mathcal{K}}} G_{r'}(\theta) F_{r'}(\delta|\theta) \frac{\partial S^{\theta,\delta}_{j'r'}(P_{r'})}{\partial P_{j'r'}} \end{bmatrix}. \tag{18}$$

That is, to identify cost differences across buyers, one needs shifters of the composition of marginal buyers across products with the same cost functions, where this shifters must affect cost only through the composition of buyers. As before, natural examples of these shifters are differences in competing products and differences in the composition of potential buyers.

5 Estimation Results

5.1 Demand estimates

Logit models without unobserved heterogeneity. After aggregating the individual-level choices in 200 different demographic groups (interactions of age, income, and household size), I estimate logit and nested logit models with different sets of of fixed-effects and controls. The estimating equation follows directly from Berry (1994):

$$\ln(S^\theta_{jr}) - \ln(S^\theta_{0r}) = -\alpha^\theta P^\theta_{jr} + \beta z^\theta_{jr} + \chi^{under50}_{n(j)r} + \chi^{over50}_{n(j)r} + \epsilon^\theta_{jr} \;\; [+\lambda \ln(S^\theta_{jr}|j \neq 0)], \tag{19}$$

where P^θ_{jr} is the (average) premium of plan j in region r for households in group θ, z^θ_{jr} describes the generosity of coverage (advertised actuarial value, or a vector of deductible, maximum out-of-pocket, and primary-care-physician visit), and χ_{jr} is a region-insurer fixed-effect, where I distinguish between over and under-50. In the most conservative specification, I include a full interaction of region-insurer-tier fixed-effect, and also distinguish between over and under-50. In the nested logit specifications, I include on the right-hand side the share of plan j in r conditional on choosing a plan, $\ln(S^\theta_{jr}|j \neq 0)$, and λ is the corresponding nesting parameter.[23]

[23] λ must vary between $[0, 1]$: as $\lambda \to 0$ one obtains the standard logit, so that the substitution between inside goods and the outside option is not different from the substitution between other goods; conversely, as $\lambda \to 1$, the

The resulting estimates under alternative sets of fixed-effects and controls for plan generosity are reported in Table 6.[24] All premium coefficients are statistically significant, and robust across specifications. The pattern is clear: the willingness-to-pay for a given plan is increasing in age, household size, and household income. Quantitatively, the ratio of the estimated α^θ between singles aged 20-40 and those aged 50-64 implies that the latter group is willing to pay approximately twice as much for the same insurance plan.

Estimates of the nesting parameter λ vary between 0.4 and 0.5, implying that there is a significant degree of heterogeneity across households in the willingness to consider a plan in the exchange.[25] Focusing on the nested logit specifications, the ratio β/α^θ corresponds to the willingness-to-pay for an increase in coverage generosity. This is estimated to be, on average, \$317 per-year for a 20% increase in actuarial value (approximately \$2,000 reduction in annual deductible) among subsidized single buyers aged 20-40. This same quantity doubles for subsidized single buyers aged 50-64.

In Table 7, I report the extensive margin semi-elasticities implied by the nested logit demand system in column (7), the specification with insurer-region fixed-effects. For different demographic and income groups, the table shows the average percent change in the probability of purchasing a plan in the exchange if all premiums increase by \$100 per-year. Reflecting the estimated heterogeneity in willingness-to-pay and price sensitivity, the estimated semi-elasticity is also decreasing (in absolute magnitude) as households get older, richer, and more numerous.

I estimate that for households without children who are younger than 40 the probability of purchasing (any) coverage drops by almost 3% if all premiums were raised by \$100/year. On the other hand, this number drops to less than 1%, on average, for households without children with average age greater than 50. Both unsubsidized households and households with children are less sensitive to premium increases.[26]

A second important property of the demand system is the substitution across different tiers when only the premiums of certain products change. To explore this, Table 8 shows the cross-

---

only choice consumers make is between inside goods and outside goods, while other distinctions between different alternatives become irrelevant.

[24]For further robustness, the Online Appendix contains estimates from the same specification where I use two years of enrollment data and alternative waves of the American Community Survey to calculate the number of potential buyers in each demographic and geographic cell.

[25]To account for this heterogeneity is important; omitting it, the simple logit estimates imply that generosity of coverage has zero (or negative) value for the average household. This is to reconcile the assumption of similar substitution patterns between all options (inside and outside of the exchange), a large share of households choosing not to purchase in the exchange, and other households who pay additional premiums to increase their coverage. Once allowing for the more flexible substitution pattern between exchange plans and outside option, the coefficient on the generosity of coverage becomes positive, statistically significant, and robust across specifications.

[26]The decline in probability of purchase among the unsubsidized when premiums increase by \$100 is between 2.2-2.6% for the younger groups, but 0.2-0.6% for households with average age greater than 50. For households with children, I estimate a semi-elasticity of 1.4% among the subsidized and 0.7% among the unsubsidized.

tier substitution patterns for different demographic groups. Specifically, each cell indicates the percentage change in the probability of purchasing a plan in a given metal tier (column) when premiums of all plans in another metal tier (row) increase by \$100/year. For the outside option, an increase in premium amounts to a decrease in premium of all the inside products, and —assuming the outside option is equivalent to uninsurance— this is conceptually equivalent to an increase of a \$100 in the mandate penalty.

The heterogeneity in response to premium changes across different groups is again very stark. The own-product semi-elasticity varies widely across groups, from 4-5%, on average, for young households without children, to less than 2% for other households. The nested logit specification also captures the important difference between the substitution across tiers and the much smaller substitution between exchange plans and outside option. Even with the flexibility of the nested logit, however, the substitution between different coverage tiers within the market is still very restricted by the IIA assumption imposed across (inside) plans within each demographic group.[27]

Mixed logit with unobserved heterogeneity. In the richest demand system, a buyer is represented by a vector of parameters $\delta = (\alpha, \beta, \chi)$, including a premium coefficient (α), a coefficient on (advertised) actuarial value (β), and a set of insurer dummies (χ^n; distinguishing between Anthem, Blue Shield, HealthNet, Kaiser, and grouping together other minor insurers).

In region r, δ can take four different values $\delta_r^1, ..., \delta_r^4$, with the normalization $\beta_r^1 = 0$. The probability that a buyer with observable demographics θ has demand parameters δ_r^k in r is

$$F_r(\delta_r^k|\theta; \mu_r) = \frac{\exp\left(\mu_r^k \theta\right)}{\sum_{k'=1}^{4} \exp\left(\mu_r^{k'} \theta\right)}, \tag{20}$$

and the likelihood of a buyer with characteristics θ choosing j in region r is then

$$\ell_{jr}^{\theta}\left(\left(\delta_r^k, \mu_r^k\right)_{k=1}^{4}\right) = \sum_{k=1}^{4} F_r(\delta_r^k|\theta; \mu_r) \frac{\exp\left(-\alpha^k P_{jr}^{\theta} + \beta^k av_{jr}^{\theta} + \chi^{k,n(j)}\right)}{1 + \sum_{j' \in J_r} \exp\left(-\alpha^k P_{j'r}^{\theta} + \beta^k av_{j'r}^{\theta} + \chi^{k,n(j')}\right)}. \tag{21}$$

Using the individual level choice data, I estimate the parameters $\left(\delta_r^k, \mu_r^k\right)_{k=1}^{4}$ separately, region-by-region, using the Expectation-Maximization algorithm (Train, 2008).[28]

The key advantage of this specification is to relax the cross-demographic and cross-region restrictions imposed in the logit and nested logit models. There are no restrictions in how preferences

[27] Indeed, the (small) asymmetries in the substitution between Bronze and Silver and, say, Bronze and Platinum within each panel of Table 8 result from variation in observed demographics within group (e.g. differences between 20 and 40 year olds, all grouped in panel (a)). The mixed logit demand system that I consider next relaxes these restrictions, allowing for more flexible substitution patterns across all options, and for richer heterogeneity in the way in which specific characteristics of products affect the willingness-to-pay of different households.

[28] This can be easily implemented using the command *lclogit* in Stata®, see Pacifico (2013).

can vary across regions. Within a given region, households can be one of four preference types, each with a separate willingness-to-pay for coverage and for the insurer-region specific characteristics (provider networks). Rather than directly determining preferences, demographics affect the probability that a given household has a specific preference type according to (20).

Table 9 reports a summary of the parameters (α, β, χ), varying across regions and demographic groups. As expected, there is significant variation across households in terms of willingness-to-pay for generosity of coverage, as well as significant differences in the taste for a specific insurer. The variation in the estimated parameters, however, is not alone very informative about the economic properties of the demand system. To illsutrate these properties, I start by looking at the implied distribution of wilingness-to-pay for a 20% increase in actuarial value of a plan; this is simply $WTP \equiv 20 \times (\beta/\alpha)$. This corresponds to an upgrade from Silver to Platinum for unsubsidized buyers, or from Bronze to Silver for buyers who benefit of cost-sharing subsidy (see Table 2).

In Figure 5, I plot the prior distribution of WTP across different types of household based on the mixed logit estimates. This shows the dispersion within and across demographic groups. Focusing on subsidy-eligible singles and couples without children (representing over 75% of exchange enrollment, and therefore driving most of my empirical findings), the difference in the WTP distribution between under-50 and older households is very stark. The mean prior WTP among under-50 is \$566/year, and increases to \$1,371 among the over-50. Both distributions are skewed to the left, with medians of \$695 and \$1,404, respectively. Within each group, the interquantile range is wider than \$600/year, going from \$433 to \$987 among the under-50, and from \$1,056 to \$1,773 among the over-50. Using the mixed logit estimates, one can also compute the posterior distribution of preferences conditional on observed choice. This is particularly important here, since it allows me to distinguish the composition, in terms of preferences (WTP) between those choosing not to purchase, those choosing low-coverage options, and those choosing to pay higher premium for more comprehensive coverage. Ultimately, in a selection market one expects these groups to have different risks and imply different costs for the insurer. Figure 6 shows the resulting distinction between the preferences of buyers and the preferences of those choosing not to purchase coverage in the exchange. The shift to the right of the WTP distribution conditioning on purchasing a plan in the exchange is large: the estimated median WTP among those purchasing coverage is \$600 higher (from \$1,000 to \$1,600), and the 25th percentile moves from \$185 to \$1,000.

To design subsidies, a key property of demand is the extensive margin semi-elsticity. Figure 7 illustrates how different age groups respond differently to a \$100 per-year increase in all premiums, and the dispersion within each group. (Here I focus on subsidized households, covering 95% of total enrollment; the Online Appendix reports results for all groups.) Similar to the logit and nested logit cases, the average response to a \$100 increase in the level of premiums (equal to a \$100 reduction in subsidies) is much larger for the relatively younger among those without children. The average change in total demand is equal to approximately -3.5% among those younger than 45, and it increases to -1.5% among those older than 55. The driver of these differences in average demand responses is the presence of a long left tail in price sensitivity among comparatively younger households, which progressively vanishes as age increases. As evident from the figure, the first quartile of extensive margin semielasticity is -6% among those aged 30, while only -2.75% for 50 or older.

The distributions of WTP and price sensitivity map directly into the demand curves that will be relevant for my counterfactuals. These are represented in Figure 8 (overimposed to raw data on average premium and choice probability) for the two groups that I will focus on: over and under-50 households without children and eligible for subsidies. The figure shows how average demand of the older group lies above the one of the younger group at any given premium. Moreover, this average difference is generated from a left tail of low WTP among young households: at an annual premium of $1,000, the bottom 25% WTP among under-50 households buy coverage with probability 0.4, while the bottom 25% WTP among over-50 buy with probability higher than 0.7.

Lastly, a major advantage of the mixed logit model is to capture substitution patterns across different metal tiers relaxing the IIA assumption. Table 10 reports the within-group median change in probability of purchasing a given tier (column) in reponse to a $100 increase in the premium of another tier (row). As expected, these substitution patterns now capture how the probability of switching between adjacent coverage levels is higher than the probability of switching, for example, from Platinum to Bronze plans. Moreover, as a major difference from the nested logit, the own-premium semielasticity is lower for higher coverage plans (among under 50, -7% and -10% for Silver and Platinum, as opposed to -12% for Bronze). As before, younger households without children are significantly more price sensitive than the other groups.

5.2 Cost estimates

Cost functions. I consider a cost function in which I allow the logarithm of (expected) cost to vary linearly across regions, insurers, and metal tiers, but also —within a given plan— with a buyer's age, household size, income, and estimated willingness-to-pay for more generous coverage (β/α). The estimating equation is:

$$\ln(C_{jr}) = \sum_{n'} \kappa^{n'} \mathbf{1}\left\{n(j) = n'\right\} + \sum_{r'} \kappa^{r'} \mathbf{1}\left\{r = r'\right\} + \sum_{z'} \kappa^{z'} \mathbf{1}\left\{z_{jr} = z'\right\} + \gamma \overline{X}_{jr} + \omega_{jr}, \tag{22}$$

where κ^n, κ^r, and κ^z are coefficients on indicators for carrier n, region r, and metal tier z, respectively. The vector $\overline{X}_{jr}$ collects average characteristics of buyers choosing plan j in r. This includes age, income, and household size, which are directly calculated from enrollment data. Additionally, $\overline{X}_{jr}$ also includes the average posterior β/α among buyers of j in r; this is derived directly from the mixed logit demand system.

Absent of selection, buyers with different characteristics (hence preferences) do not imply different costs for insurer, and one would have $\gamma = 0$. Instead, estimating $\gamma \neq 0$ indicates that cost varies across buyers. In particular, a nonzero coefficient on the average β/α is the evidence of cost varying with preferences even after controlling for demographics.

Cost parameters are reported in Table 11. Across different specifications, I find that expected cost increases with the generosity of coverage, with Silver plans implying 30% higher claims, on average, than Bronze plans, and Platinum plans implying expected claims twice as large as Bronze plans. I also find strong evidence of selection along both observable demographics and unobserved preferences for insurance. Interpreting the parameters, a 10-year increase in age increases expected claims by 20%. After controlling for demographics, above median WTP implies expected claims 25% higher than when the enrollees' WTP is in the bottom 20%.

Figure 9 shows the predictions of the cost model in dollar amounts. I estimate that expected cost grows from less than $2,000/year for enrollees of Silver and Bronze plans who are younger than 30 to over $3,000 ($4,000 for Silver) for those older than 50. The differences between coverage levels are substantial, and also increase in age. The expected cost for zero-deductible Platinum plans is equal to $3,000 for the young buyers, and increases to over $10,000 for the older ones.. Importantly, I find that the cost-ratio between 64 and 21 is between 4:1 and 5:1, depending on the tier of coverage. This is consistent with other estimates (see also Orsini and Tebaldi, 2017), the health expenditure patterns in the MEPS, and many industry and policy experts who advocate for a 5:1 age-adjustment to replace the ACA-mandated 3:1 (Blumberg and Buettgens, 2013). The heterogeneity in preferences within demographics maps directly into differences in cost. Evident from panel (b) of Figure 9, conditioning on age the interquantile range of expected cost for a Silver plan is approximately $1,000 for younger buyers, increasing to more than $2,000 for the over-50.

Joint distribution of preferences and cost. The cost estimates together with the mixed logit demand system allow me to explore the joint distribution of preferences and cost. This is the key primitive one needs to study a selection market from both a positive and normative perspective (Einav, Finkelstein, and Cullen, 2010). Even in this paper, this joint distribution is the main ingredient for my policy counterfactuals and its properties directly affect the quantitative comparison of different subsidy designs.

To represent the relationship between preferences and cost, and how this varies across different demographic groups, Figure 10 shows the joint distribution between the probability of purchasing a plan if all premiums were $3,000/year and expected cost when enrolling in a Silver plan offered by Anthem (the largest insurer in Covered California, and present in all 19 regions). I distinguish between under-50 households eligible for subsidies and without children in the left panel, and over-50 households with the same characteristics on the right. The Online Appendix reports the figure for all groups. The figure highlights the nature of selection in this market: comparing young to old, the latter group is significantly more likely to purchase coverage for the same premium (something evident from the raw data), and more than three times as costly to cover. Overall, the unconditional correlation between probability of purchase at $3,000 and expected cost is approximately 0.44. Even within a demographic group, buyers with higher demand for insurance have higher expectected cost. Indeed, the variation in preferences after fully controlling for demographics and region maps directly into differences in cost, and the conditional correlation increases to 0.81.

5.3 Alternative cost estimates and Nash-in-prices assumption

Before moving to my counterfactual analysis, I carry on two heuristic exercises to explore to which degree a Nash-in-prices assumption can help provide normative statements about the effect of different subsidy designs on market outcomes. First, I explore the relationship between marginal revenue predicted by the demand model (as a function of base prices after applying ACA regulations) and marginal cost predicted combining cost estimates from claims data with demand estimates (again as a function of base prices after applying ACA regulations yet ignoring risk-adjustment). For this comparison, it is important to recall that so far I have not imposed any assumptions on the optimality of insurers' behavior to obtain the demand-cost estimates.

Panel (a) of Figure 11 plots the relationship between marginal revenue and cost (see Appendix B1 for formal derivations), where I focus on the case in which insurers internalize the extent to which —when they offer the second-cheapest Silver plan— they can directly affect the value of subsidies. (See Online Appendix for the case in which they do not internalize this.) Overlaying a linear fit with the 45-degree line, the figure shows that there is a tight relationship between the marginal revenue and cost (the linear coefficient is 0.8 including a constant term, and 0.97 otherwise). This suggests that studying counterfactuals in which these two quantities are imposed to be equal, product-by-product, can provide a good benchmark model to analyze this market.

An alternative approach is to estimate marginal cost functions imposing Nash-in-prices, and then compare the resulting estimates to the cost-model estimated with claims data. This is shown in panel (b) of Figure 11. Similarly, I find a close relationship between the cost inferred under the assumption that insurers set base prices in a Nash-in-prices equilibrium (knowing rating regulations and internalizing the distortion induced by price-linked subsidies), and the cost estimated directly from claims data without supply-side assumptions. Here the coefficient in the linear fit is equal to 1.22 when including a constant, and 1.01 constraining the intercept to be zero.

The above discussion *does not* provide a formal test for Nash-in-prices, and the assumption carries on several limitations; for instance, it does not account for insurers' uncertainty about demand and cost in the early years of the market, and in my formulation it omits some regulatory details (e.g. the budget-neutral risk-adjustment implemented under the ACA, see discussion in Section 2). Nevertheless, at least in the case of Covered California (with peculiarities such as active purchasing and standardized plan characteristics), the assumption seems to provide a good fit to the observed data. Also considering its standard adoption in empirical industrial organization and its attractiveness for equilibrium calculations, I impose Nash-in-prices adapted to ACA regulations to predict changes to policy-relevant outcomes under alternative subsidy designs.

6 Equilibrium under Alternative Subsidy Designs

I compare Nash-in-prices equilibria under three situations:

(a) Price-linked subsidy under the ACA;

(b) Voucher of amount equivalent to the equilibrium subsidy under the ACA;

(c) Age-adjusted voucher.

In scenario (a) the subsidy scheme is the one described in Section 2: insurers set base prices b_r, premiums are age-adjusted, and the subsidy for a household with characteristics θ —$V(\theta, b_r)$ using my notation— is computed to ensure that the second-cheapest Silver plan in the region is affordable at a premium lower than $\overline{P}^{yh}$, where y and h are household income and size, respectively. I solve for the vector of equilibrium base-price $b_r^{*,ACA}$ zeroing the FOC necessary conditions for optimal pricing and imposing medical loss ratio constraints similar to those imposed by the ACA (see Appendix B2).[29]

6.1 From price-linked subsidies to vouchers

Taking $b_r^{*,ACA}$ as input, in (b) I set a fixed voucher for a household of type θ in region r equal to $V_r(\theta) = V(\theta, b_r^{*,ACA})$. That is, I provide the household with a discount equal to the one received (in equilibrium) under the ACA. I apply this discount, but do not adjust it endogenously to changes in base prices.

The key difference between the two designs is that fixed vouchers are less distortionary than price-linked subsidy. Under vouchers, insurers know that a \$1 increase in pre-subsidy premium corresponds to a \$1 increase in subsidized premium, so price increases are fully passed through to consumers. Instead, under the price-linked scheme adopted by the ACA, insurers can increase the premium of the second-cheapest Silver plan without changing the amount that buyers have to pay for this plan (yet changing premium differences from other plans).

[29] Multiplicity concerns are natural, since mine is a model with multi-product firms, mixed logit demand, and selection. These are even more severe when considering price-linked subsidies, as emphasized in Jaffe and Shepard (2017). For my counterfactuals I always use the base prices in the data as starting point, and iterate FOC until convergence. I then use 1,000 random starting point distributed uniformly in a \$100×\$100 square centered around the data and verify that the resulting equilibria do not differ from the previous equilibrium (up to simulation error, where the largest deviations are found among base prices of plans with very low market share).

For this reason, as argued in Jaffe and Shepard (2017), price-linked subsidies provide additional incentives for insurers to raise premiums, and lead to equilibria with higher markups than under fixed vouchers of equivalent amounts. Vouchers are then a good theoretical benchmark to explore alternative designs. Nevertheless, implementation of a voucher system requires a significant amount of knowledge by the regulator, and it could be (politically) easier to guarantee a premium ceiling $\overline{P}^{yh}$, and then have public finances absorb the risk of premium fluctuations across markets and over time. This argument is used widely, but the computations in Jaffe and Shepard (2017) suggest that for a wide range of regulator's uncertainty, the resulting "wrong" vouchers still lead to higher welfare, in a Nash-in-prices equilibrium, than a price-linked subsidy. Additional studies focusing on manipulation of endogenous subsidies by carriers deliver a similar point (see e.g. Decarolis, 2015; Decarolis, Polyakova, and Ryan, 2015).[30]

In Table 12, I compare equilibrium under scenarios (a) and (b). Consistently with Jaffe and Shepard (2017), I find that under fixed vouchers average markups drop, on average, by 11% (from \$952 to \$844 per-year), while the markup of the second-cheapest Silver plan drops by as much as 40%, corresponding to a reduction in base price of \$200/year. The medical loss ratio increases from 77.9% to 79.6%, reflecting a decrease in average premium of \$230/year against a decrease in average cost of \$63/year. Changes to prices correspond to changes in enrollment, coverage levels, and public spending (see also Table 13). I find that switching from the price-linked scheme to a fixed voucher leads to a 2% increase in total enrollment (+20,000), a 3% increase in total consumer surplus (+\$114 million), and at the same time a 1% reduction in public spending (-\$51 million).

6.2 Age-adjustments to the generosity of subsidies

My second counterfactual is more novel from a theoretical perspective, and it amounts to verify whether, by adjusting the generosity of subsidies to a buyer's age, it is possible to make all buyers better off without increasing public spending. The mechanism is intuitive: more generous discounts for young buyers and less generous discounts for older ones, imply that at any given base price the enrollment pools are comparatively younger. This corresponds to lower average (and marginal) cost, higher aggregate elasticity of demand, and a consequent downward pressure on prices.

Strategic complementarities between prices ensure that the new equilibrium prices are lower (see e.g. Vives, 1990), and this is enhanced by the extent to which, as is the case in my empirical model, new buyers entering the market are comparatively cheaper. If this mechanism is sufficiently strong —ultimately an empirical question I am trying to address— it is possible that the post-subsidy premium paid by all buyers are lower, and even the older households who receive a lower discount are better off than without the age adjustments (despite paying more than their younger

[30] A second, more operational, advantage of vouchers is that multiplicity issues in equilibrium computation are mitigated; see Jaffe and Shepard (2017).

counterparts who have similar income).

Mechanism. To see this more formally, consider the one-product case, with each plan j choosing a base prices to maximize expected profits taking other base prices as given.[31] One can rewrite the profit function as (see Appendix B for details):

$$\Pi_{jr}(b_r) = Q_{jr}\left(b_{jr} Af_{jr} - AC_{jr}\right), \quad \text{where} \tag{23}$$

$$Q_{jr} = \sum_{\theta,\delta} G_r(\theta) F_r\left(\delta|\theta\right) \mathcal{S}_{jr}^{\theta,\delta}\left(hA^{\tau} b_r - V_r(\theta)\right) \tag{24}$$

$$Af_{jr} = Q_{jr}^{-1} \sum_{\theta,\delta} G_r(\theta) F_r\left(\delta|\theta\right) \mathcal{S}_{jr}^{\theta,\delta}\left(hA^{\tau} b_r - V_r(\theta)\right) hA^{\tau} \tag{25}$$

$$AC_{jr} = Q_{jr}^{-1} \sum_{\theta,\delta} G_r(\theta) F_r\left(\delta|\theta\right) \mathcal{S}_{jr}^{\theta,\delta}\left(hA^{\tau} b_r - V_r(\theta)\right) \phi(\theta, \delta, z_{jr}; \kappa). \tag{26}$$

In words, applying the rating regulations and region-θ specific vouchers $V_r(\theta)$, Q_{jr} is the total enrollment in jr, Af_{jr} is the average adjustment to age and household size (so that $b_{jr}Af_{jr}$ is the average revenue for plan j), and AC_{jr} is average cost.

Taking the necessary FOC for Nash-in-prices and rearranging leads to:

$$b_{jr} = \frac{AC_{jr}}{Af_{jr}} - \left(\frac{Q'_{jr}}{Q_{jr}}\right)^{-1}\left(1 + \frac{b_{jr}Af'_{jr} - AC'_{jr}}{Af_{jr}}\right), \tag{27}$$

where Q_{jr}, AC_{jr}, and Af_{jr} are seen as functions of b_{jr}. This expression states that equilibrium base prices are the sum of (a) the ratio between average cost and average age adjustments, (b) a markup term that decreases with the aggregate elasticity of demand to base prices, and varies with a "selection term" that depends on AC'_{jr} (change in average cost as price changes) and on the slope of the average rating adjustment term Af'_{jr}.

Fixing b_r, a change in vouchers such that buyers older (younger) than 50 receive $V_r(\theta) - \Delta^{50+}$ ($V_r(\theta) + \Delta^{50-}$) affects all terms on the right-hand side of (27). First, since age adjustments (up to 3:1) do not compensate for differences in expected cost (above 4:1 in my estimates) across age, the new vouchers lower the term $\frac{AC_{jr}}{Af_{jr}}$, as more young buyers enter the market while some older buyers leave. Second, younger buyers are more price-sensitive, so that the average semi-elasticity $-\frac{Q'_{jr}}{Q_{jr}}$ is higher. As long as the selection term in parentheses does not invert these effects, insurers have incentives to lower base prices.

Letting b_r^* be the equilibrium under vouchers $V_r(\theta)$, the new equilibrium under age-adjusted vouchers —$b_r^{*,\Delta}$— is such that $b_r^* > b_r^{*,\Delta}$. This guarantees that under-50 buyers pay lower premiums,

[31] The multi-product case is a natural extension, and the intuition is the same, yet it does not provide a clean analytical solution. I discuss this case in more detail in Appendix B.

and it might be possible —ultimately an empirical question— to find $(\Delta^{50-}, \Delta^{50+})$ for which even among over-50 buyers $A^\tau b_r^* - V_r(\theta) \geq A^\tau b_r^{*,\Delta} - V_r(\theta) + \Delta^{50+}$, or simply

$$b_r^* - b_r^{*,\Delta} \geq \frac{\Delta^{50+}}{A^\tau}. \tag{28}$$

If this is the case, even over-50 buyers pay a lower post-subsidy premium. Thus, they would prefer to receive a lower discount —and pay higher premium than younger consumers with similar income— to provide additional incentive for the participation of the young.

Simulations. To investigate this mechanism, and the possibility to make all buyers better off, I simulate equilibria varying the subsidies to over and under 50-year-olds, without children, and eligible for cost-sharing subsidies (FPL<250). I change monthly vouchers of each group by up to \$50/month, in steps of \$10 ($\Delta^{50-}, \Delta^{50+} \in \{-50, -40, ..., 0, 10, ..., 50\}$); this impllies that for each of the 19 regions I simulate equilibrium under 121 different deviations in vouchers centered around the level of discounts provided under the ACA. I then aggregate outcomes across regions as quantity-weighted averages for per-person outcomes (e.g. average premium, average cost, average markup), and overall sums for total outcomes (e.g. total enrollment, total spending).

Figure 12 shows my main result, representing how enrollment among over and under-50 changes as the generosity of subsidies varies by age.[32] Additionally, Figure 13 provides a two-dimensional representation which includes other equilibrium outcomes.[33] Following the differences in demand and cost across age groups, I find that the effect of changing generosity of subsidies by age is very stark. On the one hand, enrollment of the under-50 is highly responsive to the lower premiums (higher subsidies), as showed by the steep surface in Figure 12. Conversely, the over-50, who are much less price sensitive, do not respond strongly when their subsidies are lowered.

My emphasis here is on the interaction between the two groups, driven by the combination of heterogeneity in the demand responses and cost differences. As the figure highlights (see also Figure 13), enrollment of the older group increases in the voucher of the younger group; vice versa, as subsidies for the older group increase, coverage of the young decreases. This follows the intuition from above: as younger buyers enter the market, average cost, markups, and hence premiums decrease, and buyers respond accordingly.

It is then possible to simultaneously increase the voucher of the young and lower the voucher of the over-50, and achieve an equilibrium in which both groups are better off, the per-person annual subsidy is lower, and overall spending (driven by the increase in total enrollment) increases by less than 2% (or \$50 million/year). The specific example highlighted in Figure 12 is one in which: (a) subsidies for under-50 are increased by \$40/month, (b) subsidies for over-50 are lowered by

[32] I use a quadratic spline to interpolate over the simulation grid.

[33] A more comprehensive collection of outcomes of interest can be found in Figure H in the Online Appendix.

\$30/month, (c) under-50 (subsidized, no children) enrollment is 4% higher, (d) over-50 (subsidized, no children) enrollment is 0.2% higher, (e) enrollment of subsidized with children is 2.5% higher, (f) enrollment of the unsubsidized is 2.8% higher, (g) average cost is 1.13% lower, (h) markups are 5% lower, (i) per-person subsidy is 0.5% lower, (j) total spending is 2% higher. This is only one of many examples, since Figure 13 shows that there is a region of vouchers for which the equilibrium comparative statics leads to qualitatively similar conclusions.[34]

Subsidy design under different objectives. In my last exercise, I allow voucher adjustments for all households without children, distinguishing between 6 age-income groups.[35] Calling $\Delta = \{\Delta^\theta\}$ the set of all possible voucher deviations (up to ±\$50/month), one for each group, I find the Δ's that in equilibrium:

1. Maximize total consumer surplus[36] subject to all households being better off and per-person public spending being lower than under $V_r(\theta)$;

2. Maximize total consumer surplus subject to total spending being lower than under $V_r(\theta)$;

3. Maximize total enrollment subject to total spending being lower than under $V_r(\theta)$.

Problem 1 picks one of the deviations discussed above and higlighted in Figures 12; problem 2 takes a more utilitarian approach, dropping the constraints that all buyers must be better off, and imposing a constraint on total spending (rather than per-person subsidies); lastly, problem 3 is similar to problem 2, but rather than consumer surplus the target is total enrollment (which in some contexts could be more relevant from a policy perspective).

Table 13 summarizes equilibrium outcomes (and corresponding vouchers) at the solutions to these problems, as well as the comparison between ACA price-linked subsidy and the equivalent vouchers.[37] The table consistently shows that increasing vouchers for the young, low-income households, while lowering vouchers for comparatively older, higher-income households leads to desirable

[34]Moreover, to be conservative, here I only consider small, local deviations in monthly vouchers, so the set of different vouchers combinations making all buyers better off without increasing the overall (average) generosity of subsidies is potentially larger than what I find here.

[35]These are (i) $FPL \in [130, 200)$, age<45; (ii) $FPL \in [200, 250)$, age <45; (iii) $FPL \in 250, 400)$, age<45; (iv) $FPL \in [130, 200)$, age≥45; (v) $FPL \in [200, 250)$, age ≥<45; (vi) $FPL \in 250, 400)$, age≥45.

[36]This is the usual consumer surplus in the mixed logit framework (Small and Rosen, 1981), with my specification:

$$CS = \sum_r \sum_{\theta,\delta} G_r(\theta) F_r(\alpha, \beta, \chi|\theta) \ln\left(1 + \sum_{j \in J_r} \exp\left(\alpha P_{jr}^\theta + \beta av_{jr}^\theta + \chi^{n(j)}\right)\right). \tag{29}$$

[37]In interpreting these findings, it is important to emphasize how I restrict myself to local deviations in voucher values up to \$50/month. This is sufficient for my goal here, which is to highlight how the generosity of subsidies under the ACA leaves room for deviations that can be significant from a welfare perspective. This is not indended to provide a global solutions to the subsidy design problem over the entire space of possible household-specific vouchers.

equilibria. For problem 1 (column (2) in the table), the solution indicates to provide an increase in subsidy for all the under-45 who are subsidy eligible, with maximum increase of $50/month for those with income lower than 200% of FPL or between 250-400% of FPL, and $10/month for the middle group with income between 200-250% of FPL. At the same time, the over-45 vouchers should be lowered, by $25/month, for those with income between 130-200% and 250-400% of FPL, while left unaltered for the middle income group. In terms of outcomes, compared to the vouchers equivalent to the ACA amounts, enrollment is larger by 129,000 (+11.2%), consumer surplus by $378 million (+10.9%), total spending by $13 million (+0.4%), and average cost is lower by $252 (-7.7%).

In problems 2 and 3 (columns (3) and (4) in Table 13), I allow instead for redistribution across groups. For problem 2, I find deviations in monthly vouchers for which, compared to the equilibrium in column (1), enrollment would be higher by 86,000, consumer surplus by $242 million, while total spending would be lower by $7 million and average cost by $180. The solution to problem 3 highlights the difference between maximizing consumer surplus as opposed to enrollment, where in the latter case the objective does not depend on coverage levels (which enter consumer surplus calculations), nor makes distinctions across buyers in terms of their observed characteristics or willingness-to-pay.[38] The solution to this problem implies deviations in vouchers that would lead equilibrium enrollment to be 181,000 higher than in column (1) (almost +16%), with yet a decline in consumer surplus of $335 million (-9.7%). This shows the large difference between a utilitarian measure of welfare and a raw measure of quantity that ignores who buys and what is being bought. Consequently, specific policy recommendations can differ widely under different policy objectives, while generally speaking age-adjustments to subsidy generosity can be beneficial to welfare in this market.

7 Conclusions

Summary. I use a combination of individual level data on enrollment, survey data on uninsurance, and plan-level claims data, to estimate a demand-cost model of the Californian ACA health insurance marketplace. My estimates show a large degree of heterogeneity in preferences, with comparatively older households significantly more willing to pay for individual health insurance. Moreover, within demographic groups there is a large degree of dispersion in preferences, particularly among younger households. In terms of cost, expected claims increase in both the age and the willingness-to-pay of a buyer, so that this market features adverse selection in the sense of Einav, Finkelstein, and Cullen (2010).

I use these estimates with an oligopoly model adapted to ACA regulations to study Nash-

[38] The suggested deviations in vouchers imply higher subsidies for under-45 of income up to 250% of the FPL, and for over-45 with FPL between 250-400, while subsidies are lower for other households.

in-prices equilibria under different subsidy designs. The main counterfactual shows that one can improve upon the current scheme by providing more generous subsidies to young buyers, with the possibility of lowering at the same time the subsidy to the over-50 without making them worse off. Additional results measure the distortion on insurers' incentives induced by price-linked subsidies (Jaffe and Shepard, 2017), and find subsidies that increase consumer surplus and enrollment with lower public spending than under the ACA.

Limitations and extensions. This paper considers the effect of different subsidy designs on premium and enrollment within a static, partial equilibrium framework in which insurers only set prices while the remaining market variables are held fixed. Specific features of the Californian exchange such as active purchasing, standardized contracts, and the stability of the market in its first three years, all support the ability of this framework to teach normative lessons. Nevertheless, it is important to interpret my findings through the lenses of the model I adopt here, and to consider the policy implications of this article as short-run effects within the current institutional environment. In particular, my findings apply less to exchanges in which insurers set contracts freely, or in which there is (still) a large degree of uncertainty about participating carriers, demand, and cost. In these instances it is important for future work to consider what is the "right" model of insurer behavior to study counterfactual policies.

Aspects of the market that I do not model include the setting of provider networks and price negotiations between hospitals and insurers (Ho and Lee, 2017; Shepard, 2017), and dynamic aspects such as buyers' inertia and switching costs when re-enrolling over multiple years Handel (2013). Incorporating changes to provider networks and negotiated prices does not seem a primary concern in the context of ACA exchanges, since they represent only a very small fraction of the business operations of large insurers. However, in larger government-sponsored markets such as Medicare Advantage, Medicare Part D, or if ACA exchanges were to increase in size, different subsidy designs could affect not only pricing incentives, but also (in complex ways) the insurer-provider relationships, and thus affect healthcare costs and ultimately price through this different channel. Similar considerations apply to the omission of market dynamics. In the current configurations of ACA exchanges, these are unlikely to be first order, since very few buyers are enrolled for multiple years (life events such as job transitions or marriages can shift a buyer from the exchange to a different segments of the US health insurance market). But once again, in a market where multiple years of coverage are the norm (e.g. all the segments of private Medicare), one should study how altering subsidies could impact pricing incentives when insurers take into account buyers' choice behavior over time.

Lastly, while my focus here is on subsidies, the market design problem in government-sponsored insurance (Einav and Levin, 2015) includes other regulatory tools such as risk-adjustment (particularly when not budget neutral; Mahoney and Weyl, 2014), limits to price discrimination (Handel, Hendel, and Whinston, 2015; Ericson and Starc, 2015; Orsini and Tebaldi, 2017), and insurance mandates (Hackmann, Kolstad, and Kowalski, 2015; Azevedo and Gottlieb, 2017). An exciting agenda for future research is to move beyond studies of how each of these can independently affect market outcomes, and explore how different tools interact in determining the behavior of insurers and buyers.

References

Abraham, J., C. Drake, D. W. Sacks, and K. Simon (2017): "Demand for health insurance marketplace plans was highly elastic in 2014–2015," *Economics Letters*.

Aizawa, N. (2015): "Health Insurance Exchange Design in an Empirical Equilibrium Labor Market Model," *manuscript, University of Minnesota*.

Akerlof, G. A. (1970): "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 84, 488–500.

Anthony, C. H., B. Blom, M. Booth, J. Holland, D. Fried, C. Whalen, M. Carroll, A. Lerner, A. Marcellino, and P. Greene (2015): "The budget and economic outlook: 2015 to 2025," Congressional Budget Office.

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

Berry, S., M. Carnall, and P. T. Spiller (1996): "Airline Hubs: Costs, Markups and the Implications of Customer Heterogeneity," Working Papers 5561, NBER.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–90.

Berry, S. T. (1994): "Estimating Discrete-Choice Models of Product Differentiation," *The RAND Journal of Economics*, 25, pp. 242–262.

Berry, S. T. and P. Haile (2014): "Identification in Differentiated Products Markets Using Market Level Data," *Econometrica*, 82, 1749–1797.

Bhat, C. R. (1997): "An endogenous segmentation mode choice model with an application to intercity travel," *Transportation science*, 31, 34–48.

Bindman, A. B., D. Hulett, T. P. Gilmer, and J. Bertko (2016): "Sorting Out the Health Risk in California's State-Based Marketplace," *Health Services Research*, 51, 129–145.

Blumberg, L. J. and M. Buettgens (2013): "Why the ACAs Limits on Age-Rating Will Not Cause Rate Shock: Distributional Implications of Limited Age Bands in Nongroup Health Insurance," Tech. rep., Robert Wood Johnson Foundation.

Bresnahan, T. F. (1981): "Departures from marginal-cost pricing in the American automobile industry : Estimates for 1977-1978," *Journal of Econometrics*, 17, 201–227.

Bundorf, M. K., J. Levin, and N. Mahoney (2012): "Pricing and Welfare in Health Plan Choice," *American Economic Review*, 102, 3214–48.

Bundorf, M. K. and M. V. Pauly (2006): "Is health insurance affordable for the uninsured?" *Journal of Health Economics*, 25, 650 – 673.

Chamberlain, G. (1980): "Analysis of Covariance with Qualitative Data," *The Review of Economic Studies*, 47, 225–238.

Clemens, J. (2015): "Regulatory Redistribution in the Market for Health Insurance," *American Economic Journal: Applied Economics*, 7, 109–34.

Colombo, F. and N. Tapay (2004): "Private health insurance in OECD countries," .

Curto, V., L. Einav, J. Levin, and J. Bhattacharya (2014): "Can Health Insurance Competition Work? Evidence from Medicare Advantage," Working Paper 20818, National Bureau of Economic Research.

Dafny, L., J. Gruber, and C. Ody (2015): "More Insurers Lower Premiums," *American Journal of Health Economics*.

Decarolis, F. (2015): "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?" *American Economic Review*, 105, 1547–80.

Decarolis, F., M. Polyakova, and S. P. Ryan (2015): "The Welfare Effects of Supply-Side Regulations in Medicare Part D," Working Paper 21298, National Bureau of Economic Research.

DeLeire, T., A. Chappel, K. Finegold, and E. Gee (2016): "Do Individuals Respond to Cost-Sharing Subsidies in their Selections of Marketplace Health Insurance Plans?" .

Dickstein, M. J., M. Duggan, J. Orsini, and P. Tebaldi (2015): "The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the Affordable Care Act," *American Economic Review*, 105, 120–25.

Duggan, M. and T. Hayford (2013): "Has the Shift to Managed Care Reduced Medicaid Expenditures? Evidence from State and Local-Level Mandates," *Journal of Policy Analysis and Management*, 32, 505–535.

Duggan, M., A. Starc, and B. Vabson (2014): "Who Benefits when the Government Pays More? Pass-Through in the Medicare Advantage Program," Working Paper 19989, National Bureau of Economic Research.

Einav, L., A. Finkelstein, and M. R. Cullen (2010): "Estimating Welfare in Insurance Markets Using Variation in Prices," *The Quarterly Journal of Economics*, 125, 877–921.

Einav, L., A. Finkelstein, S. P. Ryan, P. Schrimpf, and M. R. Cullen (2013): "Selection on Moral Hazard in Health Insurance," *American Economic Review*, 103, 178–219.

Einav, L. and J. Levin (2015): "Managed Competition in Health Insurance," *Journal of European Economic Association, forthcoming.*

Ericson, K. M. and A. Starc (2012a): "Designing and Regulating Health Insurance Exchanges: Lessons from Massachusetts," *INQUIRY: The Journal of Health Care Organization, Provision, and Financing*, 49, 327–338.

——— (2012b): "Heuristics and heterogeneity in health insurance exchanges: Evidence from the Massachusetts connector," *The American Economic Review*, 493–497.

——— (2013): "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange," Working Paper 19527, National Bureau of Economic Research.

——— (2014): "Measuring Consumer Valuation of Limited Provider Networks," Working Paper 20812, National Bureau of Economic Research.

——— (2015): "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange," *The Review of Economics and Statistics, forthcoming*, 97, 667–682.

Finkelstein, A., N. Hendren, and M. Shepard (2017): "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts," .

Fox, J. T., K. il Kim, and C. Yang (2016): "A simple nonparametric approach to estimating the distribution of random coefficients in structural models," *Journal of Econometrics*, 195, 236 – 254.

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

Geruso, M. (2016): "Demand Heterogeneity in Insurance Markets: Implications for Equity and Efficiency," Tech. rep., National Bureau of Economic Research.

HACKMANN, M. B., J. T. KOLSTAD, AND A. E. KOWALSKI (2015): "Adverse Selection and an Individual Mandate: When Theory Meets Practice," *American Economic Review*, 105, 1030–66.

HANDEL, B., I. HENDEL, AND M. D. WHINSTON (2015): "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk," *Econometrica*, 83, 1261–1313.

HANDEL, B. R. (2013): "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review*, 103, 2643–82.

HANDEL, B. R., I. HENDEL, AND M. D. WHINSTON (2017): "The Welfare Effects of Long-Term Health Insurance Contracts," Working Paper 23624, National Bureau of Economic Research.

HO, K. AND R. S. LEE (2017): "Insurer competition in health care markets," *Econometrica*, 85, 379–417.

HO, K. AND A. PAKES (2014): "Hospital Choices, Hospital Prices, and Financial Incentives to Physicians," *American Economic Review*, 104, 3841–84.

JAFFE, S. AND M. SHEPARD (2017): "Price-Linked Subsidies and Health Insurance Markups," .

KAUTTER, J., G. C. POPE, M. INGBER, S. FREEMAN, L. PATTERSON, M. COHEN, AND P. KEENAN (2014): "The HHS-HCC risk adjustment model for individual and small group markets under the Affordable Care Act," *Medicare & Medicaid research review*, 4.

KRINN, K., P. KARACA-MANDIC, AND L. A. BLEWETT (2015): "State-Based Marketplaces Using Clearinghouse Plan Management Models Are Associated With Lower Premiums," *Health Affairs*, 34, 161–169.

LAYTON, T., E. J. MONTZ, AND M. SHEPARD (2017): "Health Plan Payment in US Marketplaces: Regulated Competition with a Weak Mandate," Tech. rep., National Bureau of Economic Research.

LONG, S. K., A. YEMANE, AND K. STOCKLEY (2010): "Disentangling the Effects of Health Reform in Massachusetts: How Important Are the Special Provisions for Young Adults?" *The American Economic Review*, 297–302.

LUSTIG, J. (2010): "Measuring welfare losses from adverse selection and imperfect competition in privatized Medicare," *Manuscript. Boston University Department of Economics.*

MAHONEY, N. (2015): "Bankruptcy as Implicit Health Insurance," *American Economic Review*, 105, 710–46.

MAHONEY, N. AND E. G. WEYL (2014): "Imperfect Competition in Selection Markets," Tech. rep., National Bureau of Economic Research.

Manning, W. G., J. P. Newhouse, N. Duan, E. B. Keeler, and A. Leibowitz (1987): "Health insurance and the demand for medical care: evidence from a randomized experiment," *The American economic review*, 251–277.

Matzkin, L. R. (2008): "Identification in Nonparametric Simultaneous Equations Models," *Econometrica*, 76, 945–978.

Newey, W. K. and D. McFadden (1994): "Large sample estimation and hypothesis testing," *Handbook of econometrics*, 4, 2111–2245.

Orsini, J. and P. Tebaldi (2017): "Regulated Age-Based Pricing in Health Insurance Exchanges," *Available at SSRN 2464725*.

Pacifico, D. (2013): "lclogit: A Stata command for fitting latent-class conditional logit models via the expectation-maximization algorithm," *The Stata journal*, 13, 625–639.

Panhans, M. T. (2017): "Adverse Selection in ACA Exchange Markets: Evidence from Colorado," .

Pauly, M. V. (1970): "Efficiency in the provision of consumption subsidies," *Kyklos*, 23, 33–57.

Roehrig, C. S. (1988): "Conditions for Identification in Nonparametric and Parametric Models," *Econometrica*, 56, pp. 433–447.

Rosse, J. N. (1970): "Estimating Cost Function Parameters Without Using Cost Data: Illustrated Methodology," *Econometrica*, 38, pp. 256–275.

Rothschild, M. and J. E. Stiglitz (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 90, 630–49.

Ruggles, S., K. Genadek, R. Goeken, J. Grover, and M. Sobek (2015): "Integrated Public Use Microdata Series: Version 6.0 [dataset]," *Minneapolis: University of Minnesota*.

Sacks, D. W. (2017): "How Do the Mandate and the Premium Tax Credit Affect the Individual Insurance Market?" .

Saltzman, E. (2017): "Demand for Health Insurance: Evidence from the California and Washington ACA Marketplaces," .

Scheffler, R. M., D. R. Arnold, B. D. Fulton, and S. A. Glied (2016): "Differing Impacts Of Market Concentration On Affordable Care Act Marketplace Premiums," *Health Affairs*, 35, 880–888.

Shepard, M. (2017): "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," .

Small, K. A. and H. S. Rosen (1981): "Applied welfare economics with discrete choice models," *Econometrica: Journal of the Econometric Society*, 105–130.

Smith, J. C., C. Medalia, et al. (2014): *Health insurance coverage in the United States: 2013*, US Department of Commerce, Economics and Statistics Administration, Bureau of the Census.

Somaini, P. (2011): "Competition and interdependent cost in highway procurement," Tech. rep., Mimeo, Stanford University.

——— (2015): "Identification in Auction Models with Interdependent Costs," Tech. rep., Mimeo, Stanford University.

Starc, A. (2014): "Insurer pricing and consumer welfare: evidence from Medigap," *The RAND Journal of Economics*, 45, 198–220.

Stiglitz, J. E. (1987): "Chapter 15 Pareto efficient and optimal taxation and the new new welfare economics," in *Handbook of Public Economics*, Elsevier, vol. 2 of *Handbook of Public Economics*, 991 – 1042.

Summers, L. H. (1989): "Some simple economics of mandated benefits," *The American Economic Review*, 79, 177–183.

Thompson, T. S. (1989): "Identification of Semiparametric Discrete Choice Models," Tech. rep., Center for Economic Research, Department of Economics, University of Minnesota.

Train, K. E. (2008): "EM algorithms for nonparametric estimation of mixing distributions," *Journal of Choice Modelling*, 1, 40–69.

Uryas'ev, S. (1994): "Derivatives of probability functions and integrals over sets given by inequalities," *Journal of Computational and Applied Mathematics*, 56, 197 – 223.

Vives, X. (1990): "Nash equilibrium with strategic complementarities," *Journal of Mathematical Economics*, 19, 305–321.

Wagstaff, A. and E. van Doorslaer (2000): "Chapter 34 Equity in health care finance and delivery," *Handbook of Health Economics*, 1, 1803 – 1862.

Waldfogel, J. (2003): "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets," *RAND Journal of Economics*, 34, 557–68.

Weyl, E. G. and A. Veiga (2014): "The Leibniz Rule for Multidimensional Heterogeneity," *Available at SSRN 2344812*.

Appendix A:
Nonparametric identification of cost from equilibrium assumptions

In this appendix I provide conditions for nonparametric identification of the distribution of willingness to pay and of cost conditional on willingness to pay, assuming that observables consists of choices, prices, and products' characteristics.

For this I use a model that is not tailored to my specific application, omitting subsidies and other regulations. This allows me to focus on, and highlight, the novel aspect of the identification argument, which is to use equilibrium assumptions and variation in the preferences of marginal buyers to identify cross-buyer cost heterogeneity. I provide a positive result for the case of single-plan insurers (or plan-level pricing decisions), an important simplification that leaves open questions for future work. In fact, multi-product pricing decisions introduce several complications, with the need of additional conditions, a different constructive proof, or specific functional form assumptions (e.g. those in the empirical application in this paper, or in Lustig, 2010).

A1. Model and observables

I start by adopting the model of demand used in Berry and Haile (2014) (BH), and then model supply allowing costs to vary with buyers' willingness to pay, and assuming that a Nash-in-prices equilibrium realizes in each market.

Demand (adapted from BH). Each consumer i in market r chooses a plan (or product) from a set $\mathcal{J} = \{0, 1, ..., J\}$. A market consists of a continuum of consumers in the same choice environment (e.g. geographic region). Formally a market r for the J products is a tuple $\chi_r = (x_r, p_r, \xi_r)$, collecting characteristics of the products or of the market itself. Observed exogenous characteristics are represented by $x_r = (x_{1r}, ..., x_{Jr})$, where each $x_{jr} \in \mathbb{R}^K$. The vector $\xi_r = (\xi_{1r}, ..., \xi_{Jr})$, with $\xi_{jr} \in \mathbb{R}$, represents unobservables at the level of the product-market. Finally, $p_r = (p_{1r}, ..., p_{Jr})$, with each $p_{jr} \in \mathbb{R}$, represents (endogenous) prices.

Consumer preferences are represented with a random utility model quasilinear in prices (Section 4.2 in BH). Consumer i in market r derives (indirect) utility $u^i_{jr} = v^i_{jr} - p_{jr}$ when purchasing j, with the usual normalization $v^i_{0r} = 0$, for all i, all r. Given prices, the choice of each buyer is then determined by the vector $v^i_r = (v^i_{1r}, ..., v^i_{Jr})$. For each buyer in market r, v^i_r is drawn i.i.d. from a continuous density $f_r(v)$. This satisfies the following:

D1. *BH Demand structure*: There is a partition of x_{jr} into $(x^{(1)}_{jr}, x^{(2)}_{jr})$, where $x^{(1)}_{jr} \in \mathbb{R}$, such that given indexes $\delta_r = (\delta_{1r}, ..., \delta_{Jr})$, with $\delta_{jr} = x^{(1)}_{jr} + \xi_{jr}$, $f_r(v) = f(v|\delta_r, x^{(2)}_r)$.

Therefore, assuming that $\arg\max_{j\in J} u^i_{jr}$ is unique with probability one in all markets, choice probabilities (market shares) are defined by

$$s_{jr} \quad = \quad \sigma_j(\chi_r) = \int_{\mathcal{D}_j(p_r)} f(v|\delta_r, x^{(2)}_r)\, \mathrm{d}v, \;\; j = 0, 1, ..., J, \tag{30}$$

$$\mathcal{D}_j(p_r) \quad = \quad \{v : v_j - v_k \geq p_j - p_k, \text{ for all } k \neq j\}\,. \tag{31}$$

Observables. Let $z_r = (z_{1r}, ..., z_{Jr})$, $z_{jr} \in \mathbb{R}^L$, denote a vector of cost shifters excluded from the demand model. The econometrician observes $(p_{jr}, s_{jr}, x_{jr}, z_{jr})$ for all r and all $j = 1, 2, ..., J$.

Supply. Let $w_{jr} = (\xi_{jr}, x_{jr}, z_{jr}) \in \mathbb{R}^{K+L+1}$ collect characteristics (observable and unobservable) and cost shifters of product j in r. When purchasing j, a buyer i with valuations $v^i = v$ in market r increases the total expected cost for the insurer by $\psi_j(v, w_{jr})$, $\psi_j : \mathbb{R}^J \times \mathbb{R}^{K+L+1} \to \mathbb{R}$.

The function $\psi_j(\cdot, w_{jr})$ is continuous and bounded for all j, and describes how the expected cost of covering the buyer varies with her vector of valuations after conditioning on w_{jr}.

At the prices p_r the seller of j realizes profits in market r equal to

$$\Pi_{jr}(\chi_r) = p_{jr} \cdot \sigma_j(\chi_r) - \int_{\mathcal{D}_j(p_r)} \psi_j(v, w_{jr}) \cdot f(v|\delta_r, x_r^{(2)})\,\mathrm{d}v. \tag{32}$$

I assume that in each market prices are set in a complete information Nash equilibrium in pure-strategies. To formalize this, the set of marginal buyers of product j can be described by

$$\begin{aligned} \partial\mathcal{D}_j(p_r) &= \{v : v_j - v_k = p_{jr} - p_{kr} \text{ for some } k \neq j\} \qquad (33)\\ &= \lim_{\varepsilon \downarrow 0} \left\{ \mathcal{D}_j(p_r) \cap \left(\mathbb{R}^J \setminus \mathcal{D}_j(p_{jr} + \varepsilon, p_{-jr}) \right) \right\}. \qquad (34) \end{aligned}$$

Then, following Uryas'ev (1994); Weyl and Veiga (2014), quasilinearity of indirect utility with respect to price implies that, in equilibrium, in every market r:

S1. *Equilibrium*: For all $j = 1, ..., J$, $mr_{jr} = mc_{jr}$, where

$$\begin{aligned} mr_{jr} &= \sigma_j(\chi_r) - p_{jr} \cdot \int_{\partial\mathcal{D}_j(p_r)} f(v|\delta_r, x_r^{(2)})\,\mathrm{d}v, \qquad (35)\\ mc_{jr} &= -\int_{\partial\mathcal{D}_j(p_r)} \psi_j(v, w_{jr}) \cdot f(v|\delta_r, x_r^{(2)})\,\mathrm{d}v. \qquad (36) \end{aligned}$$

From S1, marginal revenues are equal to marginal costs, which must be true in a Nash-in-prices equilibrium. The integrals in mr_{jr} and mc_{jr} are well defined because $f(\cdot|\delta_r, x_r^{(2)})$ and $\psi_j(\cdot, w_{jr})$ are both continuous and bounded functions of v.

A2. Conditions for identification

Identification is defined as in Roehrig (1988); Matzkin (2008): if the unobservables differ (almost surely), then the distribution of observables differ (almost surely), where probabilities and expectations are defined with respect to the distribution of (χ_r, s_r, z_r) across markets.

My result is obtained combining conditions for identification of demand provided in BH — yielding to identification of ξ_r and then of $f(v|\delta_r, x_r^{(2)})$ — with a constructive proof to identify ψ_j which I adapted from Somaini (2011, 2015).[39] To simplify notation without loss of generality, as in BH I condition on $x_r^{(2)}$ — which unlike $x_r^{(1)}$ can affect the distribution of preferences quite arbitrarily — and suppress it.

Beside the demand and supply assumptions D1 and S1, I will use the following conditions:

C1. *BH Exogeneity of cost shifters*: For all $j = 1, ..., J$, $E[\xi_{jr}|z_r, x_r] = E[\xi_{jr}] = 0$.

C2. *BH Completeness*: For all functions $B(s_r, p_r)$ with finite expectations, if $E[B(s_r, p_r)|z_r, x_r] = 0$ with probability one, then $B(s_r, p_r) = 0$ with probability one.

C3. *Large support*: For every j, supp $v_r|\delta_r, w_{jr} \subset$ supp $p_r|\delta_r, w_{jr} \subset P$, with P bounded.

Condition C1 is a standard exclusion restriction, requiring mean independence between demand instruments and the structural erros ξ_{jr}. Condition C2 is a completeness assumption, requiring instruments to move market shares and prices sufficiently to distinguish between different functions of these variables through the exogenous variation in these instruments. C3 is a large support assumption, requiring cost shifters excluded from ψ_j to move prices

[39] This highlights the parallelism between auctions with interdependent costs and selection markets. In the former case (expected) marginal costs depend on the competitors' signals, varying with differences of bids between competitors. In a selection market (expected) marginal costs depend on the preferences of buyers choosing the plan, varying with differences of prices between competitors.

in a set that covers the support of (conditional) valuations. This is a stronger requirement than the large support assumption sufficient to identify the distributions $f(v|\delta_r)$, which would only require supp $v_r|\delta_r \subset$ supp $p_r|\delta_r$. The stronger condition in C3 allows to prove that cost functions ψ_j are also identified.

One then has:

THEOREM **1** *Under D1, S1, C1, C2, C3, ξ_r, $f(v|\delta_r)$, and ψ_j are identified.*

Proof of Theorem 1. Condition C3 implies supp $v_r|\delta_r \subset$ supp $p_r|\delta_r$, and demand is identified:

LEMMA **1** (Berry and Haile, 2014) *Under D1, C1, C2, ξ_r is identified, and $f(v|\delta_r)$ is also identified if, additionally, supp $v_r|\delta_r \subset$ supp $p_r|\delta_r$.*

Proof. Follows from Theorem 1 and Section 4.2 in BH.□

Similarly to Somaini (2011, 2015), the rest of the proof amounts to approximating for every j, every w_{jr}, and every $\widehat{v} \in$ supp $v_r|\delta_r, w_{jr}$, the integral of cost conditional on $\mathcal{D}_j(\widehat{v})$:

$$\Psi_j(\widehat{v}; w_{jr}, \delta_r) = \int_{\mathcal{D}_j(\widehat{v})} \psi_j(v, w_{jr}) \cdot f(v|\delta_r)\, \mathrm{d}v. \tag{37}$$

The mixed-partial J-1 derivative with respect to $\widehat{v}_{-j}$ yields then identification of the unknown cost function ψ_j, since

$$\frac{\mathrm{d}^{J-1}\Psi_j(\widehat{v}; w_{jr}, \delta_r)}{\mathrm{d}\widehat{v}_{-j}} = \psi_j(\widehat{v}, w_{jr}) \cdot f(\widehat{v}|\delta_r) \tag{38}$$

and $f(\widehat{v}|\delta_r)$ is identified by Lemma 1. This exploits the fact that price enters linearly in buyers' indirect utility, hence the set $\mathcal{D}_j(\widehat{v})$ is described by a set of inequalities which defines a cone in $\mathbb{R}^J$ with vertex $\widehat{v}$. The boundary of this cone is the set $\partial\mathcal{D}_j(\widehat{v})$ defined in (33); see also Figure 1 in BH.

To approximate $\Psi_j(\widehat{v}; w_{jr}, \delta_r)$, fix j, w_{jr}, and $\widehat{v} \in$ supp $v_r|\delta_r, w_{jr}$. Consider then a parametric curve $\eta : \mathbb{R}_+ \to \mathbb{R}$, with $\eta(\ell) = \widehat{v}_j + \ell$, and with this define the function $\widehat{\Psi}_j(\ell) = \Psi_j((\eta(\ell), \widehat{v}_{-j}); w_{jr}, \delta_r)$. Differentiating $\widehat{\Psi}_j(\ell)$ (and using again Uryas'ev, 1994; Weyl and Veiga, 2014) yields

$$\frac{\mathrm{d}\widehat{\Psi}_j(\ell)}{\mathrm{d}\ell} = -\int_{\partial\mathcal{D}_j((\eta(\ell), \widehat{v}_{-j}))} \psi_j(v, w_{jr}) \cdot f(v|\delta_r)\, \mathrm{d}v. \tag{39}$$

The function $\phi_j(\ell) \equiv \frac{\mathrm{d}\widehat{\Psi}_j(\ell)}{\mathrm{d}\ell}$ is bounded and continuous, and hence Riemann integrable over $[0, T]$, where by C3 the upper bound T can be chosen to be such that $\widehat{\Psi}_j(T) = 0$. Therefore,

$$\Psi_j(\widehat{v}; w_{jr}, \delta_r) = \widehat{\Psi}_j(0) = -\int_0^T \phi_j(\ell)\, \mathrm{d}\ell. \tag{40}$$

The integral in (40) can be approximated with arbitrary precision. For this, one can choose a sequence $\{\ell^n\}_{n=0}^N$ for which $0 = \ell^1 < \ell^2, ..., < \ell^{N-1} < \ell^N = T$, and using C3 build a corresponding sequence $\{\chi_r^n\}_{n=0}^N \in$ supp $\chi_r|\delta_r, w_{jr}$, such that $p_r^n = (\eta(\ell^n), \widehat{v}_{-j})$. Then, as $\max_n\{\ell^n - \ell^{n-1}\}$ becomes arbitrarily small

$$\sum_{n=0}^{N-1} \phi_j(\ell^n)(\ell^{n+1} - \ell^n) \approx \int_0^T \phi_j(\ell)\, \mathrm{d}\ell, \tag{41}$$

where all the elements in the Riemann sum are identified since by S1 each $\phi_j(\ell^n)$ can be replaced by

$$mr_{jr}^n = \sigma_j(\chi_r^n) - p_{jr}^n \cdot \int_{\partial\mathcal{D}_j(p_r^n)} f(v|\delta_r^n)\, \mathrm{d}v, \tag{42}$$

which is identified by Lemma 1.■

Appendix B:
Nash-in-prices under the ACA

Combining the notation of the institutional details described in Section 2 with the one of the econometric model in Section 3, the profit of plan j offered in region r when the base prices in the region are collected in b_r is:

$$\Pi_{jr}(b_r) = \sum_{\theta,\delta} G_r(\theta) F_r(\delta|\theta) \mathcal{S}^{\theta,\delta}_{jr}(hA^{\tau} b_r - V(\theta, b_r))(hA^{\tau} b_{jr} - \phi(\theta, \delta, z_{jr}; \kappa)), \quad (43)$$

where in this expression A^{τ} is the age-adjustment factor where τ is the average age of the household members, and h is the size of the household. The subsidy $V(\theta, b_r)$ is defined in equation (2) as a function of base prices and household income and demographics. The total profit for insurer n' is then:

$$\Pi_{n'r}(b_r) = \sum_{j:n'(j)=n} \Pi_{jr}(b_r).\text{[40]} \quad (44)$$

In a Nash-in-prices equilibrium:

$$b_{nr} \in \arg\max \Pi_{nr}(b_{nr}, b_{-nr}) \quad \text{for all } n. \quad (45)$$

B1: MR=MC under Nash-in-prices

In Section 5 I discussed how, to support the Nash-in-prices assumption, one can either (a) check the extent to which first-order optimality conditions under Nash-in-prices are met by the estimated demand-cost model, or (b) estimate cost imposing these optimality conditions directly, and then compare the resulting cost estimates to those obtained using claims data.

Both exercises require one to use the equilibrium assumption through the corresponding necessary first-order conditions:

$$\frac{\partial \Pi_{nr}(b_r)}{\partial b_{jr}} = 0 \quad \text{for all } j \quad \text{and all } n, \quad (46)$$

which substituting terms is, for all j and all n:

$$MR_{jr}(b_r) = \sum_{\theta} G_r(\theta)\left(S^{\theta}_{jr} hA^{\tau} + \sum_{j':n(j')=n(j)} \frac{\partial S^{\theta}_{j'r}}{\partial P_{jr}}\left(hA^{\tau} - \frac{\partial V(\theta, b_r)}{\partial b_{jr}}\right) hA^{\tau} b_{j'r}\right) = \quad (47)$$

$$= \quad MC_{jr}(b_r;\kappa) = \sum_{\theta,\delta}\left(G_r(\theta)F_r(\delta|\theta) \sum_{j':n(j')=n(j)} \frac{\partial S^{\theta,\delta}_{j'r}}{\partial P_{jr}}\left(hA^{\tau} - \frac{\partial V(\theta, b_r)}{\partial b_{jr}}\right)\phi(\theta, \delta, z_{j'r}; \kappa)\right). \quad (48)$$

For approach (a), I construct both, MR_{jr} and MC_{jr} for all jr, using demand estimates and the cost parameters $\widehat{\kappa}^{Claims}$ estimated using jr level average claims data. In Figure 11 in Section 5 I analyze the resulting relationship between $MR_{jr}(b_r)$ and $MC_{jr}(b_r; \widehat{\kappa}^{Claims})$.

For approach (b), I use the relationship $MR_{jr} = MC_{jr}(b_r; \kappa)$ as estimating moment to obtain the parameter κ, say $\widehat{\kappa}^{FOC}$. I then analyze the relationship between $MC_{jr}(b_r; \widehat{\kappa}^{FOC})$ and $MC_{jr}(b_r; \widehat{\kappa}^{Claims})$.

[40] Including risk-adjustment would imply defining profits as

$$\Pi_{n'r}(b_r) = \sum_{j:n'(j)=n} \Pi_{jr}(b_r) + RA\left(\{f_{n'r}(\theta,\delta), f_{-n'r}(\theta,\delta)\}_{\theta,\delta}\right),$$

where $RA(\cdot)$ is a risk-adjustment transfer computed as a function of the composition across θ, δ (and thus expected risk) of insurer n' enrollment, $\{f_{n'r}(\theta,\delta)\}_{\theta,\delta}$, and of the same measure among other insurers ($\{f_{-n'r}(\theta,\delta)\}_{\theta,\delta}$).

B2: Equilibrium under different subsidy designs

In my counterfactuals I compare equilibrium under three types of subsidy design:

(a) Price-linked subsidy under the ACA;

(b) Voucher of amount equivalent to the equilibrium subsidy under the ACA;

(c) Age-adjusted voucher.

Each design corresponds to a different function $V(\theta, b_r)$, and equilibrium b_r^* in a region r corresponds to the solution of $MR_{jr} = MC_{jr}$ subject to $\Pi_{jr} \geq 0$ and MLR constraints (see below).

For (a), one has the subsidy function specified under the ACA (equation (2) in Section 2). At the corresponding equilibrium base prices, $b_r^{*,ACA}$, a household with characteristics θ receives a subsidy equal to $V_r^{*,ACA}(\theta) = V(\theta, b_r^{*,\theta})$. Equilibrium base prices solve (48).

In (b), I eliminate the dependence of subsidies from base prices and set constant vouchers $V(\theta, b_r) = V_r^{*,ACA}(\theta)$ for all b_r. The necessary FOC becomes:

$$0 = \sum_{\theta,\delta} \left(G_r(\theta) F_r(\delta|\theta) \left(S_{jr}^{\theta,\delta} hA^\tau + \sum_{j':n(j')=n(j)} \frac{\partial S_{j'r}^{\theta,\delta}}{\partial P_{jr}} hA^\tau \left(hA^\tau b_{j'r} - \phi(\theta, \delta, z_{j'r}; \kappa) \right) \right) \right), \tag{49}$$

where the key difference from (48) is that the term $\frac{\partial V(\theta, b_r)}{\partial b_{jr}}$ is absent. This term is what captures insurers incentives to charge higher markups under price-linked subsidies (Jaffe and Shepard, 2017). Under this type of scheme, since the government increases subsidies if all base prices are higher (indeed if the base price of the second-cheapest Silver —benchmark— plan is higher), the slope of demand curves as perceived by insurers for pricing purposes is lower than the true slope of demand:

$$\left| \frac{\partial S_{j'r}^{\theta}}{\partial P_{jr}} \left(hA^\tau - \frac{\partial V(\theta, b_r)}{\partial b_{jr}} \right) \right| \leq \left| \frac{\partial S_{j'r}^{\theta}}{\partial P_{jr}} hA^\tau \right|. \tag{50}$$

with strict inequality for the benchmark plan.

In (c) I modify the voucher $V_r^{*,ACA}(\theta)$ by lowering the discount for the over-50 (without children and FPL<250) by Δ^{50+}, and simultaneously increasing the discount for the under-50 (without children and FPL<250) by Δ^{50-}. When $\Delta^{50+} < 0$ and $\Delta^{50-} > 0$ the composition of buyers at any given b_r is different than under $V_r^{*,ACA}(\theta)$: over-50 face higher premium, and purchase less, while under-50 face lower premium, and purchase more. This lowers average and marginal cost, and increase demand elasticity, at any given b_r. Therefore, at the new equilibrium base prices are lower. If this mechanism is sufficiently strong (ultimately an empirical question), base prices can drop enough to compensate (after age-adjustment) the reduction in subsidy for the older group, so that their final post-voucher premium is lower despite receiving a less generous subsidy.

Medical loss ratio. The ACA requires each insurer to have a medical loss ratio (MLR) of at least 0.8; insurers must issue rebates to consumers upon violating this limit. The MLR under the ACA is computed as

$$\text{MLR} = \frac{\text{Health Care Claims + Quality Improvements Expenses}}{\text{Premiums} - \text{Taxes, Licensing \& Regulatory Fees}}. \tag{51}$$

This is different from the traditional medical loss ratio, equal simply to the ratio between claims and premiums. Since quality improvements, taxes, and fees are not measured precisely, this formulation of the MLR leaves significant flexibility for insurers to not violate the constraint.

In my equilibrium simulations I cannot impose the MLR constraint precisely, as my model does not contain measures of quality improvements, nor of the adjustments to the denominator in (51). However, I do impose a relaxed constraint requiring insurer-level (traditional) medical loss ratio —expected total claims over total premiums— to be at least 0.7. As showed in Table 12, at the equilibrium solutions this constraint is (on average) not binding. Nevertheless, the constraint still affects equilibrium calculations by removing from the set of feasible strategies global deviations which would imply that per-buyer profits are too high, violating the MLR limit.

Figures and Tables

Figure 1: Premiums by plan, region, and household observables

Note: the figure represents, for every region, how premiums vary across plan (insurer-metal tier pair) and household demographics. Each point on the x-axis corresponds to a plan (e.g. Kaiser, HMO, Bronze), and each dot in the scatter plot corresponds to the annual premium paid for the plan by a specific type of household (e.g. 33-year-old, single, with income equal to 230% of the FPL). Products are ordered by metal tier, from left to right, with colors indicating Catastrophic coverage (blue), Bronze (maroon), Silver (green), Gold (yellow), and Platinum (gray).

Figure 2: Subsidy eligible households (ACS) and Covered CA enrollment

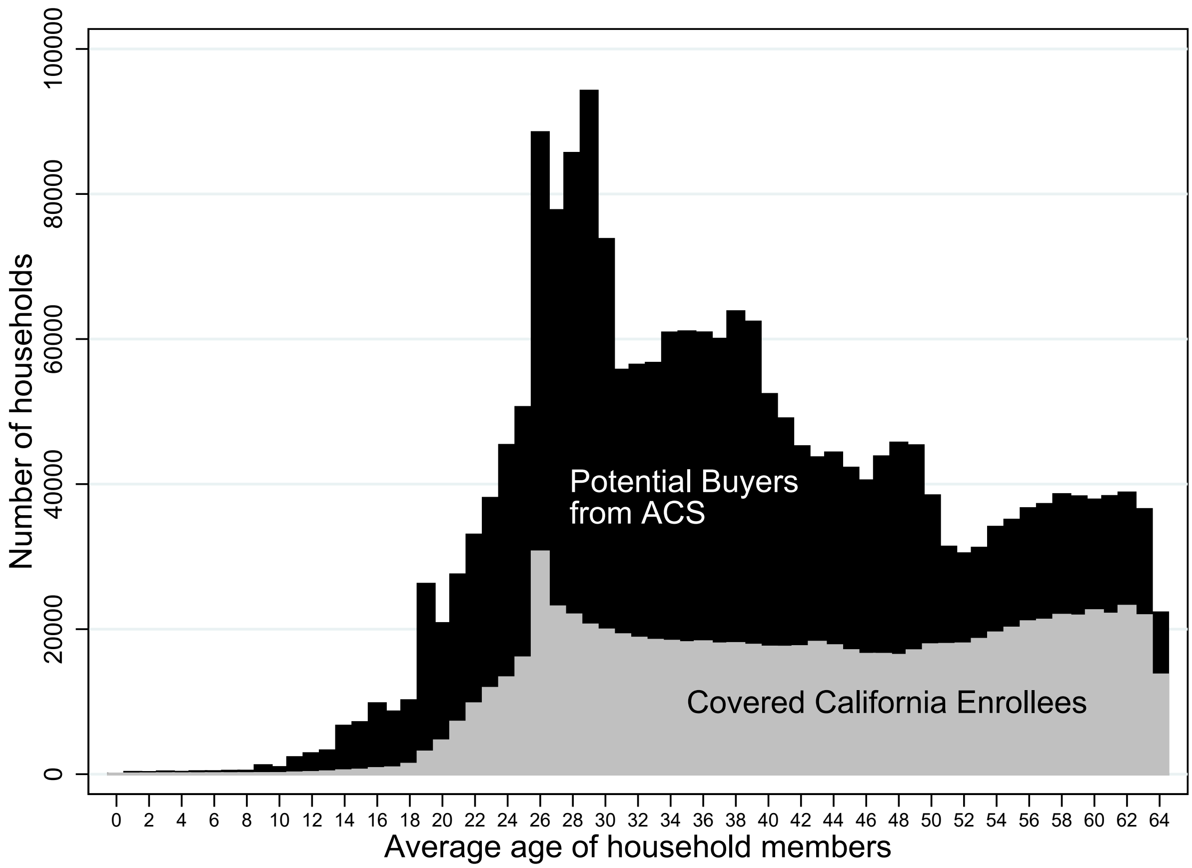

Note: the figure shows the age composition of potential buyers eligible for subsidies in the 2013 ACS and the age composition of subsidized buyers in the 2014 California exchange. Potential buyers are defined as uninsured or purchasing individual coverage. Subsidy eligibility is calculated as a function of annual income and household composition in the ACS, and it is observed in the enrollment data.

Figure 3: Average premium paid and exchange participation

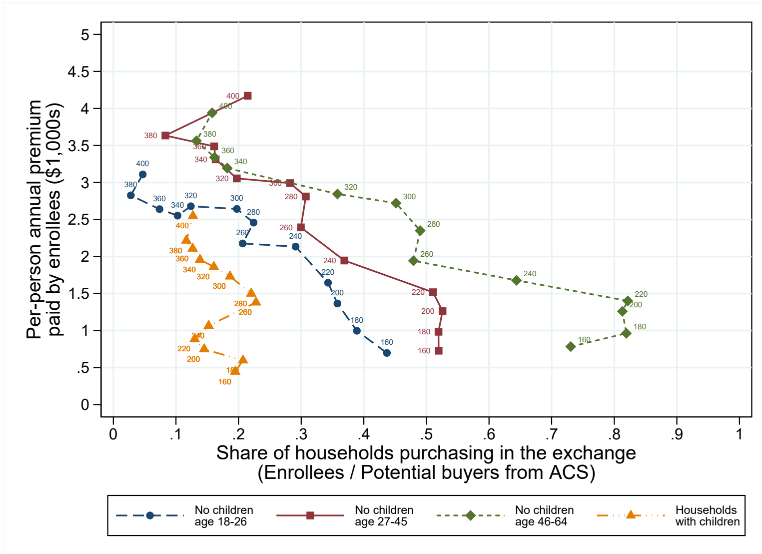

Note: the figure shows average premium paid (y-axis) and average probability of coverage (x-axis) across income and demographic groups. Each line connects different income bins (lower bound labels every point, expressed as % of FPL) for a different age and hosuehold size combination.

Figure 4: Average premium and choice probability, difference between high and low deductible

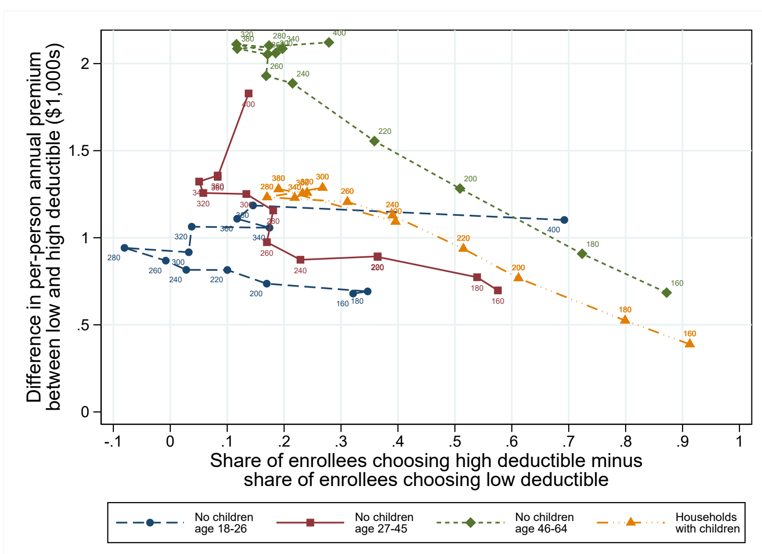

Note: the figure plots, across income and demographic groups, the relationship between average difference between Bronze (high-deductible) premium and premium of higher coverage plans (y-axis), and difference in corresponding market shares. Each line connects different income bins (lower bound labels every point, expressed as % of FPL) for a different age and hosuehold size combination.

Figure 5: Average prior willingness-to-pay across households

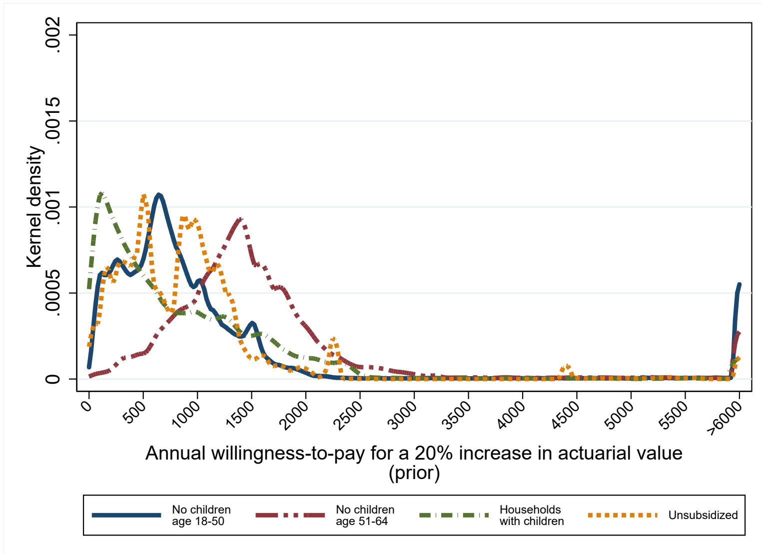

Note: kernel density of prior WTP (ratio β/α) based on mixed logit estimates, distinguishing between different demographic groups.

Figure 6: Average posterior willingness-to-pay across households conditional on choice

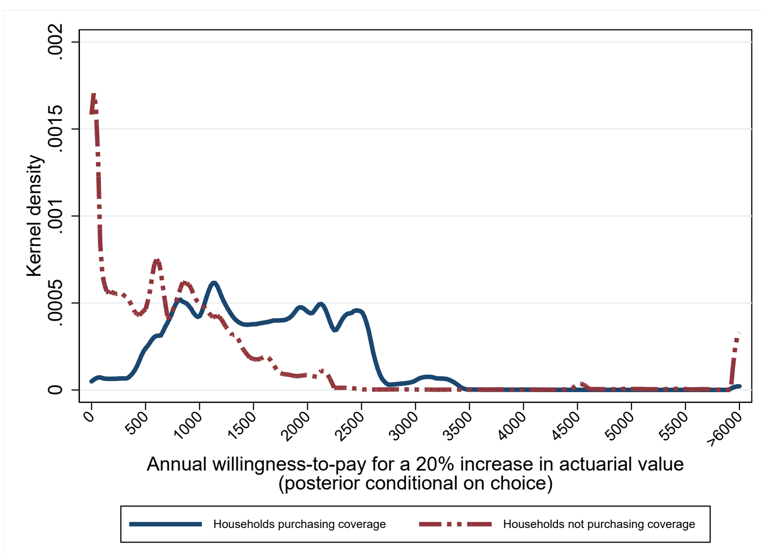

Note: kernel density of posterior WTP (ratio β/α), conditional on observed choice, distinguishing between enrollees and households who do not purchase coverage.

Figure 7: Percent change in enrollment if all premiums increase by $100/year

(a) Subsidized households without children

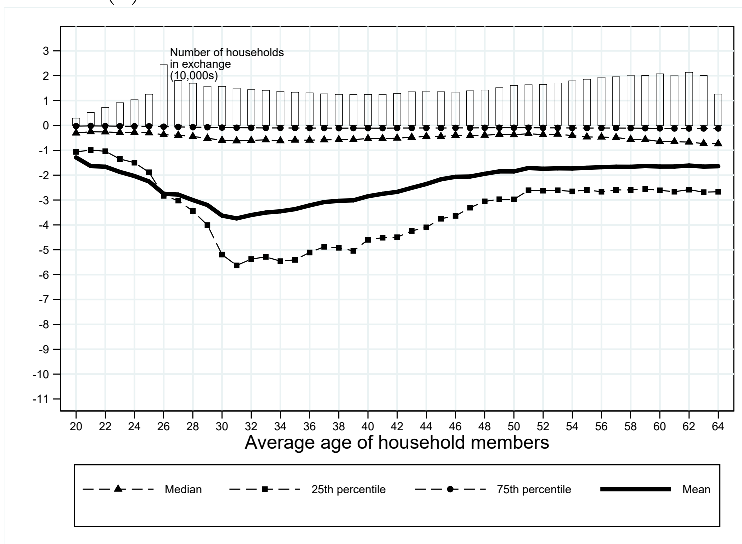

(b) Subsidized households with children

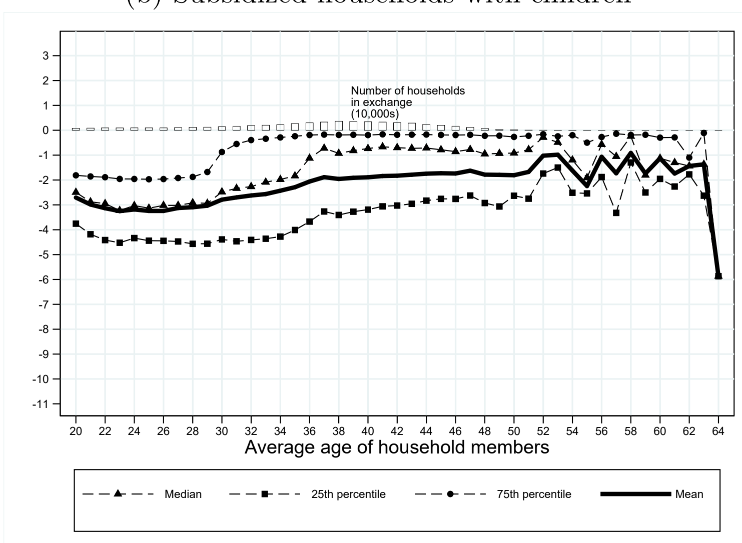

Note: extensive margin semi-elasticity by age and type of household based on mixed logit estimates and number of enrollees in the exchange. Each panel shows (for subsidized households without children and subsidized households with children, respectively) the mean, median, and 25th and 75th percentiles of percentage drop in coverage if all premiums increase by $100 per-year. The bars indicate the number of enrollees in the exchange in 10,000s. Graphs including unsubsidized households in Online Appendix.

Figure 8: Demand curves from mixed logit estimates

(a) Subsidized, no children, under 50

(b) Subsidized, no children, over 50

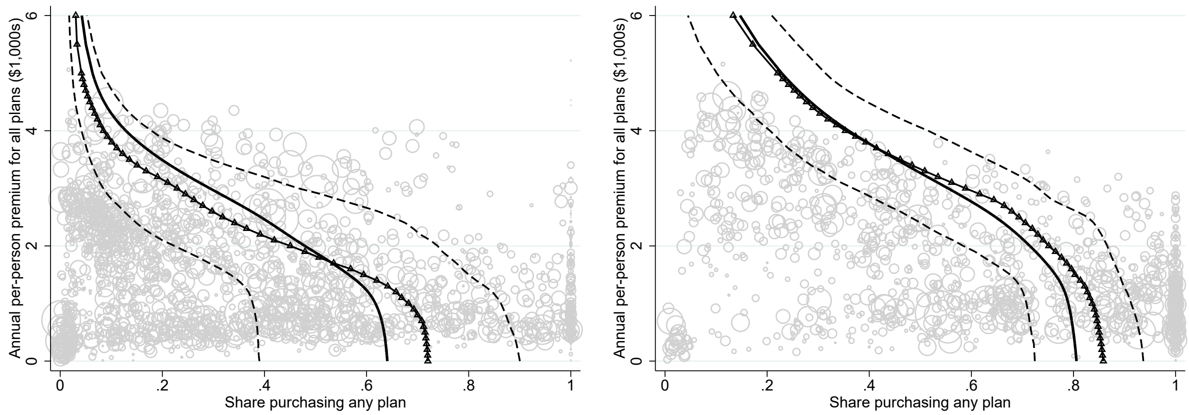

Note: demand curves from mixed logit estimates are derived computing, for each household, the probability of purchasing any plan in the exchange (x-axis) if all premiums are set equal to the amount on the y-axis. Each graph shows the 25th and 75th percentile (dashed lines), the median (connected triangles), and mean of the probability of purchase at a given premium. The shaded scatter plot corresponds to the combinations of average premium paid and coverage probability for a specific age-income-region combination as observed in the data; circles are proportional to the number of potential buyers in the age-income-region cell. Figures for all groups in Online Appendix.

Figure 9: Cost estimates

(a) Average cost by age, tier

(b) Cost dispersion by age, Silver plans

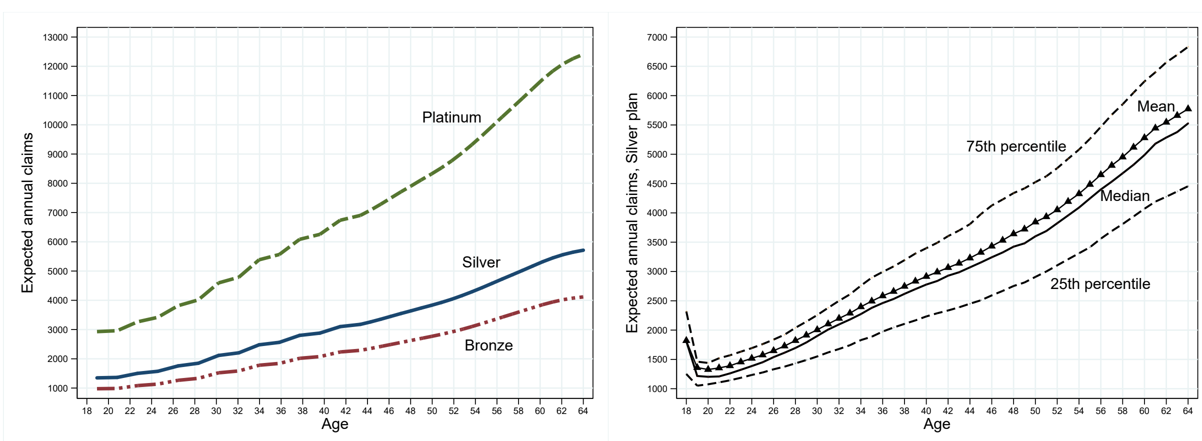

Note: estimates of expectected insurer cost by age and tier of coverage. Panel (a) shows average expected cost by age, distinguishing between Bronze, Silver, and Platinum plans. Panel (b) shows the distribution of the expected cost for a buyer upon enrolling in a Silver plan offered by Anthem (largest insurer in Covered California).

Figure 10: Joint density of preferences and cost

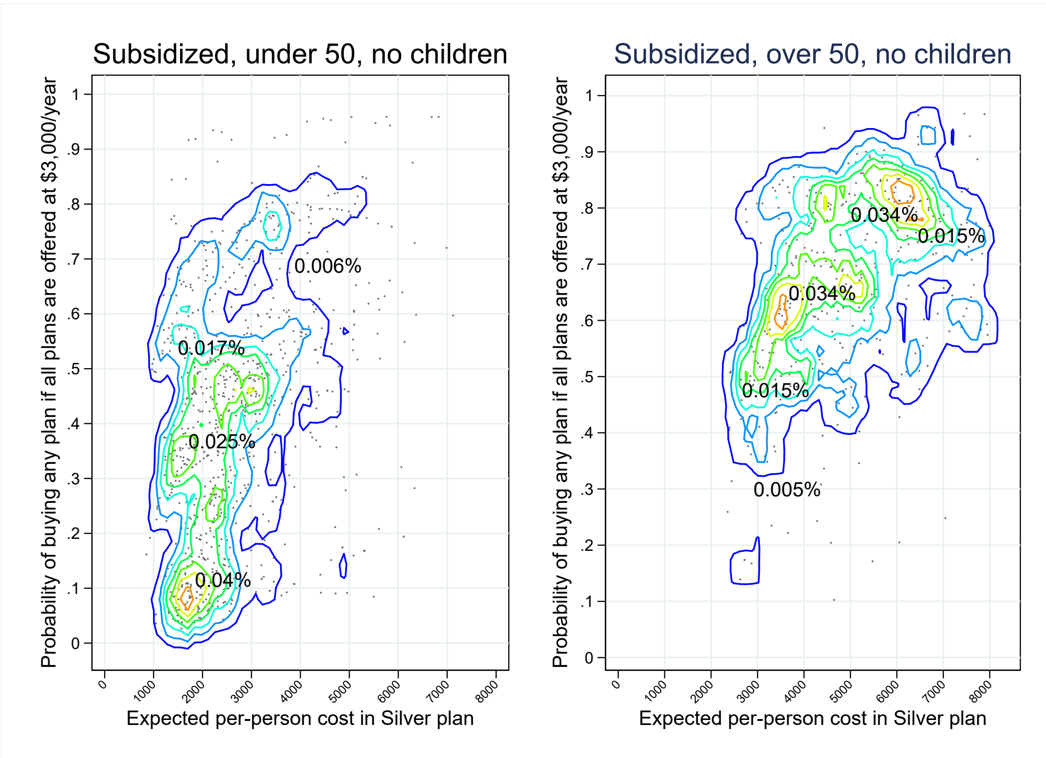

Note: based on mixed logit demand estimates and cost estimates, the figure shows countour lines of the joint kernel density of the probability of buying any plan if all premiums are $3,000/year (y-axis) and the expected cost upon enrolling in a Silver plan offered by Anthem (x-axis). In each panel the joint density overlays a scatter plot of a 1% sample of individuals from the corresponding demographic group. Graphs including households with children and unsubsidized households in Online Appendix.

Figure 11: Nash-in-prices assumption and marginal cost estimation

(a) Marginal revenue vs. marginal cost under Nash-in-prices

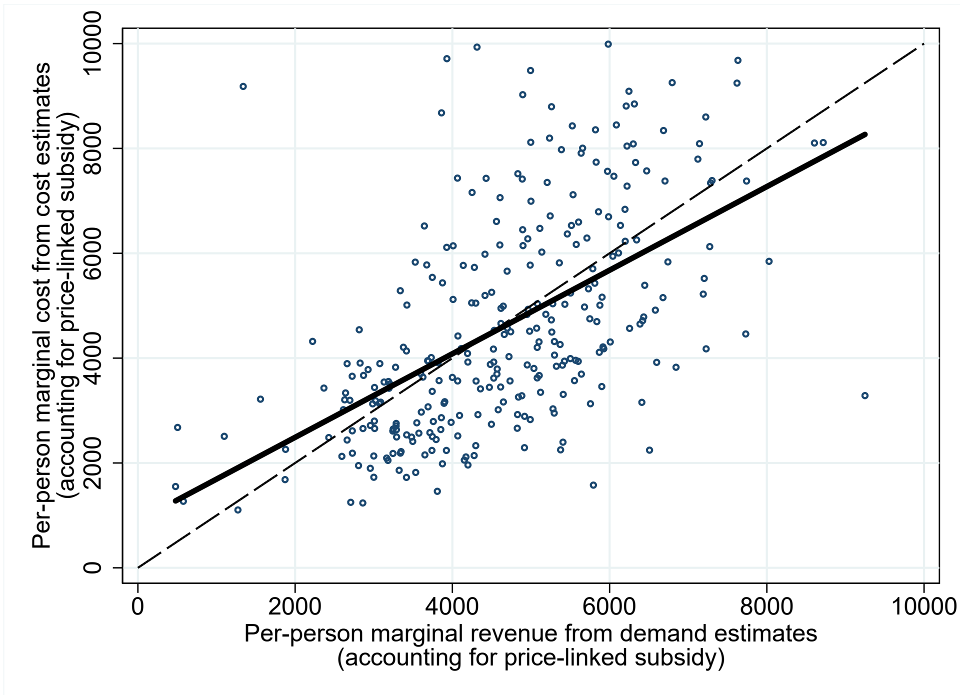

(b) Marginal cost from claims vs. marginal cost from Nash-in-prices

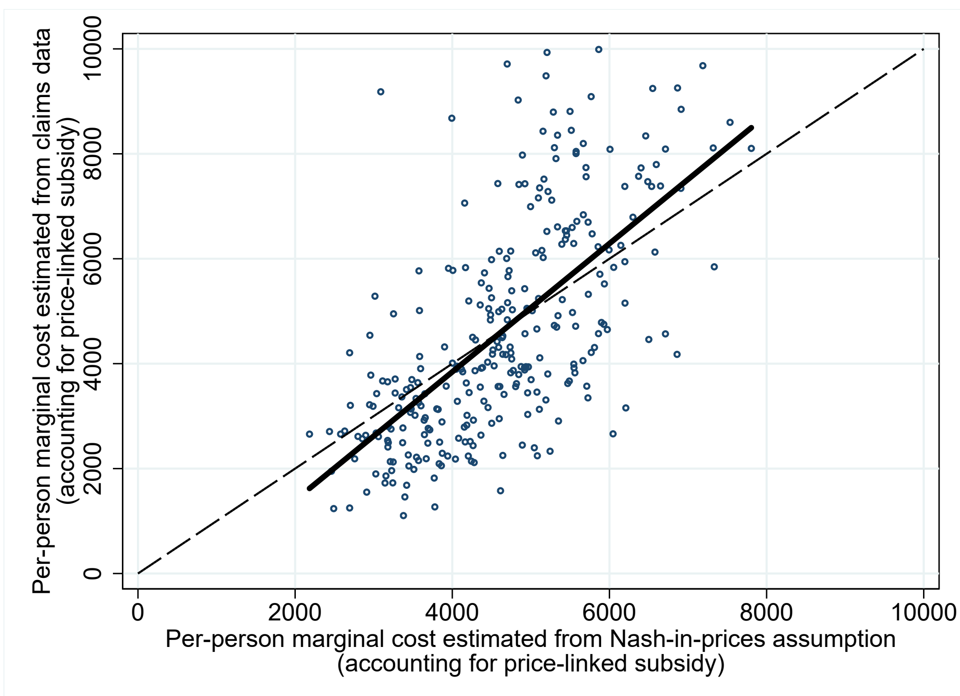

Panel (a): relationship between per-person marginal revenue and marginal cost based on mixed logit and cost estimates from average claims. Each quantity is calculated imposing age-rating adjustments and price-linked subsidies. Details in Appendix B1.

Panel (b): relationship between per-person marginal cost estimated using average claims and per-person marginal cost estimated using Nash-in-prices assumptions. Each quantity is calculated imposing age-rating adjustments and price-linked subsidies. Details in Appendix B1.

The same comparisons using fixed vouchers instead of price-linked subsidies can be found in Online Appendix.

Figure 12: Equilibrium enrollment varying monthly subsidies by age

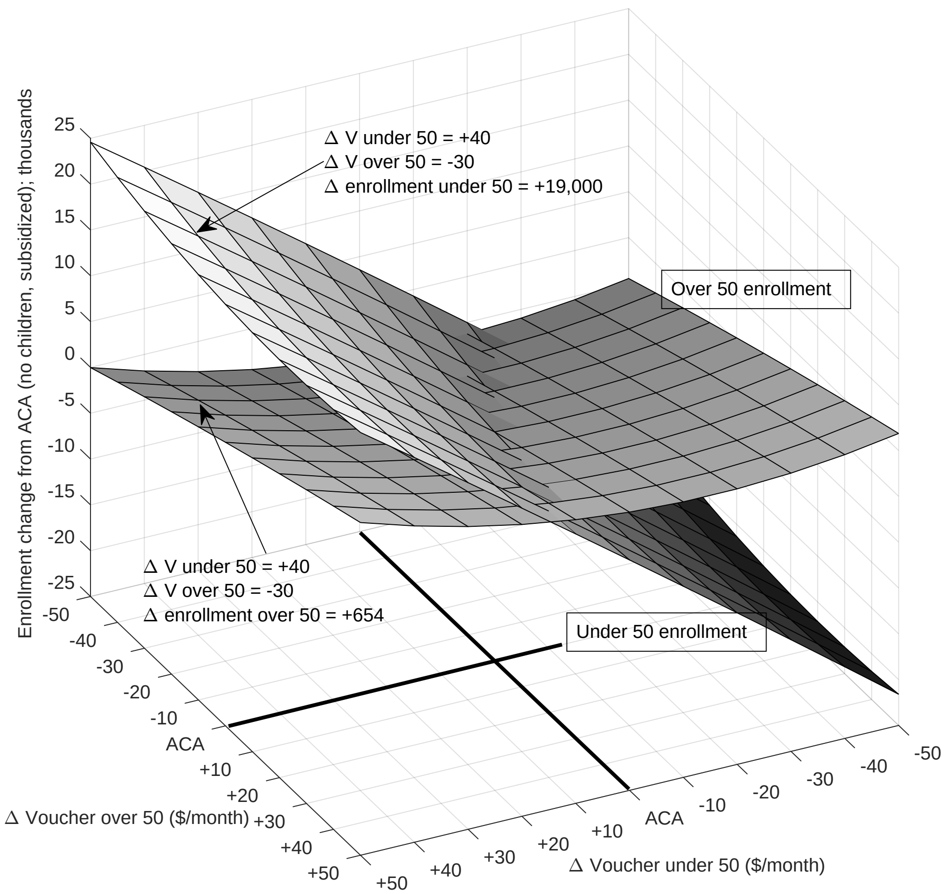

Note: the two surfaces represent the change in enrollment (in 1,000s) for under-50 and over-50, respectively, as the monthly vouchers of each age group for income up to 250% of the FPL vary around the ACA subsidies from -$50 to $50 in steps of $10. The arrow pointers highlight a combination of increase in subsidiy of under-50 and decrease in subsidy of over-50 that leads to higher enrollment in both groups. For each point in the grid of vouchers, the starting point is the base prices in the data, and equilibria are then computed zeroing FOC under a medical loss ratio of at least 0.7 and a positive expected profit constraint (see Appendix B for details).

Figure 13: Equilibrium outcomes varying monthy subsidies by age

Enrollment, subsidized with no children:

(a) Under 50 (ACA = 480,042)

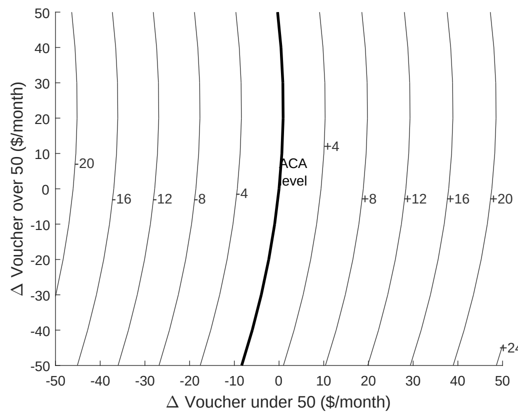

(b) Over 50 (ACA = 358,104)

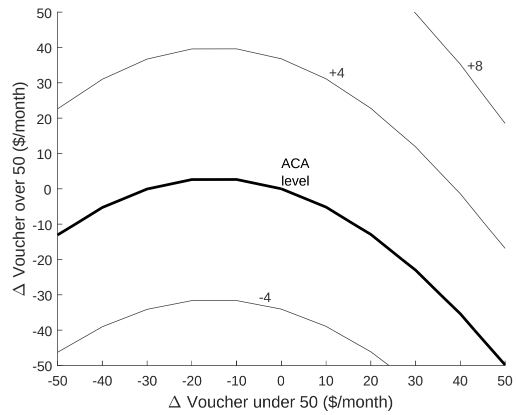

Spending:

(c) Per-person annual subsidy (ACA = $3,675)

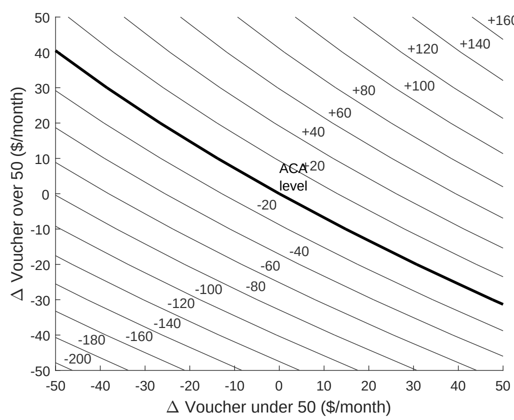

(d) Total subsidy spending (ACA = $3,396 million)

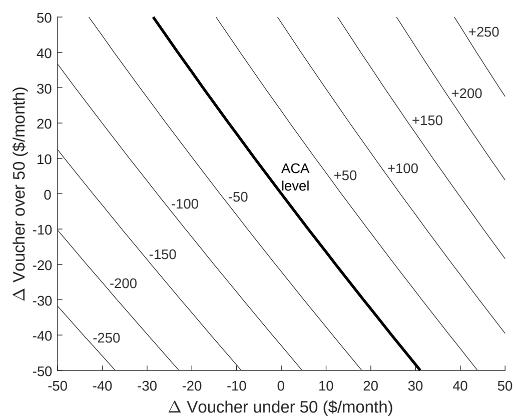

Cost and revenue:

(e) Average cost (ACA = $3,522)

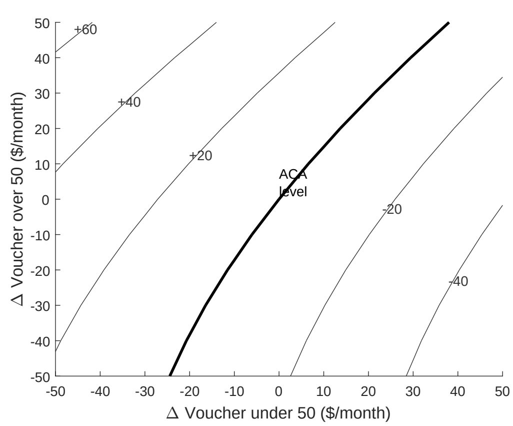

(f) Average revenue (ACA = $4,571)

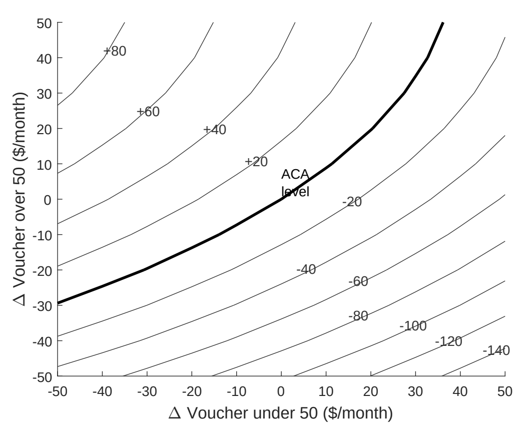

Note: Level curves of equilibrium outcomes as functions of the changes (from ACA level) in voucher for the under-50 (x-axis) and over-50 (y-axis). Level curves correspond to changes in 1,000 individuals for enrollment, change in $ for per-person subsidy, average cost, and average revenue, and $M for total spending. For each point in the grid of vouchers, the starting point is the base prices in the data, and equilibria are then computed zeroing FOC under a medical loss ratio of at least 0.7 and a positive expected profit constraint (see Appendix B for details).

Table 1: Price caps for subsidy calculation in ACA marketplaces

Income as % of FPL	up to 150%	150-200%	200-250%	250-400%
Max % of income to buy 2^{nd} cheapest Silver	4%	6.3%	8.05%	9.5%
Price cap of 2^{nd} cheapest Silver (single)	$684	$1,452	$2,416	$4,368
Price cap of 2^{nd} cheapest Silver (couple+1 child)	$1,164	$2,472	$4,008	$7,392

Note: The table shows, as a function of income, the maximum amount that can be spent on the second cheapest Silver plan in the region. For each age-income pair, the subsidy is computed as the difference between the premium of this product (after age adjustment) and the corresponding share of annual income for the buyer. The bottom row shows the corresponding price cap on monthly price for the second cheapest Silver plan in the region for singles and households of three.

Table 2: Standardized financial characteristics in Covered California

	Annual decuctible	Maximum out-of-pocket	Primary care visit	Emergency Room	Specialist visit	Preferred drugs	Advertised coverage*
Catastrophic	n.a.	$6,600	n.a.(1)	n.a.(1)	n.a.(1)	n.a.(1)	n.a.(3)
Bronze	$5,000	$6,250	$60	$300(2)	$70(2)	$50(2)	60%
Silver (>250% FPL)	$2,250	$6,250	$45	$250	$65	$50	70%
Silver (200-250% FPL)	$1,850	$5,200	$40	$250	$50	$35	74%
Gold	$0	$6,250	$30	$250	$50	$50	79%
Silver (150-200% FPL)	$550	$2,250	$15	$75	$20	$15	88%
Platinum	$0	$4,000	$20	$150	$40	$15	90%
Silver (100-150% FPL)	$0	$2,250	$3	$25	$5	$5	95%

Source: http://www.coveredca.com/PDFs/2015-Health-Benefits-Table.pdf
(1): Pay the necessary fee (negotiated between carrier and provider) until the maximum out-of-pocket is met. (2): After deductible is met, before pay the necessary fee (negotiated between carrier and provider). (3): Pay the full cost until maximum out-of-pocket is met *: These percentages are displayed to buyers when comparing products.

Table 3: Buyers, population, enrollment, and claims

	Individual-level purchases and potential buyers									
	(a): **Covered California 2014**					(b): **ACS 2013 potential buyers**				
Variable	Obs	Mean	Std. Dev.	Min	Max	Obs	Mean	Std. Dev.	Min	Max
Age	877,365	43.19	12.98	0	64	3,392,942	41.34	13.02	0	64
Household size	877,365	1.47	0.77	1	5	3,392,942	1.80	0.98	1	5
Income as % of FPL (if subsidized)	826,484	215.20	63.17	100	400	2,231,013	231.65	80.76	100	400

	Plan-level enrollment and claims									
	(c): **Enrollment in 2014 Covered CA**					(d): **2014 per-person claims ($1,000); CMS filings**				
Coverage level	N. plans	Mean	Std. Dev.	Min	Max	N. plans	Mean	Std. Dev.	Min	Max
Catastr.	80	180	225	1	1,055	78	1.969	2.088	0.283	6.432
Bronze	140	2,427	2,907	1	16,671	138	2.303	1.411	0.366	6.432
Silver	90	10,303	12,477	41	57,363	88	3.415	1.386	1.206	6.784
Gold	90	988	1,265	7	5,511	88	4.514	1.855	1.707	12.268
Platinum	90	889	1,033	8	5,555	84	8.707	5.231	1.812	20.310

Note: Panel (a) contains summary statistics of household characteristics within Covered California enrollees. Panel (b) contains summary statistics of household characteristics within ACS potential buyers (uninsured or individually insured). Panel (c) contains summary statistics of plan-level enrollment by metal tier. Panel (d) contains summary statistics of plan-level average claims from CMS filings by metal tier.

Table 4: Prices by age and income

	Average premium for subsidized buyers (200-220% FPL, $1,000)									
	(a): **Singles**					(b): **Non-singles**				
Average age of household members	Coverage level Catastr.	Bronze	Silver	Gold	Platinum	Coverage level Catastr.	Bronze	Silver	Gold	Platinum
10-30	n.a.	1.137	1.661	2.068	2.440	n.a.	0.922	2.542	3.967	5.205
30-50	n.a.	0.758	1.691	2.383	3.027	n.a.	0.472	2.525	4.246	5.701
50-64	n.a.	0.370	1.788	3.132	4.346	n.a.	0.340	2.333	4.664	7.208
	Average insurer revenue (premium for non-subsidized buyers; $1,000)									
	(c): **Singles**					(d): **Non-singles**				
Average age of household members	Coverage level Catastr.	Bronze	Silver	Gold	Platinum	Coverage level Catastr.	Bronze	Silver	Gold	Platinum
10-30	1.706	1.922	2.473	2.921	3.298	3.748	7.593	9.347	10.915	12.000
30-50	2.355	3.165	4.111	4.846	5.492	4.922	9.862	13.285	13.991	15.276
50-64	n.a.	6.031	7.825	9.275	10.548	n.a.	12.081	15.750	18.014	20.247

Note: Panel (a) contains average premium ($1,000s) by age group and metal tier for single buyers with income between 200-220% of the FPL. Panel (b) contains average premium ($1,000s) by age group and metal tier for non-single buyers with household income between 200-220% of the FPL. Panel (c) contains average premium ($1,000s) by age group and metal tier for single buyers with income above 400% of the FPL (unsubsidized). Panel (d) contains average premium ($1,000s) by age group and metal tier for non-single buyers with household income above 400% of the FPL (unsubsidized).

Table 5: Regional enrollment and market shares

		Share of enrollees choosing:										
Region	Enrollment (1)	Anthem (2)	Blue (3)	CCH (4)	Contra C (5)	HealthNet (6)	Kaiser (7)	LAC (8)	Molina (9)	Sharp (10)	Valley (11)	Western (12)
1	47,729	0.907	0.088				0.005					
2	48,673	0.314	0.169			0.037	0.429					0.050
3	65,623	0.345	0.253				0.379					0.023
4	38,218	0.167	0.226	0.296		0.057	0.254					
5	36,489	0.069	0.384		0.029	0.019	0.499					
6	58,797	0.286	0.291				0.423					
7	60,382	0.582	0.119			0.048	0.221				0.030	
8	24,492	0.152	0.234	0.107		0.050	0.457					
9	28,291	0.564	0.341			0.094						
10	58,178	0.682	0.118			0.016	0.184					
11	26,190	0.450	0.404				0.146					
12	57,358	0.487	0.433				0.080					
13	5,682	0.499	0.489				0.012					
14	16,314	0.362	0.472			0.044	0.122					
15	163,381	0.110	0.376			0.359	0.089	0.062	0.003			
16	202,185	0.230	0.232			0.330	0.101	0.097	0.011			
17	112,764	0.165	0.306			0.318	0.160		0.052			
18	124,043	0.209	0.372			0.332	0.087					
19	116,425	0.177	0.214			0.328	0.168		0.005	0.108		

Note: Column (1) contains total region enrollment (individuals), while columns (2)-(12) report, for each insurer, its share of regional enrollees. Carriers are, from left to right: Anthem Blue Cross, Blue Shield of CA, Chinese C.H., Contra Costa, HealthNet, Kaiser, LA Care, Molina, Sharp, Valley, and Western

Table 6: Parameter Estimates from Logit and Nested Logit

	Logit					Nested logit				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Premium coefficient, single subsidized:										
Premium ($1000) × Age ∈ [10,20)	-0.405***	-0.391***	-0.364***	-0.285***	-0.221***	-0.295***	-0.301***	-0.272***	-0.283***	-0.294***
	(0.0274)	(0.0284)	(0.0278)	(0.0301)	(0.0380)	(0.0286)	(0.0289)	(0.0285)	(0.0313)	(0.0391)
Premium ($1000) × Age ∈ [20,30)	-0.478***	-0.460***	-0.434***	-0.356***	-0.287***	-0.330***	-0.336***	-0.308***	-0.319***	-0.335***
	(0.0205)	(0.0212)	(0.0208)	(0.0232)	(0.0301)	(0.0209)	(0.0211)	(0.0211)	(0.0233)	(0.0301)
Premium ($1000) × Age ∈ [30,40)	-0.455***	-0.437***	-0.414***	-0.339***	-0.273***	-0.305***	-0.311***	-0.286***	-0.297***	-0.314***
	(0.0206)	(0.0208)	(0.0206)	(0.0226)	(0.0291)	(0.0212)	(0.0212)	(0.0216)	(0.0234)	(0.0295)
Premium ($1000) × Age ∈ [40,50)	-0.395***	-0.379***	-0.360***	-0.289***	-0.230***	-0.259***	-0.264***	-0.244***	-0.254***	-0.270***
	(0.0171)	(0.0170)	(0.0169)	(0.0187)	(0.0245)	(0.0167)	(0.0167)	(0.0172)	(0.0188)	(0.0248)
Premium ($1000) × Age ∈ [50,60)	-0.271***	-0.255***	-0.244***	-0.183***	-0.132***	-0.145***	-0.149***	-0.137***	-0.145***	-0.160***
	(0.0135)	(0.0138)	(0.0139)	(0.0170)	(0.0218)	(0.0142)	(0.0145)	(0.0150)	(0.0165)	(0.0207)
Premium ($1000) × Age ∈ [60,64]	-0.195***	-0.183***	-0.289***	-0.240***	-0.195***	-0.0926***	-0.0967***	-0.204***	-0.210***	-0.224***
	(0.0120)	(0.0122)	(0.0147)	(0.0163)	(0.0202)	(0.0117)	(0.0121)	(0.0142)	(0.0152)	(0.0201)
Difference in premium coefficient, other household types:										
Premium ($1000) × hh size=2	0.0697***	0.0594***	0.0671***	0.0451***	0.0161	0.00570	0.00708	0.0145	0.0177	0.0248
	(0.0132)	(0.0133)	(0.0136)	(0.0137)	(0.0168)	(0.0127)	(0.0128)	(0.0132)	(0.0136)	(0.0165)
Premium ($1000) × hh size=3	0.224***	0.210***	0.200***	0.165***	0.120***	0.120***	0.123***	0.112***	0.117***	0.129***
	(0.0199)	(0.0198)	(0.0194)	(0.0197)	(0.0240)	(0.0188)	(0.0189)	(0.0186)	(0.0194)	(0.0234)
Premium ($1000) × hh size=4	0.320***	0.304***	0.289***	0.247***	0.195***	0.200***	0.203***	0.188***	0.194***	0.207***
	(0.0234)	(0.0240)	(0.0231)	(0.0238)	(0.0289)	(0.0233)	(0.0236)	(0.0229)	(0.0239)	(0.0284)
Premium ($1000) × hh size=5	0.359***	0.340***	0.323***	0.276***	0.216***	0.215***	0.219***	0.200***	0.207***	0.221***
	(0.0234)	(0.0237)	(0.0229)	(0.0241)	(0.0290)	(0.0238)	(0.0239)	(0.0233)	(0.0245)	(0.0287)
Premium ($1000) x Non-subsidized	0.0799***	0.0853***	0.0978***	0.0840***	0.0945***	0.0876***	0.0889***	0.101***	0.102***	0.0996***
	(0.0145)	(0.0148)	(0.0142)	(0.0151)	(0.0155)	(0.0160)	(0.0160)	(0.0155)	(0.0154)	(0.0156)
Plan characteristics:										
Actuarial value (%)	-0.000913	-0.00279	-0.00298*			0.00567***	0.00492***	0.00476***		
	(0.00185)	(0.00175)	(0.00171)			(0.00143)	(0.00144)	(0.00139)		
Deductible($1000)				0.137***					-0.0394**	
				(0.0163)					(0.0156)	
Nesting Parameter:						0.522***	0.493***	0.496***	0.503***	0.412***
						(0.0162)	(0.0150)	(0.0154)	(0.0159)	(0.0240)
Observations	25,019	25,019	25,019	25,019	25,019	25,019	25,019	25,019	25,019	25,019
R-squared	0.393	0.421	0.439	0.449	0.514	0.510	0.514	0.533	0.533	0.546
Income group FE	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Control for Max-OOP and PCP visit	N	N	N	Y	N	N	N	N	Y	N
Insurer FE	Y	N	N	N	N	Y	N	N	N	N
Region FE	Y	N	N	N	N	Y	N	N	N	N
Insurer x Region	N	Y	N	N	N	N	Y	N	N	N
Insurer x Region x Over 50	N	N	Y	Y	N	N	N	Y	Y	N
Insurer x Region x Over 50 x Tier	N	N	N	N	Y	N	N	N	N	Y

Each observation is a unique (insurer-region-tier-demographic group). The 200 demographic groups are full interactions of age group (bins of 10 years), gender, annual income group (bins of $10,000), and household size. Standard errors in parentheses clustered at the region level (19 clusters); *** p<0.01, ** p<0.05, * p<0.1.

Table 7: Extensive Margin Semi-Elasticities from Nested Logit

	Age	Subsidized	Non-subsidized
Single	[20, 30)	-2.828	-2.577
	[30, 40)	-2.635	-2.313
	[40, 50)	-2.267	-1.837
	[50, 60)	-1.248	-0.687
	[60, 64]	-0.815	-0.166
Couples, no children	[20, 30)	-2.924	-2.439
	[30, 40)	-2.681	-2.193
	[40, 50)	-2.267	-1.732
	[50, 60)	-1.224	-0.578
	[60, 64]	-0.759	-0.052
Couples with children	—	-1.366	-0.701

Each cell reports the average percent change in probability of purchasing a plan in the exchange if all premiums increase by $100 per-year. Calculations are based on the nested logit estimates from column (7) in Table 6.

Table 8: Cross-Tier Semi-Elasticities from Nested Logit

Panel (a): Subsidized, no children, under 50					
$100 per-year increase in	% change in probability of choosing				
	Outside	Bronze	Silver	Gold	Platinum
Mandate penalty*	-0.521	2.446	2.624	2.406	2.412
Bronze premiums	0.187	-4.367	1.362	1.317	1.304
Silver premiums	0.172	1.016	-4.893	1.019	1.034
Gold premiums	0.0870	0.592	0.610	-5.015	0.593
Platinum premiums	0.0574	0.382	0.392	0.387	-5.215

Panel (b): Subsidized, no children , over 50					
$100 per-year increase in	% change in probability of choosing				
	Outside	Bronze	Silver	Gold	Platinum
Mandate penalty*	-0.310	0.961	1.104	1.052	1.056
Bronze premiums	0.0881	-1.938	0.506	0.502	0.498
Silver premiums	0.106	0.455	-2.211	0.545	0.551
Gold premiums	0.0655	0.317	0.365	-2.324	0.359
Platinum premiums	0.0450	0.218	0.251	0.245	-2.440

Panel (c): Subsidized, with children					
$100 per-year increase in	% change in probability of choosing				
	Outside	Bronze	Silver	Gold	Platinum
Mandate penalty*	-0.228	1.073	1.365	1.007	0.941
Bronze premiums	0.0826	-1.811	0.661	0.580	0.534
Silver premiums	0.105	0.497	-2.296	0.486	0.463
Gold premiums	0.0250	0.177	0.202	-2.125	0.171
Platinum premiums	0.0117	0.0823	0.0945	0.0834	-2.077

Panel (d): Unsubsidized					
$100 per-year increase in	% change in probability of choosing				
	Outside	Bronze	Silver	Gold	Platinum
Mandate penalty*	-0.0200	1.182	0.962	0.888	0.935
Bronze premiums	0.00806	-1.702	0.384	0.335	0.352
Silver premiums	0.00503	0.261	-1.616	0.230	0.233
Gold premiums	0.00352	0.156	0.160	-1.572	0.170
Platinum premiums	0.00298	0.126	0.133	0.143	-1.664

Note: the table shows the cross-tier semi-elasticity across demographic groups, based on the nested logit estimates in column (7) of Table 6. For each group, a cell contains the percentage change in the total enrollment of plans in the column tier if the all plans in the row tier increase annual premium by $100 per-year.

Table 9: Summary of mixed logit estimates

		Annual premium ($1,000s)					
		Mean	10th perc.	25th perc.	Median	75th perc.	90th perc.
Subsidized	no children, under-50	-1.580	-2.871	-2.314	-1.510	-0.849	-0.289
	no children, over-50	-1.482	-2.370	-1.802	-1.399	-1.062	-0.760
	with children	-0.442	-1.064	-0.751	-0.341	-0.046	0.077
Unsubsidized		-1.553	-2.827	-2.218	-1.529	-0.797	-0.232

		Actuarial value (%)					
		Mean	10th perc.	25th perc.	Median	75th perc.	90th perc.
Subsidized	no children, under-50	0.078	0.011	0.036	0.082	0.118	0.141
	no children, over-50	0.113	0.051	0.073	0.115	0.147	0.169
	with children	0.046	0.005	0.014	0.036	0.071	0.103
Unsubsidized		0.079	0.021	0.044	0.084	0.115	0.130

		Difference between Blue Cross Blue Shield and Anthem					
		Mean	10th perc.	25th perc.	Median	75th perc.	90th perc.
Subsidized	no children, under-50	0.242	-1.127	-0.374	0.452	1.140	1.720
	no children, over-50	0.240	-1.375	-0.449	0.607	1.278	1.732
	with children	0.541	-0.765	-0.276	0.953	1.601	1.877
Unsubsidized		0.202	-0.680	-0.330	0.386	1.067	1.611

		Difference between HealthNet and Anthem					
		Mean	10th perc.	25th perc.	Median	75th perc.	90th perc.
Subsidized	no children, under-50	-0.349	-2.324	-1.250	0.049	0.416	0.665
	no children, over-50	-0.453	-1.680	-1.143	-0.157	0.133	0.291
	with children	0.106	-1.541	-0.233	0.373	0.735	1.083
Unsubsidized		-0.446	-2.096	-1.183	-0.156	0.223	0.731

		Difference between Kaiser and Anthem					
		Mean	10th perc.	25th perc.	Median	75th perc.	90th perc.
Subsidized	no children, under-50	-0.108	-1.727	-0.363	0.346	0.717	1.154
	no children, over-50	0.042	-1.854	0.058	0.544	0.987	1.370
	with children	-0.161	-1.847	-0.574	0.323	0.750	1.015
Unsubsidized		-0.023	-1.736	-0.140	0.340	0.744	1.103

		Difference between other minor insurers and Anthem					
		Mean	10th perc.	25th perc.	Median	75th perc.	90th perc.
Subsidized	No children, under-50	-1.370	-2.141	-1.973	-1.721	-1.096	-0.060
	No children, over-50	-1.708	-2.387	-2.294	-2.102	-1.185	-0.455
	With children	-1.365	-2.335	-2.026	-1.656	-1.219	0.691
Unsubsidized		-1.386	-2.217	-1.994	-1.742	-0.668	-0.051

Note: the table shows summary statistics of random coefficients based on mixed logit estimates. For each parameter and demographic group, the table shows ther average of the corresponding coefficient, as well as 10th, 25th, 50th, 75th, and 90th percentiles.

Table 10: Cross-Tier Semi-Elasticities from Mixed Logit

Panel (a): Subsidized, no children, under 50					
$100 increase	Median % change in probability of choosing				
in	Outside	Bronze	Silver	Gold	Platinum
Mandate penalty*	-0.556	1.776	0.855	0.939	0.798
Bronze premiums	0.212	-12.029	1.751	1.411	0.902
Silver premiums	0.147	4.740	-6.877	4.850	3.887
Gold premiums	0.053	0.540	0.671	-11.171	0.836
Platinum premiums	0.054	0.401	0.777	0.936	-9.914

Panel (b): Subsidized, no children, over 50					
$100 increase	Median % change in probability of choosing				
in	Outside	Bronze	Silver	Gold	Platinum
Mandate penalty*	-1.503	1.014	0.764	0.717	0.672
Bronze premiums	0.672	-9.097	1.673	1.183	0.867
Silver premiums	0.392	4.572	-4.553	3.499	3.040
Gold premiums	0.116	0.407	0.447	-8.057	0.574
Platinum premiums	0.077	0.298	0.356	0.532	-7.217

Panel (c): Subsidized, with children					
$100 increase	Median % change in probability of choosing				
in	Outside	Bronze	Silver	Gold	Platinum
Mandate penalty*	-0.260	2.305	1.817	1.176	0.703
Bronze premiums	0.110	-6.065	0.849	0.570	0.446
Silver premiums	0.078	1.222	-3.867	0.813	0.577
Gold premiums	0.019	0.217	0.216	-3.871	0.149
Platinum premiums	0.008	0.130	0.126	0.112	-2.715

Panel (d): Unsubsidized					
$100 increase	Median % change in probability of choosing				
in	Outside	Bronze	Silver	Gold	Platinum
Mandate penalty*	-0.081	2.365	1.962	1.835	1.762
Bronze premiums	0.030	-2.590	0.163	0.154	0.151
Silver premiums	0.019	0.126	-2.342	0.119	0.124
Gold premiums	0.013	0.085	0.087	-2.305	0.084
Platinum premiums	0.012	0.074	0.076	0.073	-2.269

Note: the table shows the median of cross-tier semi-elasticity across demographic groups, based on the mixed logit estimates. For each group, a cell contains the median percentage change in the total enrollment of plans in the column tier if the all plans in the row tier increase annual premium by $100 per-year.

Table 11: Parameters of the cost model

	Log-annual claims						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Product characteristics:							
Bronze	(Omitted)	(Omitted)	(Omitted)	(Omitted)	(Omitted)	(Omitted)	(Omitted)
Silver	0.478***	0.399***	0.344**	0.432**	0.370***	0.306**	0.411**
	(0.132)	(0.144)	(0.140)	(0.170)	(0.141)	(0.140)	(0.170)
Gold	0.648***	0.619***	0.597***	0.608***	0.604***	0.582***	0.598***
	(0.0778)	(0.0791)	(0.0765)	(0.0850)	(0.0798)	(0.0751)	(0.0825)
Platinum	1.219***	1.171***	1.158***	1.189***	1.130***	1.113***	1.153***
	(0.0977)	(0.100)	(0.0966)	(0.107)	(0.104)	(0.0968)	(0.106)
Buyer characteristics:							
Age	0.0188**	0.0187**	0.0206**	0.0221**	0.0187**	0.0198**	0.0217**
	(0.00846)	(0.00838)	(0.00830)	(0.00911)	(0.00803)	(0.00817)	(0.00889)
FPL	0.00277	0.00243	0.00207	0.00321	0.00275	0.00227	0.00351
	(0.00251)	(0.00253)	(0.00236)	(0.00297)	(0.00245)	(0.00234)	(0.00298)
Household size	0.0870	0.0572	0.144	0.199	0.0569	0.138	0.185
	(0.112)	(0.111)	(0.131)	(0.153)	(0.115)	(0.135)	(0.157)
Low WTP (<1700)		(Omitted)	(Omitted)	(Omitted)			
High WTP (>1700)		0.113	0.172**	0.123			
		(0.0778)	(0.0766)	(0.0890)			
WTP$\in [0, 1100]$					(Omitted)	(Omitted)	(Omitted)
WTP $\in (1100, 1400]$					0.0778	0.0964	0.0938
					(0.105)	(0.0956)	(0.116)
WTP$\in (1400, 1700]$					0.133	0.124	0.0992
					(0.109)	(0.0997)	(0.118)
WTP $\in (1700, 2000]$					0.260**	0.291***	0.243*
					(0.115)	(0.109)	(0.132)
WTP >2000					0.147	0.263**	0.173
					(0.126)	(0.121)	(0.137)
Region FE	Y	Y	Y		Y	Y	
Insurer FE	N	N	Y		N	Y	
Insurer-Region FE	N	N	N	Y	N	N	Y
Observations	338	338	338	338	338	338	338
R-squared	0.463	0.467	0.572	0.597	0.472	0.575	0.600

Note: Parameters of the cost model estimated from equation (22). Each observation is an insurer-region-tier triplet, where I exclude Catastrophic coverage since it is not available for subsidized enrollees. After this exclusion, claims data used for estimation cover over 90% of enrollment, with two missing carriers (Contra Costa and Valley). Buyer characteristics are computed as average across enrolees of the plan, where WTP is the posterior of the ratio β/α, conditional on observed choice, based on mixed logit estimates. Robust standard errors in parentheses; *** p<0.01, ** p<0.05, * p<0.1

Table 12: From price-linked subsidies to vouchers of equivalent amount

	ACA price-linked subsidy scheme	Voucher equal to ACA subsidy	Relative change
Enrollment (1,000s):	1,126	1,146	+2%
Share choosing high-deductible (Bronze):	34.6%	27.1%	-21.6%
Average markup ($/year):	952	844	-11%
Average medical loss ratio:	77.9%	79.6%	+2.2%
Second-cheapest Silver (benchmark) plan:			
Base price ($/year)	2,988	2,808	-6%
Markup ($/year)	1,650	998	-40%

Note: Comparison of equilibrium outcomes between ACA price-linked subsidy and fixed vouchers of equivalent amount. The first column contains aggregate and quantity-weighted outcomes from simulating equilibrium in each region starting from the base prices in the data, and then zeroing FOC including the distortion from price-linked subsidy (see Appendix B2 for details), under a medical loss ratio of at least 0.7 and a positive expected profit constraint. For the second column, I consider fixed voucher equal to the subsidies resulting in the equilibrium in the first column. I then calculate equilibrium imposing FOC, under a medical loss ratio of at least 0.7 and a positive expected profit constraint. The third column shows percentage diferences in outcomes between the two equilibria.

Table 13: Equilibrium outcomes under alternative subsidy designs

	ACA baseline (price-linked to 2^{nd}-cheapest Silver)	Difference from ACA baseline under design (1) vouchers equal to ACA subsidies	(2) age-adjusted vouchers with all buyers better off	(3) argmax consumer surplus s.t. lower spending	(4) argmax total enrollment s.t. lower spending
Equilibrium outcomes:					
Tot enrollment (1,000s)	1,126	+20 (+2%)	+149 (+13%)	+106 (+9%)	+201 (+18%)
Tot consumer surplus ($ million)	3,348	+114 (+3%)	+492 (+15%)	+356 (+11%)	-221 (-7%)
Tot subsidy spending ($ million)	3,447	-51 (-1%)	+64 (+2%)	-58 (-2%)	-137 (-4%)
Average cost ($/year)	3,355	-63 (-2%)	-315 (-9%)	-243 (-7%)	-260 (-8%)
Monthly subsidy:					
FPL 130-200, no children, age<45	200	0	50	50	10
FPL 130-200, no children, age>45	442	0	-25	10	-50
FPL 200-250, no children, age<45	154	0	10	30	30
FPL 200-250, no children, age>45	395	0	0	-50	-25
FPL 250-400, no children, age<45	76	0	50	-25	-25
FPP 250-400, no children, age>45	318	0	-25	-50	20

Note: Equilibrium outcomes under ACA price-linked subsidy (left column), and changes in outcomes under fixed vouchers (column (1)), age-income-adjusted vouchers that maximize consumer surplus subject to per-person spending being lower than under the ACA (column (2)), age-income-adjusted vouchers that maximize consumer surplus subject to total spending being lower than under the ACA (column (3)), age-income-adjusted vouchers that maximize total enrollment subject to total spending being lower than under the ACA (column (4)). Equilibria are calculated starting from the base prices in the data, and then imposing FOC under a medical loss ratio of at least 0.7 and a positive expected profit constraint.

SUPPLEMENTARY APPENDIX

FOR ONLINE PUBLICATION ONLY

Additional Figures and Tables

Figure A: Age adjustment factors in ACA marketplaces

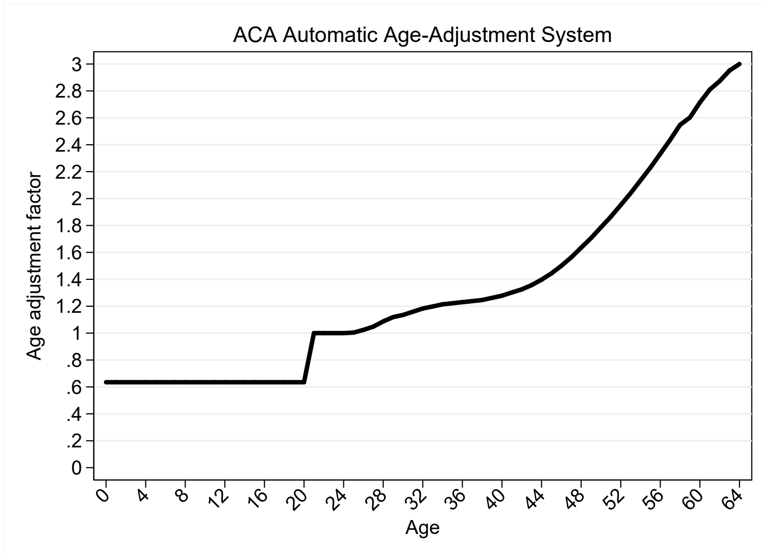

Note: For every age τ, the line in the figure shows the corresponding factor A^{τ}, which is used to compute the price of a high-income, unsubsidized buyer (equation (1)). If the base price of the product is b_{jr}, $P^{\tau}_{jr} = A^{\tau} \cdot b_{jr}$. This is also equal to the total amount that the insurer receives when a subsidized buyer purchases the plan.

Figure B: Cost sharing subsidies and share of households choosing a Silver plan

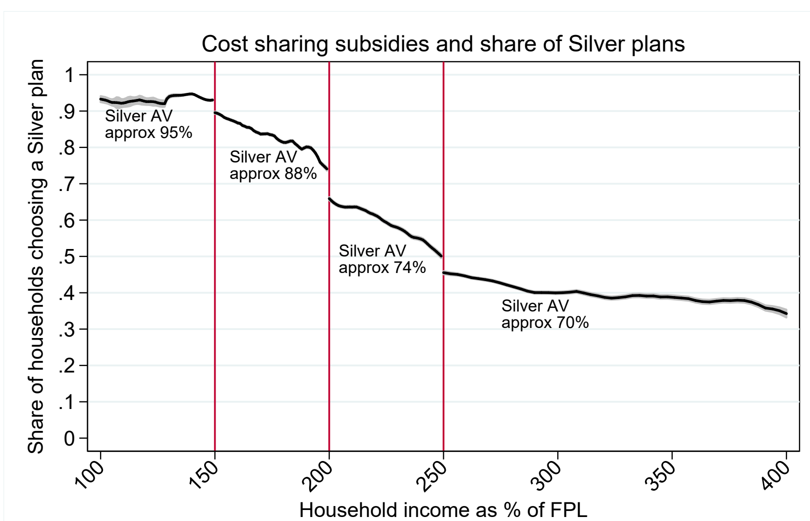

Note: The lines show the local polynomial smoothing of the share of household choosing a Silver plan as a function of the household income as % of the FPL.
All subsidized buyers in Covered California 2014-15, N=1,203,721.

Figure C: Rating regions in Covered California

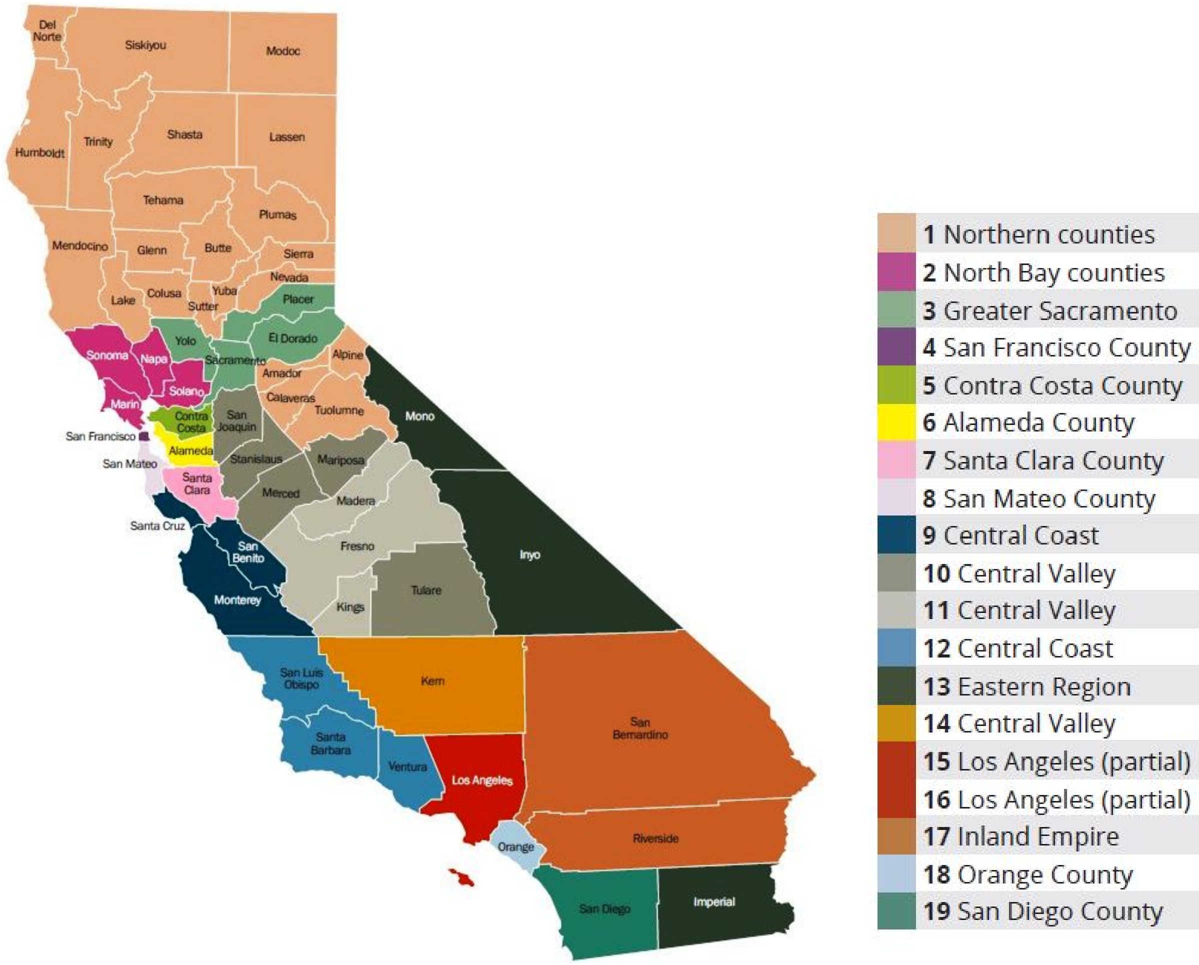

Source: www.CoveredCA.com

The map shows the 19 rating regions in Covered California. In each region, every spring insurers can announce their participation in the following year's open enrollment. The marketplace needs to authorize entry, and requires the insurer to offer five coverage levels with pre-determined financial characteristics (Table 2). In the summer insurers set one base price for every level of coverage in every region where they entered, prices and subsidies are then calculated from base prices applying ACA regulations (Section 2).

Figure D: Percent change in enrollment if all premiums increase by $100/year

(a) Subsidized households without children

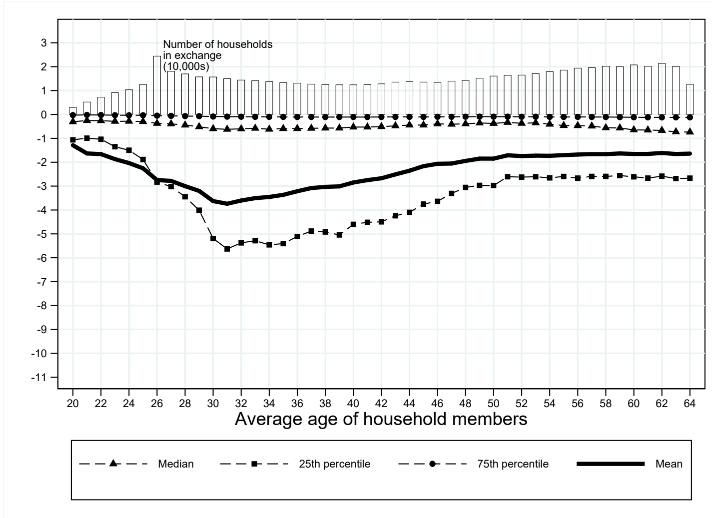

(b) Subsidized households with children

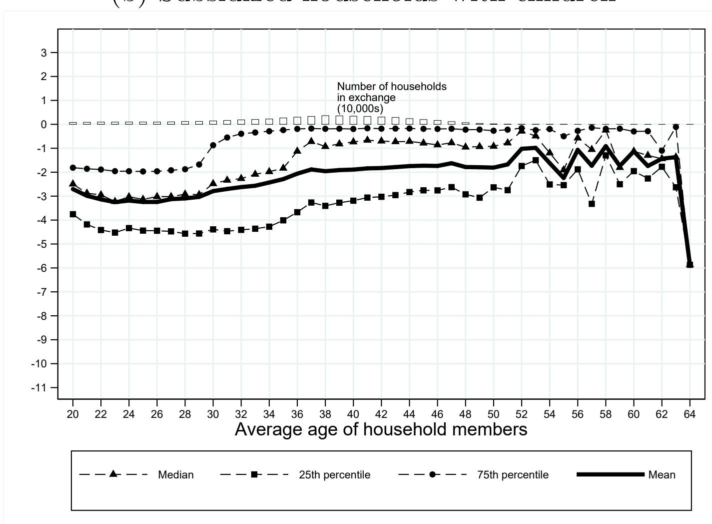

(c) Unsubsidized households

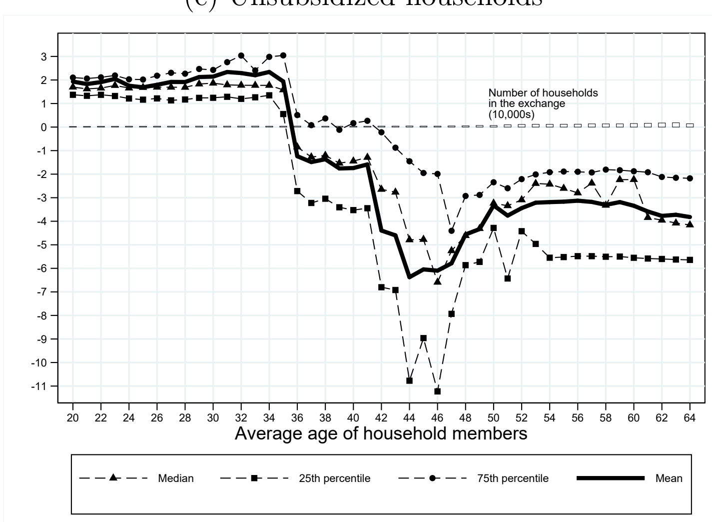

Note: extensive margin semi-elasticity by age and type of household based on mixed logit estimates and number of enrollees in the exchange. Each panel shows (for subsidized households without children and subsidized households with children, respectively) the mean, median, and 25th and 75th percentiles of percentage drop in coverage if all premiums increase by $100 per-year. The bars indicate the number of enrollees in the exchange in 10,000s.

Figure E: Demand curves from mixed logit estimates

(a) Subsidized, no children, under 50

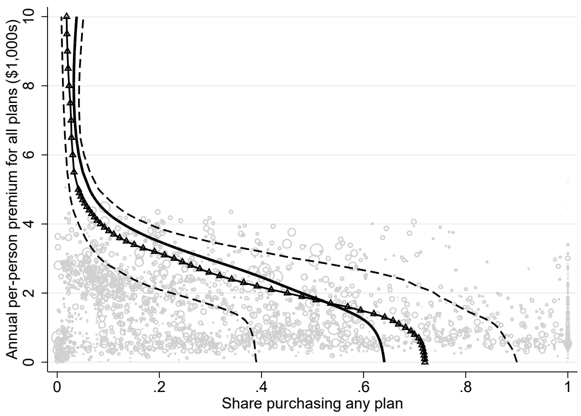

(b) Subsidized, no children, over 50

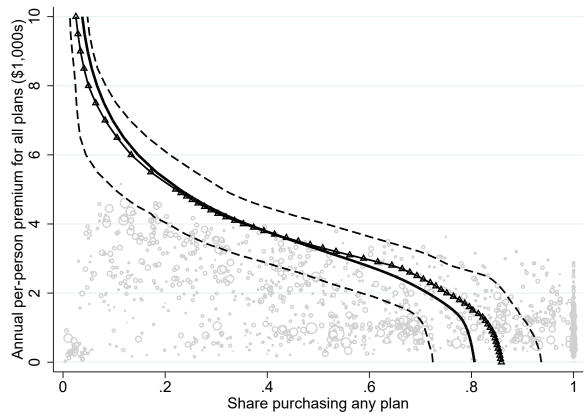

(c) Subsidized, with children

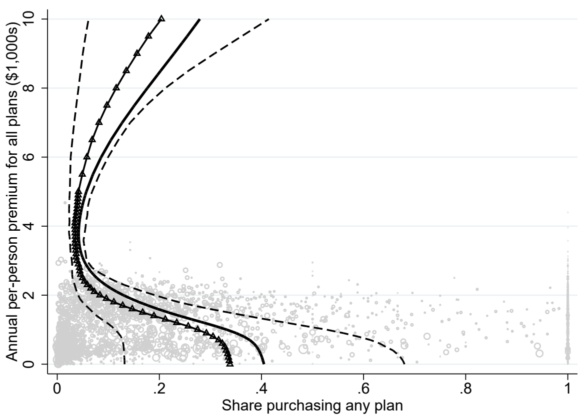

(d) Unsubsidized

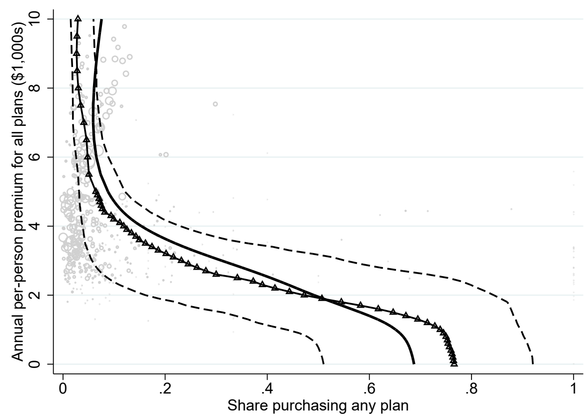

Note: demand curves from mixed logit estimates are derived computing, for each household, the probability of purchasing any plan in the exchange (x-axis) if all premiums are set equal to the amount on the y-axis. Each graph shows the 25th and 75th percentile (dashed lines), the median (connected triangles), and mean of the probability of purchase at a given premium. The shaded scatter plot corresponds to the combinations of average premium paid and coverage probability for a specific age-income-region combination as observed in the data; circles are proportional to the number of potential buyers in the age-income-region cell.

Figure F: Joint density of preferences and cost

(a) Subsidized households without children

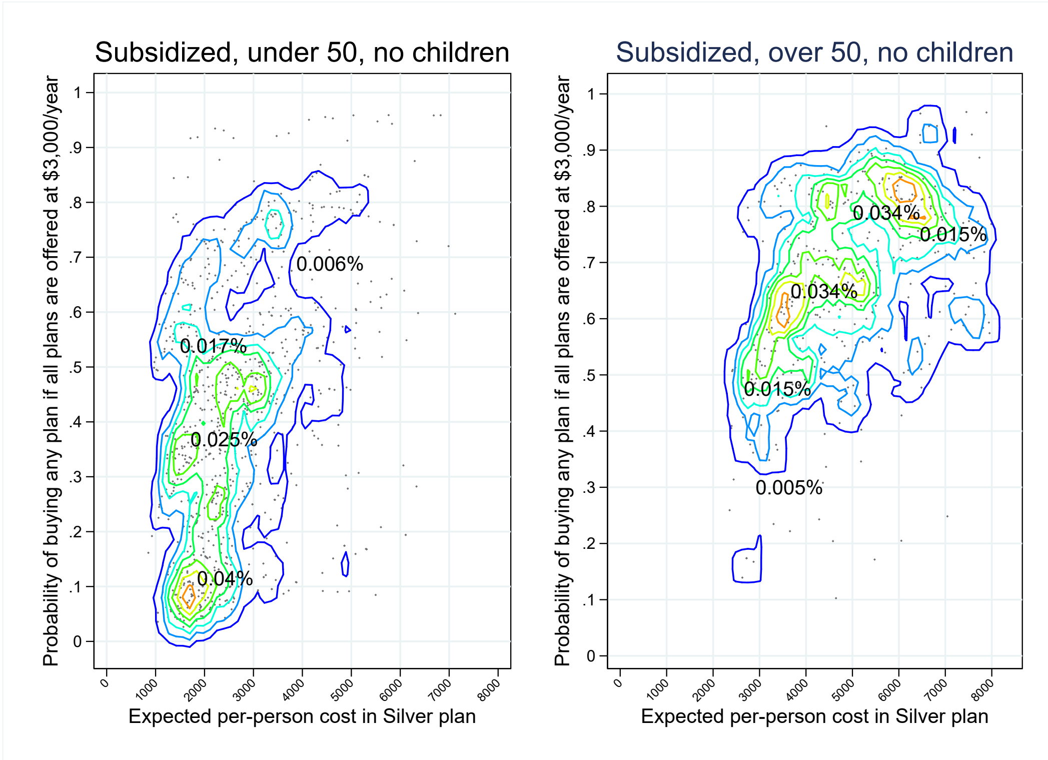

(b) Subsidized households with children and unsubsidized

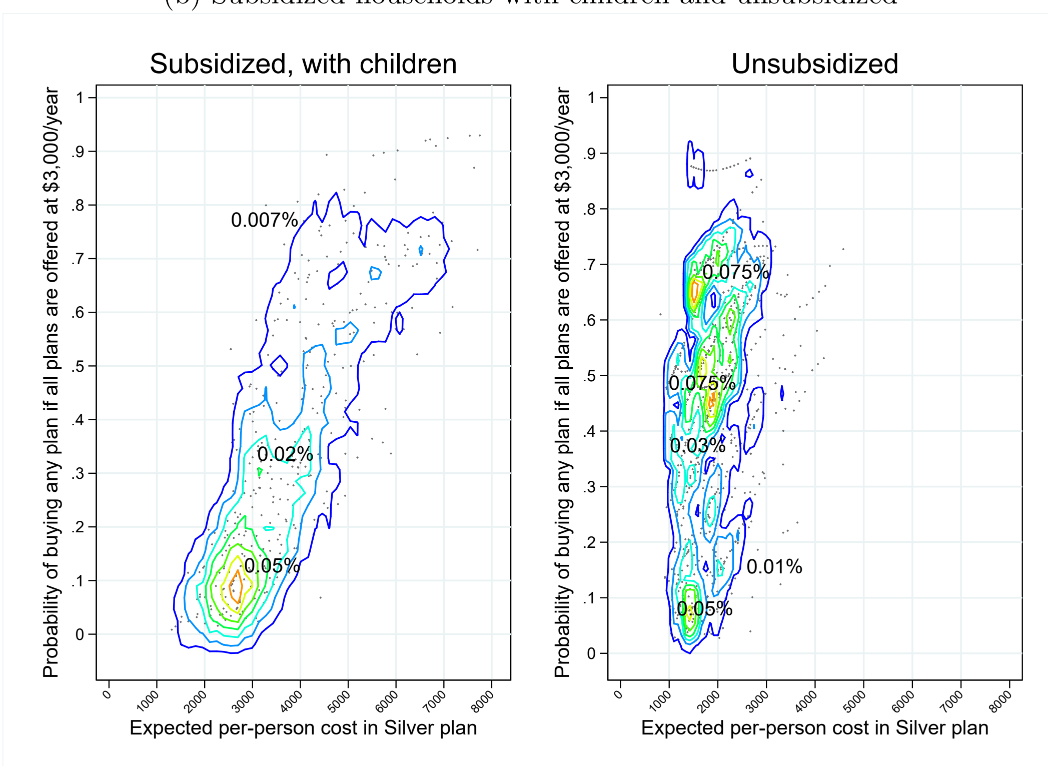

Note: based on mixed logit demand estimates and cost estimates, the figure shows countour lines of the joint kernel density of the probability of buying any plan if all premiums are $3,000/year (y-axis) and the expected cost upon enrolling in a Silver plan offered by Anthem (x-axis). In each panel the joint density overlays a scatter plot of a 1% sample of individuals from the corresponding demographic group.

Figure G: Nash-in-prices assumption and marginal cost estimation

Accounting for price-linked subsidy:

(a) MR vs. MC under Nash-in-prices

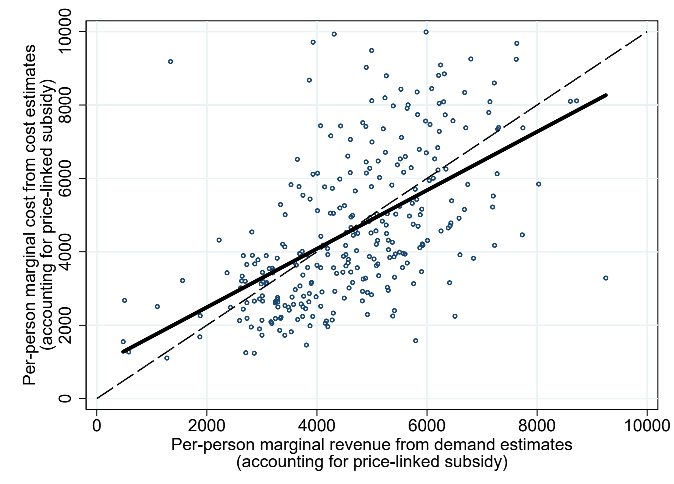

(b) MC from claims vs. MC from Nash-in-prices

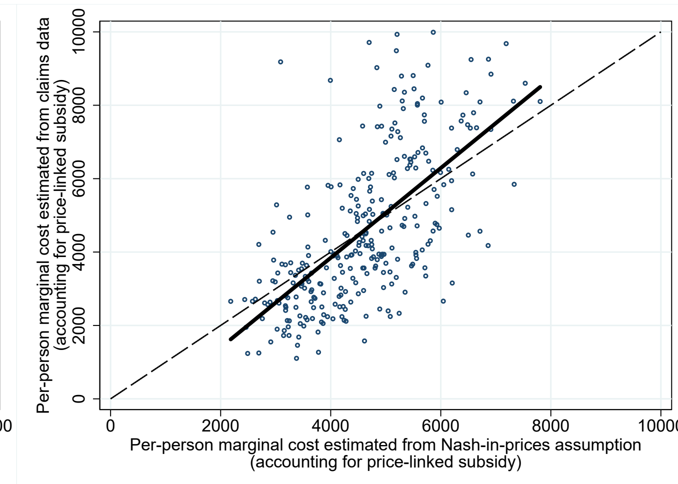

Ignoring price-linked subsidy:

(c) MR vs. MC under Nash-in-prices

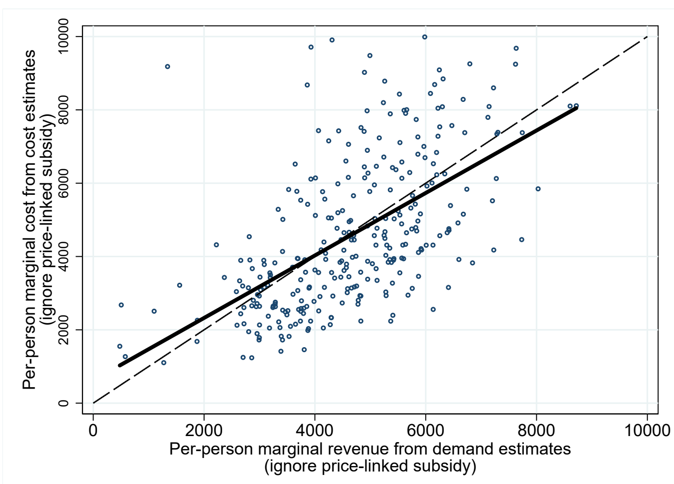

(d) MC from claims vs. MC from Nash-in-prices

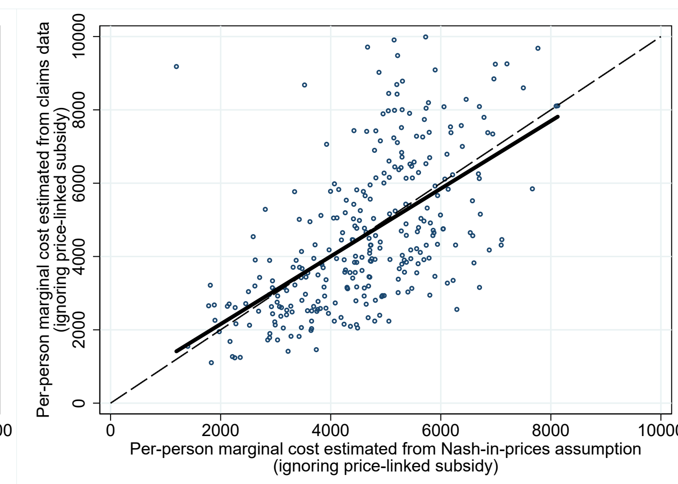

Panel (a): relationship between per-person marginal revenue and marginal cost based on mixed logit and cost estimates from average claims. Each quantity is calculated imposing age-rating adjustments and price-linked subsidies. Details in Appendix B1.

Panel (b): relationship between per-person marginal cost estimated using average claims and per-person marginal cost estimated using Nash-in-prices assumptions. Each quantity is calculated imposing age-rating adjustments and price-linked subsidies. Details in Appendix B1.

Panel (c): relationship between per-person marginal revenue and marginal cost based on mixed logit and cost estimates from average claims. Each quantity is calculated imposing age-rating adjustments while taking subsidies as fixed vouchers. Details in Appendix B1.

Panel (d): relationship between per-person marginal cost estimated using average claims and per-person marginal cost estimated using Nash-in-prices assumptions. Each quantity is calculated imposing age-rating adjustments while taking subsidies as fixed vouchers. Details in Appendix B1.

Figure H: Equilibrium outcomes varying monthy subsidies by age

(a) Under 50, subsidized, no children (ACA = 480,042)

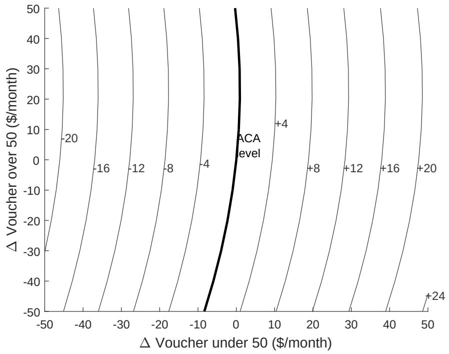

(b) Over 50, subsidized, no children (ACA = 358,104)

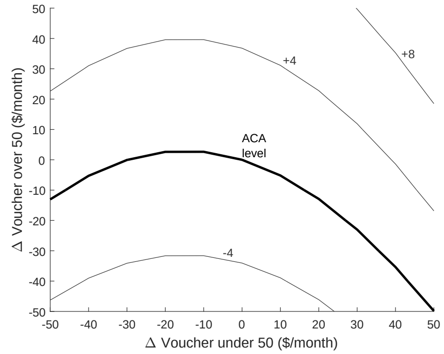

(c) Subsidized with children (ACA = 234,947)

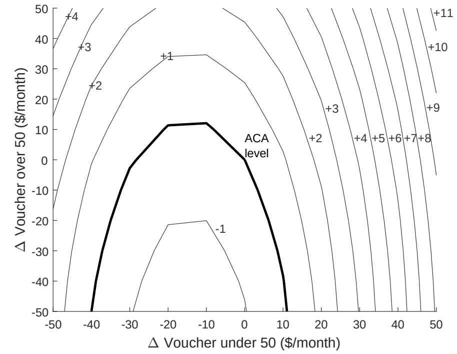

(d) Unsubsidized (ACA = 71,042)

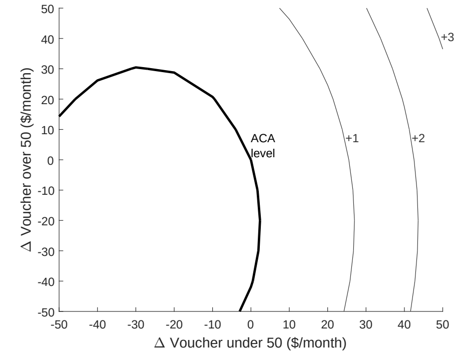

(e) Total enrollment (ACA = 1,144,135)

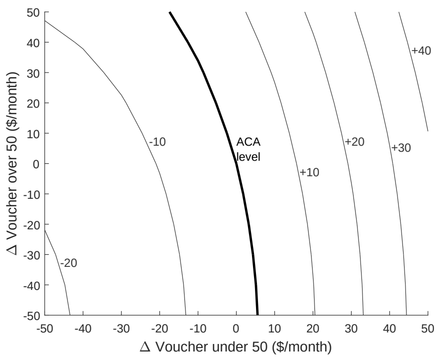

(f) Total subsidy spending (ACA = $3,396 million)

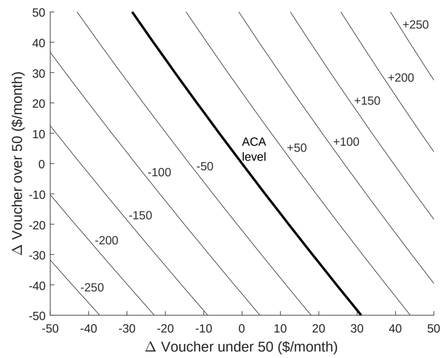

(g) Total profits (ACA = $1,001 million)

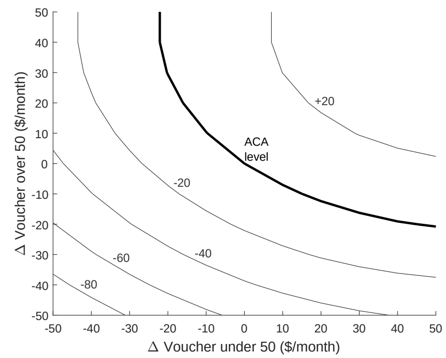

(h) Consumer surplus + profits (ACA = $4,464 million)

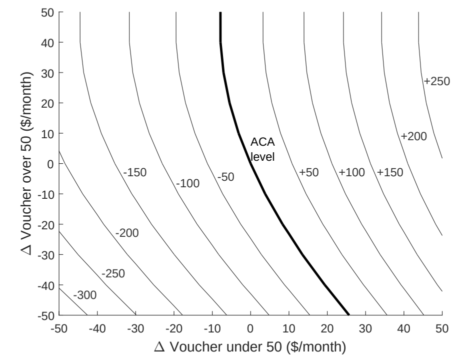

(i) Average cost (ACA = $3,522)

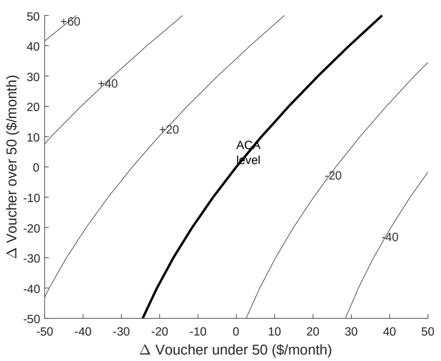

(j) Average revenue (ACA = $4,571)

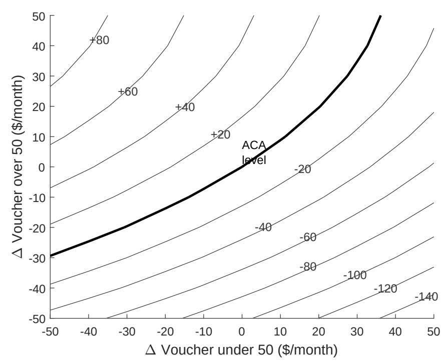

(k) Average markup (ACA = $1,049)

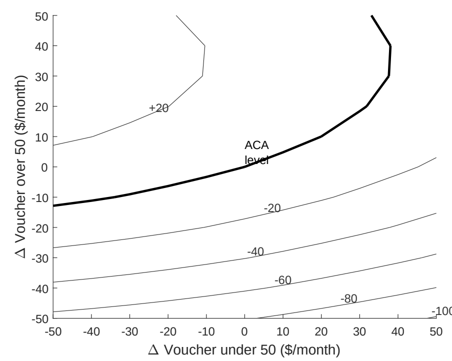

(l) Per-person subsidy (ACA = $3,675)

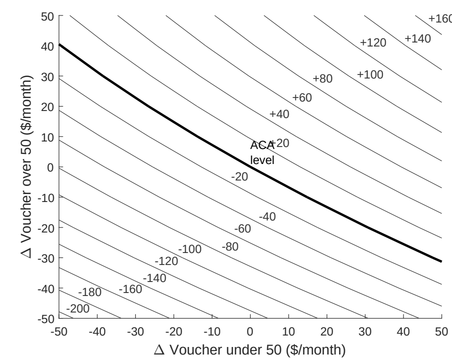

Note: Level curves of equilibrium outcomes as functions of the changes (from ACA level) in voucher for the under 50 (x-axis) and over 50 (y-axis). Level curves correspond to changes in 1,000 individuals for enrollment, change in $ for per-person subsidy, average cost, average revenue, and markup, and $M for total spending, profits, and consumer surplus.

Table A: Logit and nested logit estimates using different enrollment years and waves of the ACS to define potential buyers

	$ln(s_{jr}^{\tau,y,h}) - ln(s_{0r}^{\tau,y,h})$									
	Logit					Nested logit				
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Enrollment data:	2014	2015	2014 & 2015	2015	2014 & 2015	2014	2015	2014 & 2015	2015	2014 & 2015
Potential buyers:	ACS 2013	ACS 2013	ACS 2013	ACS 2014	ACS 2013 &2014	ACS 2013	ACS 2013	ACS 2013	ACS 2014	ACS 2013 &2014
Premium($1000) × Age ∈[10,20)	-0.285***	-0.265***	-0.275***	-0.281***	-0.285***	-0.283***	-0.253***	-0.268***	-0.261***	-0.273***
	(0.0301)	(0.0273)	(0.0277)	(0.0298)	(0.0268)	(0.0313)	(0.0278)	(0.0281)	(0.0285)	(0.0263)
Premium($1000) × Age ∈[20,30)	-0.356***	-0.328***	-0.341***	-0.337***	-0.346***	-0.319***	-0.283***	-0.300***	-0.288***	-0.303***
	(0.0232)	(0.0232)	(0.0220)	(0.0276)	(0.0236)	(0.0233)	(0.0240)	(0.0223)	(0.0276)	(0.0234)
Premium($1000) × Age ∈[30,40)	-0.339***	-0.304***	-0.321***	-0.300***	-0.319***	-0.297***	-0.253***	-0.274***	-0.246***	-0.270***
	(0.0226)	(0.0221)	(0.0213)	(0.0217)	(0.0195)	(0.0234)	(0.0226)	(0.0216)	(0.0221)	(0.0197)
Premium($1000) × Age ∈[40,50)	-0.289***	-0.266***	-0.277***	-0.263***	-0.276***	-0.254***	-0.221***	-0.237***	-0.214***	-0.233***
	(0.0187)	(0.0160)	(0.0162)	(0.0188)	(0.0156)	(0.0188)	(0.0153)	(0.0157)	(0.0183)	(0.0151)
Premium($1000) × Age ∈[50,60)	-0.183***	-0.160***	-0.171***	-0.185***	-0.184***	-0.145***	-0.118***	-0.131***	-0.143***	-0.144***
	(0.0170)	(0.0171)	(0.0164)	(0.0164)	(0.0144)	(0.0165)	(0.0157)	(0.0152)	(0.0156)	(0.0133)
Premium($1000) × Age ∈[60,64]	-0.240***	-0.225***	-0.232***	-0.223***	-0.231***	-0.210***	-0.195***	-0.203***	-0.190***	-0.200***
	(0.0163)	(0.0170)	(0.0155)	(0.0123)	(0.0124)	(0.0152)	(0.0173)	(0.0150)	(0.0113)	(0.0114)
Premium($1000) × Age ∈ hh size=2	0.0451***	0.0306**	0.0375***	0.0303**	0.0378***	0.0177	0.00570	0.0115	0.00637	0.0122
	(0.0137)	(0.0124)	(0.0127)	(0.0132)	(0.0128)	(0.0136)	(0.0122)	(0.0124)	(0.0133)	(0.0126)
Premium($1000) × Age ∈ hh size=3	0.165***	0.133***	0.148***	0.131***	0.148***	0.117***	0.0854***	0.101***	0.0845***	0.101***
	(0.0197)	(0.0185)	(0.0186)	(0.0183)	(0.0178)	(0.0194)	(0.0180)	(0.0181)	(0.0181)	(0.0172)
Premium($1000) × Age ∈ hh size=4	0.247***	0.217***	0.231***	0.206***	0.226***	0.194***	0.164***	0.178***	0.153***	0.173***
	(0.0238)	(0.0237)	(0.0233)	(0.0203)	(0.0205)	(0.0239)	(0.0235)	(0.0231)	(0.0203)	(0.0204)
Premium($1000) × Age ∈ hh size=5	0.276***	0.246***	0.260***	0.256***	0.267***	0.207***	0.177***	0.192***	0.188***	0.198***
	(0.0241)	(0.0233)	(0.0231)	(0.0182)	(0.0195)	(0.0245)	(0.0240)	(0.0234)	(0.0181)	(0.0195)
Premium($1000) × Non-subsidized	0.0840***	0.104***	0.0948***	0.104***	0.0941***	0.102***	0.122***	0.113***	0.118***	0.110***
	(0.0151)	(0.0136)	(0.0134)	(0.00948)	(0.0103)	(0.0154)	(0.0160)	(0.0147)	(0.0113)	(0.0112)
Non-subsidized	-1.917***	-2.158***	-2.044***	-2.233***	-2.074***	-2.252***	-2.554***	-2.413***	-2.584***	-2.421***
	(0.163)	(0.186)	(0.165)	(0.153)	(0.143)	(0.162)	(0.207)	(0.174)	(0.172)	(0.149)
Deductible($1000)	0.137***	0.176***	0.157***	0.170***	0.154***	-0.0394**	-0.00759	-0.0231	-0.000558	-0.0198
	(0.0163)	(0.0146)	(0.0135)	(0.0157)	(0.0135)	(0.0156)	(0.0155)	(0.0142)	(0.0148)	(0.0130)
Max out-of-pocket($1000)	-0.138***	-0.124***	-0.131***	-0.118***	-0.128***	-0.00791	-0.0142	-0.0117	-0.0208	-0.0149
	(0.0154)	(0.0163)	(0.0138)	(0.0170)	(0.0133)	(0.0121)	(0.0119)	(0.00995)	(0.0126)	(0.00980)
Primary care visit($100)	-0.414**	-0.561***	-0.492***	-0.614***	-0.520***	0.0696	0.0506	0.0600	-0.0622	0.000864
	(0.165)	(0.133)	(0.131)	(0.121)	(0.129)	(0.150)	(0.137)	(0.130)	(0.129)	(0.126)
Nesting Parameter						0.503***	0.475***	0.488***	0.466***	0.484***
						(0.0159)	(0.0148)	(0.0143)	(0.0209)	(0.0174)
Observations	25,019	24,953	49,972	23,347	48,366	25,019	24,953	49,972	23,347	48,366
R-squared	0.449	0.432	0.440	0.440	0.445	0.533	0.511	0.522	0.514	0.524
Year x Insurer x Region x Over 50	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y

Estimates from specifications (4) and (9) of Table 6 in the paper using different enrollment years and waves of the ACS to define potential buyers. Columns (3) and (8) pool both years of enrollment using the 2013 ACS to define potential buyers. Columns (5) and (10) are panel specifications where the 2013 ACS defines 2014 potential buyers and the 2014 ACS defines 2015 potetntial buyers. Each observation is a unique (insurer-region-tier-demographic group). The 200 demographic groups are full interactions of age group (bins of 10 years), gender, annual income group (bins of $10,000), and household size. Standard errors in parentheses clustered at the region level (19 clusters).
*** p<0.01, ** p<0.05, * p<0.1