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

Pietro Tebaldi+

November 4, 2016

Abstract. To design premium subsidies in a health insurance market it is necessary to estimate consumer demand and study how different subsidy schemes affect insurers' incentives. Combining data from the Californian ACA marketplace with a model of insurance demand and insurers' competition, I identify and estimate demand and cost primitives, and assess equilibrium outcomes under alternative subsidy designs. I find that vouchers are less distortionary than subsidies calculated from market premiums, and — given age-heterogeneity in demand and cost — tailoring subsidies to age leads to an equilibrium 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. A special thanks goes to Liran Einav, Tim Bresnahan, Matt Jackson, and Mark Duggan for their support and research advise, along with four anonymous referees for their constructive feedback. I also thank Sofronis Clarides, Mark Shepard, and Jonathan Williams for the helpful discussions, and I want to acknowledge the helpful comments I received from Matt Backus, Lanier Benkard, Stephane Bonhomme, Marika Cabral, Jeff Clemens, Francesco Decarolis, Michael Dickstein, Michael Dinerstein, Dave Donaldson, Matthew Gentzkow, Daniel Gottlieb, Josh Gottlieb, Michael Greenstone, Daniel Grodzicki, Han Hong, James Heckman, Caroline Hoxby, Jon Levin, Greg Lewis, Davide Malacrino, Joe Orsini, Ali Hortacsu, Ariel Pakes, Luigi Pistaferri, Maria Polyakova, Ilya Segal, Paulo Somaini, Alonso Villacorta, Glen Weyl, Michael Whinston, as well as participants at the Stanford IO, Theory, Public, Applied, and Econometrics workshops, the Bocconi University Applied Micro seminar, the Microsoft Research Economics Colloquium, the 2015 ESWC, the 2015 EARIE, 2016 NBER Insurance Workshop, 2016 IIOC, 2016 EIEF Micro Conference, and seminars at Rice University, University of Washington, University of Chicago, Yale University, Northwestern University, MIT Sloan, University of Pennsylvania, New York University, Georgetown University, University of Wisconsin - Madison, and Penn State University. All remaining mistakes are mine.

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

1 Introduction

Despite its growing importance in the provision of government-sponsored health insurance (**?**), the response of market outcomes to the decision of how to set premium subsidies in a health insurance market is still largely unexplored. A recent, large-scale example of such market design decision is found in the low-income subsidy introduced by the 2010 US health care reform (Patient Protection and Affordable Care Act; ACA). Since 2014, under this program the federal government transferred approximately $40 billion per-year to private insurers, providing discounts on health insurance premiums to more than 10 million citizens.[1] Knowledge of the relationship between subsidy design and policy-relevant outcomes such as coverage levels and public spending is critical to evaluate the success of this reform, 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 low-income buyers, insurers' price competition, and selection — correlation between a buyer's willingness to pay and ex-ante expected health cost. Characteristics of demand determine the extent to which subsidies increase buyers' participation. Pricing incentives and market power of imperfectly competitive insurers react to these changes in demand, but also to corresponding changes in average cost driven by differences in the composition of enrollment pools.

To account for these effects, and compare different subsidy designs, I combine data from the first year of the Californian ACA marketplace — in which 90% of the 1.4 million buyers receive federal subsidies — with a model of insurers' competition customized to include subsidies and other ACA regulations. I discuss identification of demand and supply primitives exploiting variation in the regional composition of buyers, and assuming equilibrium pricing. I then use these estimates and the model to carry on quantitative comparisons of different designs of subsidy programs in terms of equilibrium prices, enrollment, markups, and public spending. Within this framework my results suggest that the ACA subsidy scheme leaves room for improvements that are quantitatively significant and consistent with theoretical predictions. The alternatives I consider can potentially reduce insurers' market power, and increase incentives for the participation of young buyers that directly affects average cost — and thus prices and public spending — in the newly created state-level marketplaces.

The paper makes three main contributions. First, from a theoretical perspective I show how, in a market with a group of buyers who are cheaper to cover and more price sensitive than others — adverse selection — , tailoring the generosity of subsidies to favor this group can lead to an equilibrium where all groups are better off and public spending is lower. Intuitively, shifting subsidy generosity from the high cost, high demand group to the low cost, low demand group changes the composition of enrollment pools, hence lowers average cost and increases aggregate elasticity. This

[1] See for example the 2015 report from the Congressional Budget Office, **?**.

lowers equilibrium prices, and it can increase quantity purchased for all groups while also reducing public spending. Importantly, since also the group receiving a lower subsidy can be made better off, the benefits of heterogeneous subsidization can be achieved while avoiding redistributive concerns.

Second, a key contribution in my study is to provide estimates of demand and cost primitives using detailed data from an ACA marketplace. Other work that evaluates the role of different ACA regulations using post-reform data (**????**) exploits cross-sectional variation in outcomes across states, or state-level variation over time, without adopting any explicit demand-supply model. Here, instead, I am able to use more granular variation in prices and enrollment within a single state-level marketplace, and develop an empirically tractable model that allows me to estimate demand in the subsidized population (largely undocumented in previous studies), and also account for imperfect competition and equilibrium implications of different parts of the ACA regulatory framework.[2]

The third contribution comes from a more methodological perspective. Many papers in the empirical literature on selection markets identify heterogeneity in risk and preferences relying on the availability of cost data, or using external surveys (see e.g. **??????**). Here, instead, with an approach that is similar in spirit to the one of **?** — although our estimations and implementations differ — , I present conditions under which the well-known inversion of first-order optimal pricing conditions (see e.g. **???**) can be extended to models of regulated insurance supply to identify costs varying with buyers' preferences when cost heterogeneity is not observed.

The structure of my analysis is as follows. I start in Section **??** with a stylized model that highlights the theoretical implications of the two design decisions I consider in the paper: The first is whether or not price discounts should be fixed ("vouchers"), or calculated as a function of market prices ("price-linked", see also **?**); the second is whether discounted prices should be equal across all buyers with the same income, or adjusted to buyers' age. Both decisions are closely related to the ACA design, which features price-linked subsidies by which, given income, discounted prices do not vary by age. The model highlights that vouchers are less distortionary than price-linked subsidy, and that targeting subsidies in favor of younger buyers can make all buyers better off and reduce per-person public spending. The theoretical discussion also highlights how the quantitative impact of different subsidy designs depends on the primitives of the market, and particularly on the intensity of price competition and the heterogeneity in price elasticity of demand and cost across different groups.

In Section **??** I present regulatory details and data on prices and enrollment from the first year of the Californian marketplace. This is the largest among ACA marketplaces, where subsidy eligible buyers can choose between different coverage options offered by a set of participating insurers.

[2]Here I focus on a limited number of subsidy design alternatives, but my model could also be used to evaluate, for example, changes to age-rating rules, the implementation of the mandate penalty, or changes to the set of contracts that insurers' must offer and/or subsidized buyers are allowed to choose from.

This provides me with an attractive setup to estimate demand for coverage among the low-income uninsured, since the pre-reform period did not offer a reliable, large-scale environment in which buyers from this segment of the population were making active choices.[3] I combine this novel dataset with a discrete-choice model of insurance demand and ACA-regulated price competition in Section **??**; this combination allows me to identify and estimate primitives under the status-quo, and simulate equilibrium after changing the subsidy scheme while keeping other regulations fixed.

In Section **??** I discuss how, in the context of ACA marketplaces, existing regulations point directly toward a suitable instrumental variable strategy that can be followed to identify demand. The ACA only allows insurers to set one base price for every insurance plan in a given market (group of counties), and pre-determined pricing schedules are used to translate this base price to the premium faced by buyers of different age and income. Because this regulation links profits across heterogeneous buyers to the same univariate decision, when setting base prices insurers must consider the composition of buyers.[4] Indeed, age-income composition is a strong predictor of cross-market variation in prices in my data. Consequently, assuming that individual preferences are independent from market composition — e.g. willingness to pay among 21-year-olds is independent from the number of uninsured 64-year-olds — I use age-income composition to instrument for price and obtain consistent estimates of demand for various age and income groups.

For identification of insurers' costs I rely instead on the combination of equilibrium assumptions with existing pricing regulations. To gain intuition, suppose that insurers cannot vary price by age. Once I obtain demand estimates, I construct marginal revenues for every observed insurance contract. Equilibrium pricing implies that marginal revenues are equal to marginal cost, which results from a combination of unobserved cost across different age groups. Hence, assuming that age is the only determinant of insurers' expected cost, to estimate age-driven cost heterogeneity I compare marginal revenues/cost between contracts whose marginal buyers are predominantly "old" and contracts whose marginal buyers are instead predominantly "young". More precisely, variation in age-composition of marginal buyers — which can be derived after estimating age-specific demand curves — can be used to identify the impact of buyers' age on insurers' expected costs. The general identification result is stronger, providing conditions for identification when buyers' cost may also differ with preferences along unobservable dimensions; this follows from results in **?** combined with a constructive proof partially adapted from **??**.

In Section **??** I present the resulting estimates. For demand I find that price sensitivity decreases

[3] Before the ACA, the most similar environment in which subsidies were provided to buyers with income between 100-300% of the federal poverty level was the Massachusetts' Connector. However, under the Commonwealth Care Program these low-income buyers were not making active choices, but instead automatically assigned to a plan based on their income (see e.g. **?**).

[4] As we discuss in **?**, the automatic age-rating adjustments mandated by the ACA apply the same age gradient — or standard age-rating curve — across different markets. As long as this constant factors cannot perfectly capture age-differences in prices that would result if age-based price discrimination was allowed, regional variation in the age-composition of potential buyers becomes a relevant determinant of constrained prices.

in age and income, in line with common wisdom and existing literature on insurance demand; (e.g. ??????). For example, own-price elasticity varies from 3 for low-income young adults to 1 for high-income over 45. More importantly, a $100 increase in all prices implies a drop in overall coverage of 12% among low-income under 45, 4% among low-income over 45, and less than 3% among high-income. Estimated cost heterogeneity across age groups is also sizable, and consistent with existing evidence from the Medical Expenditure Panel Survey. On average, the expected cost of covering a buyer older than 45 on a zero deductible plan is approximately $8,100 per-year, dropping to $2,460 for younger buyers. Estimated costs also vary across insurers, by up to $600 per-person, and across different regions of the state.

These demand and cost estimates are the main inputs for my counterfactual analysis in Section ??, where I start by comparing fixed vouchers to the price-linked discounts adopted under the ACA. The current scheme calculates discounts from market prices to ensure that buyers do not pay more than a predetermined amount. As a consequence, if all prices increase by $1, low-income buyers do not face any price change, while government expenditure increases by $1 per buyer. This creates incentives for insurers to set higher prices than if discounts were instead predetermined vouchers. Indeed, my simulations imply that equilibrium insurers' markups are 15% lower if the current subsidy scheme is replaced by a voucher; under this alternative discounts are of equal amount to those ACA ones but are not adjusted to insurers' decisions. This reduction in markups corresponds to a lower per-buyer spending by the government of approximately 5%, while coverage among low-income buyers is 10% higher than under the status quo. The second design decision I consider is whether the subsidized price that a buyer pays should vary not only with her income, but also her age. Consistently with the aforementioned theoretical predictions, I find that, by raising vouchers to under 45 by $400 and lowering those to over 45 by $200, enrollment among the young raises by 50% while enrollment of the older is approximately unchanged. The reason is that, in equilibrium, a higher share of young enrollees reduces average cost by 15%, and markups by more than 20% (young buyers have also higher elasticity). Hence prices are lower, and this offsets the $200 reduction in the subsidy to the older group. At the same time, per-insured government expenditure is also reduced by approximately 15% ($600 per-buyer-year).

In addition to the papers discussed earlier, my work relates to numerous studies that use pre-reform data to study regulations that closely resemble those introduced by the ACA. Many of them use data from the Massachusetts' health exchange, a setting similar to ACA marketplaces.[5,6] In

[5] See also ??????.

[6] Other studies combine theoretical results and simulations: the impact of insurance mandates and minimum coverage provisions is studied by ?; the relationship between risk-adjustment and insurers' competition by ?; the long-run welfare impact of community rating rules by ?; the interaction between exchange design and labor markets by ?. Differently from my work, these papers abstract away from various aspects of market structure and imperfect competition observed in the US insurance market (see e.g. ??), and quantifications use data or estimates from employer-sponsored insurance (e.g. ??).

this context, **?** measure the welfare effect of an insurance mandate, and **?** develop and estimate a demand-supply model to study several ACA-like regulations, with main focus on the effect of age-based price adjustments for high-income buyers. My work complements theirs by focusing on the low-income segment of the population — not making active choices in the MA context —, and precisely studying the subsidy program for these buyers. More recently, **?** studied the welfare impact of adopting vouchers as opposed to a price-linked subsidy in a situation in which the government can be uncertain about market primitives, and particularly about insurers' costs. This is a relevant addition to my comparison of different subsidy designs, since in my counterfactuals vouchers are considered as a theoretical benchmark but I do not discuss implementation issues. Lastly, even outside the ACA setting there is a growing literature adapting industrial organization techniques to analyze the interaction between regulations and supply behavior of private health insurers. The cases of Medicare Advantage, Medicare Part D, Medigap, and Medicaid are the main focus of, among many others, **?**, **?**, **?**, **?**, **?**, and **?**. Notably, **?** shows distortions of insurers' decisions due to the design of subsidies in Medicare Part D, and in a broad review of this literature **?** explicitly discuss the importance of properly accounting for market power when designing these programs.

2 Subsidy design in health insurance

2.1 Setup

Consider a market with J health insurers. For now, each offers just one insurance plan $j = 1, ..., J$, with $j = 0$ denoting the outside option. I relax this assumption and consider multi-plan insurers later in the context of my application. Non-price characteristics of each plan are fixed, and their generosity of coverage is the same, so differences in demand across plans are driven by brand preferences and attributes of the provider networks.

Buyers are of one of two types, say young and old, denoted by $\tau = Y, O$, and I will use $G(\tau) \in [0, 1]$ to denote the fraction of type τ buyers in the market. Different types of buyers may have different health status (and thus cost for the insurer) and demand for insurance coverage. In particular, when selling coverage to a type τ buyer, insurer j expects to incur a cost equal to C_j^τ. Demand is instead defined as follows: Each individual buyer i has willingness-to-pay for product j equal to v_j^i, and the vector $v^i = (v_1^i, ..., v_J^i)$ is drawn *i.i.d* from the c.d.f. $F(v|\tau)$, conditionally on the buyer's type. Demand can then be represented by $\sigma_j(P, \tau)$, a function — derived from $F(v|\tau)$ — indicating the probability that a buyer of type τ chooses j when prices are $P = (P_1, ..., P_J)$. I assume that $F(v|\tau)$ is such that $\sigma_j(P, \tau)$ is strictly decreasing in P_j, and that it is continuous and differentiable. I also use $\eta_{jk}^\tau \equiv \left|\frac{\partial \sigma_j(P,\tau)}{\partial P_k}\right| / \sigma_j(P, \tau)$ to denote the semi-elasticity of demand for j by type τ buyers with respect to the price of plan k; in this section this is treated as constant.

As an extreme example of limits to price discrimination (similar to those mandated by the ACA, see Section **??**), insurers cannot vary prices by τ, so each sets a single price P_j that applies to all buyers. Expected profits are then a weighted average of profits across the two types:

$$\Pi_j(P_j, P_{-j}) = G(Y) \cdot \left[\sigma_j(P, Y) \cdot (P_j - C_j^Y)\right] + G(O) \cdot \left[\sigma_j(P, O) \cdot (P_j - C_j^O)\right]. \tag{1}$$

In this sense, heterogeneity in demand and cost across types is used to model selection: Even if observable, the type of a buyer is not priced, and neither the average nor the marginal cost curves of a product are necessarily constant functions of the corresponding pricing decision.

I assume complete information, and that prices form a Nash Equilibrium. That is, each insurer j sets its price P_j to maximize $\Pi_j(P_j, P_{-j})$ taking P_{-j} as given. I also maintain the assumption that (primitives are such that) prices are strategic complements, so that equilibrium comparative-static results from **?** can be applied.

2.2 Subsidies

To consider the effect of a subsidy program, let a subsidy design be a function $S(P, \tau) > 0$, such that type τ buyers face the discounted price vector $P - S(P, \tau) = (P_1 - S(P, \tau), ..., P_J - S(P, \tau))$. This will change demand by both types, from $\sigma_j(P, \tau)$ to $\sigma_j(P - S(P, \tau), \tau)$, and will have a corresponding effect on profits, equilibrium prices, and government expenditure.

I consider a situation in which the government tries to increase coverage for both young and old individuals, and, given coverage levels, to minimize public spending per-buyer. Therefore, the key inputs to the subsidy design problem are the effects of different $S(\cdot, \cdot)$ on equilibrium coverage levels and on expected spending; shedding light on these effects is the primary goal of this paper.

For every chosen design the equilibrium price vector (pre-subsidy) — $P^{*,S}$ — is such that, for every product, the price is the sum of average cost and markup. Characteristics of demand and cost across different types of buyers determine how average cost and markup depend on the chosen subsidy design. To see this formally, let $\alpha_j^S(P)$ be the share of young buyers of plan j when the subsidy design S is adopted and prices are P:

$$\alpha_j^S(P) = \frac{G(Y)\sigma_j(P - S(P, Y), Y)}{G(Y)\sigma_j(P - S(P, Y), Y) + G(O)\sigma_j(P - S(P, O), O)}. \tag{2}$$

I use this to define the corresponding average cost — $AC_j^S(\cdot)$ — and markup — $MK_j^S(\cdot)$ — functions under S (see supplementary Appendix S1 for detailed derivations):

$$AC_j^S(P) = C_j^Y \cdot \alpha_j^S(P) + C_j^O \cdot (1 - \alpha_j^S(P)), \tag{3}$$

$$MK_j^S(P) = \frac{1 - \left[\alpha_j^S(P)\left(1 - \alpha_j^S(P)\right)\left(\eta_{jj}^O - \eta_{jj}^Y\right)\left(C_j^Y - C_j^O\right)\right]}{\eta_{jj}^Y \alpha_j^S(P) + \eta_{jj}^O(1 - \alpha_j^S(P))}. \tag{4}$$

Then, $P^{*,S}$ is an equilibrium under S if for each j

$$P_j^{*,S} = AC_j^S(P^{*,S}) + MK_j^S(P^{*,S}), \tag{5}$$

so that the problem of the government — if cost and demand primitives are known — amounts to choosing S knowing that prices will then satisfy this equilibrium condition, and coverage and spending will respond accordingly. For this I also let g^S be the per-insured public spending in equilibrium corresponding to a subsidy design S.[7]

Targeted or non-targeted subsidies? A first comparison is between a non-targeted subsidy, for which the subsidy does not depend on τ, so $S(P,\tau) = S(P)$, in contrast to a targeted subsidy that does. The key result here is that, in a market with adverse selection, tailoring the generosity of subsidies to favor the cheaper-to-cover *and* more-price-sensitive group can lead to an equilibrium where all groups are better off, and the government spends less to subsidize a buyer.

The following proposition states that a voucher that is more generous for the low-cost, high-elasticity group leads to lower prices than a non-targeted voucher. Moreover, if the price reduction is sufficiently large (where this possibility depends on the heterogeneity in demand and cost across groups), all buyers are better-off and the government spends less per-insured buyer:

PROPOSITION **1** *If prices are strategic complements and $C_j^Y < C_j^O$, $\eta_{jj}^Y > \eta_{jj}^O$ for all j:*

(a) if S is a non-targeted voucher scheme $S(P,\tau) = V$, for which, at the equilibrium prices $P^{*,S}$, $\alpha_j^S\left(P^{*,S}\right) < \frac{1}{2}$ for all j, then there exists a $\Delta > 0$ such that, for the targeted voucher scheme $\widehat{S}(P,Y) = V + \Delta$, $\widehat{S}(P,O) = V - \Delta$, $P^{*,\widehat{S}} < P^{*,S}$;

(b) if, moreover, $P_j^{*,S} - P_j^{*,\widehat{S}} > \Delta$ for all j, then equilibrium quantities purchased under $\widehat{S}$ are higher in both groups, and $g^{\widehat{S}} < g^S$.

I prove this in Appendix A, but the intuition is as follows: The starting situation is one in which the government uses a non-targeted voucher, fixing an amount $V > 0$ such that $S(P,\tau) = V$, while the alternative is a targeted voucher for which $\widehat{S}(P,\tau) = \widehat{V}^\tau$, with $\widehat{V}^Y > V > \widehat{V}^O$. For any given P, the composition of buyers under the two schemes differs, with $\alpha_j^S(P) < \alpha_j^{\widehat{S}}(P)$ for all j because of quasi-linearity of preferences. Hence, since young buyers are cheaper to cover *and* more price sensitive, $\widehat{S}$ implies lower average cost *and* markups: $AC_j^S(P) > AC_j^{\widehat{S}}(P)$, $MK_j^S(P) > MK_j^{\widehat{S}}(P)$, for all P, all j. Hence, replacing S with $\widehat{S}$ induces an equilibrium with lower prices, $P^{*,\widehat{S}} < P^{*,S}$ (part (a)). Moreover, if these price reductions are larger than $V - \widehat{V}^O$ (the amount by which discounts to old buyers are lowered under $\widehat{S}$), coverage is higher for all buyers, and spending per-buyer is lower (part (b)).

[7] $g^S = \dfrac{G(Y)\,\sigma_j\left(P^{*,S} - S\left(P^{*,S},Y\right),Y\right)S\left(P^{*,S},Y\right) + G(O)\,\sigma_j\left(P^{*,S} - S\left(P^{*,S},O\right),O\right)S\left(P^{*,S},O\right)}{G(Y)\,\sigma_j\left(P^{*,S} - S\left(P^{*,S},Y\right),Y\right) + G(O)\,\sigma_j\left(P^{*,S} - S\left(P^{*,S},O\right),O\right)}$.

Importantly, Proposition **??** entails a possibility result, and this depends — in a complex way — on the interaction between demand and cost primitives in the market. This primitives become then key estimands for any quantitative evaluation of this subsidy design decision in a specific context.

Price-linked discounts or vouchers? A second relevant design decision is whether subsidies should be ex ante fixed by the regulator, or computed ex post as a function of market prices, as it is currently done under the ACA (see Section **??**). Practically, one can consider a scheme S with price-linked discounts, for which $\frac{\partial S(P,\tau)}{\partial P_j} > 0$ (for some j), or a voucher program where instead $\frac{\partial S(P,\tau)}{\partial P_j} = 0$ for all j and all P. This is also the main focus of independent work by **?**, where they also discuss several welfare consequences and critical implementation issues of this policy choice.

Price-linked discounts may be desirable if the government — not knowing demand and cost primitives — is unable to predict price. Adjusting subsidies to prices reduces then the possibility that discounts are too low (or too high) than what would be necessary to induce a target coverage level. However, adjusting subsidies to prices can distort insurers' incentives, and lead to an equilibrium with higher prices and higher spending by the government than what would result if subsidies were ex ante fixed. The intuition is straightforward and clearly resembles the difference between lump-sum as opposed to proportional taxes: If price increases are partly covered by discount adjustments, insurers maximize profits as if buyers were less price sensitive, and thus have additional incentives to set higher prices. The magnitude of this distortion decreases with the intensity of price competition and with the degree of horizontal differentiation in the market. I formalize this in Proposition **??** in the supplementary Appendix S1.

Relevant primitives. The previous discussion highlighted qualitative effects of different subsidy design decisions. Heterogeneity in cost and preferences across buyers, and how this interacts with pricing restrictions, are key estimands to determine the direction and magnitude of these differences. Moreover, the number of insurers and the corresponding intensity of price competition are also essential to predict how changes to a given design may imply a different equilibrium in a specific institutional context.

To study these design trade-offs in the context of ACA marketplaces, I go on to introduce relevant institutional details and data from the Californian marketplace. I then adapt the model to this setting, and obtain estimates of demand and cost primitives, which I then use to quantify the above theoretical predictions.

3 ACA marketplaces

3.1 Institutional context and federal regulations

As of 2013, 17 percent of US citizens younger than 65 did not have health insurance coverage (**?**). Affordability of the annual premium was a prominent reason why those uninsured did not purchase coverage in the private market (**?**), and this was one of the main motivations for the ACA. 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. Indeed, in the first two years of their operation, approximately 90 percent of buyers on these marketplaces received premium subsidies,[8] associated with annual government disbursements of approximately $40 billion.[9]

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 zipcodes — defining the level at which decisions by buyers and insurers take place. Every spring, insurers announce their interest in offering plans in each region in the subsequent calendar year. Entrants undergo a certification process,[10] 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 increased coverage options, and are ordered by an estimate of the actuarial value of their coverage: 60% for Bronze, 70% for Silver, 80% for Gold, and 90% or more for Platinum plans.

Federal law requires each participating insurer to offer at least one Silver and one Gold plan, but does not specify the precise combination of deductible and co-pays that lead to these coverage levels. Certain states, including California which is my focus later in this paper, have stricter regulations. They specify the full coverage details and mandate insurers to offer all five coverage levels in each region at which the insurer offers plans. In these marketplaces, insurers are differentiated only in their brand name, the structure of their provider networks, and the associated premiums.

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 each fall. Coverage then lasts for the subsequent calendar year.

[8] See e.g. https://aspe.hhs.gov/sites/default/files/pdf/83656/ib_2015mar_enrollment.pdf.

[9] See **?**.

[10] In California, for example, it is required that the providers include ten essential categories: (1) ambulatory patient services, (2) emergency services, (3) hospitalization, (4) maternity and newborn care, (5) mental health and substance use disorder services, (6) prescription drugs, (7) rehabilitative services and devices, (8) laboratory services, (9) preventive services / chronic disease management, and (10) pediatric services.

Pricing regulations. One important provision of the ACA is that insurers are not allowed to arbitrarily vary prices depending on buyers' observable characteristics. The only characteristic that affects annual premiums is the buyer's age, but even this adjustment is done in a pre-specified way. That is, each plan j offered in region r is associated with a single base price $b_{jr} > 0$, which is then translated to age-specific premiums using given age adjustment factors A^τ, equal for all products:

$$P_{jr}^\tau = A^\tau \cdot b_{jr}. \tag{6}$$

Age adjustments vary between 1 (for 21-year-old buyers) and 3 (for 64-year-old buyers), with details for all ages shown in Figure **??**.

This form of automatic age adjustment implies that — as long as age adjustment factors do not *perfectly* match age-differences in premiums that would result absent this regulation — base prices (hence premiums for all buyers) should vary with the composition of buyers in a market. For example, if buyers can only be 21 or 64-year-old, and the ratio between unconstrained premiums at 21 and 64 is, say, 1:5 (as approximately resulting from the Medical Expenditure Panel Survey, the pre-ACA data analyzed in **?**, and my own estimates in this paper) in response to the regulation premiums should be higher in markets with a higher share of 64-year-olds.

Subsidies. Although P_{jr}^τ is the premium received by the seller if a τ-year-old buyer enrolls in plan j in region r, subsidies are provided for all buyers whose household's annual income is below four times the federal poverty level (FPL; approximately $47,000 in 2014). For this, the law establishes a cap on the premium amount each individual has to pay for the second-cheapest Silver plan (benchmark plan) in each region. This cap is a function of the individual's household income (see Table **??**), ranging from $684 for the lowest income group to $4,368 for the highest income group (among subsidized individuals). Importantly, given income this cap amount does not vary by age.

This defines a premium discount for each age (τ) and income (y) in region r equal to

$$S^{\tau,y}(b_r) = \max\left\{A^\tau \cdot b_r^* - \overline{P}^y, 0\right\}, \tag{7}$$

where $\overline{P}^y$ is the premium cap for individuals with income y, and b_r^* is the base price of the benchmark plan in the region. Since the law establishes that each buyer must pay at least $1, the price of plan j for an individual with age τ and income y in region r is equal to

$$P_{jr}^{\tau,y} = \max\{P_{jr}^\tau - S^{\tau,y}(b_r), 1\}. \tag{8}$$

Figure **??** summarizes how age adjustments and subsidies transform base prices into prices faced by specific individuals, based on their age and income.

It is important to clarify the different role played by subsidy design as opposed to pricing

regulations — and in particular age-rating restrictions as studied in **??**. On the one hand, rating regulations determine payments to insurers as a function of pricing decisions *given* a selection of buyers into plans. Differently, subsidies directly determine such selection into enrollment pools (both size and composition) by altering the final discounted price for a buyer of a given age and income.

Cost-sharing reductions. A last important regulation in ACA marketplaces is the provision of cost-sharing reductions, available for buyers purchasing a Silver plan and whose income is lower than 2.5 times the FPL. For them, the federal government covers part of deductible and out-of-pocket expenses, so that the actuarial value of Silver plans increases from 70% to approximately 87%.[11] Cost-sharing reductions do not directly affect prices, but make Silver plans more attractive. Moreover, although the insurer covers approximately 70% of the health expenses, buyers' utilization when enrolled in a Silver plan is as if the plan covered 87% of expenses, and therefore likely to be higher (c.f. "moral hazard" in health insurance, see e.g. **??**).

3.2 The case of Covered California

With its 1.4 million enrollees in 2014, the Californian marketplace (Covered California) is the largest among ACA marketplaces,[12] and provides a useful setup to quantify the effect of alternative subsidy designs. The state is divided into 19 rating regions (see map in Figure **??**), with the number of insurers active in each region varying between 3 and 6, for a total of 11 participants.

Strengthening federal regulations, in Covered California financial details of different levels of coverage are fully standardized (see Table **??**), and when selling plans in a region each insurer must offer one plan in each level. Moreover, while the federal law allows some premium adjustments for tobacco use, these are not allowed in Covered California.

3.3 Data and summary statistics[13]

My main data is an extract of the official records of Covered California, obtained via Public Records Act (CA Gov §6250). I combine this source with the 2012-2013 Area Health Resources File[14] to construct the distribution of potential buyers across different age and income groups in the 19

[11] Specifically, for income levels between 100-150% of the FPL the deductible is reduced to $0, and maximum out-of-pocket is $2,250; for income levels between 151-200% of the FPL the deductible increases to $550, and maximum out-of-pocket is also $2,250; while for income levels between 201-250% of the FPL the deductible is $1,850, and the maximum out-of-pocket increases to $5,200.

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

[13] Note: The Californian marketplace and the Census Bureau regularly provide updated and increasingly granular data that could be use for a similar analysis. What follows is updated to data releases up to January 2015, but future analyses may be enriched with the use of additional information.

[14] See http://ahrf.hrsa.gov/.

Covered California rating regions. A detailed description of the original data files and of the construction of the final dataset can be found in the supplementary Appendix S2.

In every rating region r I observe the complete list of insurer-tier combinations (j), and the corresponding base price (b_{jr}). The Covered California dataset contains 401 products (insurer-tier-region) and corresponding pricing decisions.

I also observe market shares by 6 age-income groups: combinations of age between 20-29, 30-44, and 45-64, and low-income subsidized (LI) or not (HI). For each age-income pair (τ, y), and each region r, $q_{jr}^{\tau,y}$ is the observed number of enrollees in plan j, and $G_r(\tau, y)$ is the number of potential buyers. The ratio $s_{jr}^{\tau,y} = q_{jr}^{\tau,y}/G_r(\tau, y)$ is then the observed share of buyers choosing j conditional on the corresponding age-income group.

Knowing the ACA regulations on price adjustments, subsidies, and cost-sharing support, for each age-income pair and each product I then know the price paid by the buyer $P_{jr}^{\tau,y}$, the price received by the insurer P_{jr}^{τ}, and product characteristics z_{jr}^{y} that are relevant for buyers' decisions (see Table **??**).[15] Thus, corresponding to the 401 products and base prices — one pair for every (region, insurer, tier) — , I observe a total of 2,160 combinations — one for every (region, insurer, tier, age, income) — of prices (for the buyer and for the seller), product's characteristics, and market share within the specific age-income group.

Prices across age and income. Premiums that are relevant for buyers' decisions are summarized in Table **??**, which highlights the effect of age adjustments and subsidies.

Premiums for the high-income older than 45 are 3 times larger than those for high-income younger than 29. Because of the subsidy formula, however, this monotonicity does not hold for low-income buyers. For them, the subsidy design implies that Silver plans are available for approximately the same amount for all ages (the second cheapest Silver for *exactly* $1,452 for any age). For lower coverage, premiums decrease in age, while the opposite is true for Gold and Platinum plans, a pattern that is mechanically implied by (**??**) and (**??**), and can be seen directly in Figure **??**.

Market composition and subsidies. The size and composition of potential buyers across the 19 regions — important for insurers' pricing decisions and resulting enrollment — is summarized in Table **??**. In terms of market size, the overall number of potential buyers varies from less than 34,000 for the easter region (Mono, Inyo, and Imperial counties) to approximately 1,000,000 in each of the two regions covering Los Angeles. Importantly, moreover, an average of 70% of potential buyers are eligible for ACA subsidies.

[15] To calculate subsidies and cost-sharing reductions I treat subsidized buyers as if their annual income was 200% of the FPL. This income level was approximately the mean (and median) among subsidized buyers. In particular, the distribution of 2014 subsidized enrollees was 19% in 100-150% FPL, 36% in 151-200% FPL, 19% in 201-250% FPL, and 25% in 250-400% FPL. Source: http://hbex.coveredca.com/data-research/ . See Appendices B for further details.

The age-income composition of the relevant population is also quite heterogeneous across regions, as already summarized in Table **??** but also depicted in Figure **??**. For example, the San Francisco region has the minimum share of potential buyers older than 45 (36%), and a share of potential buyers eligible for subsidies of just 58%. At the opposite extreme, the northern region[16] has maximum shares of both, over 45 (49%) and subsidy eligible (80%). As evident from Figure **??** many intermediate cases (rich *and* old, or poor *and* young regions) are also observed.

Given the ACA subsidy design, the price of the second-cheapest Silver plan for low-income buyers is fixed. However, as the base price of the benchmark plan varies across regions, the corresponding per-buyer transfer from the government to insurers also varies. As reported in Table **??**, the base price of the benchmark plan is on average equal to $3,044, ranging from $2,338 to $3,690. This variation translates into a discount for a 21-year-old low-income buyer (paid to sellers directly by the government) equal to $1,592 in the average region, but this ranges from $886 to $2,238. For the oldest group of low-income buyers, the same variation in base prices of benchmark plans implies an average discount equal to $7,680, ranging from $5,562 to $9,618.

Enrollment and market structure. Given prices and market composition, to estimate demand I will use the corresponding purchase decisions across the 19 regions in the marketplace. Average per-plan enrollment across age-income groups and different levels of coverage is summarized in Table **??**, while cross-region averages of the decision of enrolling in Covered California are summarized in Figure **??**. A first salient fact is that, due to the subsidies, participation among low-income buyers is higher than among high-income unsubsidized. While they are 70% of potential buyers, subsidy recipients make up for 88% of adult enrollment in Covered California.[17] A second relevant fact emphasized in Figure **??** is that, among the subsidized buyers, those who are older than 45 are more likely to purchase insurance. Since subsidies keep the level of prices almost invariant across ages, this heterogeneity in participation decisions is informative on how older buyers are more willing to pay for coverage when compared to the younger group. Lastly, the vast majority of subsidized buyers chose either Bronze or Silver coverage. This concentration in plan selection can be largely explained by the availability of cost-sharing support, which makes Silver coverage comparable to Gold and Platinum while significantly cheaper.

Beside prices, population, and enrollment, another dimension of regional variation that I will exploit is variation in the combination of participating insurers and resulting distribution of market shares. The number of insurers goes from three to six, for a total of eleven participants. Four are

[16] Region 1: Alpine, Amador, Butte, Calaveras, Colusa, Del Norte, Glenn, Humboldt, Lake, Lassen, Mendocino, Modoc, Nevada, Plumas, Shasta, Sierra, Siskiyou, Sutter, Tehama, Trinity, Tuolumne and Yuba counties.

[17] My analysis and the main results are driven almost exclusively by estimates about this group, and – since my main focus is precisely on the subsidy program – little would change if I were to exclude the high-income unsubsidized. The exclusion of all children and young adults under the age of 20 is instead a significant limitation of my analysis, which might be resolved in the future as additional data are released.

large players — Anthem, Blue Shield, HealthNet, and Kaiser —, operating almost everywhere in the state. The remaining seven are smaller, local insurers offering coverage only in a small number of regions. Insurers are differentially attractive in different regions and for different income groups, as I summarize in Table **??**. Among the four largest carriers, each captures on average between 15-36% of subsidized buyers. Yet these shares range from a minimum of 10% to a maximum of 30% for HealthNet, 50% for Kaiser and Blue Shield, and over 90% for Anthem, the largest insurer in Covered California. The comparison between market shares within subsidized to those within high-income buyers suggests that preferences for insurers may vary by income. For example, while Anthem has, on average, a share of high-income buyers of approximately 43%, within subsidized this drops to 36%. In contrast, a similar difference is "gained" in the opposite direction by Blue Shield, whose share within low-income, subsidized buyers is 6% higher than within high-income. Part of these differences in the success of large insurers can be explained by the role played by small local insurers. For instance, Chinese Community Health Plan captures 33% of low-income buyers in San Francisco region 4, and 12% in San Mateo region 8; within high-income, however, this number drops to less than 7% in both regions. The opposite pattern is also observed, with Sharp — a local insurer in San Diego — enrolling 9% of low-income buyers, while 23% of buyers who pay their premiums in full.

4 Econometric model

The simple model introduced in Section **??** is now extended, adapted to nest ACA regulations, and later combined with Covered California data to study the effect of different subsidy designs in this particular market.

4.1 Primitives

Markets, insurers, and products. There are R geographic markets (regions), indexed by $r = 1, ..., R$. In each r, a population of individuals is offered J health insurance plans by N insurers, indexed by $n = 1, ..., N$.[18] For each n, $J_n \subset J$ is the set of products offered by n, and with a slight abuse of notation $n(j)$ will denote the seller of product j.

Individual buyers. A potential buyer i in region r is defined by a tuple (τ^i, y^i, v^i, c^i); superscripts are used throughout to index buyers, while subscripts index regions, insurers, and products.

The pair $(\tau^i, y^i) \in \mathcal{T} \times Y$ (both finite sets) denotes age and income of buyer i, and it is observed by all agents. Preferences of a buyer are instead unobserved by sellers, and are described by the vector $v^i = (v^i_1, ..., v^i_J) \in \mathbb{R}^J$. This collects i's willingness to pay for each of the J products relative

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

to the outside option $j = 0$. For buyers of age-income (τ, y), $P_r^{\tau,y} = (P_{1r}^{\tau,y}, ..., P_{Jr}^{\tau,y})$ denotes the "price vector" of differences between the price of each j and the price of the outside option (e.g. tax penalty for lack of insurance). Hence, i chooses j when $v^i \in \mathcal{D}_j(P_r^{\tau^i,y^i})$, where

$$\mathcal{D}_j(P_r^{\tau^i,y^i}) = \left\{ v \in \mathbb{R}^J : \underset{k \in J}{\text{argmax}} \left\{ v_k - P_{kr}^{\tau^i,y^i} \right\} = j \text{ , and } v_{jt} \geq P_{jr}^{\tau^i,y^i} \right\}. \tag{9}$$

Because it is not relevant to my analysis, I abstract away from how each v^i could be derived from a more primitive model of choice under uncertainty (see **?**, for an example of this derivation with CARA preferences).

The last element characterizing i, $c^i = (c_1^i, ..., c_J^i) \in \mathbb{R}_+^J$, collects the costs that each insurer expects to bear if i enrolls in a given product. That is, c_j^i is equal to the amount that the seller of j expects to spend to reimburse the health services of i under insurance policy j during the coverage period. Differences in c_j^i across j may reflect different underlying contracts with health providers, differences in administrative costs, differences in the generosity of coverage, or differences in expected utilization of health services.

Population. In every region, insurers share a common belief that age and income among potential buyers is distributed as $G_r(\tau, y)$, with $G_r(\tau, y) \geq 0$, and $\sum_{\tau,y} G_r(\tau, y) = 1$ for all r. Additionally, insurers know that, conditional on a pair (τ, y), preferences and cost of buyers in region r are distributed according to the continuous density $h_r(v^i, c^i|\tau, y)$.

Rather than the entire joint distribution h_r, the two relevant primitives that I will focus on throughout are the marginal density of preferences conditional on age and income:

$$f_r(v^i|\tau, y) = \int_{\mathbb{R}_+^J} h_r(v^i, c^i|\tau, y) \, \mathrm{d}c^i, \tag{10}$$

and the vector of expected costs for each j in region r conditional on age, income, and buyer's preferences:

$$\psi_r(v, \tau, y) \equiv \int_{\mathbb{R}_+^J} c^i \cdot \frac{h_r(v, c^i|\tau, y)}{f_r(v|\tau, y)} \, \mathrm{d}c^i. \tag{11}$$

In words, $\psi_{jr}(v, \tau, y)$ is the insurer's expected cost when covering under plan j a buyer with age-income (τ, y) and preferences v in region r.

Demand. The function $\sigma_{jr}^{\tau,y}(P_r^{\tau,y})$ denotes the probability that a buyer with age-income (τ, y) purchases j when the prices are $P_r^{\tau,y}$. Using the above notation this can be expressed as

$$\sigma_{jr}^{\tau,y}(P_r^{\tau,y}) \equiv \int_{\mathcal{D}_j(P_r^{\tau,y})} f_r(v^i|\tau, y) \, \mathrm{d}v^i. \tag{12}$$

4.2 Regulations, profits, and equilibrium

Insurers do not adjust prices across age and income. Instead, for every product j, in region r the seller chooses only a univariate base price b_{jr}. A regulation then determines both, the price paid by a buyer as function of age and income, and the price received by the seller as a function of age only.

Formally, $P_{jr}^{\tau} = R^{\tau}(b_{jr})$ is the revenue for the seller if a buyer of age τ purchases j. Of this amount, the buyer only pays the subsidized price $P_{jr}^{\tau,y} = P_{jr}^{\tau} - S^{\tau,y}(b_r)$, where the discount $S^{\tau,y}(b_r)$ is covered by the government, and it is computed via the subsidy formula (**??**) as a function of base prices in the region.

The expected profits for insurer n in region r are then a function of b_r:

$$\Pi_{nr}(b_r) = \sum_{j \in J_n} \sum_{\tau,y} G_r(\tau, y) \cdot \underbrace{\left(\sigma_{jr}^{\tau,y}(P_r^{\tau,y}) \cdot P_{jr}^{\tau} - \int_{\mathcal{D}_j(P_r^{\tau,y})} \psi_{jr}(v^i, \tau, y) \cdot f_r(v^i|\tau, y)\, \mathrm{d}v^i \right)}_{\text{Expected profits from buyers with age-income } (\tau,y)}. \tag{13}$$

With this as payoff function of the pricing game, base prices are set in each region as a Nash equilibrium, with each insurer maximizing $\Pi_{nr}(b_r)$ taking base prices of other insurers as given.[19]

5 Identification

A critical challenge for identification of the model's primitives is that, differently from a market without selection, costs for insurers may not only vary by product, but also with the characteristics of the buyer. For this, my approach here relies on supplementing choice and price data with supply-side equilibrium assumptions, applying to this context the "inversion of first-order conditions" widely adopted by the industrial organization literature since **?** and **?**. In terms of estimation, this idea is already present in **?**, but my main point here is to highlight sufficient conditions on observables that ensure that demand and cost functions are identified.

Observables. In every region, and for every product $j =$ insurer-tier, observables consist of the collection

$$\mathcal{O}_{jr} \equiv \left(b_{jr}, \left\{ P_{jr}^{\tau,y}, P_{jr}^{\tau}, s_{jr}^{\tau,y}, z_{jr}^{y}, G_r(\tau, y) \right\}_{\tau \in \mathcal{T}, y \in Y} \right), \tag{14}$$

along with functions $R^{\tau}(\cdot)$ and $S^{\tau,y}(\cdot)$ which determine price adjustments, and subsidies. Importantly, entry and product characteristics are fixed and exogenous throughout.

[19] Since this is a model of imperfect competition between insurers offering vertically and horizontally differentiated plans, and the distribution of individual preferences and cost admits a continuous density, profit functions are continuous. This reduces concerns for lack of equilibrium existence emerging from discontinuity in payoffs, a first-order issue in models of insurance supply under perfect competition between vertically differentiated plans (see **?**, for a rich discussion). An existence result in a model similar to mine is provided by **?**.

5.1 Overview of nonparametric identification

Nonparametric identification is formalized in Theorem **??** in Appendix B, but here I provide a brief summary before introducing a parametric version of the model and the corresponding identification argument that I use for estimation.

First, $f_r(\cdot|\tau, y)$ is identified for each age-income pair by inversion of the relationship between observed and model-predicted market shares: $s_{jr}^{\tau,r} = \sigma_{jr}^{\tau,y}(P_r^{\tau,y})$. Sufficient conditions for this are provided by **?** (BH), and the essential requirement is observability of "price instruments" that do not directly affect choices, yet affect pricing decisions. In my model, base prices (and thus prices faced by buyers) depend on the population composition in terms of age and income, since this affects the shape of payoff functions in (**??**). If unobservables relevant for individual choices do not depend on population composition, the collection $\{G_r(\tau, y)\}_{\tau\in\mathcal{T},y\in Y}$ can be used as a set of valid price instruments (this is similar to **?**).[20] BH then show that if the support of prices (conditional on instruments) is sufficiently large, there exists a unique density $f_r(\cdot|\tau, y)$ yielding $\sigma_r^{\tau,y}(\cdot)$ consistent with the observed market shares $s_{jr}^{\tau,y}$.

Then, given that $\{f_r(\cdot|\tau, y)\}_{\tau\in\mathcal{T},y\in Y}$ is identified, expected cost functions (ψ_r) are the remaining unknowns in the equilibrium FOC (here simplified to single-product sellers)

$$MR_{jr} = \underbrace{\sum_{\tau,y} G_r(\tau, y) \cdot \frac{\partial}{\partial b_{jr}} \left(\int_{\mathcal{D}_j(P_r^{\tau,y})} \psi_{jr}(v^i, \tau, y) \cdot f_r(v^i|\tau, y)\, \mathrm{d}v^i \right)}_{\text{Marginal cost}}, \tag{15}$$

where the left-hand side is equal to marginal revenues and does not depend on ψ_r. This is an underidentified system, with more unobserved unknowns than observed pricing decisions. Yet, leveraging on quasi-linearity of preferences,[21] restrictions on how costs may vary across markets, and a constructive proof analogous to the one in **??** yield sufficient conditions for identification of ψ_r. An essential requirement for this is that the same variation in prices used to identify demand induces variation in the composition of marginal buyers (in terms of v, τ, and y) across observed products. In addition to the main result in Appendix B, in the supplementary Appendix S3 I provide an additional discussion and present alternative conditions for identification that do not rely on a large support assumption.

[20] Identification of demand may also rely on different strategies, varying across contexts and data structures. For example, regression-discontinuity is used in **?**, while a mix of BLP-style instruments (other products' characteristics, c.f. **?**) and Hausman-style instruments (prices in different geographic markets) are used by **?**. See also **?**.

[21] Because preferences are quasi-linear in price, the set $\mathcal{D}_j(P_r^{\tau,y})$ is a cone in $\mathbb{R}^J$, and the set of marginal buyers of product j in r is the limit of an "L-shaped" set in $\mathbb{R}^J$ (limit for $\varepsilon \to 0$ of $\mathcal{D}_j(P_r^{\tau,y}) \cap \left(\mathbb{R}^J \setminus \mathcal{D}_j(P_{jr}^{\tau,y}, P_{-jr}^{\tau,y} - \varepsilon)\right)$). **??** can then be followed to construct expected cost functions. (See Appendix B for details).

5.2 Parametric identification using Covered California data

Sufficient conditions for identification hold under the following combination of parametric assumptions which I use for estimation using Covered California data.

Demand. Each individual i with age-income $(\tau^i, y^i) = (\tau, y)$ has a (money-metric) indirect utility from purchasing j in r equal to

$$U^i_{jr} = \underbrace{-\alpha^{\tau,y} \cdot P^{\tau,y}_{jr} + (\beta^{\tau,y})' z^y_{jr} + \delta^{\tau,y}_{n(j)}}_{\mu^{\tau,y}_{jr}} + \xi^{\tau,y}_{jr} + \epsilon^i_{jr} \,, \tag{16}$$

where ϵ^i_{jr} is *iid* distributed as a standard extreme value type I. $\xi^{\tau,y}_{jr}$ is instead the (endogenous) mean-zero error term capturing how unobserved characteristics of j in region r affect choice by buyers with age-income (τ, y).

The unknown parameters for every observed age-income combination are the price coefficient $\alpha^{\tau,y}$, the coefficients on product characteristics (deductible and out-of-pocket maximum) $\beta^{\tau,y}$, and insurers' fixed-effects $\delta^{\tau,y} = \left(\delta^{\tau,y}_1, ..., \delta^{\tau,y}_N\right)$.[22]

The standard inversion for multinomial logit discrete choice (see e.g. **?**) implies then that $\sigma^{\tau,y}_{jr}(P^{\tau,y}_r) = \exp\left(\mu^{\tau,y}_{jr} + \xi^{\tau,y}_{jr}\right) / \left(1 + \sum_{k \in J} \exp\left(\mu^{\tau,y}_{kr} + \xi^{\tau,y}_{kr}\right)\right)$, and $s^{\tau,r}_{jr} = \sigma^{\tau,y}_{jr}(P^{\tau,y}_r)$ yields to

$$\xi^{\tau,y}_{jr} = \ln(s^{\tau,y}_{jr}) - \ln(s^{\tau,y}_{0r}) - \mu^{\tau,y}_{jr}. \tag{17}$$

The parameters in (**??**) are identified if, next to standard rank conditions requiring population composition to shift prices conditional on insurer and product characteristics,[23] for all (τ, y)

$$E\left[\xi^{\tau,y}_{jr} \middle| z^y_r, \{G_r(\tau, y)\}_{\tau \in \mathcal{T}, y \in Y}\right] = 0; \tag{18}$$

requiring conditional mean independence between unobservables in the demand from a specific age-income group and population composition. Simply put, this assumes for example that unobservables (network of providers) affecting demand by 21-year-olds do not depend on the number of 64-year-olds in the same market.

If this is the case, since insurers' profit functions (equation (**??**)) depend on the age-income composition of the region, variation in this composition as the one shown in Figure **??** can induce variation in observed base prices, which can then be used to identify demand. Figures **??** and **??** show that this variation is present in Covered California, highlighting the response of base prices

[22] Estimates from alternative demand specifications including a random-coefficients model à la BLP (**??**) are presented in Appendix C.

[23] $E\left[\left[z^y_r, \{G_r(\tau, y)\}_{\tau \in \mathcal{T}, y \in Y}\right]' \left[z^y_r, P^{\tau,y}_{jr}\right]\right]$ must be full-column rank.

(in natural logarithm) to variation in age-composition, and income-composition, respectively, after controlling for insurer and level of coverage fixed-effects.[24] From Figure **??**, the base price is higher in regions with a larger number of potential buyers older than 45. This is consistent with the ex-ante expected cost of covering old buyers being higher than three times the one of covering young buyers, and under the age rating regulations insurers set higher prices when relatively more buyers are over 45.[25] Instead, as shown in Figure **??**, regions with a higher percentage of subsidy-eligible buyers have lower prices. This is consistent with both, low-income buyers being more price sensitive, and with the market being larger (insurers are able to operate more efficiently, or pay lower rates to providers).

One possible concern with this approach is that insurers could have altered provider networks or marketing strategies taking into consideration market composition. If these features enter the error term $\xi_{jr}^{\tau,y}$, this would *not* be independent from $\{G_r(\tau,y)\}_{\tau\in\mathcal{T},y\in Y}$, violating (**??**). To partially address this concern, in Appendix C I discuss robustness to an alternative strategy, where demand is estimated simultaneously across all age-income pairs with insurer-region fixed-effects. These fixed-effects capture omitted characteristics which are constant across plans offered by the same carrier in a region, including provider networks. The remaining endogenous error ξ_{jt}^{τ} is then largely driven by functional form mis-specification (e.g. differences in the money-metric valuation of provider networks across buyers of different age). From a practical perspective, this alternative approach yields similar estimates of the elasticities relevant for my final counterfactual.

Cost. Expected cost of coverage varies by age of buyer, tier of coverage, insurer, and region:[26,27]

$$\psi_{jr}(\tau) = \phi^{\tau} + \phi_{\text{tier}} + \phi_{n(j)} + \phi_r. \tag{19}$$

Taking the estimated $\widehat{\sigma}_{jr}^{\tau,y}(\cdot)$ as given, I then assume that base prices are set in equilibrium, and that insurers' FOCs must hold up to a mean-zero error term. I interpret the source of this error as discrepancy between the population composition as assumed by insurers when setting base prices

[24] I discuss first-stage estimates together with demand estimates in Section **??**.

[25] In **?** we provide a *pseudo* placebo test for this, considering a large sample of counties across the entire US. While before the ACA there was no significant relationship between price of young buyers and number of old buyers in the county, this is now a positive and highly significant effect.

[26] Such coarseness in the heterogeneity of cost across different buyers is dictated by data limitations: As evident from the conditions of Theorem **??** in Appendix B, and Proposition **??** in supplementary Appendix S3, the dimensionality of cost heterogeneity is bounded above by the dimensionality of (estimated) demand heterogeneity.

[27] From a practical perspective, the presence in the ACA exchanges of several risk-adjustment mechanisms (risk corridors, re-insurance, and standard risk adjustment; see e.g. http://kff.org/health-reform/issue-brief/explaining-health-care-reform-risk-adjustment-reinsurance-and-risk-corridors/) mitigate possible concern that large heterogeneity in cost conditional on age, tier, insurer, and region plays an important role in insurers' incentives. Concerns however are still present, since risk-adjustment is being phases in slowly, and it is unlikely to completely remove unobservable risk heterogeneity across buyers. In future work it will be important to estimate demand models with richer heterogeneity in preferences, which will then allow to estimate a richer heterogeneity in buyers' costs.

and the age-income composition G_r as observed in my data.

Formally, $\widehat{\eta}_{jr}$ is the value of marginal profits evaluated using the observed G_r:

$$\begin{aligned}\widehat{\eta}_{jr} &= \frac{\partial \widehat{\Pi}_{n(j)r}(b_r)}{\partial b_{jr}} \\ &= \widehat{MR}_{jr} - \widehat{MC}_{jr} \\ &= \widehat{MR}_{jr} - \sum_{k \in J_{n(j)}} \sum_{\tau,y} G_r(\tau, y) \cdot \frac{\partial \widehat{\sigma}_{kr}^{\tau,y}}{\partial b_{jr}} \cdot \psi_{kr}(\tau),\end{aligned} \tag{20}$$

where marginal revenues $\widehat{MR}_{jr}$ are constructed from demand estimates and base prices applying age rating and subsidies. As long as the difference between the observed G_r and the population composition assumed by insurers is independent from agents' choices,

$$E\left[\widehat{\eta}_{jr} \middle| z_r, b_r, s_r\right] = 0, \tag{21}$$

where conditioning variables exclude the population composition G_r which is correlated with $\widehat{\eta}_{jr}$.

Identification of tier, insurer, and region parameters of the individual cost function (**??**) is standard, relying on variation of marginal revenues (equal to cost) across tiers, insurers, and regions, respectively. Instead, identification of cost differences for the same product across different buyers is peculiar to my context, and in general to selection markets.

For this, parametric restrictions such those in (**??**) are needed, since otherwise the number of unknowns is larger than the number of observed pricing decisions.[28] Moreover, as it is evident from equation (**??**), insurers' risk-neutrality implies that marginal profits are linear in expected costs across buyers of different age, and this delivers the condition for identification of ϕ^τ.

Specifically, to distinguish cost differences between two age-groups τ, τ', the composition of marginal buyers across the two groups must vary across observed products, requiring variation in the ratio

$$\frac{\text{marginal buyers of age-income } \tau, y}{\text{marginal buyers of age-income } \tau', y} = \frac{G_r(\tau, y) \cdot \partial \widehat{\sigma}_{kr}^{\tau,y} / \partial b_{jr}}{G_r(\tau', y) \cdot \partial \widehat{\sigma}_{kr}^{\tau',y} / \partial b_{jr}}. \tag{22}$$

This is a special case of the much more general result presented in Appendix B, and the main intuition is summarized with an abstract example in Figure **??**, in which I consider two age groups, over and under 45. Demand estimates and population composition can be used to calculate per-buyer marginal revenues (vertical axis), and percentage of marginal buyers younger than 45 (horizontal axis). This can be done for every observed product (panel (A)), assuming that cross-insurer, cross-tier, and cross-region differences have been controlled for. The co-variation between marginal revenues and composition of marginal buyers is observed, but can also be predicted by the model

[28] This differs from the case of standard consumption goods, where assuming equilibrium pricing one identifies marginal cost nonparametrically for every observed product without additional assumptions (see e.g. **???**).

for any given pair of parameters $\phi^{20\text{-}44}$ and $\phi^{45\text{-}64}$ (panel (B)). This follows the linear relationship:

$$\widehat{MR}_{jr} = \widehat{MC}_{jr} = (\text{under 45 marginal buyers}) \cdot \phi^{20\text{-}44} + (\text{over 45 marginal buyers}) \cdot \phi^{45\text{-}64}. \quad (23)$$

As long as the age-composition of marginal buyers varies, there exists a unique pair of cost parameters $\widehat{\phi}^{20\text{-}44}$ and $\widehat{\phi}^{45\text{-}64}$ for which the model-predicted marginal cost is (in expectation) equal to the one implied by demand estimates and equilibrium pricing conditions (panel (C)).

6 Estimates from Covered California data

6.1 Demand by age and income

First-stage: population composition and base prices. Estimated coefficients from a hedonic regression of base prices on product characteristics and age-income composition are reported in Table **??**. As anticipated, the effect of population composition on pricing decision is large in size and highly significant. Looking at specifications in levels and natural logarithms, the *Share 45-64* and *Share 100-400% FPL* – both excluded from individual demand equations – explain approximately 20% of the variation in base prices across markets.

After controlling for insurer and contract details, the base price of a plan increases by 8.3% (+$243 per year) if the share of over 45 among potential buyers increases by 10%. The share of subsidy eligible buyers also matters: a 10% increase in the size of this group relative to all potential buyers induces an estimated drop in base price of 6.6% (−$205 per year).

Demand estimates. Demand estimates across the six age-income groups in Covered California data are obtained exploiting this variation in prices, as well as cross-market differences in the set of participating insurers.

Coefficients from equation (**??**) are reported in Table **??**, and are obtained correcting for endogeneity of $P_{jr}^{\tau,y}$ using a control function. This is convenient here because the subsidy design implies that the *level* of prices of low-income buyers is maintained approximately constant (at least for the second cheapest Silver plan) at $1,452. Therefore, the instruments vary mostly the *dispersion* of prices for subsidized buyers around this value. With a control function, however, I can exploit the rating regulations and capture unobservables affecting rating decisions by adding to the indirect utility the residuals of the hedonic regression in the first stage. Robustness to different specifications and methods to correct for endogeneity of premiums — including standard 2SLS estimates and the alternative specification with insurer-region fixed-effects discussed in the previous section — is presented in Appendix C. These alternative strategies yield very similar elasticity estimates. The bottom rows of Table **??** report averages of (age-income specific) own-price elasticities (% drop in plan enrollment if the annual premium increases by 1%), and extensive margin semi-elasticities

(% drop in overall coverage in the region if all prices increase by $100).

Commenting on the results: Preferences for insurance are estimated to be quite heterogeneous across different age groups, and between subsidized and unsubsidized buyers. This is largely driven by differences in product choice as relative prices across tiers and across insurers vary for different groups and across rating regions. Magnitudes of the estimated coefficients on annual premium are decreasing in income and age, and, more relevantly, implied own-price elasticities go from -3% for the 20-44 low-income people, to -1.5% for 45-64 low-income. This value drops to -1% (approximately constant in age) among the high-income.

Extensive margin semi-elasticities are also highly heterogeneous, a result that is largely driven by the variation in average participation decisions as summarized in Figure **??**. If all premiums increase by $100, approximately 12% of buyers who are subsidized and whose age is between 20-44 would opt out of Covered California. Instead, for the over 45 subsidized, a uniform $100 increase in all prices would reduce participation by less than 4%. The same extensive margin is even smaller among high-income unsubsidized. These estimates are quantitatively aligned with what estimated by **?** using quasi-experimental variation in the Massachusetts exchange.

Beside price, also the estimated effect of product characteristics on buyers' choices are consistent with intuition:[29] buyers like lower deductible and lower out-of-pocket expenditure. In particular, comparing the two more demanded contracts in Covered California, a low-income younger than 30 would be willing to pay, on average, an additional $1,000 to be covered under a Silver plan ($550 deductible and $2,250 maximum out-of-pocket) rather than under a Bronze plan ($5,000 deductible and $6,250 maximum out-of-pocket). For the same increase in coverage, a low-income buyer older than 45 would be willing to pay more than $2,500, showing again significant age-heterogeneity in the valuations of different contracts. Lastly, insurers' fixed-effects are aligned with differences in market shares already highlighted in Table **??**. Anthem is the most preferred, followed closely by Blue Shield, Kaiser, and HealthNet. When active, certain local insurers are comparable to the four large players (e.g. Chinese Community Health plan, Sharp, or Valley), although this is not the case for all participants, with Contra Costa, Molina, LA Care, and Valley estimated to be inferior for all buyers.

6.2 Expected cost by age, tier, insurer, region

Variation in composition of marginal buyers. In Figure **??** I reproduce using Covered California data and demand estimates the illustration of cost identification in Figure **??** (Section **??**). I plot the co-variation between estimated marginal revenues and age composition of marginal buyers. I exploit this variation to estimate cost, varying by buyer's age, contract details, insurer, and

[29] For this discussion I focus on low-income buyers since estimates for the unsubsidized group are much less precise, and not relevant for my results. These represent the minority (less than 10%) of buyers in the Covered California data, and are not directly affected by the subsidy program which I will alter in the counterfactual.

region.

As shown in the figure, the larger the fraction of under 45 buyers among those reacting to changes in b_{jr}, the lower the individual marginal revenue predicted by the demand model, equal to cost under my equilibrium assumption. At the same time, more generous plans (left vs. right panel for, respectively, Silver vs. Gold coverage) have a higher marginal revenue/cost, and – as highlighted with different symbols – after controlling for level of coverage and age composition of marginal buyers regional differences directly impact how cost estimates vary in different regions. (A similar intuition holds for differences between insurers.)

Cost estimates. Driven by this variation, cost estimates are obtained imposing the moment condition (**??**), and reported in Table **??**. Columns (1)-(3) show estimates allowing only for product, insurer, and region heterogeneity, as in a model without selection. Then, in columns (4)-(6) the same specifications are estimated allowing expected costs to vary also by age. For this, I distinguish between over and under 45-year-old, and use only the products that are available to all buyers (Bronze, Silver, Gold, and Platinum).[30]

Estimated expected cost increases in generosity of coverage and age. Looking at the more comprehensive specification in column (6) – which I will also use for my counterfactuals – to enroll a young buyer in a Bronze plan ($5,000 deductible and $6,250 out-of-pocket maximum) costs Anthem approximately zero in expectation ($200, not statistically significant). For the same plan the carrier expects instead to spend on average $5,860 if the buyer is older than 45. As coverage increases, expected cost increases as well. Enrolling a buyer in a Silver plan costs the insurer $1,600 if the buyer is under 45, and over $7,000 if older. This is the contract chosen by the vast majority of subsidized buyers (see Table **??**), hence these estimates are the most important for my final results, and the more robust across different specifications.[31] This implies a cost-ratio of approximately 4.5 between the two groups, so that age rating restrictions yield to higher prices for young buyers than what they would be charged if age-based price discrimination was unrestricted.[32]

The second relevant dimension of cost-heterogeneity is across insurers. Anthem and Blue Shield of California are estimated to have almost identical cost. For HealthNet, instead, I estimate a lower

[30] This choice is dictated by the necessity of observing sufficient variation in composition of marginal buyers across products (after conditioning on insurer and market averages). In the Covered California data, subsidized buyers are the most price sensitive, and make up for 90% of enrollment. Moreover, buyers younger than 45 have very similar preferences, so that the main variation in composition of marginal buyers and enrollment pools relevant for pricing decisions is variation in the relative share of under 45 among marginal buyers, and this is what I focus on.

[31] Importantly, for Silver plans expected cost for the insurer is driven by two factors. First, insurers expect buyers to behave as if deductible and co-pays were lower because of cost-sharing reductions, and this may affect their utilization. Given utilization, however, the insurer expects claims to reflect a contract with a $2,250 deductible and $6,250 out-of-pocket (the difference is paid for by the government).

[32] In **?** we also show that, before the ACA, the ratio of unrestricted prices of identical products between 64 and 21 was equal to 4.023. Similar numbers are obtained using the Medical Expenditure Panel Survey as shown in Appendix C.

individual cost (by $500, only significant at the 20% level), while for Kaiser Permanente a higher cost ($600, highly significant). This is driven by differences in prices and/or estimated attractiveness. Kaiser, in particular, has higher prices in most markets, and differences in elasticity do not alone explain these price differences. Hence the model rationalizes higher prices with a higher average cost for the carrier.[33]

Lastly, regional effects are also relevant to cost, and are jointly significant at any conventional level.[34] Moreover, inspecting how estimates vary across different markets, I find that cost tend to be lower in urban areas, and particularly in regions with a high number of health providers per-capita.[35] This is in line with the fact that cost of specific procedures when covered by the insurer varies with the structure of the market of health providers (see also **??**).

7 Equilibrium under different subsidy designs

Building on the theoretical insights from Section **??**, and using the richer model with estimates from Covered California, I can now compare equilibrium under different designs of the subsidy program.

7.1 Vouchers vs. price-linked discounts

My first comparison is between the ACA subsidy design and an "equivalent" voucher program, where buyers receive fixed discounts equal to those resulting in equilibrium under the current scheme. To start, recall that under the ACA design if base prices are b_r^{ACA} buyers receive price-linked discounts $S^{\tau,y}(b_r^{ACA})$ computed via equation (**??**). I compare this to the alternative design $\widehat{S}$, under which buyers receive a fixed discount ("equivalent voucher") $\widehat{S}_r^{\tau,y} = S^{\tau,y}(b_r^{ACA})$, which is *not* adjusted as insurers' pricing decisions b_r vary. Because young buyers are more price-sensitive and cheaper to cover, insurers' marginal profit functions under $\widehat{S}$ lay below marginal profit functions

[33] This fact might be surprising, as this large integrated HMO is known for providing high quality health services while able to contain its cost. The difference in prices has been explained in various, alternative ways, which I briefly discuss here although to take a specific stand favoring one explanation over the other is beyond my purpose. First of all, Kaiser's vertical integration implies that its marginal cost might include more "cost items" than those considered by a non-integrated insurer, for which individual cost is mostly driven by expected utilization and negotiated reimbursements to physicians (also see **?**). Moreover, in the public discussion it has been highlighted how Kaiser has capacity constraints, and so its strategy to post high-prices in ACA marketplaces might in fact be a way to limit the influx of new patients into its hospitals and clinics during the first years after the reform. Lastly, Kaiser's official position is that its providers' networks were not adjusted ad-hoc to the new regulations, while other insurers allegedly selected a small set of their providers to be covered by plans offered in Covered California.
(See http://articles.latimes.com/2013/jun/12/business/la-fi-kaiser-health-rates-20130613)

[34] I can only consider the joint significant of these estimates, as I do not assume independence across plans within a given market. Thus, the variance on each individual market fixed-effect cannot be estimated.

[35] For example, Los Angeles county, San Francisco, and Alameda, where according to the Dartmouth Atlas (http://www.dartmouthatlas.org/) there are approximately 3 (acute care) hospital beds per-1,000 residents, insurers' expected cost is estimated to be lower than in San Mateo, or in the Eastern region, where hospital beds per-1,000 residents are less than 2.

under S (discussion in Section **??**). Therefore, even in this richer model insurers set lower base prices under the voucher scheme $\widehat{S}$ than under the ACA design, and relevant equilibrium quantities respond accordingly.

To quantify these differences I start off by computing equilibrium base prices b_r^{ACA} under the status quo subsidy design. For this, I use demand and cost estimates from Covered California, and assume that insurers use the observed age-income composition G_r to compute expected profits in each region.[36] I then carry on the same equilibrium computation under the equivalent voucher scheme $\widehat{S}$ and compare the two equilibria in terms of enrollment by age group, average cost, average markups (difference between total per-person amount received by the insurer and average cost), and subsidy expenditure.

Table **??** reports the results of this comparison, showing that differences between the two designs are indeed sizable. In particular, under fixed vouchers equilibrium markups are 15% lower (approximately $200 per-year). This is driven by insurers setting lower base prices, yielding to lower pre-discount premiums. At these lower premiums, because vouchers are not adjusted discounted prices are lower than under the ACA design, and enrollment is higher (+7% among under 45, and +9% for the older group). With these changes in enrollment the age-composition of buyers remains almost unaltered, and this reflects into average cost being approximately the same in the two equilibria (+1%). Markups reductions are then largely explained by a lower per-person amount received by insurers, combination of lower average subsidy provided by the government (−5%), and lower contribution paid by buyers (−11%).

7.2 Age-adjusted vouchers

Next I quantify the difference in equilibrium outcomes induced by variations in the voucher amount across different age groups; this allows me to explore the potential gains from age-targeted subsidies. I consider fixed vouchers analogous to $\widehat{S}$, where instead of setting voucher values equal to the discounts provided under the ACA scheme, for any age group τ there is an amount $V^\tau > 0$ for which $\widehat{S}^\tau = V^\tau$. (I omit y from the notation since I observe only one income level in my data.) The discussion in Section **??** highlights that, within this class of subsidy schemes, reducing vouchers for the old and increasing those for the young can yield to lower average cost, higher coverage, and lower per-buyer spending. Intuitively, increasing the relative share of young buyers in enrollment pools lowers average cost and markups (by raising average elasticities), thus induces insurers to set lower prices. If this price reduction exceeds the amount by which the voucher of the older group was raised, all buyers are better off. Additionally, if the enrollment increase is sufficient to compensate the lower per-enrollee markup, also total profits are higher, so that insurers are not worse off either.

To investigate this mechanism, I simulate and report equilibrium outcomes for different pairs

[36] In practice, I find base prices that solve the equilibrium FOC precisely, rather than with the error $\widehat{\eta}_{jr}$.

of voucher values ($V^{20\text{-}44}, V^{45\text{-}64}$), varying these over a two-dimensional grid of \$100 increments. In Figure **??** I plot level curves (in the space of voucher values) for equilibrium outcomes that would likely enter the government's objective: average cost, average premium received by insurers, enrollment for different age groups, and subsidy expenditure (both per-insured person and aggregate). To facilitate comparisons, for each outcome I highlight the curve corresponding to the equilibrium level under the ACA design, and each other curve corresponds to a 10% increase (or decrease) from this level. Panel (c) and (d) show how equilibrium enrollment in the two age groups varies as a function of vouchers. Because they are less price sensitive, for old buyers level curves of enrollment are significantly sparser than those of the young, corresponding to a flatter surface. In practice, a \$100 increase in $V^{20\text{-}44}$ yields to a much larger increase in enrollment among the young than the drop in enrollment among the old implied by a \$100 reduction in $V^{45\text{-}64}$. The effect of age adjustments to vouchers follows, as it is shown directly by the downward sloping curves in panel (d): one can increase $V^{20\text{-}44}$, reduce $V^{45\text{-}64}$, and obtain higher coverage for both groups, with lower cost (panel (a)) and lower per-person public spending (panel (e)).

In Table **??** I compare with more precision equilibrium outcomes along two level curves depicted in panel (d), with voucher values for which enrollment of the over 45 is approximately constant at either 1.08 or 1.02 times the equilibrium level under the ACA. In both situations a \$400 increase in $V^{20\text{-}44}$ and a simultaneous \$200 reduction $V^{45\text{-}64}$ maintain enrollment of the older group approximately invariant. At the same time, however, enrollment among under 45 increases by approximately 60%, and average cost is 10-15% lower. Since young buyers are more price-sensitive markups are also lowered, up to a 25% drop from ACA levels (approximately \$448 per-person-year), and per-person public expenditure is reduced by more than 15% (or \$600 per-person-year). Lastly, the large increase in enrollment compensates the reduction in per-person markup, and insurers are also better off when the markets has more young buyers; total profits go from \$1.76 to \$1.89 billion.

7.3 ACA price-linked discounts with age-specific price caps

My results thus far suggest that the use of age-adjusted vouchers might be preferable to the current ACA scheme (price-linked discounts not tailored by age). However, price-linked discounts have the important advantage of ensuring the final price for the buyer. The government needs to know less about the determinants of insurers' and buyers' decisions, and can avoid the risk of setting vouchers that are either too high or too low. Nevertheless, even if only price-linked discounts are feasible, age adjustments might still be desirable due to age-heterogeneity in demand and cost.

In my last counterfactual I consider this, comparing the ACA scheme (where price caps used to compute discounts are equal to \$1,400 for all ages) to a scheme with price-linked discounts but with premium caps varying by age. Practically, I consider a scheme $\widehat{S}$ such that

$$\widehat{S}^{\tau,y}(b_r) = \max\left\{A^{\tau} \cdot b_r^* - \overline{P}^{\tau,y}, 0\right\}, \tag{24}$$

in which $\overline{P}^{20\text{-}44,y} < \overline{P}^{45\text{-}64,y}$, and b_r^* is the second-cheapest base price of Silver plans in the region. This is then the same scheme S as implemented under the ACA (equation (**??**)), with the only difference being that price caps vary also by age (under S one has $\overline{P}^{20\text{-}44,y} = \overline{P}^{45\text{-}64,y} = \overline{P}^{y}$). Practically, one can also think of this as a scenario in which the government provides — on top of the current subsidy program — additional incentives for the participation of young buyers with different fiscal instruments.

The results of this comparison are reported in Table **??**, where I show how equilibrium outcomes respond to progressively lower price caps for the under 45. Higher generosity of the subsidy scheme for young buyers yields again to large increases in their participation, and a corresponding reduction in average cost and per-buyer subsidy outlays. Yet, gains relative to the ACA scheme are smaller than under age-adjusted vouchers, since markup reductions are lower (from over 20% to less than 10%), a consequence of the additional distortions induced by price-linked discounts.

7.4 Summary of equilibrium comparative statics under different designs

I summarize the above comparisons in Figure **??**, where I express differences between the ACA design and possible alternatives in head-counts (for enrollment) and dollars (for average cost, markup, and subsidy).

The figure highlights the different role played by the two aspects of subsidy design that I focused on. On the one hand, using vouchers instead of price-linked subsidies increases price competition, and thus lowers markups and the cost for the government of covering a low-income buyer. On the other hand, tailoring discounted prices to age does not imply a redistributive trade-off, but rather is a powerful tool to affect the market average cost, insurers' markups, and per-buyer government expenditure. Importantly, the two mechanisms co-exist and either complement or offset each other under different design alternatives.

8 Conclusions

The recent changes to the US health care system, primarily induced by the 2010 national reform, opened many questions for regulators and economists. Within the growing body of work on regulation of private insurance, in this paper I used an empirically tractable model of imperfect competition between insurers to study (static) equilibrium under different designs of the ACA subsidy program. The applied contribution is two-folded. First, I estimated the model in the post-reform status quo, regulated by the ACA and in which the vast majority of buyers are low-income and thus subsidy eligible. Second, my counterfactuals highlighted that the ACA subsidy scheme leaves room for possible improvements that are quantitatively significant and consistent with theoretical predictions. Additionally, from a methodological perspective, I formalized sufficient conditions to

identify (unobserved) cost primitives in a selection market, relying on the combination of equilibrium assumptions with the variation in prices and choices exploited to identify demand.

In my application I used data from the first year of operations of the Californian marketplace to obtain estimates of demand and cost. Importantly, the buyers in this market belong to a segment of the population that was largely underrepresented in previous studies, and I found a large degree of heterogeneity (both in demand and cost) across buyers of different age. These estimates are the drivers of my counterfactual results. Price-linked subsidies as mandated by the ACA increase insurers' market power, implying higher markups, lower coverage, and higher spending when compared to a mechanism where low-income discounts are not adjusted to prices. I quantified this distortion to be approximately \$200 per-person-year (15%). A second result is that age-adjustments of subsidized premiums within a given income level might lead to better outcomes, both in terms of enrollment levels and efficiency in the use of public funds. This alternative might be easier to implement than exogenous vouchers, and the gains follow directly from the heterogeneity in cost across buyers: Raising the participation of "young invincibles" generates a positive externality on the entire market, reducing costs, prices, and the public spending for every insured buyer. Supporting intuition, my simulations suggested that potential gains are sizable: One can maintain enrollment among buyers who are older than 45 approximately unaltered, increase enrollment among younger buyer by more than 50%, while reducing per-insured public spending by \$600 per-year. With the government intent to increase coverage while limiting spending, a modification of the subsidy scheme to allow for age-specific premium even among low-income buyers could then improve upon the current regulation.

Looking at ACA marketplaces, here I focused on the design of the subsidy program, with a partial equilibrium analysis that holds other parts of the regulations unchanged. Subsidies are indeed only one piece of the large regulatory innovation mandated by the reform, and different rules complement each other in generating market outcomes. Versions of the model I employed here could also be used to study other ACA regulations, such as age-rating restrictions, cost-sharing support, risk-adjustment mechanisms, or tax-penalties for the uninsured.

Pricing decisions can be arguably seen as a natural starting point to analyze insurers' competition in a new institutional context. However, pricing incentives are just one part of the complex puzzle that economists and regulators need to consider. Dynamic incentives of buyers — e.g. choice inertia (**?**) — and insurers — e.g. entry and network-setting decisions (**?**) — could also play a relevant role in determining outcomes in this market. As the marketplaces are now in the third year of operations, in the future it will be possible to use richer models and data to account for these additional dimensions. It is somewhat comforting for my analysis — and necessary in order to take seriously my cost estimates and counterfactuals — that the Californian market has been a very stable environment thus far: Prices and market shares of individual plans did not vary significantly

from year one to year two, and no relevant entry or exit episodes occurred.[37,38] Moreover, recently released data on statewide average per-enrollee claims incurred in 2014 by Anthem Blue Cross — the largest insurer in Covered California — confirmed that my cost estimates matched closely the realized cost of this carrier. The average yearly incurred claim filed by Anthem was $4,632, I estimated an expected average cost for Anthem equal to $4,616.[39] This also provides evidence that the use of supply-side first-order conditions to estimate sellers' costs can be a viable approach to study supply in selection markets for which cost data are not available.

To conclude, the study of the design of the ACA and other government-sponsored insurance programs remains challenging. This defines a rich agenda for future work using multiple years of prices and enrollment, detailed information on networks of providers and utilization, individual-level demographics matched to plan choice, and detailed claims information. This could lead to more precise estimates of relevant primitives and quantifications of the corresponding policy implications.

[37] The autocorrelation in plan-level base prices between 2014 and 2015 was 0.98, while the autocorrelation in plan-level market shares between 2014 and 2015 was 0.96. No entry occurred, and the only insurer leaving the marketplace in 2015 was Contra Costa, the smallest insurer in 2014 with less than 1,000 enrollees.

[38] My modeling assumptions would be harder to justify for those states that in the early years of the ACA marketplaces saw significant churn in prices, market shares, and insurers' participation decisions.

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

APPENDIX

Appendix A. Proof of Proposition ??

I prove part (a), and part (b) follows immediately by simple algebra.
Note that for any τ the probability that a buyer chooses j at the prices P is

$$\sigma_j(P,\tau)=\int_{\{v_j\geq P_j\}\cap\{v_j-v_k\geq P_j-P_k\forall k\}}\mathrm{d}F(v|\tau). \tag{25}$$

Then, when all prices are lowered by Δ for Y buyers:

$$\sigma_j(P-\Delta,Y)=\int_{\{v_j\geq P_j-\Delta\}\cap\{v_j-v_k\geq P_j-P_k\forall k\}}\mathrm{d}F(v|Y)>\sigma_j(P,Y). \tag{26}$$

Symmetrically, when all prices are increased by Δ for O buyers:

$$\sigma_j(P+\Delta,O)=\int_{\{v_j\geq P_j+\Delta\}\cap\{v_j-v_k\geq P_j-P_k\forall k\}}\mathrm{d}F(v|O)<\sigma_j(P,Y). \tag{27}$$

It follows that for all P $\alpha_j^S(P)<\alpha_j^{\widehat{S}}(P)$. From this, and recalling that $C_j^Y<C_j^O$, $AC_j^{\widehat{S}}(P)<AC_j^S(P)$. Moreover, since $\eta_{jj}^Y>\eta_{jj}^O$, the denominator in $MK_j^S(P)$ is smaller than the denominator in $MK_j^{\widehat{S}}(P)$. Then, if for some $\Delta>0$, the numerator in $MK_j^S(P)$ is larger than the numerator in $MK_j^{\widehat{S}}(P)$ — at least in a neighborhood of $P^{*,S}$ —, then the right-hand side of (**??**) is lower under $\widehat{S}$ than under S, and since prices are strategic complements **?** implies $P^{*,\widehat{S}}<P^{*,S}$. To obtain the result it is then sufficient to consider the equilibrium point $P^{*,S}$, and setting $\widetilde{\alpha}_j^{*,S}(\Delta)\equiv\alpha_j^{\widehat{S}}\left(P^{*,S}\right)$ check that the function

$$\Phi_j(\Delta)=1-\left(C_j^Y-C_j^O\right)\left(\eta_{jj}^O-\eta_{jj}^Y\right)\widetilde{\alpha}_j^{*,S}(\Delta)\left(1-\widetilde{\alpha}_j^{*,S}(\Delta)\right)$$

is decreasing in $\Delta\in\left[0,\overline{\Delta}\right]$ for some $\overline{\Delta}>0$. Since $\widetilde{\alpha}_j^{*,S}(0)<\frac{1}{2}$, $\frac{\mathrm{d}\widetilde{\alpha}_j^{*,S}(0)}{\mathrm{d}\Delta}>0$, one has

$$\Phi_j'(0)=-\left(C_j^Y-C_j^O\right)\left(\eta_{jj}^O-\eta_{jj}^Y\right)\frac{\mathrm{d}\widetilde{\alpha}_j^{*,S}(0)}{\mathrm{d}\Delta}<0,$$

and the result follows.■

Appendix B. Nonparametric identification

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.[40]

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 **?**).

B1. Model and observables

I start by adopting the model of demand used in **?** (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 $\delta_{jr} = x^{(1)}_{jr} + \xi_{jr}$, $f_r(v) = f(v|\delta_r, x^{(2)}_r)$.

[40] This is the main difference between my approach and existing work on identification of demand and cost in selection markets, where the observability of costs (e.g. ex-post claims) has been assumed (c.f. **??**, and many others). One exception is **?**, who like me estimates costs using equilibrium pricing conditions. While our estimators are similar, my point here is to formalize which variation is sufficient for identification.

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

$$s_{jr} = \sigma_j(\chi_r) = \int_{\mathcal{D}_j(p_r)} f(v|\delta_r, x_r^{(2)})\,\mathrm{d}v, \;\; j = 0, 1, ..., J, \tag{28}$$

$$\mathcal{D}_j(p_r) = \{v : v_j - v_k \geq p_j - p_k, \text{ for all } k \neq j\}. \tag{29}$$

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}.[41]

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

$$\Pi_{jr}(\chi_r) = p_{jr} \cdot \sigma_j(\chi_r) - \int_{\mathcal{D}_j(p_r)} \psi_j(v, w_{jr}) \cdot f(v|\delta_r, x_r^{(2)})\,\mathrm{d}v. \tag{30}$$

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

$$\partial\mathcal{D}_j(p_r) = \{v : v_j - v_k = p_{jr} - p_{kr} \text{ for some } k \neq j\} \tag{31}$$

$$= \lim_{\varepsilon \downarrow 0} \left\{\mathcal{D}_j(p_r) \cap \left(\mathbb{R}^J \setminus \mathcal{D}_j(p_{jr} + \varepsilon, p_{-jr})\right)\right\}. \tag{32}$$

Then, following **??**, 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

$$mr_{jr} = \sigma_j(\chi_r) - p_{jr} \cdot \int_{\partial\mathcal{D}_j(p_r)} f(v|\delta_r, x_r^{(2)})\,\mathrm{d}v, \tag{33}$$

$$mc_{jr} = -\int_{\partial\mathcal{D}_j(p_r)} \psi_j(v, w_{jr}) \cdot f(v|\delta_r, x_r^{(2)})\,\mathrm{d}v. \tag{34}$$

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.

[41] As in Section **??**, ψ and $f(v)$ could be derived from a more primitive joint distribution $h(v, c)$ over individual preferences and cost, as shown in equations (**??**) and (**??**). If this joint distribution admits a continuous density, the resulting ψ_j is continuous in v, which here is a maintained assumption.

B2. Conditions for identification

Identification is defined as in **??**: 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 **??**.[42] 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 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. In supplementary Appendix S3 I discuss conditions for identification that do not require C3, hence — although imposing stronger restrictions on the model — are more operational in many applications.

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 ??. Condition C3 implies supp $v_r|\delta_r \subset$ supp $p_r|\delta_r$, and demand is identified:

LEMMA **1** (**?**) *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 $p_r|\delta_r$.*

[42] 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.

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

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

$$\Psi_j(\widehat{v}; w_{jr}, \delta_r) = \int_{\mathcal{D}_j(\widehat{v})} \psi_j(v, w_{jr}) \cdot f(v|\delta_r)\, \mathrm{d}v. \tag{35}$$

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

$$\frac{\mathrm{d}^{J-1}\Psi_j(\widehat{v}; w_{jr}, \delta_r)}{\mathrm{d}\widehat{v}_{-j}} = \psi_j(\widehat{v}, w_{jr}) \cdot f(\widehat{v}|\delta_r) \tag{36}$$

and $f(\widehat{v}|\delta_r)$ is identified by Lemma **??**. 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 (**??**); see also Figure 1 in BH.

To approximate $\Psi_j(\widehat{v}; w_{jr}, \delta_r)$, fix j, w_{jr}, and $\widehat{v} \in \text{supp } v_r|\delta_r, w_{jr}$. 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 **??**) yields

$$\frac{\mathrm{d}\widehat{\Psi}_j(\ell)}{\mathrm{d}\ell} = -\int_{\partial\mathcal{D}_j((\eta(\ell),\widehat{v}_{-j}))} \psi_j(v, w_{jr}) \cdot f(v|\delta_r)\, \mathrm{d}v. \tag{37}$$

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,

$$\Psi_j(\widehat{v}; w_{jr}, \delta_r) = \widehat{\Psi}_j(0) = -\int_0^T \phi_j(\ell)\, \mathrm{d}\ell. \tag{38}$$

The integral in (**??**) 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 $\{\chi_r^n\}_{n=0}^N \in \text{supp } \chi_r|\delta_r, w_{jr}$, such that $p_r^n = (\eta(\ell^n), \widehat{v}_{-j})$. Then, as $\max_n\{\ell^n - \ell^{n-1}\}$ becomes arbitrarily small

$$\sum_{n=0}^{N-1} \phi_j(\ell^n)(\ell^{n+1} - \ell^n) \approx \int_0^T \phi_j(\ell)\, \mathrm{d}\ell, \tag{39}$$

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

$$mr_{jr}^n = \sigma_j(\chi_r^n) - p_{jr}^n \cdot \int_{\partial\mathcal{D}_j(p_r^n)} f(v|\delta_r^n)\, \mathrm{d}v, \tag{40}$$

which is identified by Lemma **??**.■

C Robustness and consistency with external cost data

C1. Robustness of price elasticities of demand to different specifications

Table **??** reports average own-price elasticity by age groups (for low-income buyers) using different demand specifications:

- In column (1) I report elasticities from the baseline specification, corresponding to the last row of Table **??**;
- In column (2) I report elasticities from the same specification, where instead of a control-function I use a standard 2SLS logit-demand for every age-income observed in the data;
- In column (3) I report elasticities obtained from a mixed-logit as in **?**, with random coefficients on price with age-income specific mean. The indirect utility from purchasing j in r for a buyer i with age-income (τ, y) is

$$U^i_{jr} = -(\alpha^i + \alpha^{\tau,y}) \cdot P^{\tau,y}_{jr} + (\beta^{\tau,y})' z^y_{jr} + \delta^{\tau,y}_{n(j)} + \xi^{\tau,y}_{jr} + \epsilon^i_{jr}, \tag{41}$$

  where $\alpha^i \sim \mathcal{N}(0, \sigma^{\tau,y})$, and as in the baseline model ϵ^i_{jr} is *iid* distributed as a standard extreme value type I. The estimated $\sigma^{\tau,y}$ is negligible.
- In column (4) I report elasticities obtained from the joint-estimation (N=2,160) of the six demand system (one for every age-income) via simple OLS with insurer-region-income fixed effects and restricting the coefficients on deductible and out-of-pocket limits to vary only by income. This is meant to capture the unobserved network of providers, varying across regions for the same insurers, but equal within a given insurer-region pair. Formally the specification is:

$$U^i_{jr} = -(\alpha^\tau + \alpha^y) \cdot P^{\tau,y}_{jr} + (\beta^y)' z^y_{jr} + \delta^y_{n(j)r} + \xi^{\tau,y}_{jr} + \epsilon^i_{jr}. \tag{42}$$

- In column (5) I report elasticities (coefficient estimates) obtained from a log-quantity on log-price 2SLS regression, a specification that departs from the multinomial choice model:

$$\ln(q^{\tau,y}_{jr}) = -(\alpha^y + \alpha^\tau) \cdot \ln(P^{\tau,y}_{jr}) + (\beta^y)' z^y_{jr} + \delta_{n(j)} + \delta^y + \gamma \ln(M_r) + \nu^{\tau,y}_{jr}, \tag{43}$$

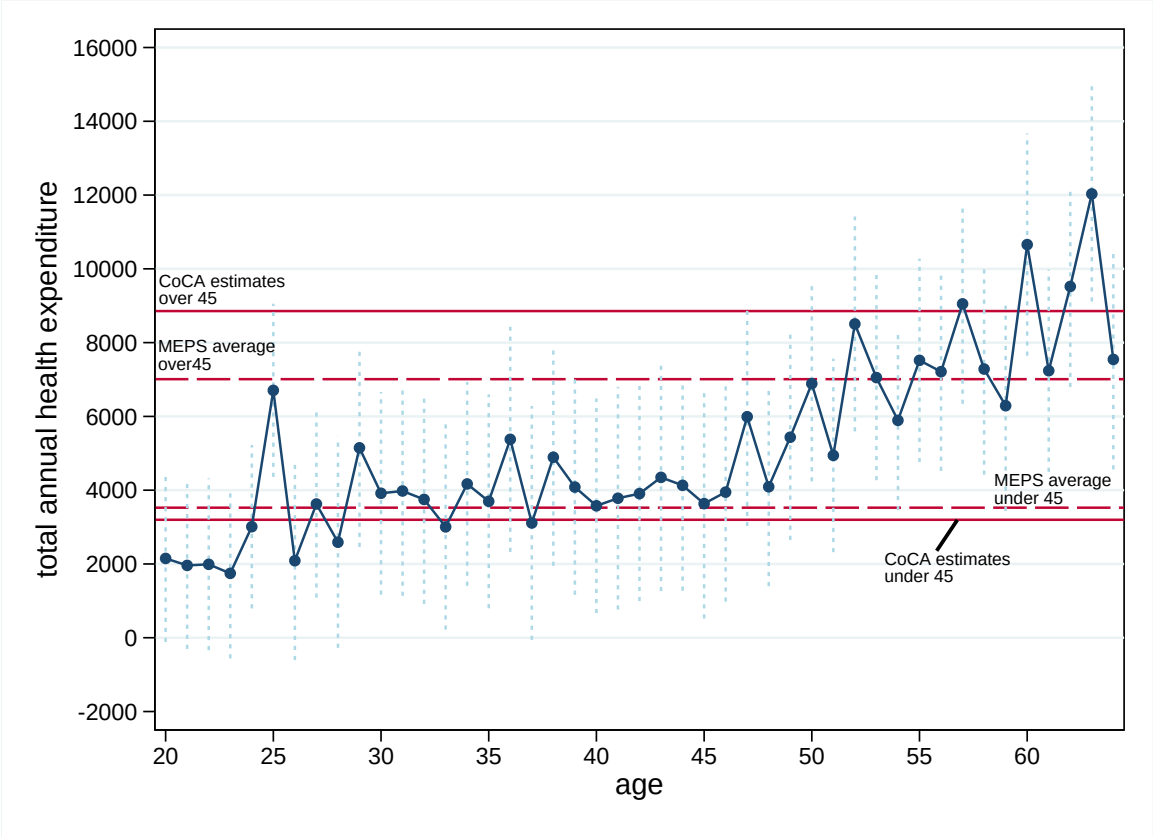
  where the instruments for α^y and α^τ are the age-income composition of the region, and M_r is the total number of potential buyers in the market.
- In columns (6)-(7) I report elasticities from the baseline specification where I deflate (inflate) the total number of potential buyers in all regions — constructed with error from the Area Health Resources file — by 40%.

C2. External cost data

In the figure below I report a comparison between my cost estimates — obtained assuming equilibrium pricing and using the demand estimates from Covered California — to the distribution of annual health expenditure by age in the Medical Expenditure Panel Survey.

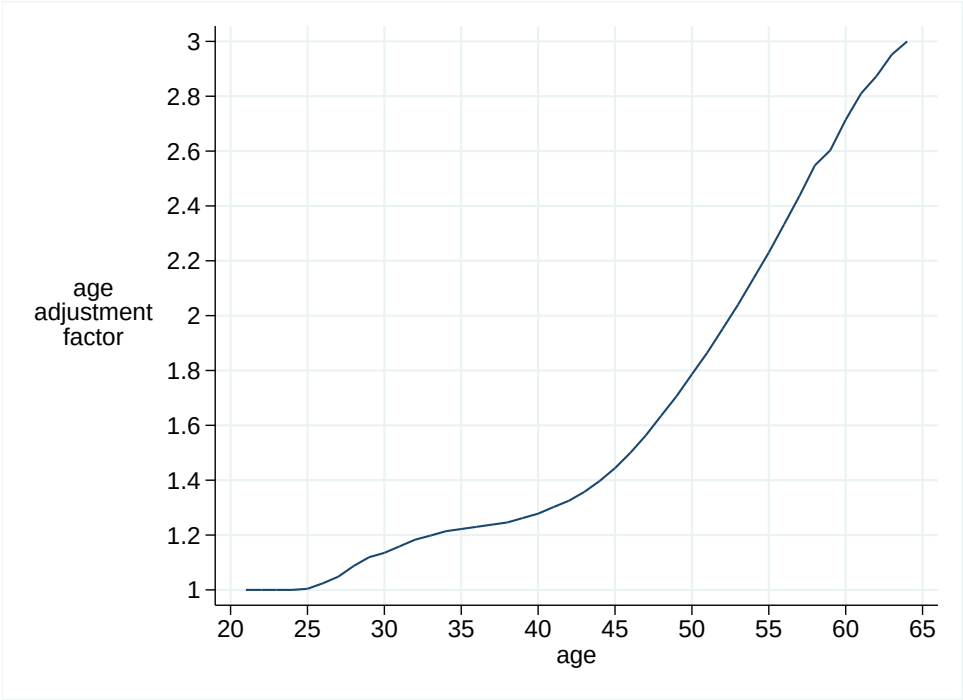
To be more specific, using the survey I regress 2011 annual health expenditure (sum of insured claims and individual contributions) of privately insured buyers between 20-64, with income lower than $44,000, on age fixed-effects. These are plotted with the corresponding 95% confidence intervals, and compared to my cost estimates for the average $0 deductible, Platinum plan.

Annual health expenditure in the MEPS and estimated expected cost



FIGURES AND TABLES

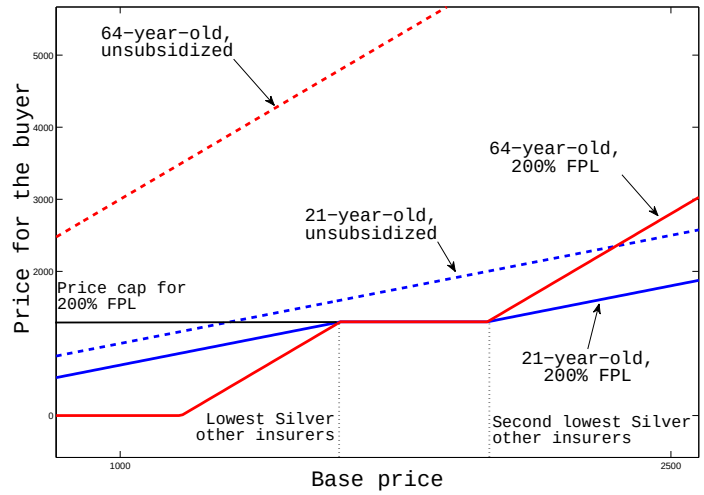
Figure 1: Age adjustment factors in ACA marketplaces



Note: Age adjustment factors under the ACA. 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 (**??**)). 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. This curve was suggested by the Center for Medicare and Medicaid Services, and adopted by 48 states, including California. See also **?**.

Figure 2: Transformation of base price in premium according to ACA regulations.

(a) Silver



(b) Non-Silver

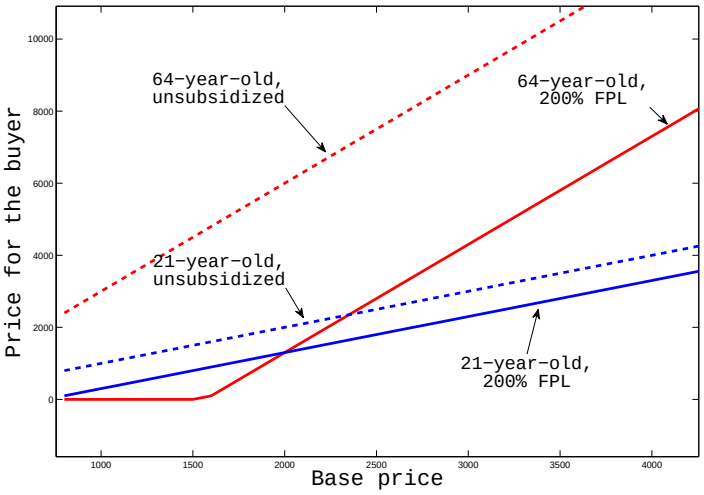
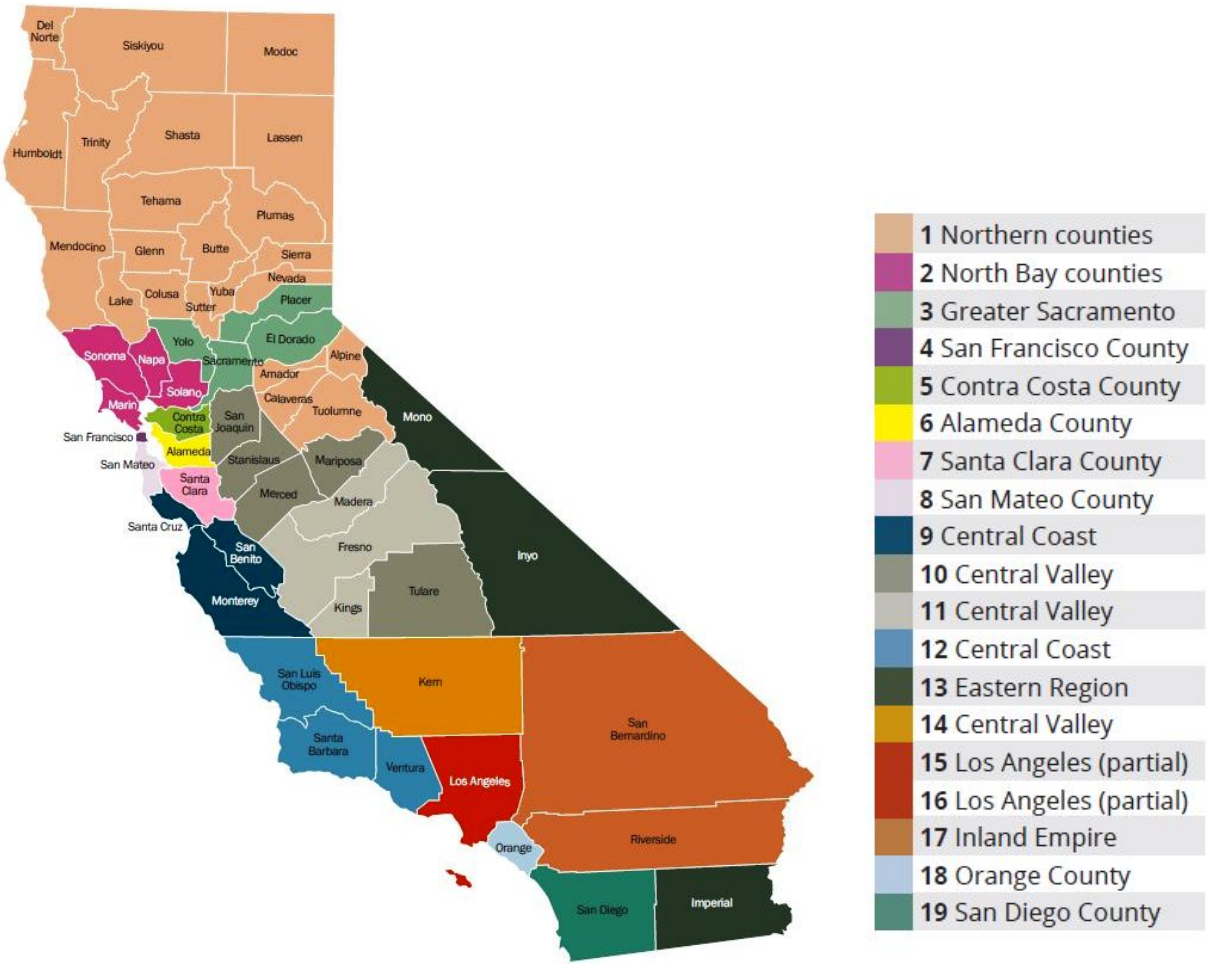


Figure **??** summarizes the transformation of base price in the price paid by 21 and 64-year-old, unsubsidized and subsidized with income at 200% of the FPL.

As the figure highlights, these transformations are different for Silver and non-Silver plans. The second-lowest base price of Silver plans is in fact used to determine subsidies, ensuring that the price paid by the buyer is lower than the price cap as determined by the law (panel (a)). This determines the size of the discounts, which are then applied to all plans (vertical distance between dashed and solid lines in panel (b)).

Note: In both panels: second lowest Silver = \$2,000, lowest Silver = \$1,600, price cap = \$1,300.

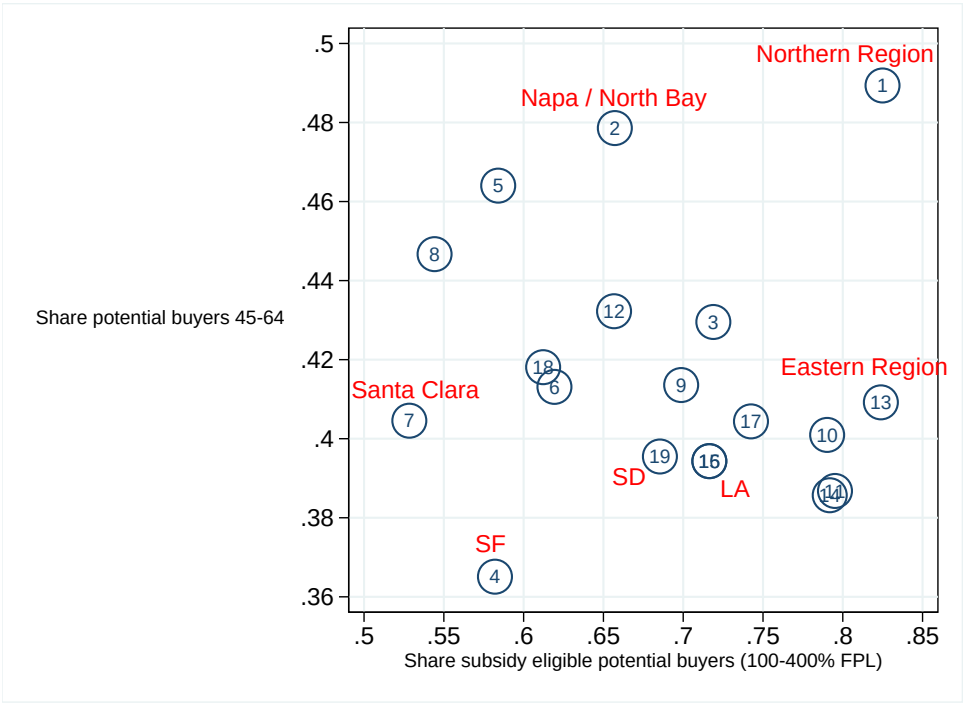
Figure 3: Rating regions in Covered California



Source: www.CoveredCA.com

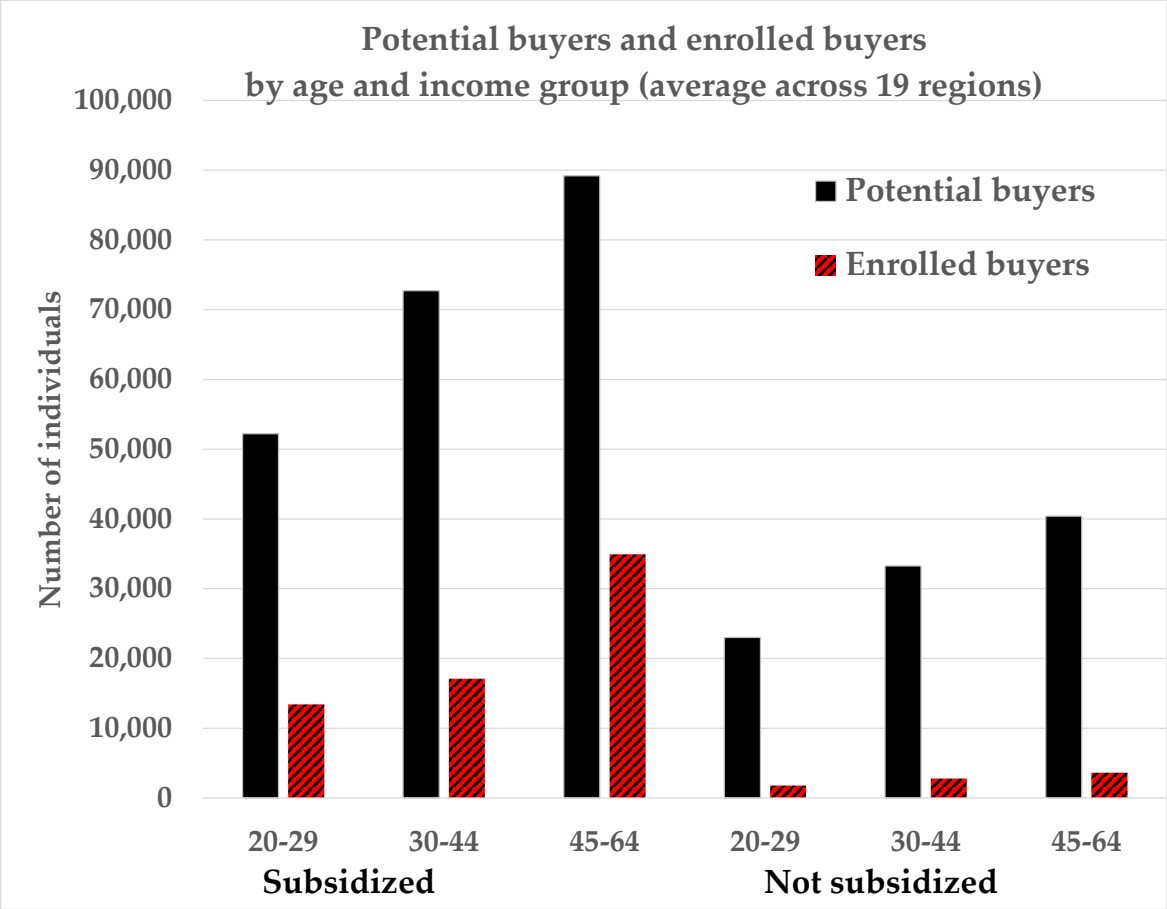
The map shows the 19 rating regions in Covered California. In each region, every spring insurers can announce their participation in the following year's open enrollment. The marketplace needs to authorize entry, and requires the insurer to offer five coverage levels with pre-determined financial characteristics (Table **??**). 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 **??**).

Figure 4: *Share over 45* and *Share subsidy eligible* across regions in Covered California



Note: Each point in the scatter is a rating region in Covered California, with the region number inside the circle. Six regions, including the three main metropolitan areas of Los Angeles (LA), San Francisco (SF), and San Diego (SD), are labeled. This shows the variation in the instruments which exploited to estimate demand. Because of rating regulations and subsidies, insurers set different base prices in markets with different age-income composition. Assuming that unobservables affecting individual choice do not depend on market composition, variation in prices across regions in this scatter, and the corresponding change in market shares, are used to estimate demand parameters.

Figure 5: Average market size and participation by age and income



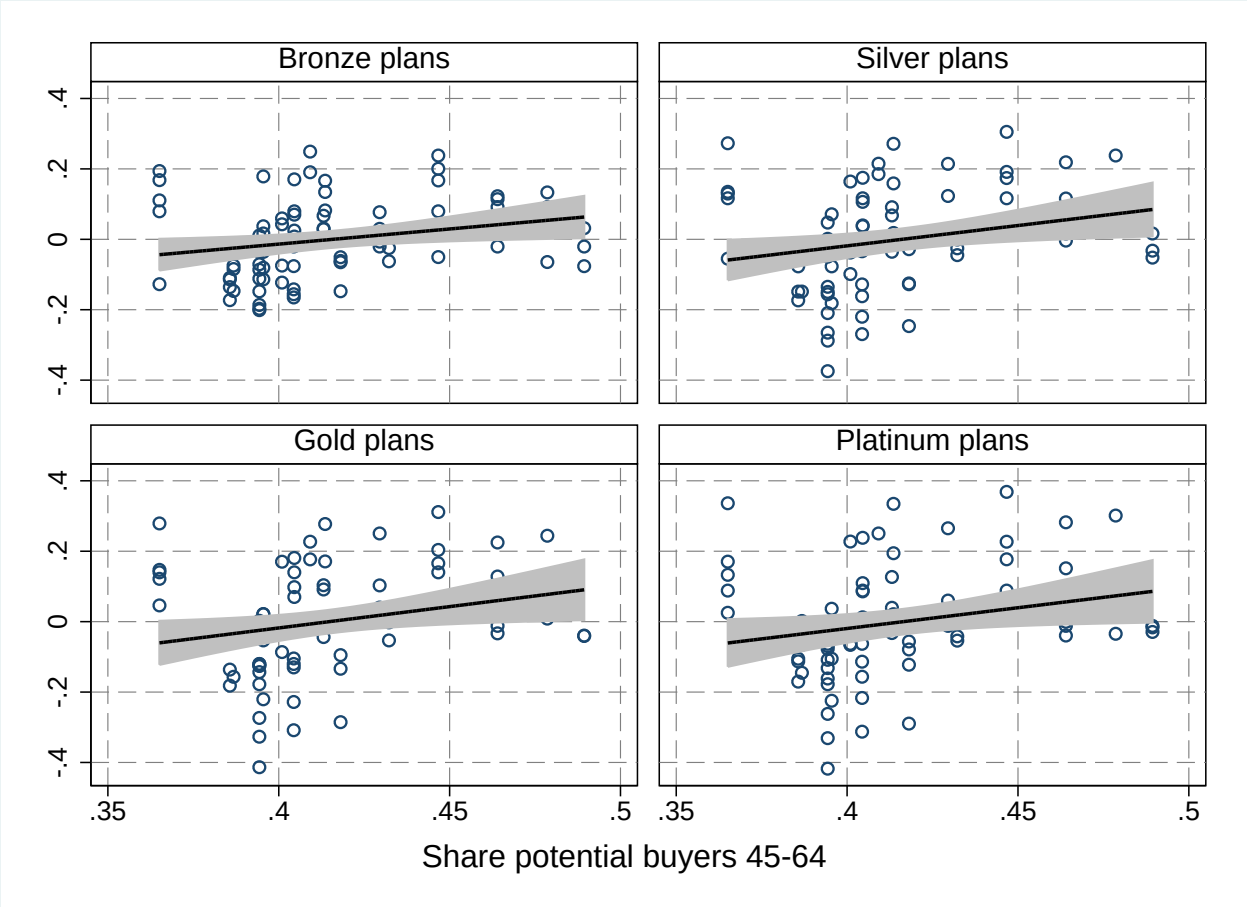
Note: The bars represent the cross-region average number of potential buyers as constructed from the Census (see supplementary Appendix S2) and the cross-region average number of enrollees.

Figure 6: First-stage: base prices vs. share of over 45

OLS residuals of

$$\log(b_{jr}) = \text{ insurer fixed-effect} + \text{tier fixed-effect}$$

plotted against the share of over 45 in the region, one graph for each coverage level



Note: Each panel contains, for the corresponding tier of coverage, the scatter plot of (x-axis) the share of potential buyers older than 45 in the region, and corresponding (y-axis) OLS residuals from the hedonic regression of (log) base prices in Covered California on insurer and tier fixed-effect; N=401.

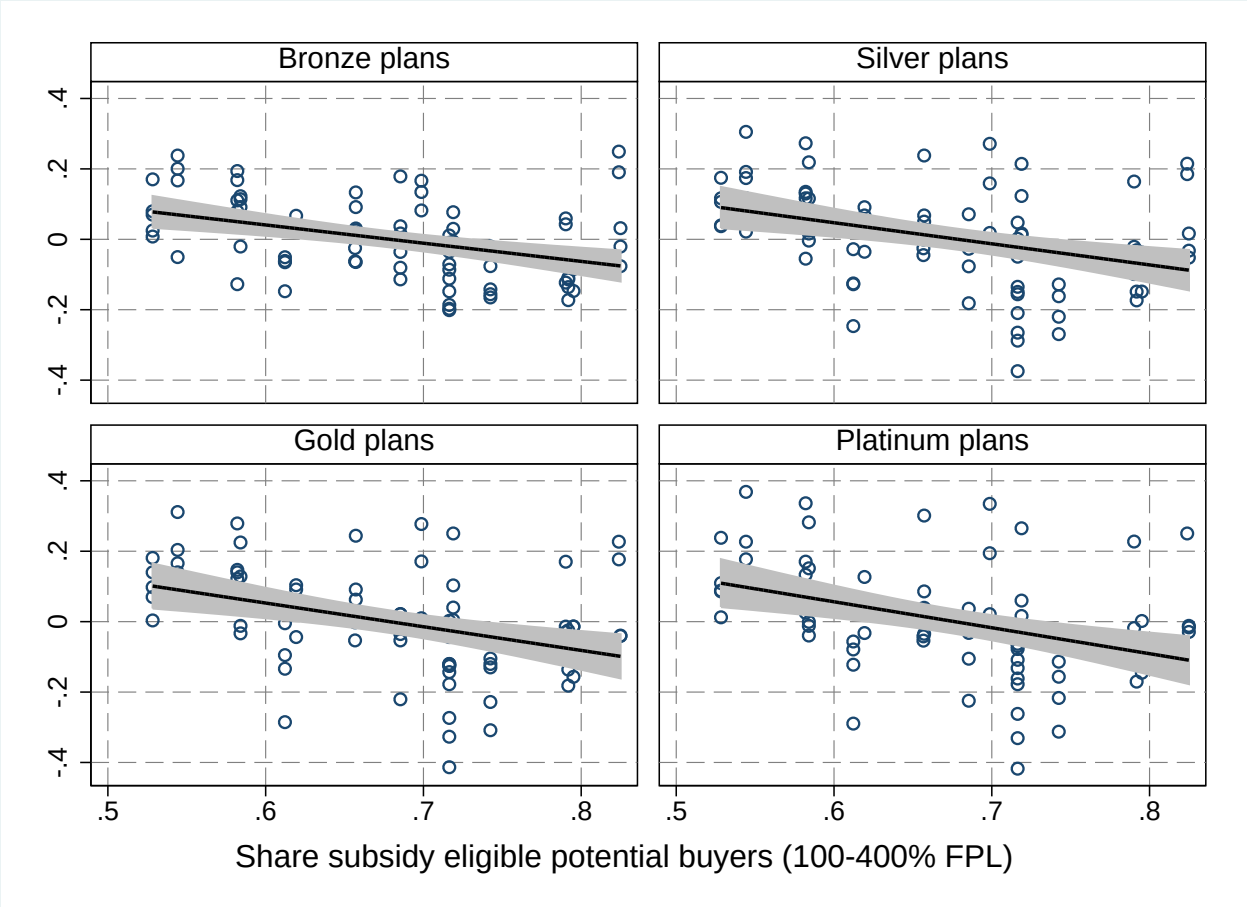
Identification of demand elasticities in the baseline specification uses the variation highlighted in the figure as source of regional variation in prices, maintaining the assumption that the unobserved age-income specific error $\xi_{jr}^{\tau,y}$ in in equation (**??**) does not depend on the share of 45 or older in the market.

Figure 7: First-stage: base prices vs. share subsidy eligible

OLS residuals of

$$\log(b_{jr}) = \text{ insurer fixed-effect} + \text{tier fixed-effect}$$

plotted against the share of subsidy eligible in the region, one graph for each coverage level

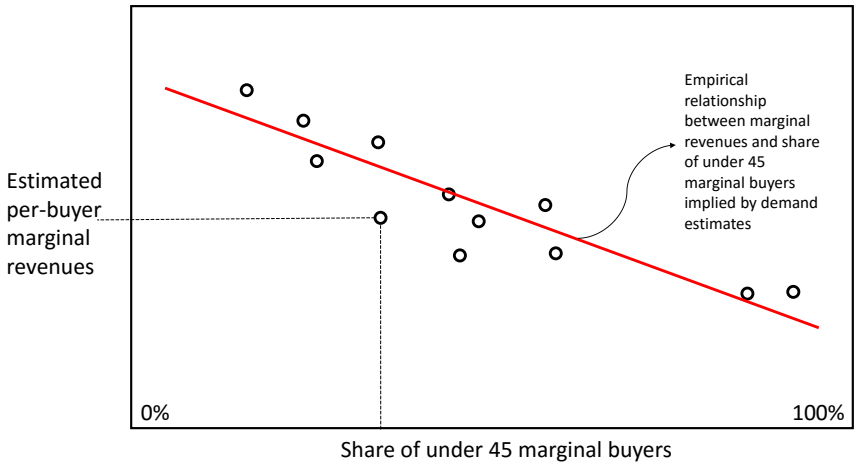


Note: Each panel contains, for the corresponding tier of coverage, the scatter plot of (x-axis) the share of potential buyers who are subsidy eligible, and corresponding (y-axis) OLS residuals from the hedonic regression of (log) base prices in Covered California on insurer and tier fixed-effect; N=401.

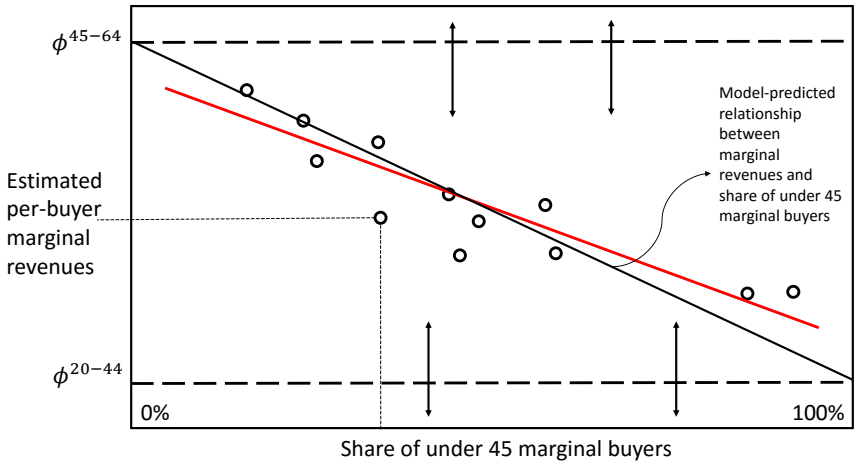
Identification of demand elasticities in the baseline specification uses the variation highlighted in the figure as source of regional variation in prices, maintaining the assumption that the unobserved age-income specific error $\xi_{jr}^{\tau,y}$ in in equation (**??**) does not depend on the share of subsidy eligible in the market.

Figure 8: Intuition for identification of differences in expected cost across age

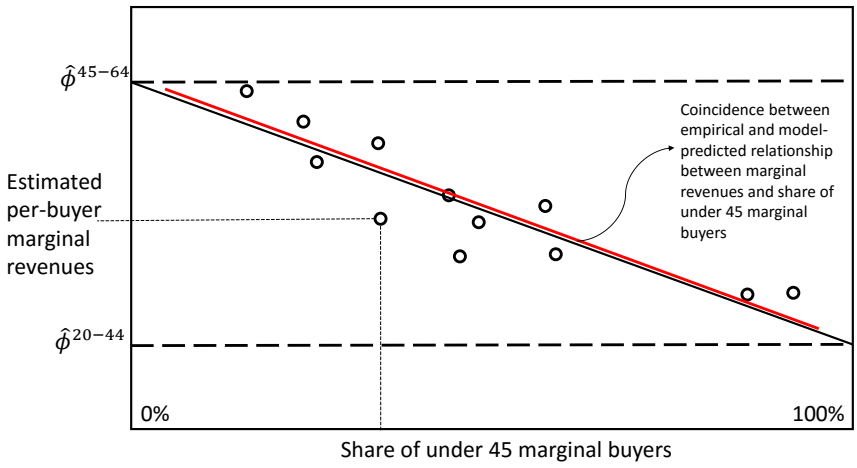
(A)



(B)

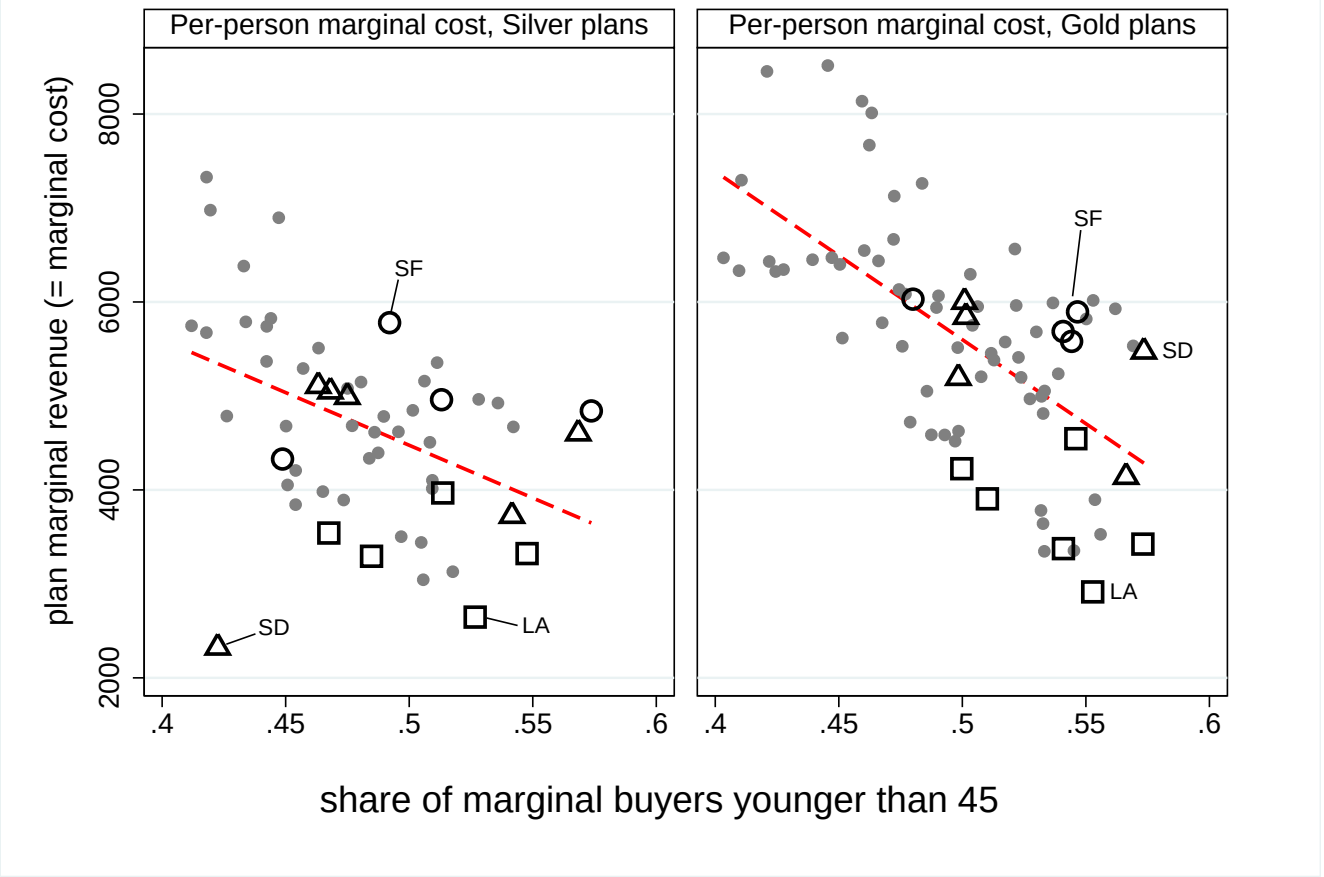


(C)



The figure illustrates an abstract of how variation in the composition of marginal buyers is used to identify cost differences across age. After estimating demand, and controlling for insurer and market differences, I consider the empirical relationship between per-buyer marginal revenues – in equilibrium equal to marginal cost – and the share of marginal buyers younger than 45 (panel (A)). For any pair of parameters $\phi^{20\text{-}44}$ and $\phi^{45\text{-}64}$ the model also implies a relationship between marginal revenues/cost and share of under 45 (panel (B)). Panel (C) shows a combination of $\widehat{\phi}^{20\text{-}44}$ and $\widehat{\phi}^{45\text{-}64}$ for which the model-predicted variation corresponds to the one inferred from demand estimates.

Figure 9: Individual marginal revenue/cost and composition of marginal buyers

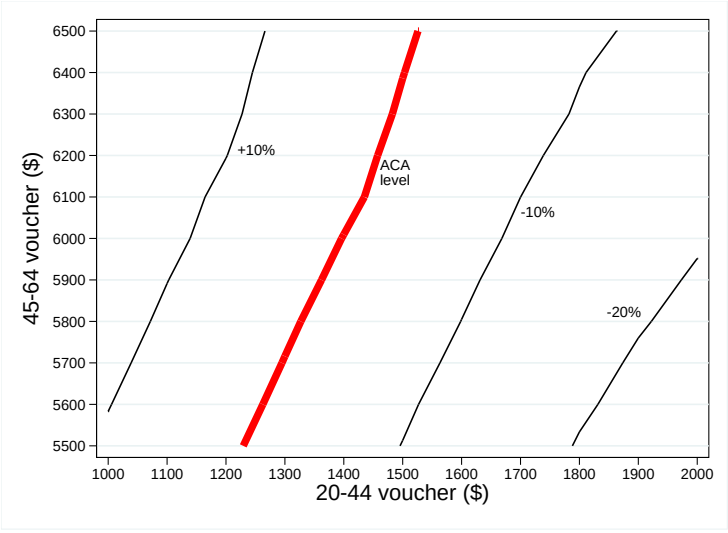


For Silver (left panel) and Gold (right panels) plans, the figure plots the relationship between average per-buyer marginal revenue predicted by the demand model (on the y-axis) and share of marginal buyers younger than 45 (x-axis); each point (or symbol) corresponds to an insurer-region pair in Covered California in the corresponding tier. Different symbols are used for the three urban areas of Los Angeles, San Francisco, and San Diego.

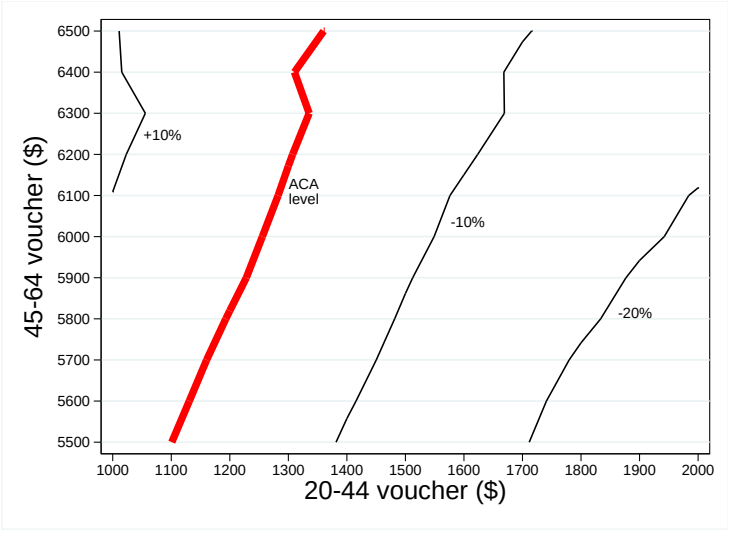
This exemplifies the variation I use to obtain cost estimates after assuming equilibrium, and thus that marginal revenues on the y-axis are equal to marginal cost: The slope of the linear fits leads to the estimated heterogeneity across ages (see also Figure **??**), the vertical distance between the two lines to the variation in contract details, and the variation between regions to the market-fixed effects.

Figure 10: Equilibrium outcomes: ACA scheme vs. age-specific vouchers

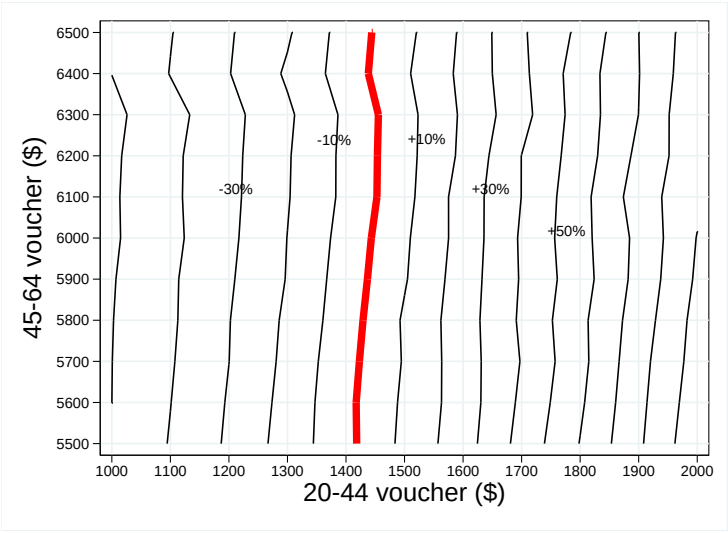
(a) Average cost (ACA level = $4,061)



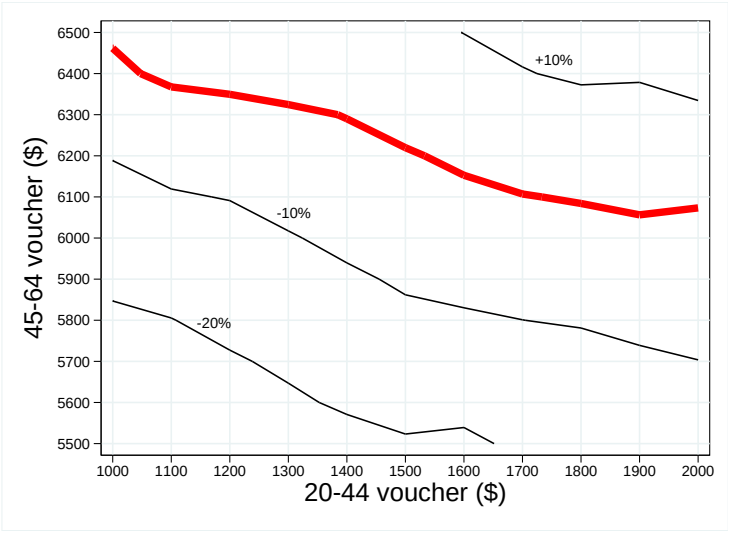
(b) Average premium (ACA level = $5,814)



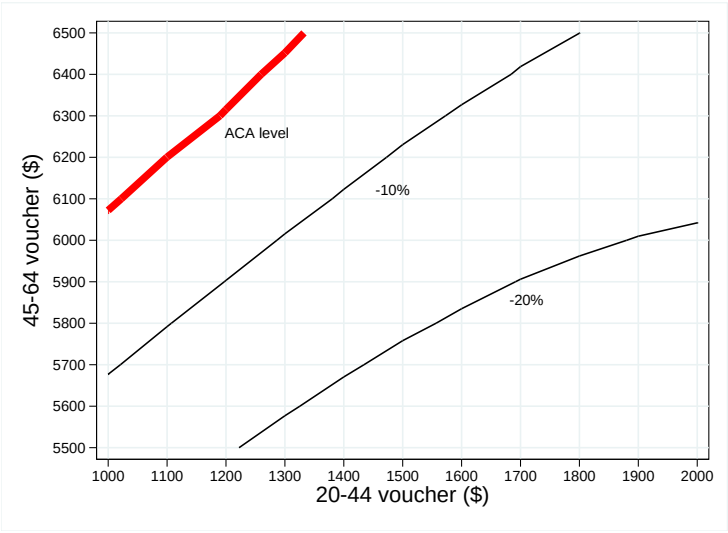
(c) Enrollment 20-44 (ACA level = 460,423)



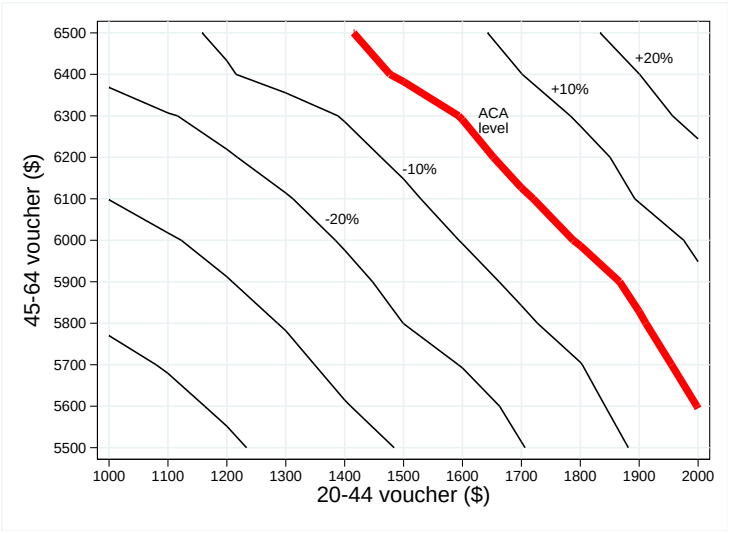
(d) Enrollment 45-64 (ACA level = 543,029)



(e) Average subsidy (ACA level = $3,944)



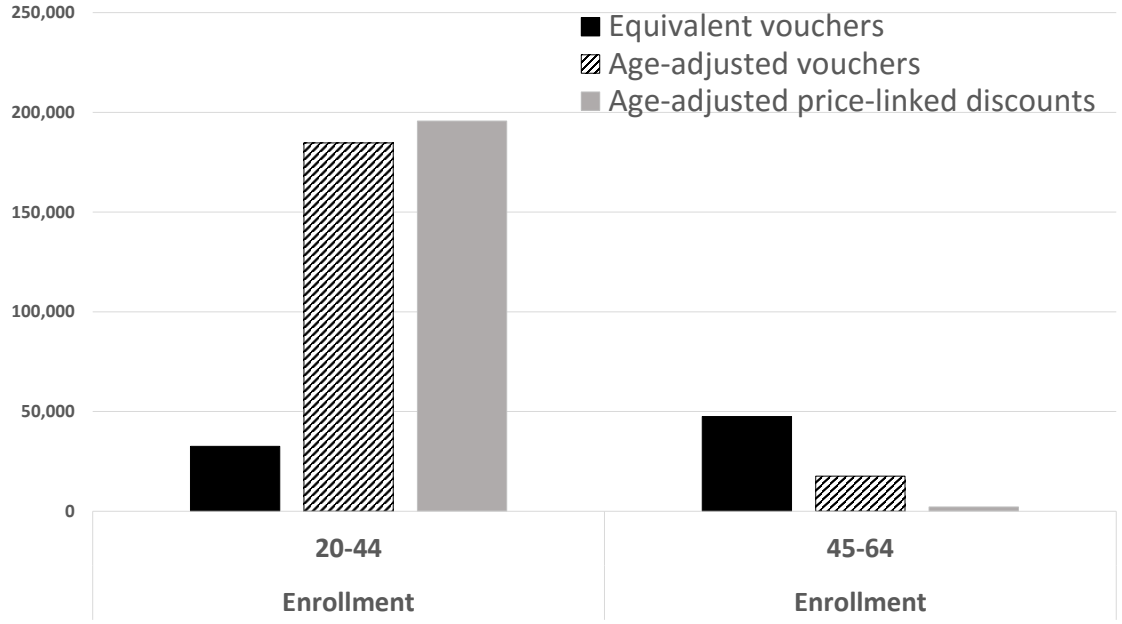
(f) Total subsidy expenditure (ACA level = $4.45 billion)



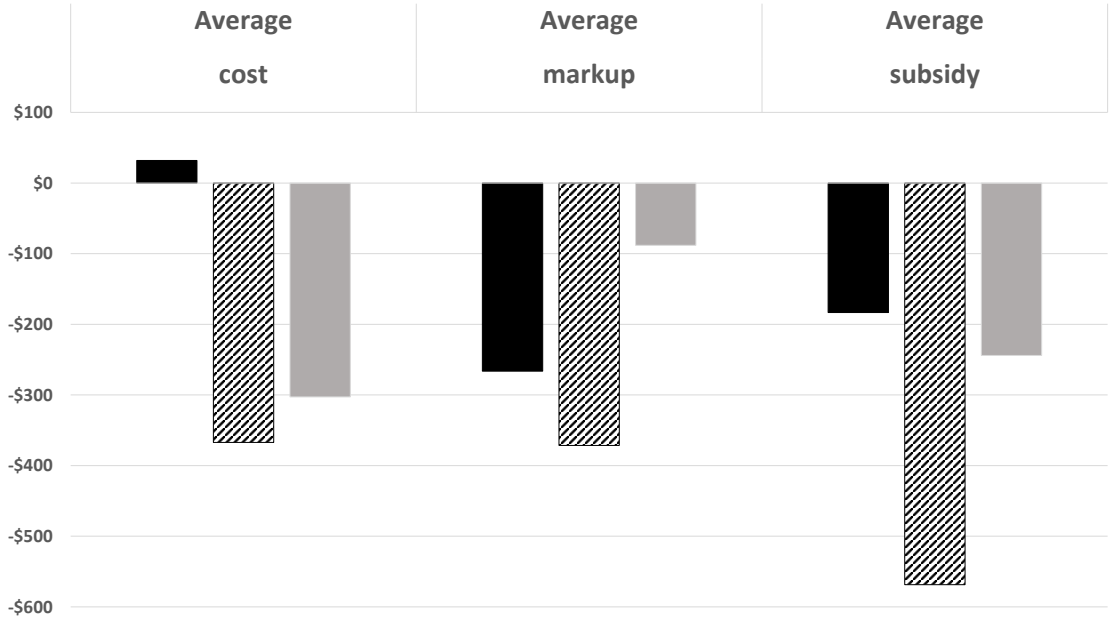
The figure shows level curves of equilibrium outcomes resulting under an age-adjusted voucher as functions of the voucher for the under 45 (x-axis) and voucher for the over 45 (y-axis). The level corresponding to the baseline ACA equilibrium (model fit) is highlighted in red, and every level curve corresponds to a 10% increase (decrease) for that level. The graph is obtained simulating equilibrium base prices — zeroing first-order conditions — as vouchers value vary over the grid in $100 increments.

Figure 11: Summary of differences between ACA scheme and counterfactual alternatives

(a) Differences in enrollment from price-linked discounts without age adjustments



(b) Differences in cost, markup, and subsidy from price-linked discounts without age adjustments



The figure summarizes the comparison between the ACA subsidy design and the three alternatives I consider in my counterfactuals: fixed vouchers with amounts equivalent to the discounts in the equilibrium under the ACA — black bars — , age-adjusted vouchers chosen to be such that over 45 enrollment is 3% higher than under the ACA — crosshatched bars — , and age-adjusted price-linked discounts where the price ceiling on Silver coverage for the under 45 is raised to have the same enrollment as under age-adjusted vouchers — shaded bars.
The bars are derived from Tables **??**, **??**, and **??**.

Table 1: Price caps for subsidy calculation in ACA marketplaces, different levels of income

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	\$684	\$1,452	\$2,416	\$4,368

Note: The table shows, as a function of a buyer's 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.

Table 2: Standardized financial characteristics for the five levels of coverage in *Covered California*

	Annual deductible	Maximum out-of-pocket	Primary care visit	Emergency Room	Specialist visit	Preferred drugs	Advertised coverage*
Minimum	n.a.	\$6,600	n.a.(1)	n.a.(1)	n.a.(1)	n.a.(1)	0%(3)
Bronze	\$5,000	\$6,250	\$60	\$300(2)	\$70(2)	\$50(2)	60%
Silver	\$2,250	\$6,250	\$45	\$250	\$65	\$50	70%
Gold	\$0	\$6,250	\$30	\$250	\$50	\$50	80%
Platinum	\$0	\$4,000	\$20	\$150	\$40	\$15	90%

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: Average premium across plans in *Covered California*

	Subsidized (200% FPL)			Unsubsidized		
	20-29	30-44	54-64	20-29	30-44	54-64
Minimum coverage	n.a.	n.a.	n.a.	2,105	2,751	6,349
Bronze	798	607	82	2,342	3,045	7,027
Silver	1,512	1,530	1,642	3,068	3,989	9,195
Gold	2,105	2,306	3,394	3,647	4,752	10,963
Platinum	2,603	2,949	4,872	4,147	5,405	12,461

Note: Minimum coverage is not available to subsidized buyers. Average premiums are annual, and computed as a simple average across plans with at least 5 enrollees in 2014, using the corresponding baseline rates. All subsidized buyers are treated as if their income was 200% FPL, and age adjustments are: $A^{[20,29]} = 1$, $A^{[30,44]} = 1.3$, $A^{[45,64]} = 3$.

Table 4: Cross-region variation: potential buyers, number of insurers, and 2^{nd} lowest Silver premium

	Mean	St. Dev.	Min	Max
Total potential buyers	310,813	293,403	33,803	1,001,100
Share potential buyers 45-64	0.417	0.033	0.365	0.489
Share subsidy eligible potential buyers (100-400% FPL)	0.689	0.093	0.528	0.825
Number of insurers	4.316	1.108	3	6
Base price of benchmark plan	3,044	409	2,338	3,690
Observations	19			

Note: Summary statistics across the 19 regions of *Covered California*. Composition of potential buyers is constructed from marketplace estimates and Census data (see supplementary Appendix S2). Number of insurers and 2014 price of second lowest Silver plan for the 21-year-old unsubsidized are directly observed in the data provided by the marketplace.

Table 5: Average enrollment across plans in *Covered California*

	Subsidized			Unsubsidized		
	20-29	30-44	54-64	20-29	30-44	54-64
Minimum coverage	n.a.	n.a.	n.a.	58	89	103
Bronze	772	982	1,932	160	254	303
Silver	2,084	2,623	5,235	121	193	221
Gold	164	204	402	54	86	97
Platinum	130	162	309	66	105	116

Note: The table reports the simple average of total enrollment in 2014 across plans in each level of coverage, split across the 6 age-income groups. Minimum coverage is not available to subsidized buyers. Subsidized buyers with income lower than 250% FPL also receive cost-sharing support, by which the deductible and co-pays of Silver plans are partially covered by the government; see Section **??** and http://www.coveredca.com/PDFs/2015-Health-Benefits-Table.pdf.

Table 6: Cross-region variation: average market shares for the 11 insurers in *Covered California*

		Market share : Subsidized buyers				Market share : Unsubsidized buyers			
Carrier name	N. regions	Mean	St. Dev.	Min	Max	Mean	St. Dev.	Min	Max
Anthem	19	0.363	0.229	0.070	0.909	0.426	0.211	0.100	0.908
Blue Shield	19	0.292	0.131	0.086	0.525	0.226	0.074	0.086	0.352
Kaiser	18	0.209	0.164	0.005	0.511	0.213	0.121	0.006	0.444
HealthNet	13	0.147	0.155	0.010	0.354	0.159	0.066	0.037	0.241
Molina	4	0.021	0.028	0.003	0.063	0.022	0.034	0.001	0.073
Chinese C.H.	2	0.223	0.151	0.116	0.330	0.052	0.025	0.034	0.070
LA Care	2	0.086	0.029	0.066	0.106	0.144	0.019	0.131	0.158
Western	2	0.029	0.017	0.017	0.041	0.075	0.033	0.052	0.098
Sharp	1	0.090	0.000	0.090	0.090	0.224	0.000	0.224	0.224
Valley	1	0.028	0.000	0.028	0.028	0.034	0.000	0.034	0.034
Conta costa	1	0.027	0.000	0.027	0.027	0.029	0.000	0.029	0.029

Note: Summary statistics of 2014 market shares for each carrier among unsubsidized and subsidized buyers. The statistics are taken across the regions in which the insurer participates.

Table 7: First-stage: OLS of baseline rates on product characteristics and population composition

	Baseline rate		Log-baseline	
	(1)	(2)	(3)	(4)
Share 45-64	**3317.678**	**2432.069**	**1.132**	**0.830**
	(593.650)	(615.024)	(0.177)	(0.180)
Share 100-400% FPL	**-2064.92**	**-2057.36**	**-0.663**	**-0.657**
	(258.363)	(267.766)	(0.079)	(0.081)
Annual deductible ($1,000)	**-248.892**	**-248.483**	**-0.088**	**-0.088**
	(9.921)	(9.335)	(0.003)	(0.003)
Maximum Out-of-Pocket ($1,000)	**-215.229**	**-212.666**	**-0.049**	**-0.048**
	(37.092)	(35.261)	(0.010)	(0.009)
Anthem		**5419.657**		**8.614**
		(394.33)		(0.112)
Blue Shield		**5371.63**		**8.620**
		(395.29)		(0.113)
Chinese Comm. Health		**5326.933**		**8.605**
		(376.98)		(0.107)
Conta costa		**5177.109**		**8.556**
		(405.32)		(0.118)
HealthNet		**5366.38**		**8.579**
		(414.97)		(0.116)
Kaiser Permanente		**5478.72**		**8.652**
		(393.26)		(0.112)
LA Care		**4756.582**		**8.387**
		(398.86)		(0.111)
Molina		**4923.06**		**8.456**
		(404.74)		(0.116)
Sharp		**5179.485**		**8.548**
		(388.14)		(0.112)
Valley		**5158.483**		**8.545**
		(366.06)		(0.105)
Western		**5625.458**		**8.679**
		(427.04)		(0.122)
Constant	**5023.521**		**8.489**	
	(402.59)		(0.114)	
Observations (insurer-tier-region)	391	391	391	391
Adjusted R^2	0.7586	0.9821	0.7857	0.9997
First-stage partial R^2	0.1990	0.1775	0.2205	0.2016
F-statistic	56.16	41.79	67.36	49.50

Note: OLS regressions of base price, in levels (columns (1) and (2)) and taking the natural logarithm of the dependent variable (columns (3) and (4)). Column (2) and (4) include insurer fixed effects, which are also reported. Robust standard errors in parentheses. All bold coefficients are significant at the 1% level. The first-stage partial R^2 is the ratio between the reduction in sum of square residuals obtained by including *Share 45-64* and *Share 100-400% FPL* and the total sum of square residuals when these variables are omitted. The F-statistic is for the joint Wald test with null hypothesis being that the coefficients on *Share 45-64* and *Share 100-400% FPL* are both zero.

Table 8: Demand and elasticity estimates, parameters of equation (**??**)

	Subsidized			Unsubsidized		
	20-29	30-44	45-64	20-29	30-44	45-64
Product characteristics:						
Annual premium ($1,000)	**-1.781**	**-1.417**	**-0.605**	**-0.315**	**-0.338**	**-0.118**
	(0.169)	(0.124)	(0.064)	(0.121)	(0.103)	(0.040)
Annual Deductible ($1,000)	-0.097	*-0.103*	-0.018	-0.051	*-0.099*	-0.069
	(0.062)	(0.059)	(0.068)	(0.046)	(0.050)	(0.046)
Maximum Out-of-Pocket ($1,000)	**-0.344**	**-0.350**	**-0.368**	0.024	0.008	0.034
	(0.044)	(0.043)	(0.053)	(0.059)	(0.064)	(0.061)
Control function	0.159	0.123	*0.378*	-0.086	0.062	0.056
	(0.159)	(0.147)	(0.201)	(0.155)	(0.164)	(0.159)
Insurer fixed effects:						
Anthem	0.522	-0.093	*-0.504*	**-4.294**	**-3.706**	**-4.099**
	(0.357)	(0.303)	(0.261)	(0.615)	(0.679)	(0.629)
Blue Shield	0.342	-0.286	**-0.737**	**-4.949**	**-4.445**	**-4.822**
	(0.341)	(0.286)	(0.244)	(0.572)	(0.634)	(0.580)
Chinese Comm. Health	0.240	**-0.975**	-0.778	**-5.950**	**-5.428**	**-6.228**
	(0.369)	(0.485)	(0.547)	(0.829)	(0.898)	(0.848)
Conta costa	**-0.760**	**-1.731**	**-2.363**	**-6.437**	**-6.063**	**-6.416**
	(0.350)	(0.293)	(0.257)	(0.707)	(0.770)	(0.705)
HealthNet	-0.591	**-1.272**	**-1.763**	**-5.136**	**-4.562**	**-5.055**
	(0.381)	(0.329)	(0.300)	(0.609)	(0.668)	(0.623)
Kaiser Permanente	0.320	-0.375	**-1.055**	**-4.791**	**-4.245**	**-4.777**
	(0.445)	(0.382)	(0.355)	(0.608)	(0.668)	(0.615)
LA Care	**-1.996**	**-2.671**	**-3.185**	**-6.244**	**-5.615**	**-6.258**
	(0.537)	(0.514)	(0.471)	(0.666)	(0.724)	(0.673)
Molina	**-3.604**	**-4.171**	**-4.876**	**-8.529**	**-7.880**	**-8.498**
	(0.448)	(0.423)	(0.430)	(0.681)	(0.721)	(0.695)
Sharp	*-0.630*	**-1.021**	**-1.498**	**-4.660**	**-3.900**	**-4.300**
	(0.318)	(0.267)	(0.226)	(0.620)	(0.680)	(0.630)
Valley	*-0.859*	**-1.995**	**-1.716**	**-6.351**	**-5.933**	**-6.166**
	(0.346)	(0.300)	(0.272)	(0.668)	(0.734)	(0.676)
Western	**-1.166**	**-1.673**	**-2.376**	**-5.643**	**-5.011**	**-5.488**
	(0.420)	(0.369)	(0.356)	(0.677)	(0.740)	(0.689)
Observations (insurer-tier-region)	321	323	328	391	396	401
Adjusted R^2	0.953	0.962	0.927	0.966	0.961	0.964
Average premium paid	1,328	1,289	1,230	2,694	3,517	8,153
Own-price elasticity (%)	-3.076	-2.586	-1.496	-0.961	-1.342	-1.077
Change in enrollment if all prices increase by $100 (%)	-12.886	-10.744	-3.799	-2.923	-3.125	-1.091

Note: Control function logit demand regressions for each $(age, income)$ group. Robust standard errors in parentheses. Bold coefficients are significant at the 1% level, italics coefficients are significant at the 10% level. The second to last row reports the average of implied product-level own-price elasticity (% drop in enrollment in the plan if the premium increases by 1%). The bottom row reports the average of implied region-level extensive margin, showing the % change in overall enrollment in Covered California plans if all prices for the $(age, income)$ in the region increase by $100.

Table 9: Estimates of expected cost for the insurer, parameters of equation (**??**)

		No buyer heterogeneity			Age heterogeneity					
		(1)	(2)	(3)	(4)		(5)		(6)	
Product characteristics:					Age of buyer		Age of buyer		Age of buyer	
	Maximum				20-44	45-64	20-44	45-64	20-44	45-64
Deductible:	Out-of-pocket:									
$5,000	$6,250	**2,790**	**2,275**	**3,394**	-193	**5,142**	157	**3,888**	202	**5,860**
		[335]	[339]	[648]	[671]	[610]	[644]	[792]	[436]	[601]
$2250[(1)]	$6,250	**4,013**	**4,379**	**4,614**	*1,204*	**6,539**	**2,334**	**6,064**	**1,607**	**7,265**
		[288]	[249]	[350]	[703]	[609]	[714]	[691]	[533]	[461]
$0	$6,250	**4,751**	**9,230**	**5,286**	*1,727*	**7,062**	**6,786**	**10,517**	**2,458**	**8,116**
		[573]	[785]	[965]	[785]	[735]	[1,104]	[1,013]	[640]	[576]
$0	$4,000	**4,227**	**2,387**	**6,131**	*1,704*	**7,040**	656	**4,387**	**3,198**	**8,855**
		[478]	[597]	[872]	[823]	[699]	[866]	[980]	[564]	[588]
Insurer, difference from Anthem (omitted):										
Blue Shield			**-373**	-54			*-321*		-14	
			[180]	[157]			[170]		[88]	
HealthNet			**-1,289**	**-859**			**-1,048**		*-516*	
			[170]	[276]			[188]		[231]	
Kaiser Permanente			194	*371*			**384**		**635**	
			[165]	[189]			[136]		[117]	
Other, minor			111	-187			207		-174	
			[279]	[171]			[268]		[210]	
Market FE:										
1: Northern counties				-204					-311	
2: North Bay				169					200	
3: Sacramento				548					816	
4: San Francisco				-622					35	
5: Contra Costa				27					160	
6: Alameda				-852					-177	
7: Santa Clara				-311					279	
8: San Mateo				317					459	
9: Central coast 1				-70					588	
10: Central valley 1				-784					-109	
11: Central valley 2				-1,160					-293	
12: Central coast 2				-572					-224	
13: Eastern region				1,164					1,701	
14: Central valley 3				-827					-279	
15: Los Angeles 1				-1,516					-1,135	
16: Los Angeles 2				-1,030					-773	
17: Inland Empire				-893					-585	
18: Orange county				-940					-643	
19: Sand Diego				-204					292	
F-stat on market FE:				55.943					98.193	
(p-value)				(<0.0001)					(<0.0001)	

Note: The table reports the estimates of average per-buyer expected costs for different specifications of equation (**??**). Columns (1)-(3) estimate cost varying only by contract type, seller, and region, but constant across buyers. Columns (4)-(6) allow cost to vary also by age. Standard errors are block-bootstrapped drawing regions, and repeating the entire 2-step estimation procedure; 200 draws. The F-statistic on the market fixed effects tests the null hypothesis that these are all equal to zero.

Table 10: Comparison between ACA price-linked discount and vouchers

	Enrollment 20-44	Enrollment 45-64	Average cost	Average markup	Per-person subsidy	Total spending
ACA status quo levels: (endeogenous discount)	460,423	543,029	\$4,061	\$1,753	\$3,944	\$4.45 billion
Percentage change under equivalent voucher (1)	+7%	+9%	+1%	-15%	-5%	+3%

Note: Market outcomes in the simulated equilibrium under the ACA (fit of baseline model) and relative change in these outcomes in the equilibrium under a voucher program that provides, in every region, discounts equal to those resulting in equilibrium under the ACA. The average "equivalent" voucher is \$1,500 for the under 45 and \$6,500 for the over 45.

Table 11: Equilibrium outcomes under different age-adjusted vouchers

		Enrollment 20-44	Enrollment 45-64	Average cost	Average markup	Per-person subsidy	Total spending
Vouchers at which 45-64 enrollment ≈ 1.07-1.10 times the ACA level							
$V^{[20,44]}$	$V^{[45,64]}$	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)
$1500^{(1)}$	$6500^{(1)}$	1.07	1.09	1.01	0.85	0.95	1.03
1600	6500	1.22	1.10	0.98	0.83	0.93	1.08
1700	6400	1.38	1.10	0.93	0.79	0.90	1.10
1900	6300	1.70	1.07	0.87	0.79	0.85	1.16
Vouchers at which 45-64 enrollment ≈ 1.02-1.03 times the ACA level							
$V^{[20,44]}$	$V^{[45,64]}$	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)
1500	6300	1.07	1.02	0.99	0.83	0.91	0.95
1700	6200	1.40	1.03	0.91	0.79	0.86	1.03
1800	6200	1.55	1.03	0.88	0.78	0.84	1.07
1900	6100	1.75	1.02	0.84	0.76	0.82	1.10

Note: Market outcomes (relative to the ACA equilibrium in the baseline model) in the simulated equilibrium under alternative pairs of age-adjusted vouchers. Both panels show market outcomes changing as the voucher for the under 45 is raised and the voucher for the over 45 is lowered. The top panel corresponds to a level curve of over 45 enrollment equal to 102-103% of the ACA level, while the bottom panel corresponds to a level curve of over 45 enrollment equal to 107-110% of the ACA level; see also Figure **??**.

Table 12: Equilibrium outcomes under the ACA scheme with lower price caps for under 45

		Enrollment 20-44	Enrollment 45-64	Average cost	Average markup	Per-person subsidy	Total spending
$\overline{P}^{[20,44]}$	$\overline{P}^{[45,64]}$	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)	(ACA =1)
1400	1400	1	1	1	1	1	1
1300	1400	1.13	1	0.97	0.99	0.98	1.04
1200	1400	1.28	1	0.95	0.97	0.96	1.08
1100	1400	1.42	1	0.93	0.95	0.94	1.12
900	1400	1.75	1	0.88	0.91	0.91	1.22
800	1400	1.94	1	0.86	0.91	0.90	1.29

Note: Market outcomes (relative to the ACA equilibrium in the baseline model) in the simulated equilibrium under alternative pairs of age-adjusted price-linked subsidies. The table shows market outcomes changing as the price cap on the second-cheapest Silver plan in the region for the under 45 is progressively lowered.

Table 13: Robustness of average own-price elasticity (subsidized) to demand model and market size

	Baseline	Alternative demand models				Construction of potential buyers	
Age	Logit Control func. (1)	Logit 2SLS (2)	Logit BLP (3)	Logit Fixed effects (4)	log-linear 2SLS (5)	-40% market size (6)	+40% market size (7)
20-29	-3.075	-2.938	-2.980	-2.740	-2.698	-2.985	-3.100
30-44	-2.586	-2.421	-2.495	-2.394	-2.013	-2.557	-2.596
45-64	-1.496	-1.368	-1.461	-1.703	-1.224	-1.429	-1.517

Note: Robustness of own-price demand elasticities for different age groups among low-income, subsidy eligible buyers. Details for each specification can be found in Appendix C1. Column (1) are elasticities in the baseline model (see Table **??**), the model in column (2) uses IV-2SLS instead of control function, the model in column (3) allows a "BLP" random coefficient on price, the model in column (4) estimates coefficients via OLS adding insurer-times-region fixed effects (equation (**??**)), the model in column (5) departs from discrete choice and estimates elasticity regressing log-enrollment on log-premium and other controls (equation (**??**)), columns (6) and (7) show the baseline elasticities when market size in all regions is increased (decreased) by 40%.

SUPPLEMENTARY APPENDIX

FOR ONLINE PUBLICATION ONLY

S1: Supplementary Appendix to Section ??

S1.1. Derivation of equilibrium condition (??)

Given a subsidy design S, rewrite the profit function for product j as

$$\Pi_j^S\left(P_j, P_{-j}\right) = Q_j^S\left(P_j, P_{-j}\right)\left(P_j - AC_j^S\left(P_j, P_{-j}\right)\right), \tag{44}$$

where $AC_j^S(\cdot)$ is defined in (3) and $Q_j^S(P) = G(Y)\,\sigma_j\left(P - S(P,Y), Y\right) + G(O)\,\sigma_j\left(P - S(P,O), O\right)$. If P^* is an equilibrium one has

$$\begin{aligned}\frac{\partial \Pi_j\left(P_j^*, P_{-j}^*\right)}{\partial P_j} &= \frac{\partial Q_j^S\left(P_j^*, P_{-j}^*\right)}{\partial P_j}\left(P_j^* - AC_j^S\left(P_j^*, P_{-j}^*\right)\right) + Q_j^S\left(P_j^*, P_{-j}^*\right)\left(1 - \frac{\partial AC_j^S\left(P_j^*, P_{-j}^*\right)}{\partial P_j}\right)\\ &= 0;\end{aligned}$$

and rearranging terms leads to $P_j^* = AC_j^S\left(P_j^*, P_{-j}^*\right) - \frac{Q_j^S\left(P_j^*, P_{-j}^*\right)}{\frac{\partial Q_j^S\left(P_j^*, P_{-j}^*\right)}{\partial P_j}}\left(1 - \frac{\partial AC_j^S\left(P_j^*, P_{-j}^*\right)}{\partial P_j}\right)$.

To simplify this further, rewrite

$$\begin{aligned}-\frac{Q_j^S\left(P_j^*, P_{-j}^*\right)}{\frac{\partial Q_j^S\left(P_j^*, P_{-j}^*\right)}{\partial P_j}} &= \frac{Q_j^S\left(P_j^*, P_{-j}^*\right)}{G(Y)\,\sigma_j\left(P^* - S(P^*, Y), Y\right)\eta_{jj}^Y + G(Y)\,\sigma_j\left(P^* - S(P^*, O), O\right)\eta_{jj}^O}\\ &= \frac{1}{\alpha_j^S(P^*)\,\eta_{jj}^Y + \left(1 - \alpha_j^S(P^*)\right)\eta_{jj}^O};\end{aligned} \tag{45}$$

$$\begin{aligned}&\frac{\partial AC_j^S\left(P_j^*, P_{-j}^*\right)}{\partial P_j} = \left(C_j^Y - C_j^O\right)\frac{\partial \alpha_j^S(P^*)}{\partial P_j} =\\ &= \left(C_j^Y - C_j^O\right)\frac{\frac{\partial\left(G(Y)\sigma_j\left(P^* - S(P^*,Y),Y\right)\right)}{\partial P_j} Q_j^S(P^*) - G(Y)\,\sigma_j\left(P^* - S(P^*,Y),Y\right)\frac{\partial Q_j^S(P^*)}{\partial P_j}}{\left(Q_j^S(P^*)\right)^2}\\ &= \left(C_j^Y - C_j^O\right)\left(\eta_{jj}^O - \eta_{jj}^Y\right)\alpha_j^S(P^*)\left(1 - \alpha_j^S(P^*)\right).\end{aligned} \tag{46}$$

Putting together (**??**) and (**??**) gives the expression for the equilibrium price in (**??**):

$$P_j^* = AC_j^S\left(P_j^*, P_{-j}^*\right) + \underbrace{\frac{1 - \left(C_j^Y - C_j^O\right)\left(\eta_{jj}^O - \eta_{jj}^Y\right)\alpha_j^S\left(P^*\right)\left(1 - \alpha_j^S\left(P^*\right)\right)}{\alpha_j^S\left(P^*\right)\eta_{jj}^Y + \left(1 - \alpha_j^S\left(P^*\right)\right)\eta_{jj}^O}}_{MK_j^S(P_j^*, P_{-j}^*)}. \tag{47}$$

S1.2. Price-linked discounts vs. vouchers

The following proposition shows that, as long as there is a group of buyers who are more price sensitive and cheaper to cover, given price-linked subsidies it is possible to find a voucher scheme leading to an equilibrium with lower prices.

PROPOSITION **2** *Assume that prices are strategic complements, and let S be a subsidy design for which at any P, $\frac{\partial S(P,\tau)}{\partial P_j} = \delta > 0$, for some j (which may depend on P). Then, if $C_j^Y < C_j^O$, and $\eta_{0j}^Y > \eta_{0j}^O$ for all j, a voucher scheme $\widehat{S}$ such that $\widehat{S}\left(P,\tau\right) = \widehat{V}^\tau = S\left(P^{*,S},\tau\right)$ is such that $P^{*,\widehat{S}} < P^{*,S}$.*

Proof. From the equilibrium comparative static results formalized in **?**, the proof amounts to show that, for all P at which all products make weakly positive profits, setting $\widehat{S}\left(P,\tau\right) = \widehat{V}^\tau = S\left(P,\tau\right)$ implies

$$\frac{\partial \Pi_j^S\left(P_j, P_{-j}\right)}{\partial P_j} \geq \frac{\partial \Pi_j^{\widehat{S}}\left(P_j, P_{-j}\right)}{\partial P_j} \text{ for all } j, \tag{48}$$

with a strict inequality for at least one j. (This immediately implies that $\frac{\partial \Pi_j^{\widehat{S}}\left(P^{*,S}\right)}{\partial P_j} < 0$ for at least one j: At the equilibrium prices under S, if the vouchers $\widehat{S}$ are adopted, at least one insurer wants to lower its price).

To show (**??**), given that at prices P for at least one j, $\frac{\partial S(P,\tau)}{\partial P_j} = \delta > 0$, and for all k, $P_k \geq AC_k^S\left(P\right)$, one has that for any j:

$$\begin{aligned}
\frac{\partial \Pi_j^S\left(P_j, P_{-j}\right)}{\partial P_j} - \frac{\partial \Pi_j^{\widehat{S}}\left(P_j, P_{-j}\right)}{\partial P_j} &\geq \delta\left(\frac{\partial Q_j^S\left(P\right)}{\partial P_0}\left(P_j - AC_j^S\left(P_j, P_{-j}\right)\right) - Q_j^S\left(P\right)\frac{\partial AC_j^S\left(P\right)}{\partial P_0}\right) \\
&= \delta\left(\underbrace{P_j - AC_j^S\left(P_j, P_{-j}\right)}_{>0} - \underbrace{\frac{Q_j^S\left(P\right)}{\frac{\partial Q_j^S(P)}{\partial P_0}}}_{>0}\frac{\partial AC_j^S\left(P\right)}{\partial P_0}\right).
\end{aligned}$$

Hence (**??**) follows if $\frac{\partial AC_j^S(P)}{\partial P_0} < 0$: raising the price of the outside good (or equivalently the subsidy)

lowers average cost. This is indeed the case as long as $C_j^Y < C_j^O$, and $\eta_{0j}^Y > \eta_{0j}^O$:

$$\begin{aligned}
\frac{\partial AC_j^S(P)}{\partial P_0} &= \left(C_j^Y - C_j^O\right)\frac{\partial \alpha_j^S(P)}{\partial P_0} \\
&= \left(C_j^Y - C_j^O\right)\frac{\frac{\partial (G(Y)\sigma_j(P-S(P),Y))}{\partial P_0} Q_j^S(P) - G(Y)\,\sigma_j(P-S(P,Y),Y)\,\frac{\partial Q_j^S(P)}{P_0}}{\left(Q_j^S(P)\right)^2} \\
&= \left(C_j^Y - C_j^O\right)\left(\eta_{0j}^Y - \eta_{0j}^O\right)\left(1-\alpha_j^S(P)\right) < 0.
\end{aligned}$$

Therefore (**??**) holds and, by **?**, $P^{*,\widehat{S}} < P^{*,S}$. ■

S2: Details of data construction

S2.1. Enrollment by age and income

In 2014, Covered California provided me with total enrollment Q_{jr} for every product j (insurer-tier pair) in every region r, for a total of $N = 401$ products in the marketplace.

Additionally, the marketplace published for every income $y = LI, HI$, and every region r:

- $\Pr[\tau, y|r]$: the share of buyers of age τ, $\tau \in \{[20, 29], [30, 44], [45, 64]\}$, and income y in region r;
- $\Pr[n, y|r]$: the share of income y buyers of plans sold by insurer n, in region r;
- $\Pr[x, y|r]$: the share of income y buyers of plans in tier x, in region r.

Therefore, since $j = (n, x)$, I impute the enrollment $q_{jr}^{\tau,y}$ of product j in region r restricted to the age-income group (τ, y) via Bayes rule:

$$q_{jr}^{\tau,y} = Q_{jr} \cdot \Pr[\tau, y|j, r], \tag{49}$$

$$\Pr[\tau, y|j, r] = \Pr[\tau, y|n, x, r] = \frac{\Pr[\tau, y, n, x|r]}{\Pr[n, x|r]} = \frac{\Pr[n, y|r] \Pr[x, y|r] \Pr[\tau, y|r]}{\Pr[n, x|r]}, \tag{50}$$

where

$$\Pr[n, x|r] = \Pr[j|r] = \frac{Q_{jr}}{\sum_k Q_{kr}}. \tag{51}$$

This assumes that conditional on income, tier choice does not depend on the specific insurer, and that both tier and insurer choice do not depend on age.

Importantly, more granular data released in October 2015 (after the second open enrollment, see http://hbex.coveredca.com/data-research/) confirmed that the market shares conditional on age and income resulting from this construction match very closely the choices taken by buyers in the first two years of the marketplace. This data could later be incorporated in my analysis.

S2.2. Market composition

The 2012-2013 Area Health Resources file (http://ahrf.hrsa.gov/) contains, for every county c in the US, the following variables:

- U_c^τ: number of uninsured individuals of a given age τ, $\tau \in \{[20, 29], [30, 44], [45, 64]\}$;
- H_c^y: number of households with income y, $y \in \{[0, 10k), [10k, 15k), [15k, 25k), [25k, 50k), [50k, 100k), [100k, \infty)\}$.

Since subsidy eligibility for a family of two starts at \$16,000, and stops at \$70,000, I approximate the share of uninsured of a given age τ and income $y = LI, HI$ in county c as

$$G_c(\tau, LI) = \frac{U_c^\tau}{\sum_{\tau'} U_c^{\tau'}} \frac{\sum_{y \in [15k, 100k)} H_c^y}{\sum_{y \in [15k, \infty)} H_c^y}, \tag{52}$$

$$G_c(\tau, HI) = \frac{U_c^\tau}{\sum_{\tau'} U_c^{\tau'}} \left(1 - \frac{\sum_{y \in [15k, 100k)} H_c^y}{\sum_{y \in [15k, \infty)} H_c^y} \right). \tag{53}$$

Then, if $\mathcal{C}(r)$ is the set of counties in region r, the share of age-income (τ, y) potential buyers is calculated as the following population-weighted average:

$$G_r(\tau, y) = \frac{\sum_{c \in \mathcal{C}(r)} U_c G_c(\tau, y)}{\sum_{c \in \mathcal{C}(x)} U_c}, \tag{54}$$

where $U_c = \sum_\tau U_c^\tau$.

I obtain similar results when using alternative constructions based on the PUF file of the Californian Health Insurance Survey (http://healthpolicy.ucla.edu/chis/Pages/default.aspx), or using marketplace estimates of the subsidy eligible in each region before the first open enrollment period.

The latter source displays a similar cross-region variation but significantly smaller overall market size. As a result, in several regions enrollment during the first year was greater than the predicted number of eligible buyers, and since then the marketplace removed these predictions from the web. Yet, omitting the regions with negative share of the outside good, demand estimates using this source to determine potential buyers remain similar to the ones in my baseline specification.

Since the income composition measured here does not correspond exactly with subsidy eligibility under the ACA, in one robustness check in Appendix C I inflate or deflate the overall market size in each region by up to 40%, and demand estimates across age-groups are not substantially different. In fact, what is relevant for the validity of my results is that my construction captures correctly the cross-region variation in age-income shares, rather than the exact size of each market.

S3: Identification without a large support assumption

Setup.

- J products are offered in R markets.
- A buyer i in market r is a pair (z_{ir}, v_{ir}). $z_{ir} \in \mathbb{R}^Q$ is a collection of buyer-specific observables (e.g. age, income, gender, zip-code); $z_{ir} \in \mathcal{Z} = \{z^1, ..., z^T\}$, a finite set. $v_{ir} = (v_{i1r}, ..., v_{iJr}) \in \mathcal{V}$ is a vector of (money-metric) valuations collecting the buyer's willingness-to-pay for each of the J products.
- The probability that a buyer has characteristics $\widehat{z} \in \mathcal{Z}$ in market r is $\mu_r(\widehat{z}) \geq 0$; $\sum_{\widehat{z} \in \mathcal{Z}} \mu_r(\widehat{z}) = 1$.
- Conditional on $z_{ir} = \widehat{z}$, in market r v_{ir} is drawn *i.i.d* from the density $f(v_{ir}|z_{ir}, x_r, \xi_r)$, with support $\mathcal{V}$.
  $x_r = (x_{1r}, ..., x_{Jr})$ collects observed characteristics of the J products in r, with each $x_{jr} \in \mathbb{R}^K$.
  $\xi_r = (\xi_{1r}, ..., \xi_{Jr})$ collects unobserved characteristics of the J products in r affecting preferences, with each $\xi_{jr} \in \mathbb{R}$.
- Prices are equal for all buyers in a market, and are collected in the vector $p_r = (p_{1r}, ..., p_{Jr})$, $p_{jr} \in \mathbb{R}$.
  Given prices, the set of valuations of buyers choosing j is $D_j(p_r) = \{v \in \mathcal{V} : v_j - p_j \geq v_k - p_k, \, \forall k\}$.
  Given prices, the set of marginal buyers for j is $\partial D_j(p_r) = \{v \in \mathcal{V} : v_j - p_j = v_k - p_k \text{ for some } k\}$.
- The (expected) cost for the seller when a buyer with $(z_{ir}, v_{ir}) = (\widehat{z}, \widehat{v})$ purchases j in r is $\psi_{jr}(\widehat{z}, \widehat{v})$. The function $\psi_{jr} : \mathcal{Z} \times \mathcal{V} \to \mathbb{R}_+$ is not assumed to be constant, thus this is market with selection.
- Setting $\chi_r = (p_r, x_r, \xi_r)$, the probability that a buyer with characteristics $z_{ir} = \widehat{z}$ chooses j in r is
  $s_{jr}(\widehat{z}) = \sigma_j(\widehat{z}, \chi_r) = \int_{D_j(p_r)} f(\widehat{v}|\widehat{z}, x_r, \xi_r) \, \mathrm{d}\widehat{v}$. Expected profits for the seller of j in r are then
  $$\Pi_{jr}(\chi_r) = \underbrace{p_{jr} \sum_{\widehat{z} \in \mathcal{Z}} \mu_r(\widehat{z}) \sigma_j(\widehat{z}, \chi_r)}_{\text{Revenues}} - \underbrace{\sum_{\widehat{z} \in \mathcal{Z}} \mu_r(\widehat{z}) \int_{D_j(p_r)} \psi_{jr}(\widehat{z}, \widehat{v}) f(\widehat{v}|\widehat{z}, x_r, \xi_r) \, \mathrm{d}\widehat{v}}_{\text{Cost}}.$$

Observables and demand identification. Let $w_r = (w_{1r}, ..., w_{Jr})$ denote product/market-specific cost-shifters excluded from buyers' preferences (e.g. service fees for hospitals and clinics covered by a given j); each $w_{jr} \in \mathbb{R}^L$. The econometrician observes, for all r, all j, and all z, the collection $(\mu_r(z), s_{jr}(z), p_{jr}, x_{jr}, w_{jr})$. **??** provide sufficient conditions under which ξ_r and $f(\cdot|\cdot, x_r, \xi_r)$ are identified. I impose these conditions and treat these demand primitives as known

henceforth.[43]

Cost identification — two types. The remaining unknown primitives — one value for each $\psi_{jr}(\widehat{z},\widehat{v})$ — are more numerous than the observed supply decisions — say N, one value for each price p_{jr}.[44]

I start by considering the simple case in which cost only depends on observable characteristics of the buyer, i.e. $\psi_{jr}(\widehat{z},\widehat{v}) = \psi_{jr}(\widehat{z})$; this reduces the number of unknowns to $N \times T$. To make things even simpler, suppose that $T = 2$, i.e. $\mathcal{Z} = \{z^1, z^2\}$, so that there are only two unknowns $\psi_{jr}(z^1), \psi_{jr}(z^2)$ for each observed pair jr.

The first step is to assume that prices are set optimally by sellers, implying that $mr_{jr} = mc_{jr}$. In this expression: $mr_{jr} = \sum_{\ell=1}^{2} \mu_r(z^\ell) \left(\sigma_j(z^\ell, \chi_r) - p_{jr} \int_{\partial D_j(p_r)} f(\widehat{v}|z^\ell, x_r, \xi_r)\, \mathrm{d}\widehat{v} \right)$, while the right-hand side (marginal cost) is

$$mc_{jr} = \sum_{\ell=1}^{2} \mu_r(z^\ell) \int_{\partial D_j(p_r)} \psi_{jr}(z^\ell,\widehat{v}) f(\widehat{v}|z^\ell, x_r, \xi_r)\, \mathrm{d}\widehat{v} = \sum_{\ell=1}^{2} \psi_{jr}(z^\ell) \left(\mu_r(z^\ell) \int_{\partial D_j(p_r)} f(\widehat{v}|z^\ell, x_r, \xi_r)\, \mathrm{d}\widehat{v} \right).$$

I then assume that $\psi_{jr}(\cdot) = \psi(\cdot; x_{jr}, \xi_{jr}, w_{jr})$, letting cost functions vary only with product characteristics and cost shifters. Then, for any given triplet $(\overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr})$, sufficient conditions to identify $\psi(\cdot; \overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr})$ are

(i) there are at least two products jr and $\widetilde{jr}$ with $(x_{jr}, \xi_{jr}, w_{jr}) = (x_{\widetilde{jr}}, \xi_{\widetilde{jr}}, w_{\widetilde{jr}}) = (\overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr})$; and

(ii) the two products have a different composition of marginal buyers in terms of z^1 and z^2:

$$\frac{\mu_r(z^1) \int_{\partial D_j(p_r)} f(\widehat{v}|z^1, x_r, \xi_r)\, \mathrm{d}\widehat{v}}{\mu_r(z^2) \int_{\partial D_j(p_r)} f(\widehat{v}|z^2, x_r, \xi_r)\, \mathrm{d}\widehat{v}} \neq \frac{\mu_{\widetilde{r}}(z^1) \int_{\partial D_{\widetilde{j}}(p_{\widetilde{r}})} f(\widehat{v}|z^1, x_{\widetilde{r}}, \xi_{\widetilde{r}})\, \mathrm{d}\widehat{v}}{\mu_{\widetilde{r}}(z^2) \int_{\partial D_{\widetilde{j}}(p_{\widetilde{r}})} f(\widehat{v}|z^2, x_{\widetilde{r}}, \xi_{\widetilde{r}})\, \mathrm{d}\widehat{v}}.$$

If this is the case $\left[\psi(z^1; \overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr}), \psi(z^2; \overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr})\right]$ is the unique solution of the linear system describing the first-order conditions that must hold for both products:

$$\begin{bmatrix} mr_{jr} \\ mr_{\widetilde{jr}} \end{bmatrix} = \begin{bmatrix} \mu_r(z^1) \int_{\partial D_j(p_r)} f(\widehat{v}|z^1, x_r, \xi_r)\, \mathrm{d}\widehat{v} & \mu_r(z^2) \int_{\partial D_j(p_r)} f(\widehat{v}|z^2, x_r, \xi_r)\, \mathrm{d}\widehat{v} \\ \mu_{\widetilde{r}}(z^1) \int_{\partial D_{\widetilde{j}}(p_{\widetilde{r}})} f(\widehat{v}|z^1, x_{\widetilde{r}}, \xi_{\widetilde{r}})\, \mathrm{d}\widehat{v} & \mu_{\widetilde{r}}(z^2) \int_{\partial D_{\widetilde{j}}(p_{\widetilde{r}})} f(\widehat{v}|z^2, x_{\widetilde{r}}, \xi_{\widetilde{r}})\, \mathrm{d}\widehat{v} \end{bmatrix} \begin{bmatrix} \psi(z^1; \overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr}) \\ \psi(z^2; \overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr}) \end{bmatrix}.$$

In words: With two types of buyers which may imply different cost for the same product jr, it is necessary to observe two products with the same characteristics, and that variation in the characteristics of their competitors and/or market composition — and thus prices — induce variation in

[43] See conditions D1, C1, C2, C3 in Appendix B; or Theorem 1 and Section 4.2 in **?**. **?** (see in particular Section 4.2) provide a rich discussion on the advantages of rich variation in individual level data to trace out the heterogeneity in preferences in each market relaxing specific functional form and parametric assumptions.

[44] With $\psi_{jr}(\widehat{z},\widehat{v}) \equiv c_{jr}$, i.e. assuming away selection, this would not be the case, and one could use the traditional cost-identification results discussed in **??**.

the composition of their marginal buyers in terms of the two types. If the two types are YOUNG and OLD, if one product has a 50:50 ratio of marginal buyers across YOUNG and OLD, it is necessary to observe another product for which (i) one assumes the same cost function and (ii) the ratio of marginal buyers across YOUNG and OLD is not 50:50.

Cost identification — general case. This approach can be extended to the general model. What follows differs slightly from Appendix B, yet the main point is conceptually the same.[45] As above I impose:

A1. For all j and all r, $mr_{jr} = mc_{jr}$, where $mr_{jr} = \sum_{\widehat{z}\in\mathcal{Z}} \mu_r(\widehat{z}) \left(\sigma_j(\widehat{z}, \chi_r) - p_{jr} \int_{\partial D_j(p_r)} f(\widehat{v}|\widehat{z}, x_r, \xi_r)\, \mathrm{d}\widehat{v} \right)$,
and
$mc_{jr} = \sum_{\widehat{z}\in\mathcal{Z}} \mu_r(\widehat{z}) \int_{\partial D_j(p_r)} \psi_{jr}(\widehat{z}, \widehat{v}) f(\widehat{v}|\widehat{z}, x_r, \xi_r)\, \mathrm{d}\widehat{v}$.

A2: For any j, $\psi_{jr}(\widehat{z}, \widehat{v}) = \psi(\widehat{z}, \widehat{v}, x_{jr}, \xi_{jr}, w_{jr})$, where x_{jr} may include the identity of the seller.

Then I restrict the number of "cost-types" to be finite, the *key assumption* for my argument:[46]

A3: There exists a finite (disjoint) partition of $\mathcal{V}$, say $\{\mathcal{V}^1, \mathcal{V}^2, ..., \mathcal{V}^M\}$ such that, for any $\widehat{z} \in \mathcal{Z}$, if $\widehat{v}, \widetilde{v} \in \mathcal{V}^m$ for some $m = 1, ..., M$, then $\psi(\widehat{z}, \widehat{v}, x_{jr}, \xi_{jr}, w_{jr}) = \psi(\widehat{z}, \widetilde{v}, x_{jr}, \xi_{jr}, w_{jr})$.

Using $\partial\mathcal{D}_{jr}(\mathcal{V}^m, z^\ell) = \left(\mu_r(z^\ell) \int_{\partial D_j(p_r)\cap\mathcal{V}^m} f(\widehat{v}|z^\ell, x_r, \xi_r)\, \mathrm{d}\widehat{v} \right)$ to denote the density of marginal buyers for j in r with characteristics z^ℓ and valuations $\widehat{v} \in \mathcal{V}^m$ I then have the following:

PROPOSITION **3** *Under A1-A3, $\psi(\cdot, \cdot, \overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr})$ is identified if there exists a set of $H \geq T \times M$ pairs $\widetilde{jr}$ such that*
(i) *$(x_{\widetilde{jr}}, \xi_{\widetilde{jr}}, w_{\widetilde{jr}}) = (\overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr})$ for all $\widetilde{jr} = 1, ..., H$; and*
(ii) *the H-by-$(T \times M)$ matrix of marginal buyers*

[45] In Appendix B I focus only on heterogeneity in cost due to differences in preferences, ignoring observable heterogeneity across buyers. I impose cross-product restrictions, and use a large support condition on prices that is extremely demanding on the data. This stronger condition allows me to provide a constructive proof for the identification of cost functions using variation in the set of marginal buyers, without assuming that selection is limited to a finite set of possible preference types. Here I impose more assumptions, but the conditions for identification become more transparent and operational.

[46] For example, suppose I was to estimate the following parametric model:
Facing prices p_r, the indirect utility that buyer i derives from j in r is

$$u_{ijr} = -\alpha_i p_{jr} + \beta_i x_{jr} + \xi_{jr} + \epsilon_{ijr},$$

where (α^i, β^i) collects random parameters drawn from a distribution $G(\alpha, \beta|z_{ir})$ with finite support $A \times B$ (similar to the demand system in **?**).
Assumption A3 then holds by assuming that $\psi(\alpha, \beta, \epsilon, x_{jr}, \xi_{jr}, w_{jr}) = \psi(\alpha, \beta, x_{jr}, \xi_{jr}, w_{jr})$, requiring that the idiosyncratic preference shock ϵ is uninformative about the buyer's risk.

$$\begin{array}{c} h=1 \\ \vdots \\ h=H \end{array} \left[\begin{array}{cccccc} \partial\mathcal{D}_{\widetilde{jr}}(\mathcal{V}^1,z^1) & \partial\mathcal{D}_{\widetilde{jr}}(\mathcal{V}^2,z^1) & ... & \partial\mathcal{D}_{\widetilde{jr}}(\mathcal{V}^M,z^1) & \partial\mathcal{D}_{\widetilde{jr}}(\mathcal{V}^1,z^2) & ... & \partial\mathcal{D}_{\widetilde{jr}}(\mathcal{V}^M,z^T) \\ \vdots & \vdots & & \vdots & \vdots & & \vdots \\ \partial\mathcal{D}_{\widetilde{j'r'}}(\mathcal{V}^1,z^1) & \partial\mathcal{D}_{\widetilde{j'r'}}(\mathcal{V}^2,z^1) & ... & \partial\mathcal{D}_{\widetilde{j'r'}}(\mathcal{V}^M,z^1) & \partial\mathcal{D}_{\widetilde{j'r'}}(\mathcal{V}^1,z^2) & ... & \partial\mathcal{D}_{\widetilde{j'r'}}(\mathcal{V}^M,z^T) \end{array}\right]$$

is full-column rank w.p.1 with respect to the (conditional) distribution of $\left((\mu_r, x_{-jr}, \xi_{-jr}, w_{-jr})\middle|(\overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr})\right)$.

Beside the simple proof that I report below, the intuition is similar to the 2-types case analyzed earlier. To identify differences in the cost of a product j induced by differences in the type of buyer, one can use variation in the composition of marginal buyers within groups of products with otherwise (assumed) equal cost structures. This variation is induced by variation in the characteristics of opponents $(x_{-jr}, \xi_{-jr}, w_{-jr})$, or variation in the composition of potential buyers in the market $\mu_r(\cdot)$. Both induce variation in prices and choices, thus marginal buyers' composition, but do not directly affect individual cost functions for product j.

Since the number of possible "cost-types" that can be identified is bounded by the number of "preference types" distinguished by the demand system, the availability of rich individual-level observables and large variation in prices is important: Both can allow the estimation of a demand system with richer heterogeneity (see **?**), and this can then lead to the estimation of cost functions with less limits on selection.

Proof of Proposition ??

Assume that there are two functions $\psi(\cdot,\cdot,\overline{x}_{jr},\overline{\xi}_{jr},\overline{w}_{jr}) \neq \widehat{\psi}(\cdot,\cdot,\overline{x}_{jr},\overline{\xi}_{jr},\overline{w}_{jr})$ for which observables are the same with strictly positive probability. Pick for any $m = 1, ..., M$ an arbitrary $v^m \in \mathcal{V}^m$. By A1-A3, with positive probability, for all pairs jr with $(x_{jr}, \xi_{jr}, w_{jr}) = (\overline{x}_{jr}, \overline{\xi}_{jr}, \overline{w}_{jr})$

$$mr_{jr} = \sum_{m,\ell}^{M,T} \psi(z^\ell, v^m, x_{jr}, \xi_{jr}, w_{jr}) \partial\mathcal{D}_{jr}(\mathcal{V}^m, z^\ell)$$

$$mr_{jr} = \sum_{m,\ell}^{M,T} \widehat{\psi}(z^\ell, v^m, x_{jr}, \xi_{jr}, w_{jr}) \partial\mathcal{D}_{jr}(\mathcal{V}^m, z^\ell),$$

thus

$$\sum_{m,\ell}^{M,T} \left(\psi(z^\ell, v^m, x_{jr}, \xi_{jr}, w_{jr}) - \widehat{\psi}(z^\ell, v^m, x_{jr}, \xi_{jr}, w_{jr})\right) \partial\mathcal{D}_{jr}(\mathcal{V}^m, z^\ell) = 0.$$

Conditions (i) and (ii) imply then that $\psi(\cdot,\cdot,\overline{x}_{jr},\overline{\xi}_{jr},\overline{w}_{jr}) = \widehat{\psi}(\cdot,\cdot,\overline{x}_{jr},\overline{\xi}_{jr},\overline{w}_{jr})$ w.p.1, a contradiction.■