This work is distributed as a Discussion Paper by the

STANFORD INSTITUTE FOR ECONOMIC POLICY RESEARCH

SIEPR Discussion Paper No. 15-012

Estimating Equilibrium in Health Insurance Exchanges: Analysis of the Californian Market under ACA

By

Pietro Tebaldi

Stanford Institute for Economic Policy Research
Stanford University
Stanford, CA 94305
(650) 725-1874

The Stanford Institute for Economic Policy Research at Stanford University supports research bearing on economic and public policy issues. The SIEPR Discussion Paper Series reports on research and policy analysis conducted by researchers affiliated with the Institute. Working papers in this series reflect the views of the authors and not necessarily those of the Stanford Institute for Economic Policy Research or Stanford University

Estimating Equilibrium in Health Insurance Exchanges: Analysis of the Californian Market under the ACA *

Pietro Tebaldi
Stanford University†

April 30, 2015

Abstract

This paper develops and estimates a model of a regulated health insurance exchange, in which insurers' ability to adjust prices across buyers with different observed risk or preferences is restricted. I show conditions under which the joint distribution of risk and preferences is identified, even when the econometrician does not observe any information on individual risk. These primitives can then be used to simulate equilibrium under alternative regulations. I estimate the model with data from the first year of the Californian exchange under the Affordable Care Act, where age-rating restrictions and a subsidy program determine the way in which insurers' decisions translate to expected profits. For this market, I investigate alternative designs of the subsidy program. Compared to the subsidy formula mandated by the healthcare reform, the adoption of a voucher program – providing buyers with a lump-sum equal to 70-80% of their expected expenditure – would transfer welfare away from insurers, favoring consumers and/or taxpayers. Simulations of equilibrium under this alternative policy result in total coverage between 100-115% of the levels achieved by the current regulations, while also reducing government expenditure, average premiums, and markups, by 0-20%, 12-15%, and 22-27%, respectively.

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

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

---

*I thank Liran Einav, Tim Bresnahan, and Matt Jackson for the helpful research advise during the writing of this draft. I also thank Lanier Benkard, Michael Dickstein, Dave Donaldson, Mark Duggan, Egemen Eren, Daniel Grodzicki, Han Hong, Caroline Hoxby, Brad Larsen, Jon Levin, Davide Malacrino, Joe Orsini, Ariel Pakes, Luigi Pistaferri, Maria Polyakova, Ilya Segal, as well as participants at the Stanford IO and Applied Micro workshops, and Bocconi University Applied Micro seminar. All remaining mistakes are mine.

†Department of Economics. 579 Serra Mall, Stanford, CA, 94305. Email: ptebaldi@stanford.edu.

1 Introduction

The role of private insurance markets is central to the US healthcare system. These markets operate under what is called "managed competition", to emphasize how competition of private insurers is embedded in a highly regulated environment, with the intent of the law being to avoid possible market failures due to adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976), and preserve a fair access to health services for low-income individuals.

A recent, large scale example of managed competition is provided by how the recent US national reform (Patient Protection and Affordable Care Act, ACA henceforth) regulates the so-called "individual market" for health insurance: the segment of the market serving individuals and families who are not offered coverage through an employer-sponsored program, and are not beneficiaries of public insurance. Along with a universal ban on insurers' ability to reject risky buyers, the central provision of the law is the establishment of state-level health insurance exchanges. These new marketplaces have been central to the public discussion regarding the healthcare reform, joining the list of regulated health exchanges already operating in the market of supplemental insurance for citizens older than 65 (Curto et al., 2014; Starc, 2014; Decarolis, 2015). The diffusion of regulated health insurance exchanges calls for a careful analysis of how supply, market outcomes, and welfare, respond to different policies.

To carry on such analyses, it is necessary to estimate the joint distribution of buyers' risk and preferences, since the extent to which these primitives are correlated is precisely the difference between insurance markets and markets without selection. For this task, different works used a variety of data formats and econometric techniques, largely relying on *some* information about differential riskiness of buyers with different characteristics. This requires either direct observation of insurance claims (Einav et al., 2010a), or the use of external survey data (Starc, 2014). In new markets such as the ACA exchanges, however, detailed claims data might not be available, being an important piece of proprietary information left to private insurers. At the same time, available surveys might not well represent the population of buyers in a market which is meant to serve outsiders of the previous healthcare system. Our ability to analyze these new markets for policy purposes is then still an open question.

For this, my paper serves two main purposes. First, it provides an empirically tractable model of a regulated health exchange and shows conditions under which the joint distribution of risk and preferences can be estimated, even when there is no data describing risk varying across different types of buyers. Second, I apply this framework to one ACA state-level health exchange, using 2014 data from the Californian market. I use my estimates to document differences in risk and price sensitivity across buyers of different age and income, and to describe the interaction between regulations and insurers' incentives. I then study how market outcomes would vary by adopting a different design of the subsidy program, showing how replacing the ACA-mandated formula with a voucher-type mechanism could transfer welfare from insurers to consumers and/or lower the burden on the government's budget.

In the first part of the paper (Section 2), I present a theoretical model of a health insurance exchange. I focus on rating regulations, the typical aspect of managed competition that limits the ability of insurers to price discriminate between buyers with different observable risk or preferences. Taking entry decisions and the product space as given, I explicitly model the way in which limits on price discrimination imply a relationship between population characteristics (e.g. age and income distribution) and insurers' decisions.

From this fact, leveraging on a "Nash-in-prices" equilibrium assumption to obtain a pricing equation, variation of observed population characteristics across different geographic areas can be exploited to invert the mapping from the joint distribution of risk and preferences to observed market outcomes (enrollment and insurance premiums). Sections 3 and 4 formalize this, providing conditions under which, even when the econometrician cannot observe any information on different costs across different buyers, data on enrollment and prices from a regulated exchange are sufficient to estimate demand and costs allowed to be heterogenous across buyers.

To obtain intuition: suppose that insurers cannot adjust rates on observables correlated with risk, say age. In a market in which potential buyers are distributed across old and young in a 2:1 ratio, prices will be higher than in a market in which old and young buyers are in a 1:1 ratio. If preferences of young buyers do not depend on the number of old buyers in the same market, variation in prices induced by varying age-composition of potential buyers can be exploited to estimate demand via standard methods. At this point, marginal revenues at the observed prices can be constructed, while marginal costs are equal to a weighted-average of costs (still unknown) for buyers of different age, where the weights correspond to shares of marginal buyers across ages (which can also be constructed from demand estimates). Then, by imposing the equilibrium optimality conditions which equate marginal revenues to marginal costs, variation in the age-composition of marginal buyers across products in the data can suffice to estimate age-specific costs, even without *any* information on how risk evolves with a buyer's age.

This theoretical framework nests the context of state-level exchanges subject to ACA regulations, as described in Section 5, and it is applied to data from the first year of the Californian exchange; the largest among ACA exchanges, with over two million people covered since early 2014. For this market, the essential regulations determining the mapping from an insurer's pricing decision to his expected profits are (1) standardized age-based pricing (Orsini and Tebaldi, 2014), and (2) the federal subsidy program, providing discounts to buyers whose income is less than four times the federal povery level (FPL), or approximately 45,960 dollars per-year for a single adult. These two features are built explicitly in my model of insurance demand and optimal pricing. I estimate and describe the equilibrium incentives under the status-quo (Section 6), and then show how the formula adopted to calculate discounts for subsidy-eligible buyers could be ameliorated to favor consumers, and/or lower the burden on public finances (Section 7).

Starting from demand estimates, my results show that price sensitivity decreases in age and income, where I consider three age groups (20-29, 30-44, and 45-64), and distinguish between "subsidized" and "unsubsidized" buyers (below or above four times the FPL). The more price-sensitive group consists of the young adults (age 20-29) who are eligible for subsidies: for them, a 1% increase in premium (approximately 17 dollars per-year) induces on average a 3.30% drop in demand. At the other end of the spectrum, for the older adults (age 45-64) who do not receive subsidies, a 1% increase in premium (approximately 89 dollars per-year) induces on average a 1.04% drop in demand.

The second step of estimation delivers estimates of expected annual health expenditure, that are allowed to depend on which insurer the buyer decides to purchase his coverage from, and varying in the buyer's income and age. In particular, for any insurer operating in the exchange, I estimate expected costs distinguishing between subsidized and unsubsidized buyers, and between two age groups (20-44, and 45-64). Estimated health risk increases in age and decreases in income, and it is quantitatively aligned with what one would find using the Medical Expenditure Panel Survey. A preview of these estimates is shown in the right panel of the table below, which along with the left panel summarizes the co-variation in risk and preferences that I find in the population of buyers of the Californian exchange.

Preferences (left) and risk (right) as estimated in the Californian exchange

Price elasticity			
Age	20-29	30-44	45-64
Subsidized	3.3	2.7	1.5
Unsubsidized	1.5	2.1	1.0

Expected health expenditure			
Age	20-29	30-44	45-64
Subsidized	\$3,700	\$3,700	\$6,800
Unsubsidized	\$2,500	\$2,500	\$4,500

Note: Derived from results and tables presented in Section 6. Elasticities correspond to the percentage drop in demand resulting from a one percent increase in price. Health expenditure is intended for one calendar year.

These estimates are a key ingredient for the counterfactual analysis presented in Section 7, where I focus on the design of the subsidy program. This program, similarly to the one in place in the market for supplemental insurance for senior citizens (Curto et al., 2014; Duggan et al., 2014; Decarolis, 2015), significantly reduces the amount paid directly by low-income consumers to insurers, with the difference being covered through public funding. As discussed in Einav and Levin (2015), this lowers buyers' price sensitivity, thus increasing the ability of insurers to charge higher prices without facing the *true* slope of the demand schedule. As a consequence, a large fraction of the welfare gains created by these programs goes to insurers, who benefit of increased market power (Duggan et al., 2014; Einav and Levin, 2015). An important challenge for policymakers is then to find ways to transfer welfare gains away from insurers, to favor instead consumers, or to reduce the burden on taxpayers.

For this purpose, here I discuss how the subsidy formula adopted under the ACA is particularly problematic. Discounts to low-income buyers are in fact computed as an increasing function of market premiums (see details in Section 5). The higher the premium, the larger the subsidy, with buyers not experiencing any change in discounted prices. Intuitively, this mechanism reduces price competition, raising prices in equilibrium. This distortion in insurers' incentives could be avoided by adopting a voucher-type mechanism, in which low-income buyers benefit of a discount that does not increase with market prices.

In my counterfactuals, first I quantify the overall effect of the current subsidy program (whose legitimacy in numerous states is now under scrutiny of the US Supreme Court), then I compare equilibrium outcomes under the ACA subsidy scheme to the ones obtained under a voucher-type mechanism.

With the first of these exercises, comparing the current market to a situation without any subsidy, my results emphasize the importance of government support to low-income uninsured. Without the program, total demand among subsidized buyers would fall by more than 95%, as young, subsidized are hit by a 115% premium increase, while their older counterparts would face an average price increase of approximately 500%. Since the market without subsidy would have less price-sensitive, richer buyers becoming the largest portion of demanders, prices would increase by approximately 10-12% even for this group, as individual markups would raise by 20%.

Rather than this extreme scenario, I then explore a policy that replaces the ACA subsidy mechanism with a voucher program. For this exercise, the amount of the voucher becomes a key policy parameter, since there is no formula that computes discounts from realized prices. To carry on the comparison with the ACA formula without imposing any prior assumption on the "correct" discount level, I compute equilibrium outcomes for vouchers that vary from 0-100% of the average expected health expenditure of a buyer, where the average is taken across all insurance carriers active in the market. My results are promising: when providing low-income buyers with a discount equal to 70-80% of their annual expected cost for the average insurer, market outcomes are such that welfare is transferred from insurers toward consumers and/or taxpayers. More specifically, under this alternative policy total coverage remains between 100-115% of the levels achieved under the current scheme, while government expenditure, average premiums, and markups, are reduced by 0-20%, 12-15%, and 22-27%, respectively.

These analyses are chosen for their policy relevance, but are just an example of the type of equilibrium comparisons that are possible with estimates from a health exchange, which can be obtained through the econometric model presented in the first part of the paper.

Relationship to Existing Literature. This work belongs to a growing literature exploring how alternative regulations affect outcomes and welfare in private health insurance markets. The case of "managed competition" within the Medicare Advantage, Medicare Part D, and Medigap programs is the main focus of, among others, Curto et al. (2014), Duggan and Hayford (2011), Duggan et al. (2014), Decarolis (2015), and Starc (2014). Like here, also these papers emphasize how insurers' market power should be accounted for when studying alternative regulations and subsidy programs in health exchanges. Compared to these, however, the individual market under the ACA – object of my application – serves a different population, and it is subject to rules that differ in many aspects from the ones imposed on the markets for senior citizens.

Within the smaller, more recent literature studying the impact of the ACA regulatory regime on market outcomes, an important contribution is the one of Handel et al. (2015). They use a rich empirical model (estimated in the context of employer-sponsored insurance) to simulate the long-run welfare consequences of limitations to risk-based premium adjustments. They do so assuming perfect competition, and a universal participation mandate (buyers are not able to leave the market). Although my analysis has a more limited scope, a contribution is to resemble better the market structure of ACA exchanges, considering oligopolistic competition between insurers with heterogeneous characteristics, while also allowing buyers to leave the market when facing high prices. In this, my approach is similar to the one of Starc (2014) and Ericson and Starc (2015): the former studies the Medigap market, while the latter analyses ACA pricing regulations with estimates from the pre-ACA Massachusetts' exchange, and doest not focus on the subsidy program which is the main object of my counterfactuals.

Recent work by Dafny et al. (2014) also analyzes competition, and the role of market power, in ACA exchanges. Their results emphasize a negative relationship between number of insurers and prices, but do not model explicitly insurers' incentives and how these are altered by pricing regulations. Other papers can be classified along two lines: some (Hackmann et al., 2013) use pre-reform data (or data from different markets) to simulate the effect of ACA regulations, others (Orsini and Tebaldi, 2014; Dickstein et al., 2015) use post-reform data but do not craft the details of ACA regulations in a comprehensive econometric model. Here, I use data from ACA exchanges and estimate a model of the new market in which the insurers' incentives are explicitly accounted for, therefore allowing for counterfactual simulations which can be used to explore future directions for different regulatory designs.

From a methodological perspective, within the existing literature proposing welfare and regulatory analyses of insurance markets, identification of the joint distribution of risk and preferences relies on the availability of detailed cost data (Einav et al., 2010a,b; Hendren, 2013; Handel, 2013; Bundorf et al., 2012). When cost information across buyers with different characteristics was not available, authors matched demand data to representative surveys (Starc, 2014), or assumed directly a functional form relating demand and cost (Lustig, 2010). Here I show how, when studying a regulated exchange, it is possible to circumvent this problem without additional data or functional form assumptions.

2 Equilibrium Model of a Regulated Health Insurance Exchange

Setup and Notation. I consider the situation in which a population of individuals is offered J differentiated health insurance plans by N firms. For a given firm f, the set $J_f \subset J$ denotes products (coverage options) offered by f. The set of active firms, and all non-price observable products' characteristics, are taken as fixed and exogenous. This is a relevant assumption, which I hold throughout, making my analysis different from the one of Rothschild and Stiglitz (1976), while closer in spirit to the ones of Akerlof (1970), Einav et al. (2010a), and Handel et al. (2015). The plausibility of this assumption will depend on the regulatory context, as I discuss further in Sections 5 and 6. Finally, as in Bundorf et al. (2012) and Starc (2014), rather than assuming that plans only differ in their generosity (percentage of health expenditures reimbursed), here I also allow for "horizontal" differences between products. This captures differences in insurers' attractivenness within a given level of coverage, induced for example by brand loyalty, or by differences in the network of providers covered by the plan.

The population of potential buyers is represented by a distribution $G(\zeta, \tau)$, where $\zeta \in \mathcal{Z}$ is a (possibly infinite) vector of consumer characteristics that cannot be elicited by the insurer, and therefore cannot be used for rating purposes. $\tau \in \mathcal{T}$ is instead a finite vector of observable characteristics (e.g. age, or gender); absent rating regulations, these can be used for rating purposes. As in Einav et al. (2010a) and Handel et al. (2015), the precise nature of ζ and τ does not need to be specified. They each could represent, for example, consumers' risk-preferences, a detailed health history of the individual, a risk-score, age, income, or gender. What is relevant here is that I explicitly distinguish between two types: one can be used for rating purposes, while the other cannot.

The indirect utility that a buyer (ζ_i, τ_i) derives when purchasing plan j at price p_{ij} is denoted $u_{ij}(p_{ij}) = v_j(\zeta_i, \tau_i) - p_{ij}$, where $v_j : \mathcal{Z} \times \mathcal{T} \to \mathbb{R}$ describes the monetary valuation for product j, a function of the buyer's characteristics (ζ_i, τ_i).

If a given individual (ζ_i, τ_i) enrolls in plan $j \in J_f$, over the coverage period (typically a calendar year), insurer f expects to face a cost equal to $\phi_j(\zeta_i, \tau_i)$. If the same individual is charged a price p_{ij}, the expected profit for f is then $p_{ij} - \phi_j(\zeta_i, \tau_i)$.

With this formulation I avoid to make specific assumptions on how individual characteristics translate in expected cost for the insurer. The function $\phi_j : \mathcal{Z} \times \mathcal{T} \to \mathbb{R}$ implicitly captures this process: this depends on the individual's health status, but also on the set of providers covered by each specific plan, as well as on other idiosyncrasies such as, for example, different contracts between insurers and providers, or "moral hazard" considerations, with utilization varying with coverage choice. The model does not distinguish between these different aspects entering the mapping ϕ_j, as they do not affect insurers' incentives. Clearly, however, when engaging in welfare considerations, further assumptions about the mechanisms underlying ϕ_j will become necessary.

Rating and Demand. Prices for an individual (ζ_i, τ_i) depend on the observable type τ_i, also referred to as "rating type". In particular, $r_j : \mathcal{T} \to \mathbb{R}$ is a rating function indicating the price at which product j is priced across buyers with different τ. Individual (ζ_i, τ_i) can buy j at price $p_{ij} = r_j(\tau_i)$. In what follows, omitting product subscripts indicates vectors, so that individual (ζ_i, τ_i) faces the price vector $p_i \equiv (p_{i0}, p_{i1}, ..., p_{iJ}) = (r_0(\tau_i), r_1(\tau_i), ..., r_J(\tau_i)) \equiv r(\tau_i)$; one price for each product in J, adding the outside option $j = 0$. ($r_0(\tau)$ does not need to be zero; a positive price for the outside option is used to model public transfers, as it will become clear in Section 5.)

Each buyer observes the set of available products and prices before making his coverage decision. In this discrete choice environment, for individual (ζ_i, τ_i) to choose product j, it must be that, for any product $k = 0, 1, ..., J$,

$$v_j(\zeta_i, \tau_i) - v_k(\zeta_i, \tau_i) \geq p_{ij} - p_{ik}. \tag{1}$$

Let $v(\zeta_i, \tau_i)$ denote the vector $(v_j(\zeta_i, \tau_i))_{j=0}^{J}$, and Γ_j be the matrix such that, for a vector $x = (x_j)_{j=0}^{J}$, $\Gamma_j x = (x_j - x_k)_{k=0}^{J}$. (See for example Thompson (1989).) From (1), total demand for product j among potential buyers depends on the collection of rating functions $r = (r_0, r_1, ..., r_J)$. Formally, this is

$$D_j(r) = \int_{\mathcal{Z}\times\mathcal{T}} \mathbf{1}\{\Gamma_j v(\zeta, \tau) \geq \Gamma_j r(\tau)\}\, dG(\zeta, \tau), \tag{2}$$

where $\mathbf{1}\{\cdot\}$ is the indicator function, and $\geq$ is used to compare vectors element-by-element.

Supply and Profits. The N insurers choose the rating functions $r_1, ..., r_J$ to maximize expected profits *à la* Bertrand: prices are determined simultaneously, and fixed before the market opens and demand realizes.

For what follows, it is convenient to introduce the notation for τ-specific demand and costs curves: $d_j(p; \tau)$ and $c_j(p; \tau)$, respectively. Using $G(\zeta|\tau)$ to indicate the distribution of ζ conditional on a given rating type τ, these are

$$d_j(p;\tau) = \int_{\mathcal{Z}} \mathbf{1}\{\Gamma_j v(\zeta, \tau) \geq \Gamma_j p\}\, dG(\zeta|\tau), \text{ and} \tag{3}$$

$$c_j(p;\tau) = \int_{\mathcal{Z}} \mathbf{1}\{\Gamma_j v(\zeta, \tau) \geq \Gamma_j p\}\, \phi_j(\zeta, \tau)\, dG(\zeta|\tau). \tag{4}$$

For a collection of rating functions r, insurer f realizes expected profits equal to

$$\Pi_f(r) = \sum_{j\in J_f} \left(\underbrace{\int_{\mathcal{T}} r_j(\tau)\, d_j(r(\tau);\tau)\, dG(\tau)}_{\text{Revenues for product } j} - \underbrace{\int_{\mathcal{T}} c_j(r(\tau);\tau)\, dG(\tau)}_{\text{Costs for product } j} \right), \tag{5}$$

where $G(\tau)$ denotes the marginal distribution of rating types τ in the population. I mantain the assumption that the underlying primitives are such that both, $d_j(p;\tau)$ and $c_j(p;\tau)$, are continuous and differentiable functions of p, for all $\tau \in \mathcal{T}$.

Unconstrained Equilibrium. Insurers choose rating functions r^* simultaneously to maximize expected profits, as in a Bertrand-Nash equilibrium. That is, r^* is an equilibrium if, for all firms f, the collection $\left(r_j^*\right)_{j \in J_f}$ is a solution to

$$\max_{(r_j)_{j \in J_f}} \Pi_f \left((r_j)_{j \in J_f}, (r_k^*)_{k \notin J_f} \right).$$

Since profit functions are continuous and differentiable, equilibrium can be characterized via first-order optimality conditions. In particular, in this situation in which insurers are freely allowed to price discriminate across different τ, in equilibrium one will have that

$$\frac{\partial \Pi_f (r^*)}{\partial (r_j (\tau))} = 0 \quad \text{for all } f \text{, all } j \in J_f \text{, all } \tau \in \mathcal{T};$$

or equivalently

$$d_j (r^* (\tau); \tau) + \sum_{j' \in J_f} \frac{\partial d_{j'} (r^* (\tau); \tau)}{\partial p_j} - \frac{\partial c_{j'} (r^* (\tau); \tau)}{\partial p_j} = 0 \quad \text{for all } f \text{, all } j \in J_f \text{, all } \tau \in \mathcal{T}. \tag{6}$$

Notice how, in this case, the composition of the population across rating types does not affect equilibrium prices, which depend on the conditional distributions $G(\zeta|\tau)$ but not on $G(\tau)$.

Since in this model I consider imperfect competition between firms offering horizontally differentiated products, concerns for (lack of) equilibrium existence are mitigated. Although the problem of existence of pure-strategy Nash equilibrium in a multi-product oligopoly is still unsolved by the literature, product differentiation is sufficient to ensure the continuity of profit functions, whose failure is the central problem that other models of competitive insurance provision have dealt with. (See Handel et al. (2015) for a thorough discussion.)

In my model, existence could possibly fail to hold if profit functions are not quasi-concave. This, however, does not depend on the insurance nature of the market, since the same problem arises even if marginal costs are assumed to be constant in (ζ, τ). In Appendix A I discuss this, and propose a set of possible sufficient conditions, each ensuring equilibrium existence in the pricing game analyzed in this paper. Moreover, in the empirical application I am able to check numerically whether my final estimates are constistent with the existence of an equilibrium.

Rating Regulations and Constrained Equilibrium. A rating regulation consists of a set of constraints imposed on the function r_j, that is the formula that the insurer uses to discriminate across different rating types. Restrictions on rating adjustments are modeled by imposing that $r_j \in \mathcal{R}$, where $\mathcal{R} \subset \mathbb{R}^{\mathcal{T}}$ describes the (compact) subset of allowed rating functions.

Subject to the rating regulations $\mathcal{R}$, insurers set prices in equilibrium. A constrained equilibrium is a collection r^* such that, for all f, $\left(r_j^*\right)_{j\in J_f}$ solves

$$\max_{(r_j)_{j\in J_f} \text{ s.t. } r_j\in\mathcal{R}} \Pi_f\left((r_j)_{j\in J_f}, (r_k^*)_{k\notin J_f}\right).$$

This formulation is general, and can cover arbitrary constraints.[1] Here, however, I focus on the case of *automatic adjustments*: situations in which prices across τ's are adjusted according to a pre-determined rule. The same framework can be used to include more complicated constraints, but my choice here is to keep the presentation uncluttered and tightly linked to my empirical application.

Automatic rating adjustments are described by a function $a : \mathbb{R}\times\mathcal{T} \to \mathbb{R}$, and by a normalized type $\tau_0 \in \mathcal{T}$; one then has $\mathcal{R} = \mathcal{R}^a$, with

$$\mathcal{R}^a = \left\{r_j(\tau) = a\left(r_j(\tau_0), \tau\right) \text{ for all } \tau\in\mathcal{T}\right\}, \tag{7}$$

where this imposes that $a\left(r_j(\tau_0), \tau_0\right) = r_j(\tau_0)$, i.e. $a(\cdot, \tau_0)$ is the identity function. In words, once the price $r_j(\tau_0)$ for buyers of type τ_0 is determined, prices for any other $\tau\in\mathcal{T}$ are determined using the function a. When $a(\cdot,\tau)$ is the identity function for all τ, this corresponds to perfect community rating, where no price discrimination is allowed (Buchmueller and Dinardo (2002), or Sasso and Lurie (2009)). When $a(\cdot,\tau)$ varies with a pre-determined risk-score, this corresponds to *ex-ante* risk-adjustment as implemented by the Medicare Advantage program (Brown et al. (2014), or Curto et al. (2014)).

In equilibrium, each insurer will optimally choose prices for types τ_0, knowing that these will map directly into prices for all other buyers. In particular, the first-order conditions explicitly show how $\mathcal{R}^a$ links profits across different rating types τ, thus creating a link between prices for a given τ and the distribution of rating types in the population of potential buyers.

Assuming that the function $a(\cdot,\tau)$ is differentiable in its first argument for all τ, in a constrained equilibrium r^* the following must hold:

- For all f, for all $j\in J_f$, $r_j^*\in\mathcal{R}^a$, where $\mathcal{R}^a$ is defined in (7);
- For all f, for all $j\in J_f$,

$$\int\left(d_j\left(r_j^*(\tau);\tau\right) + \sum_{j'\in J_f} r_j^*(\tau)\frac{\partial d_{j'}\left(r_j^*(\tau);\tau\right)}{\partial p_j} - \frac{\partial c_{j'}\left(r_j^*(\tau);\tau\right)}{\partial p_j}\right)\frac{\partial a\left(r_j^*(\tau_0),\tau\right)}{\partial\left(r_j(\tau_0)\right)}dG(\tau) = 0. \tag{8}$$

[1] To ensure a solution to the insurer's pricing problem, $\mathcal{R}$ must be compact in $\mathbb{R}^{\mathcal{T}}$ (the space of functions from $\mathcal{T}$ to $\mathbb{R}$), which is a standard Euclidean space since $\mathcal{T}$ is finite.

Condition (8) implies that, in a constrained equilibrium, rates for each τ are affected by the entire distribution $G(\tau)$ in the population, and this was not the case without rating regulations, as shown in the equilibrium condition (6) above.

3 Econometric Model

Markets. The model introduced in the previous section refers to a single geographic market, also referred to as "rating region" in the context of health exchanges. I use m to denote a market, with J_m representing the corresponding set of products. For any firm f active in market m, $J_{fm} \subset J_m$ collects the products offered by f.

Formally, a market m corresponds to a triplet (J_m, χ_m, G_m), where $\chi_m = \left(x_m, r_m, (\xi_m^\tau)_{\tau \in \mathcal{T}}\right)$.

- The vector $x_m = (x_{jm})_{j \in J_m}$ collects exogenous products' characteristics for products sold in m.
- The vector $r_m = (r_{jm})_{j \in J_m}$ collects the rating functions in market m.
- For each $\tau \in \mathcal{T}$, $\xi_m^\tau = \left(\xi_{jm}^\tau\right)_{j \in J_m}$ collects the structural errors specific to the valuation of a buyer of type τ for product j in m. Each element ξ_{jm}^τ represents m-specific characteristics of product j, observed by buyers and insurers – thus affecting both pricing and purchase decisions – but unobserved by the econometrician.
- G_m is the m-specific distribution over (ζ, τ). For all m, $G_m(\zeta, \tau) = G_m(\tau)\, G(\zeta|\tau)$; with the distribution of characteristics ζ conditional on the rating type τ assumed constant across markets.

I assume that the m-specific $G_m(\tau)$ is observed with noise: the econometrician observes a distribution $\widetilde{G}_m(\tau) = G_m(\tau) + \eta_m(\tau)$.[2] The measure η_m is assumed to be drawn independently across markets from the space of measures over $\mathcal{T}$, and such that $\mathbb{E}[d\eta_m(\tau)] = 0$ for all τ, with the expectation taken across m. Noisy measurement of $G_m(\tau)$ is common in most applications. Often, the composition of the population of potential buyers across different τ's is constructed from external data sources, creating the measurement error η_m. The assumption that $\mathbb{E}[d\eta_m(\tau)] = 0$ for all τ means that, on average across markets, each τ is not systematically over-represented or under-represented. I assume instead that insurers have *common knowledge* of the true distribution $G_m(\tau)$, an important simplification which is largely motivated by the need to keep equilibrium analysis tractable.

Random Utility. The vector of non-τ characteristics ζ_i contains all the elements of a standard random utility model, $\zeta_i = (\beta_i, \epsilon_i)$. β_i is a vector of parameters specific to individual i, while $\epsilon_i = (\epsilon_{ij})_{j \in J_m}$ is a collection of i's idiosyncratic preference shocks, one for each product in m.

[2]For both G_m and $\widetilde{G}_m$ to be probability measures, this implicitly requires that $\int d\eta_m(\tau) = 0$ for all m.

The valuation for a product j is market-specific, and denoted by v_{jm}. In particular, for each τ one has

$$v_{jm}\left(\zeta_i, \tau\right) = \beta_i' x_{jm} + \epsilon_{ij} + \xi_{jm}^{\tau}. \tag{9}$$

After making distributional assumptions on the vector of taste-shocks ϵ_i, say $\epsilon_i \sim F(\epsilon)$, one obtains a re-writing of (3) as

$$d_{jm}\left(r_m(\tau); \tau | \chi_m\right) = \int \left(\int \mathbf{1}\left\{ \Gamma_j \left(\beta' x_m + \xi_m^{\tau} + \epsilon \right) \geq \Gamma_j r_m(\tau) \right\} dF(\epsilon) \right) dG(\beta|\tau), \tag{10}$$

where I emphasize the dependence of the model-predicted market shares on χ_m. If ϵ_i is assumed to follow a standard type I distribution, one obtains the standard expression for a mixed logit demand system, one for each rating type τ:

$$d_{jm}\left(r_m(\tau); \tau | \chi_m\right) = \int \frac{\exp\left(\beta' x_{jm} + \xi_{jm}^{\tau} - r_{jm}(\tau)\right)}{\exp\left(-r_{0m}(\tau)\right) + \sum_{j' \in J_m} \exp\left(\beta' x_{j'm} + \xi_{j'm}^{\tau} - r_{j'm}(\tau)\right)} dG(\beta|\tau),$$

where I adopted the usual normalization $v_{0m} = 0$. Notice that the price of the outside good for a given type τ, $r_{0m}(\tau)$, could be non-zero. This is used to model subsidies or tax penalties due to non-participation, as I will discuss further in Section 5.

Cost Functions. Assuming that the idiosyncratic component of preferences ϵ_i does not affect the cost of individual i, one can write $\phi_{jm}(\zeta, \tau) = \phi_{jm}(\beta, \tau)$. This function is then specified as

$$\phi_{jm}(\beta, \tau) = w_m \phi_f(\beta, \tau; x_{jm}), \tag{11}$$

where w_m is a multiplicative cost index observed in the data. This variable, normalized to unity in one market, describes differences in input costs across different geographic markets. Many measures of this kind are publicly available, since they are used to adjust reimbursements to physicians and clinics by public insurance programs. (A leading example is the county-level Medicare geographic adjustment factor used for fee-for-service payments; MaCurdy et al. (2014).)

The function ϕ_f captures the extent to which, for the same amount of insurable risk for a given individual, expected claims may differ across insurers and products; this heterogeneity can emerge from differences in the set of physicians covered by a specific insurer f, differences in the type of plan (e.g. HMO, PPO, or vertically integrated), or difference in expected utilization. With this cost structure, equation (4) of the theoretical model can be written as

$$c_{jm}\left(r_m(\tau); \tau | \chi_m\right) = w_m \int \phi_f(\beta, \tau; x_{jm}) \left(\int \mathbf{1}\left\{ \Gamma_j \left(\beta' x_m + \xi_m^{\tau} + \epsilon \right) \geq \Gamma_j r_m(\tau) \right\} dF(\epsilon) \right) dG(\beta|\tau). \tag{12}$$

Demand Equations. Upon observing, in each market m, market shares within all types $\tau \in \mathcal{T}$,

$s_m^\tau = \left(s_{jm}^\tau\right)_{j \in J_m}$, one has a set of equations corresponding to τ-specific demand curves:

$$s_{jm}^\tau = d_{jm}\left(r_m(\tau);\tau|\chi_m\right), \qquad \tau \in \mathcal{T},\ j \in J_m. \tag{13}$$

Supply Equations. Assuming that observed rates are an equilibrium in each market, and therefore that the first-order conditions (8) hold, the equations for optimal pricing depend on the regulation $\mathcal{R}^a$, and on market characteristics (J_m, χ_m, G_m), where $G_m(\tau)$ is replaced by the observed distribution $\widetilde{G}_m(\tau)$:

$$\begin{aligned} &\int \left(d_{jm}\left(r_m(\tau);\tau|\chi_m\right) + \sum_{j' \in J_{fm}} r_{j'm}(\tau) \frac{\partial d_{j'm}\left(r_m(\tau);\tau|\chi_m\right)}{\partial p_j} \right) \frac{\partial a\left(r_{jm}(\tau_0),\tau\right)}{\partial\left(r_j(\tau_0)\right)} d\widetilde{G}_m(\tau) \\ = &\int \left(\sum_{j' \in J_{fm}} \frac{\partial c_{j'm}\left(r_m(\tau);\tau|\chi_m\right)}{\partial p_j} \right) \frac{\partial a\left(r_{jm}(\tau_0),\tau\right)}{\partial\left(r_j(\tau_0)\right)} d\widetilde{G}_m(\tau), \qquad j \in J_m. \end{aligned} \tag{14}$$

Importantly, while an observation in the system of demand equations in (13) above corresponds to a pair (j,τ) in the data, in the pricing equations in (14) observations correspond to the number of rating decisions made by insurers, one for each product j in the data.

4 Identification

Identification of the model is presented in two distinct steps. I first explain the identification of demand primitives for each rating type τ, i.e. the collection $\left(\xi_m^\tau, G(\beta|\tau)\right)_{\tau\in\mathcal{T}}$. Then, I consider the identification of cost functions ϕ_f. Together, these primitives describe the joint distribution of risk and risk-preferences, needed to analyze alternative policy and welfare in health exchanges.

The reason to favor a two-step procedure over a simultaneous equations approach is twofold. First is transparency: Demand curves are identified under weak assumptions on the price-setting process, and estimated via common techniques. On the other hand, estimation of cost functions requires stronger assumptions, and employs a non-standard result. The second reason is practical implementation: The amount of flexibility in the way in which costs can vary in (β,τ) depends on conditions that must be met by the estimated demand system. These conditions can be verified only *after* estimating demand across τ's.

Demand Identification. The econometrician observes $\left(\widetilde{G}_m, J_m, x_m, w_m, r_m, (s_m^\tau)_{\tau\in\mathcal{T}}\right)$. The set of rating types $\mathcal{T}$ is a finite set, and for each τ, my setup and the discussion of demand identification is equivalent to the one in Section 4.1 and 4.2 in Berry and Haile (2014) (BH henceforth).

For any market (J_m, χ_m, G_m), considering separately each pair (J_m, χ_m^τ), my model is a subcase of BH: the assumptions on the random utility introduced in Section 3 imply their Assumptions 1, 2, and 5. These are, respectively, what they call *index restriction* – here implied by linearity in structural errors, *connected substitututes* – here corresponding to *reactiveness* of market shares to price changes, and *linearity in price* – assumed in my model. To leverage on Theorem 1 and Section 4.2 in BH, one needs to find proper instruments for prices, and observing sufficient variation in market shares and prices induced by variation in these instruments.

Under $\mathcal{R}^a$, for each τ, the price vector $p_m^\tau = r_m(\tau)$ depends on the distribution of rating types in market m. This is true in the equilibrium model presented in Section 2, but it would be true in a much broader class of supply models, including for example average cost pricing. In fact, for all pricing mechanism in which rates depend on the characteristics of the expected enrollment pool, a regulation such as $\mathcal{R}^a$ implies that prices for types τ depend on the composition of the population G_m, and thus on $\widetilde{G}_m$. (See also my discussion of estimation assuming alternative supply models in Appendix B.)

Assuming that composition of population in m – and in particular the number of potential buyers of different types – is independent from the unobservable determinants of τ's preferences, one can use $d\widetilde{G}_m(\tau')$, $\tau' \neq \tau$, as a set of valid instruments for p_m^τ. To formalize this, the assumption yielding a sufficient condition for identification of demand is:

Assumption **1** [Instruments for price] *For each $\tau \in \mathcal{T}$, there exist a set $\mathcal{I}(\tau) \subset \mathcal{T} \setminus \{\tau\}$, such that*

(a) $\mathbb{E}\left[\xi_m^\tau | \left(d\widetilde{G}_m(\tau')\right)_{\tau' \in \mathcal{I}(\tau)}, x_m\right] = 0$ *almost surely, and*

(b) *for all functions $B(s_m^\tau, p_m^\tau)$ with finite expectation,*
if $\mathbb{E}\left[B(s_m^\tau, p_m^\tau) \,|\, \left(\widetilde{G}_m(\tau')\right)_{\tau' \in \mathcal{I}(\tau)}, x_m\right] = 0$ *almost surely, then* $B(s_m^\tau, p_m^\tau) = 0$ *almost surely.*

From BH one has then the following:

Proposition **1** *If Assumptions 1 holds, then $G(\beta|\tau)$ and ξ_m^τ are identified for all τ.*

Proof. See Theorem 1 and Section 4.2 in Berry and Haile (2014).□

Discussion. Identification of demand explicitly exploits how, for a given set of rating restrictions $\mathcal{R}^a$, insurers will set optimal prices considering the distribution of the population across rating types τ, varying across markets. If individual preferences (e.g. young buyers) do not vary with how many people in a set of different types (e.g. old buyers) could potentially buy insurance in a the same market, the regulation $\mathcal{R}^a$ and a price mechanism accounting for the composition of the enrollment pool, generate together a source of exogenous variation in prices within each τ, and demand conditional on τ can is consistently estimated via standard techniques, common since Berry (1994); Berry et al. (1995), and recently formalized in Berry and Haile (2014).

Cost Identification. Consider a situation in which, for all τ, demand primitives $\widehat{G}(\beta|\tau)$ and $\widehat{\xi}_m^\tau$ have been consistently estimated. For each τ, for all products j in all markets m one can construct $\widehat{d}_{jm}(p;\tau|\chi_m)$ as in (10).

Also let B be the support of all possible random coefficient vectors β in the population. For any measurable subset $B' \subset B$, and any τ, one can define demand specific to the subset of population $\{(\beta_i, \epsilon_i, \tau_i) : \beta_i \in B', \tau_i = \tau\}$. Formally:

$$\widehat{d}_{jm}\left(p;\tau,B'|\chi_m\right) = \int_{B'} \left(\int \mathbf{1} \left\{ \Gamma_j \left(\beta' x_m + \widehat{\xi}_m^\tau + \epsilon \right) \geq \Gamma_j r_m(\tau) \right\} dF(\epsilon) \right) d\widehat{G}(\beta|\tau). \tag{15}$$

With this expression, for any finite partition of B, say $\mathcal{B} = \{B_k\}$, in disjoint, measurable sets $\{B_k\}$ (with $\cup_k B_k = B$) one can write:

$$\widehat{d}_{jm}(p;\tau|\chi_m) = \sum_k \widehat{d}_{jm}(p;\tau,B_k|\chi_m).$$

A relevant requirement for cost identification is that, for any product characteristics x_{jm}, and for any insurer f, the function $\phi_f(\beta,\tau;x_{jm})$ has to be constant within certain groups of buyers. This means assuming that the population can be divided in "minimal" sets of individual characteristics within which expected risk does not vary. Within these groups, total cost becomes the familiar product between total quantity (enrollment) and a constant marginal cost (risk). Formally, I introduce the following:

Assumption **2** [Finite cost types] *There exists a finitie, disjoint, measurable partition $\mathcal{B} = \{B_k\}$ of B, such that*

- *for all f, all τ, $\phi_f(\cdot,\tau;x_{jm})$ is constant in β for all $\beta \in B_k$, all $B_k \in \mathcal{B}$;*
- *$\widehat{d}_{jm}(p;\tau,B_k|\chi_m)$ defined in (15) is continuous and differentiable in p for all $B \in \mathcal{B}$.*

Under Assumption 2, the expression for costs within each τ in (12) can be written as

$$c_{jm}(r_m(\tau);\tau|\chi_m) = w_m \sum_k \phi_f^{\tau,k}(x_{jm})\, \widehat{d}_{jm}(r_m(\tau);\tau,B_k|\chi_m),$$

where $\phi_f^{\tau,k}(x_{jm})$ is the value taken by $\phi_f(\beta,\tau;x_{jm})$ for all $\beta \in B_k$ – assumed to be constant. Moreover, from the second part of Assumption 2, one can write

$$\frac{\partial c_{jm}(r_m(\tau);\tau|\chi_m)}{\partial p_j} = w_m \sum_k \phi_f^{\tau,k}(x_{jm}) \frac{\partial \widehat{d}_{jm}(r_m(\tau);\tau,B_k|\chi_m)}{\partial p_j}. \tag{16}$$

Through (16), marginal cost from buyers of type τ is expressed as the weighted sum of costs across subsets of buyers, with risk assumed constant within these subsets (Assumption 2). The weights in the sum are the derivatives of demand within the same subsets of buyers.

Next, I introduce two possible assumptions under which one can identify ϕ_f for all f. The first one restricts the variety of products that the same firm offers across markets.

Assumption **3** [Standard products] *For all firms f, $J_{fm} = J_f$ for all m.*

The other corresponds to the situation in which one can describe the insurance product with characteristic x_{jm} as share of total costs reimbursed – or actuarial value $AV(x_{jm})$.

Assumption **4** [Linear contracts] *For all firms f, $\phi_f(\beta, \tau; x_{jm}) = AV(x_{jm})\,\widetilde{\phi}_f(\beta, \tau)$, for some $\widetilde{\phi}_f$ and a known function $AV(x_{jm}) \in (0, 1)$.*

Depending on the regulatory context, and on the information available to the researcher, one assumption will be preferred to the other. Importantly, under Assumption 4, the type of product chosen cannot affect the individual's expected total claims $\widetilde{\phi}_f(\beta, \tau)$, ruling out "moral hazard" consideration where utilization can vary with the level of coverage.

Identification of ϕ_f relies on Assumptions 1,2,3; or on Assumptions 1,2,4. In each of the two cases, a sufficient condition for identification is provided by invertibility of a particular matrix; this requires variation in the composition of the set of marginal buyers at the observed prices across products in the data.

In what follows, I call a "dataset" the collection of all observed markets, $(J_m, \chi_m, G_m)_m$.

Under Assumptions 2, and 3, identification considers the *matrix of marginal buyers for firm f*, Ω_f:

Definition **1** *Mantain Assumptions 2 and 3.*
For a given dataset $(J_m, \chi_m, G_m)_m$, the matrix of marginal buyers for f *under the regulation $\mathcal{R}^a$ is the $(|m| \times |J_f|)$-by-$(|J_f| \times |\mathcal{T}| \times |\mathcal{B}|)$ matrix Ω_f, with generic entry*

$$\omega_f\left[(j, m), (j', \tau, B_k)\right] = \frac{\partial \widehat{d}_{j'm}(r_m(\tau); \tau, B_k|\chi_m)}{\partial p_j} \frac{\partial a(r_{jm}(\tau_0), \tau)}{\partial(r_j(\tau_0))} d\widetilde{G}_m(\tau).$$

A row is (j, m), a column is (j', τ, k): $j, j' \in J_f$, $\tau \in \mathcal{T}$, $B_k \in \mathcal{B}$.

In words, a row of Ω_f is a product (j, m) sold by f in the data. Each entry $\omega_f[(j, m), (j', \tau, k)]$ is the variation in the mass of buyers with $\tau_i = \tau$, and $\beta_i \in B_k$, purchasing a product j' sold by f in m, and reacting to small changes in the price of product j. With this in hand, the first identification result follows:

Proposition **2** *If Assumptions 1, 2, and 3 hold, rates are set in equilibrium under $\mathcal{R}^a$, and the matrix of marginal buyers Ω_f is full-column-rank almost surely for all f, then cost functions ϕ_f are identified.*

Proof. Under Assumptions 1, 2, and 3 demand is identified, and one can write the equilibrium supply equations (14) as

$$\frac{\widehat{MR}_{jm}}{w_m} = \sum_{j' \in J_f} \sum_{\tau} \sum_{k} \phi_f^{\tau,k}(x_j)\, \omega_f\left[(j,m),(j',\tau,B_k)\right], \qquad j \in J_m, \tag{17}$$

where $\widehat{MR}_{jm}$ (marginal revenues of product jm) is the left-hand side in (14), with d_{jm} replaced by the estimated $\widehat{d}_{jm}$. Let $\varepsilon_{jm}(\phi)$ be the difference between the left- and right-hand side of (17), evaluated at the cost functions $\phi = \left(\phi_f\right)_f$. For the true cost functions ϕ, since the measurement error η_m is independent from other observables, $\mathbb{E}\left[\varepsilon_{jm}(\phi)\right] = 0$. (Omitted algebra in the supplementary appendix.) Now, assume that, for a given $\widehat{\phi}$, $\varepsilon_{jm}\left(\widehat{\phi}\right) = \varepsilon_{jm}(\phi)$; that is, ϕ and $\widehat{\phi}$ are observationally equivalent. Since the left-hand side in (17) does not depend on ϕ, this implies that

$$\sum_{j' \in J_f} \sum_{\tau} \sum_{k} \left(\phi_f^{\tau,k}(x_j) - \widehat{\phi}_f^{\tau,k}(x_j)\right) \omega_f\left[(j,m),(j',\tau,B_k)\right] = 0,$$

and since Ω_f is full-column-rank almost surely, this implies that $\phi_f = \widehat{\phi}_f$, for all f almost surely. Then, ϕ_f is identified for all f.□

When products offered by the same firm vary across markets, Ω_f is not well defined. Identification of ϕ can then rely on Assumption 4, rather than 3, still exploiting variation in the composition of marginal buyers across products, although represented by a different matrix.

It is convenient to notice that, under linear contracts, insurers are covering risk units rather than head counts: when covering d buyers in product x_{jm}, the total expected risk is a share $AV(x_{jm})$ of the total risk of the d buyers. Furthermore, under Assumptions 2, and 4, the expected cost for the firm generated from insuring a given group of buyers (β,τ) with $\beta \in B_k$, can be aggregated across all the products offered by the same firm in market m, and counting the risk units sold within the group:

$$\text{Total cost from } (\beta,\tau),\ \beta \in B_k \quad = \quad \widetilde{\phi}_f(\beta,\tau) \sum_{j \in J_{fm}} \underbrace{AV(x_{jm})\, \widehat{d}_{jm}(r_m(\tau);\tau,B_k|\chi_m)}_{\text{Risk units enrolled in product } j \text{ among } (\beta,\tau),\ \beta\in B_k}.$$

Under Assumptions 2, and 4, identification considers the *matrix of marginal risk units for firm* f, $\widetilde{\Omega}_f$:

Definition 2 *Mantain Assumptions 2 and 4.*
For a given dataset $(J_m, \chi_m, G_m)_m$, *the* matrix of marginal risk units for f *under the regulation*

$\mathcal{R}^a$ is the $(\sum_m |J_{fm}|)$-by-$(|\mathcal{T}| \times |\mathcal{B}|)$ matrix $\widetilde{\Omega}_f$, with generic entry

$$\widetilde{\omega}_f\left[(j,m),(\tau,B_k)\right] = \sum_{j' \in J_{fm}} AV\left(x_{j'm}\right) \frac{\partial \widehat{d}_{j'm}\left(r_m(\tau);\tau,B_k|\chi_m\right)}{\partial p_j} \frac{\partial a\left(r_{jm}(\tau_0),\tau\right)}{\partial\left(r_j(\tau_0)\right)} d\widetilde{G}_m(\tau).$$

A row is (j,m), a column is (τ,k): $j \in J_{fm}$, $\tau \in \mathcal{T}$, $B_k \in \mathcal{B}$.

In words, a row of $\widetilde{\Omega}_f$ is a product (j,m) sold by f in the data. Each entry $\widetilde{\omega}_f\left[(j,m),(\tau,k)\right]$ is the variation in the total amount of risk units purchased from f in m by buyers with $\tau_i = \tau$, and $\beta_i \in B_k$, reacting to small changes in the price of product j.

The main difference with Ω_f is that, here, the allocation of buyers across products in J_{fm} can be aggregated through knowledge of the function AV. The matrix $\widetilde{\Omega}_f$ only describes changes in the aggregate composition of the enrollment pool across different (β,τ), where the relevant quantities are risk units rather than individuals, and is well defined even when the set of products J_{fm} varies across markets.

The second identification results requires sufficient variation in the composition of marginal buyers after aggregating costs in risk units:

Proposition **3** *If Assumptions 1, 2, and 4 hold, rates are set in equilibrium under $\mathcal{R}^a$, and the matrix of marginal risk units $\widetilde{\Omega}_f$ is full-column-rank almost surely for all f, then cost functions ϕ_f are identified.*

Proof. Under Assumptions 1, 2, and 4 demand is identified, and one can write the equilibrium supply equations (14) as

$$\frac{\widehat{MR}_{jm}}{w_m} = \sum_{\tau}\sum_{k} \widetilde{\phi}_f^{\tau,k} \widetilde{\omega}_f\left[(j,m),(\tau,B_k)\right], \qquad j \in J_m, \tag{18}$$

where $\widehat{MR}_{jm}$ (marginal revenues of product jm) is the left-hand side in (14), with d_{jm} replaced by the estimated $\widehat{d}_{jm}$, and $\widetilde{\phi}_f^{\tau,k} = \phi_f^{\tau,k}\left(x_{jm}\right)/AV\left(x_{jm}\right)$, constant under Assumption 4. As in the proof of Proposition 2, $\varepsilon_{jm}\left(\widetilde{\phi}\right)$ is the difference between the left- and right-hand side of (18), evaluated at $\widetilde{\phi}$. At the true $\widetilde{\phi}$, $\mathbb{E}\left[\varepsilon_{jm}\left(\widetilde{\phi}\right)\right] = 0$. If, for a given $\widehat{\phi}$, $\varepsilon_{jm}\left(\widehat{\phi}\right) = \varepsilon_{jm}\left(\widetilde{\phi}\right)$, this implies that

$$\sum_{\tau}\sum_{k}\left(\widetilde{\phi}_f^{\tau,k} - \widehat{\phi}_f^{\tau,k}\right)\widetilde{\omega}_f\left[(j,m),(\tau,B_k)\right] = 0,$$

and since $\widetilde{\Omega}_f$ is full-column-rank almost surely, this implies that $\widetilde{\phi}_f = \widehat{\phi}_f$, for all f almost surely. Then, $\widetilde{\phi}_f$ (and thus ϕ_f) is identified for all f.□

Discussion. After estimating demand, to estimate cost (risk) allowed to vary across different subsets of buyers, one can relax the reliance on additional cost data by making assumptions on the pricing mechanism, and imposing restrictions on the *richness* of cost heterogeneity within each rating type τ. For the pricing mechanism, here I focus on equilibrium pricing as presented in Section 2, but similar results could be derived for the case of average cost pricing; see Appendix B. Assumption 2 restricts the extent to which consumers with different risk preferences can have different risk: within any rating type τ, there needs to be a finite number of possible risk profiles, each corresponding to subsets of preference parameters β.

In particular, to apply Proposition 2, one needs to have that the number of types for which risk can vary is less than the number of markets in the data. Similarly, to apply Proposition 3, the number of types for which a different risk profile can be estimated is less than the total number of observations for a given firm. Moreover, sufficient variation in the composition of marginal buyers across the different groups must be observed, as formalized by requiring Ω_f (respectively $\widetilde{\Omega}_f$) to be full-column-rank. Economically, this requires that different products in the data are facing different compositions of marginal buyers, following changes in population composition and/or market structure across markets.

Under these assumptions, the key idea behind indentification of ϕ_f can be simplified to the following: Suppose that prices are such that marginal revenues are equal to marginal costs. Assumption 2 implies that marginal costs can be expressed as a weighted average of costs across different types of buyers, with the weights being shares of marginal buyers across different types. As long as demand for each type is estimated, the number of marginal buyers by type can be constructed for each product in the data (equation (16)). Then, variation in these weights, i.e. variation in the composition of marginal buyers across products in the data, can be used to recover costs varying across different types.

A simple example can be used to clarify ideas. Consider a case in which $\mathcal{T} = \{Young, Old\}$, so that τ represents age of the buyer, $\mathcal{R}^a$ imposes that prices must be equal across τ's, and there is no heterogeneity within τ – i.e. $\mathcal{B} = \{B\}$ in Assumption 2. To identify costs varying by τ two requirements must be met. First, it is necessary to observe two products in the data; this differs from the common framework in which observation of one product and knowledge of elasticities are sufficient to determine the product's marginal cost (e.g. Bresnahan (1981)). Second, if the first product's marginal buyers are a combination of $Young$ and Old in a 1:2 ratio, the second product must have a combination of marginal buyers across $Young$ and Old in a ratio 1:X, with $X \neq 2$. If these two requirements are verified, costs across the two groups are distinguished by the data, since there is at most one solution to the linear system defined by equilibrium first-order conditions

$$\begin{aligned} \text{Marginal revenue product 1} \quad &= \quad \text{Cost } Young + 2 \text{ Cost } Old \\ \text{Marginal revenue product 2} \quad &= \quad \text{Cost } Young + X \text{ Cost } Old. \end{aligned}$$

In what follows, I present the institutional details of health exchanges regulated by the Affordable Care Act, then go on and estimate a model of the Californian exchange with data from the first year of operations.

5 ACA Health Insurance Exchanges

Individual Market after the ACA. The ACA affected the individual market starting in late 2013. Since then, health insurers who want to offer products in this market are required to comply with many regulations common to "managed competition" in health insurance markets. These pose constraints on (i) the number and "quality" of products offered (both in terms of actuarial value and network of physicians and hospitals), (ii) possible adjustments to the product mix across different geographic locations, (iii) timing of entry/exit decisions as well as changes in prices and/or product characteristics.

An exchange is the central marketplace for the individual market within a state.[3] The territory is divided in rating regions (groups of counties), each operating as a separate market within which the regulations on entry, product mix determination, and pricing apply.

During the spring of every year, insurers announce their participation in a region and need to obtain the necessary qualification for their insurance products. In practice, the exchange verifies that the network of physicians and hospitals covered by a plan is adequate and complies with the minimum standards imposed by the federal law.

After entry decisions are made (and approved), insurers determine their prices and product characteristics. Products are "binned" in metal tiers (Catastrophic, Bronze, Silver, Gold, and Platinum) according to their projected actuarial value for a representative buyer, and every participating insurer has to offer at least two plans (Silver and Gold). In some states (e.g. California, focus of my empirical application) products are fully standardized: the exchange requires the insurer to offer one product in each of the four metal tiers (excluding Catastrophic coverage), and the financial characteristics are fixed across all products within a tier. In practice, a buyer in these states will find that all Bronze plans have the same, identical mix of deductible, co-pays, and out-of-pocket limits; this being true for all tiers. [4]

Each product in a region must be available to any buyer without further restrictions,[5] and pricing is constrained by the rating regulation that I spell out below in thorough details. Once

[3]Insurers can offer products outside the exchange, but they are subject to the same regulations about product characteristics, pricing, and minimum-loss-ratio. The key differences are that federal subsidies are only available for products in the exchange, and that the timing for entry/exit – as well as changes in product mix – are not constrained as in the exchange.

[4]The importance of product standardization when regulating health insurance markets is the main focus of Ericson and Starc (2013).

[5]This is the most dramatic change from the pre-reform era, when insurers were able to adjust the product mix, and even deny coverage, based on the individual characteristics of each buyer.

prices are determined (between late summer and early fall) they are publicly announced along with the products' details, and cannot be adjusted until the following year. Buyers can access the market only during a fixed "open enrollment" window, corresponding to the last months of the calendar year.[6] They are free to choose any product in their region and buy insurance for the following year, with coverage starting on the first day of January.

Beyond establishing these rules constraining the supply of health insurance, the government's role further extends to guarantee affordability of coverage for low-income buyers. In what follows I present together the details of the price adjustments across buyers with different characteristics implied by rating rules and the subsidy program. These will be the key rating adjustments, as captured by $\mathcal{R}^a$ in the model introduced in Section 2.

Rating Regulations and Federal Subsidies. Within a rating region (geographic market), for any insurance plan in their product mix, carriers can only set one rate, corresponding to the annual premium for a 21 years old buyer who is not eligible for federal subsidies, say P (a choose 21 as a normalization, without loss of generality). Each premium is then age-adjusted: that is, buyers of different ages will pay a premium equal to $F \times P$, where F is an "age factor" (with $F = 1$ for 21-year-old), determined by the Center of Medicare and Medicaid Services. This "standard age rating curve" is reported in Table 1. Each exchange could alter the curve, but 48 states (including California) choose to adopt the standard one.

After the age-adjustment, a buyer whose income is below four times the federal povery level (FPL) receives a discount, an amount which is paid to the insurer by the government. Under ACA regulations, this amount is determined so that the buyer can afford the second cheapest Silver plan (or "benchmark" plan) for a pre-determined fraction of her income. This varies from 4-9% of the buyer's yearly income, as summarized by Table 2. Buyers can use the discount for the purchase of any product in the region (excluding Catastrophic coverage), with prices bounded below by zero. Additional rate adjustments are based on tobacco use (up to a factor of 5) and family status, with specific rules varying across states.

Focusing on age adjustments and subsidies for individual adults, it is evident how – taking entry and product mix as fixed – the insurer's decision is reduced to a one-dimensional element: the rate for a single group of buyers (21-year-olds with income above four times the FPL). This is translated by the regulation in the price faced by all buyers of different income and age within a region.

These institutional details are nested by the model I introduced in Section 2:

- Rating types $\tau = (\text{age}, \text{income})$;
- The standardized type τ_0 is the 21-years-old, unsubsidized buyer;

[6] The year 1 open enrollment took place from October-March, in year 2 it is reduced to November-February, and in the upcoming years will be November-December.

- $\mathcal{R}^a$ is defined in (7), where the function a reduced to applying the age-factors provided by the standard age-rating curve (Table 1);

- The price of the outside option $j = 0$ is equal to zero for unsubsidized buyers, while it is equal to the discount determined under the ACA formula for subsidized buyers. For a buyer τ eligible for subsidies and with age factor F_τ, suppose that $\bar{c}_\tau$ is the maximum contribution computed using Table 2, and let $b^*(\tau_0)$ be the price of the benchmark plan for a 21-years-old. Then, the price of the outside option for τ is
$$r_0(\tau) = \max\{0, \bar{c} - F \times b^*(\tau_0)\}.$$

6 First Year of the Californian exchange

I estimate the model using data from the first enrollment period of the Californian exchange, Covered California (CoCA). This is a state-managed exchange that registered over 1.7 million enrollees in year 1, making it the largest in the US individual market after the ACA.

Representativenness. The regulatory framework complies with the ACA regulations outlined in Section 5, with two important additional restrictions that support my modelling assumptions. First, product characteristics are exogenous. State regulators determined the details of the set of products that must be offered by all participating insurers. While the ACA only requires that two products with approximate actuarial value (AV) be offered (one Silver plan – $\approx$ 70% AV, and one Gold plan – $\approx$ 80% AV), in California insurers must offer four products (Bronze, Silver, Gold, and Platinum) whose financial characteristics (deductible, out-of-pocket limits, and specific co-payments and co-insurance rates) are fixed by the regulation. Second, only age rating and subsidy are applied: in this CoCA departs from the less stringent ACA regulation by banning price adjustments based on tobacco use.

Entry and Other Dynamic Considerations. The Californian exchange established a final set of regulations with large anticipation and operated under considerably less uncertainty when compared to the experience of many other states. This reduces the concern about this market being far from the equilibrium outcomes that we would expect to observe several years from now. In particular, my choice to study rating decisions considering entry as given and publicly observed is consistent with the timeline of the Californian exchange: insurers announced their participation and obtained all the necessary certifications before March 2013, after the final set of rules was also announced. After several months the final price schedule was announced and made public (late August). Enrollment opened regularly on October 1st, and the technical difficulties preventing consumers from accessing the platform – a widespread problem in many states – did not appear as a significant issue in the Californian experience.

Of course, my analysis still suffers from the large uncertainty surrounding the market in the early years. Although exchange regulators made an effort to provide issuers with information about the potential population of buyers, with emphasis on the number of them who were expected to be eligible for subsidies, it is reasonable to expect that insurers will be able to experiment and optimize their competitive behavior over time. Fortunately, unlike in many states where the churn in participation completely redesigned market structures going from the first to the second year of the ACA implementation, in California the set of insurers active in each region stayed the same across the two periods and price adjustments from year one to year two were small.

6.1 Data

The data used for estimation consists of rating and enrollment data from year 1 of CoCA. Provided by the exchange, this data reports quantities purchased for each plan as well as composition of these quantities in terms of age (divided in three bins: 20-29, 30-44, and 45-64) and income status (subsidized or not). I complement these data with information about the set of potential buyers, also decomposed by age and income status for the same six types. Potential buyers information are derived from CoCA estimates and Census data.

The state of California was divided in 19 regions, all consisting of groups of counties (except Los Angeles County, broken down in two separate regions). In each region the number of participating insurers varies between 3 and 6, with a total of 11 carriers operating in at least one region in the state. This amounts to a total of 380 unique products: metal tier, insurer, region combinations.

Panel (a) of Table 3 shows summary statistics for the rate charged to a 21-year-old across the products available in CoCA, while panel (b) shows the average premium paid by consumers, further broken down by metal tier.[7] The standardized product characteristics are fixed across insurers within a given tier, as reported in panel (c) of Table 3.

Variation in composition of potential buyers and market structure across regions is summarized in Table 4. This represents the source of exogenous variation exploited by the model to identify demand and, in the second estimation step, costs conditional on type of buyer.

Market shares, which together with premiums are the dependent variable in the econometric model, are summarized in Table 5. This shows a great deal of heterogeneity across income status, insurer, and age. Differences across type of buyers are also emphasized in Figure 1, plotting market shares against yearly premium for the six groups of buyers observed in the data.

6.2 Parametrization

Demand. I estimate six Logit demand systems, one for each of the six groups observed in the data. Specifically, I posit that for any $\tau = (age, income)$ a buyer i in market (or region) m derives

[7] Catastrophic coverage is not available to subsidy beneficiaries.

(money-metric) utility from product $j \in \mathcal{J}_r$ equal to

$$u_{i,j,\tau,m} = -\alpha_\tau p_{j,\tau,m} + \beta'_\tau x_j + \xi_{j,\tau,m} + \epsilon_{i,j,\tau,m} \tag{19}$$

where x_j contains deductible, maximum out-of-pocket, and insurer fixed effects, $\xi_{j,\tau,m}$ is a region-age-insurer-specific unobservable – interpretable as utility from the specific network of providers, or as other unobservable factors affecting demand from the group, e.g. advertising. The idiosyncratic error term $\epsilon_{i,j,\tau,m}$ is assumed to follow a standard Type I distribution.

Denoting by $s_{j,\tau,m}$ the market share of product $j \in J_m$ among buyers of type τ, the estimating equation for demand (corresponding to equation (13) in Section 3) becomes then

$$s_{j,\tau,m} = \frac{\exp\left\{-\alpha_\tau p_{j,\tau,m} + \beta'_\tau x_j + \xi_{j,\tau,m}\right\}}{1 + \sum_{\widehat{j} \in J_m} \exp\left\{-\alpha_\tau p_{\widehat{j},\tau,m} + \beta'_\tau x_{\widehat{j}} + \xi_{\widehat{j},\tau,m}\right\}}, \tag{20}$$

which can be transformed in the usual linear form:

$$\ln\left(s_{j,\tau,m}\right) - \ln\left(s_{0,\tau,m}\right) = -\alpha_\tau p_{j,\tau,m} + \beta'_\tau x_j + \xi_{j,\tau,m}, \tag{21}$$

where $s_{0,\tau,m}$ is the fraction of potential buyers of type τ who decide not to purchase any product in the exchange.

To compute the discount received by subsidized buyers (thus the price of the outside option for each group), I consider the 150-200% bin in Table 2; this was the most numerous group among subsidized in the first year of the exchange.

Cost. The cost specification takes a simple linear form, where expected health expenditure is allowed to vary by age and income. In applying Assumption 2 in this context, however, this simple version of the model assumes that expected health expenditure does not vary within an income-age pair, where I consider a coarser age grouping: 20-44, and 45-64. In terms of the notation used in Section 4,

$$\mathcal{B} = \left\{(\text{20-44,subsidized}), (\text{45-64,subsidized}), (\text{20-44,unsubsidized}), (\text{45-64,unsubsidized})\right\}.$$

Additional data, or estimating a richer demand system, would allow me to relax this restrictions.

I estimate costs under Assumption 4, applying Proposition 3. Note, however, that Assumption 3 holds for CoCA, and I will consider costs estimated applying proposition 2 in my robustness checks. When enrolling a buyer of type τ in market m in a plan j, insurer f faces expected costs equal to $w_m \phi_f(\tau; x_j)$, where w_m is the population-weighted Medicare Geographic Adjustment Factor

(GAF) in the region (MaCurdy et al., 2014), and

$$\phi_f(\tau; x_j) = AV(x_j)\,\widetilde{\phi}_f(\tau)\,. \tag{22}$$

Here, the actuarial value function AV takes values equal to 0.5, 0.6, 0.7, 0.8, 0.9, for Catastrophic, Bronze, Silver, Gold, and Platinum plans, respectively. The target of estimation is then $\widetilde{\phi}_f$ evaluated at each element of the partition $\mathcal{B}$ introduced above. That is, for each insurer, and for each element of $\mathcal{B}$, I estimate an average expected health expenditure when buyers enroll in a plan offered by f. These are translated in realized costs by the AV function and the market-specific GAF w_m.

6.3 Estimation Results

As discussed in Section 4, my empirical strategy is broken down in a two step procedure. First, I estimate demand for each τ using a standard IV Logit approach. Second, I construct the supply equation and estimate costs. Confidence intervals for the second step estimates are constructed bootstrapping the entire estimation procedure.

First-Stage. The rating regulation is such that the premium faced by buyers of a given type τ depends on the number of potential buyers of different types $\widehat{\tau}$. This can be seen explicitly from the insurer first-order condition in equilibrium. By assuming that the unobservable error term $\xi_{j,\tau,m}$ does not depend on the number of type $\widehat{\tau}$ buyers – e.g. the number of potential buyers between 45-64 years of age in a regionof does not affect the individual demand of a 21-years-old – this provides valid instruments for premium to consistently estimate demand.

To show that population composition does affect equilibrium prices, in Table 6 I show the output of a "first-stage" OLS regression. This shows the relationship between population composition in a market m and the rate $r_{j,m}(\tau_0)$ of a product for types τ_0 (21-year-olds, unsubsidized). The premium for a young, unsubsidized buyer increases with the share of potential buyers older than 45, decreases in the share of potential buyers eligible for subsidies, as well as with the total number of potential buyers in the market. As expected, plans with higher deductible and maximum out-of-pocket expenditure are offered at cheaper annual premiums. The coefficients are robust to the inclusion of insurer fixed effects in both their magnitudes and statistical significance.

Interpreting the results, these imply that a 10% increase in the share of potential buyers over 45 corresponds to a 7.6% increase in annual premium for young buyers who do not benefit of federal subsidies.[8] Similarly, a 10% increase in the share of potential buyers eligible for subsidies is associated to a 14% decrease in annual premium for this baseline group.

[8]Consistent with this estimate from Californian data only, in Orsini and Tebaldi (2014) we find that, when looking at the 34 federally-run HIXs, under the age-rating regulation a 10% increase in the share of buyers older than 45 increases premiums of benchmark (second cheapest) Silver plans by approximately 4-6%.

This provides evidence that the relationship between market composition and insurer incentives anticipated by the theoretical model is empirically relevant. This is exploited as source of exogenous variation to estimate demand.

IV Logit Demand Estimates. The coefficient estimates from equation (21), one for each of the six groups in the data, are reported in Table 7. I find that price sensitivity decreases in age within each income group, and it is higher for beneficiaries of federal subsidies.

When comparing average "own-price elasticity", i.e. considering the percent drop in demand away from a product increasing its premium by one percentage point, the most elastic group corresponds to the 20-29 buyers eligible for subsidies, with an average elasticity equal to 3.37 percent (approximately \$17). The least price elastic group corresponds instead to over 45 buyers who are not eligible for subsidies because their income exceeds four times the FPL. In response to a one percent premium increase (approximately \$90), a product loses on average 1.01 percent of its demanders.

The demand estimates also show that subsidized buyers experience a higher disutility from increases in annual deductible and maximum out-of-pocket expenditure. The latter does not seem to significantly affect the utility of high-income buyers.

Finally, insurer-specific fixed-effects suggest that horizontal preferences across insurers (likely due to different perceptions of the network of providers and different brand/advertising effects) are different across the two income groups, while do not differ across ages within a given income group.

Inversion of FOC and Cost Estimates. With the demand estimates reported in Table 7 I can estimate cost functions using the supply equations (14) of Section 3. As discussed in Section 4, for each insurer[9] I construct the matrix of marginal risk units for each product in the data (Definition 2), and regress marginal revenues on the marginal risk units for each of the four groups in $\mathcal{B}$ (two income groups, and two age groups).

Here I allow for insurer-specific costs only for subsidized buyers (approximately 89% of enrollees). Costs for richer, unsubsidized buyers are already estimated with very large variance when not allowing them to vary by carrier. In fact, variation in marginal risk units for this group is very small, and this group consists of a negligible fraction of marginal buyers for most products in the data.

The cost estimates resulting from this estimation step are reported in Table 8. I find that subsidized buyers are on average riskier than their wealthier counterparts (although for the latter group estimates are very imprecise, with confidence intervals fully containing those obtained for subsidized buyers). Among subsidized buyers, those with age between 20-44 have expected expenditures equal to \$3,760, when averaging across insurers. This grows to \$6,808 for those in the same

[9] I consider the four largest carriers, active in at least ten markets, Anthem, BCBS, Kaiser, and HealthNet, and create a residual group which I call "Other".

income group with age between 45-64.

These estimates show some heterogeneity in expected costs across insurers, where it is important to recall that my model does not distinguish between the case in which these differences are due to different utilization (here induced by differences in the set of health providers covered by a carrier, since financial characteristics are standardized), or by differences in the specific reimbursement to providers bargained by each single insurer. Among the major four carriers in CoCA, the one with the lowest costs of covering subsidized buyers is HealthNet, while the one with the highest costs is Kaiser Permanente.

Importantly, these estimates varying by age and income of a buyer were obtained without any information on differences in costs across these groups. For this, I relied on Proposition 3, exploiting the variation in the composition of marginal buyers across products in the data, and assuming that the first-order optimality conditions characterizing equilibrium hold at the observed prices.

Before using my estimates to characterize the insurer pricing problem under ACA regulations, it is important to verify how different these are from those that one would obtain using a representative survey. For this purpose, I choose the Medical Expenditure Panel Survey, widely used in the empirical literature on health insurance markets, and focus on the variable characterizing annual health expenditure considering the sample of those who are privately insured, with age between 20-64. This comparison is reported in Table 9, and confirms that my estimates are very close to what would be obtained with external data.

The Insurer Problem: Regulations and Selection. With the estimates of demand and expected costs presented above, one obtains a full description of the problem faced by insurers when setting prices under ACA regulations. In particular, it is possible to show explicitly how changes in the rate for the 21-year-olds unsubsidized determine the number of enrollees, the composition of the enrollment pool, and average revenues and costs for each plan; these map directly in expected profits. The primitives used fort this purpose are the main input for counterfactual analyses in which one can simulate equilibrium under different policies, as I do in Section 7.

Using a simple example, in Figure 2 I show how, accounting for imperfect competition and differentiated products, different plans in the dataset face a rating problem with different characteristics.

I consider the rating region corresponding to Los Angeles County, and for two plans – a Bronze plan offered by HealthNet, and a Gold plan offered by Kaiser Permanente – I show the determinants of the plans' expected profits, where the carrier can choose the point on the horizontal axis (*Enrollment*) by altering the annual premium for 21-year-olds unsubsidized. This relationship between the choice variable ($r_{jm}(\tau_0)$ using the notation of Section 2) and total demand is omitted from the graph, but it is immediately implied by the rating regulation ($\mathcal{R}^a$, which maps $r_{jm}(\tau_0)$ in $r_{jm}(\tau)$ for all τ), demand estimates ($\widehat{d}_{jm}(\cdot;\tau)$; Table 7), and the composition of potential buyers ($\widetilde{G}_m(\tau)$;

Table 4). The graphs in the figure are constructed holding all other prices as fixed, and varying the rate for the 21-year-olds unsubsidized for the single plan.

Clearly, increases in enrollment are induced by reductions in premiums; the insurer directly reduces $r_{jm}(\tau_0)$, this in turns reduces prices for all τ's and attracts new enrollees. At the same time, the average per-enrollee revenue (vertical axis in the top panels of Figure 2) changes, since different types react differently to such price changes, and therefore the composition of the enrollment pool also changes (bottom panels of Figue 2). These changes in the composition of enrollment pools also determine directly the slope of the average cost curve for the plan. Traditional tests for adverse selection focused precisely on this slope (see e.g. Einav and Finkelstein (2011) for a review).

Looking at the top panels of Figure 2, in which the vertical dashed lines indicate the levels observed in the 2014 data, insurers choose a position on the horizontal axis (enrollment, or quantity) in order to maximize the area of the rectangular region lying below the average revenue (price) and above the average cost, as it would be in a standard "price-quantity" graph. In equilibrium, this is done simultaneously by all insurers.

From the figure one can notice how, after estimating demand, and thus constructing the average revenue curves and the composition of enrollment (bottom panels) for all of the 380 plans in the data, the goal of the second-step of estimation is to find values of expected expenditure – one for each of the different groups of buyers – such that the resulting average cost curves are consistent with simultaneous profit maximization by all insurers in all markets. In other words, cost estimation relies on altering cost curves simultaneously for all plans offered by a given carrier, imposing consistency with maximization of expected profits taking other prices as given.

From the resulting estimates, the price-quantiy graph and the enrollment composition for the two plans in my example feature different properties (and full equilibrium computation relies on 380 such graphs). In particular, the Bronze plan (left panels in Figure 2) faces an average revenue curve presenting a kink, and at this point the average cost curve changes slope, from upward ("advantageous selection") to downward ("adverse selection").

This pattern is induced by age-rating and the ACA subsidy scheme. In particular, the region on the right of the kink, in which average cost is downward sloping, corresponds to a region in which the older group of subsidized buyers faces a zero premium for the Bronze plan, while other groups face the full change in premium induced by the insurer and the age-rating adjustments. This is because, under the ACA formula, the discount is computed to ensure affordability of a Silver plan (Table 2 in Section 5). Since prices are higher for older buyers (following the age-rating curve in Table 1 of Section 5), for this group discounts are larger, and if a low-generosity Bronze plan lowers its price, the resulting prices for which these consumers are directly responsible hit the zero bound. In this region, subsidized consumers older than 45 are insesitive to price changes, becoming the least price sensitive group. At the same time, as shown in Table 8, this is also the riskier group; this correspondence between low price sensitivity and high risk induces adverse selection in this region of prices, resulting in a downward sloping average cost curve.

As the price raises, and eventually all buyers face price changes, the low-risk unsubsidized buyers become again the least price sensitive, as it would be absent the subsidy program. In this region (on the left of the kink in the left panels of Figure 2), one observes then what is referred in the literature as "advantageous selection", with low-price sensitivity buyers being also low-risk relative to others, and inducing an upward sloping average cost curve.

The pattern observed in the pricing problem of the low-coverage Bronze plan differs from that one found when looking at the more generous Gold plan offered by Kaiser Permanente. Noticeably, this plan attracts a higher share of high-income unsubsidized buyers, which are the least price sensitive and also lower-risk. As the plan reduces its premiums (moving enrollment toward the right on the horizontal axis), more and more low-income buyers choose to enter the enrollment pool (shown in the bottom right panel of Figure 2), raising average cost. Therefore, this high-coverage plan faces advantageous selection.

This is just an example of comparative static using two competing plans, while equilibrium computations as the one presented in the next section consider all the 380 products, and how they interact with each other across the 19 Californian markets. By looking at these two cases, however, it is clear that a model that accounts for the composition of enrollment (in terms of risk and preferences) and allows for imperfect competition among multi-product, differentiated insurers, is necessary to study regulation, selection, and competition in this market. As clear from Figure 2, a unique label (e.g. presence of "adverse selection") would be insufficient in characterizing the incentives at work in the determination of market outcomes.

7 ACA Subsidies or Voucher Program?

After providing a description of the structure of risk and preferences in the Californian exchange, and characterizing the incentives of insurers when setting premiums under ACA regulation, in this section I study the effect of the subsidy program, and the consequences of altering the way in which discounts to low-income buyers are determined.

My analysis is motivated by noticing that, under the subsidy scheme mandated by the ACA, described in Section 5, the subsidy levels are linked to the level of premiums. In simple words: the higher the price (of Silver plans), the more generous the subsidy, since the government established that the second cheapest plan in the Silver coverage level has to be afforded with a constant fraction of the buyer's income (Table 2 in Section 5).

In a market where marginal costs are not correlated with a buyer's preferences, i.e. a market without "selection", this way to compute a subsidy would always yield higher equilibrium prices than those realized under a constant discount. In fact, a product whose price determines the subsidy level (or any product expecting to determine the subsidy with some probability), faces an artificially lower price elasticity, since a \$1 increase in price is covered by the government (as long as the product does not become more expensive than the third cheapest in the tier). Standard comparative static results (Topkis, 1978) imply that such product would raise its price, and as long as prices are strategic complements (Vives, 1990) equilibrium prices will be higher than under a constant discount, in which the slope of demand faced by a seller is not altered by prices.

In a health insurance exchange, however, it is hard (if not impossible) to make unambiguous theoretical predictions regarding the direction taken by equilibrium outcomes under different designs of the subsidy program. Certainly, as emphasized in Einav and Levin (2015), a challenge of economists and policymakers is to come up with effective ways to support low-income buyers in the purchase of health insurance coverage, while also preserving the competitive incentives that induce insurers in charging low premiums. Without considering how subsidy programs artificially reduce price elasticity, and doing so lower the competitive pressure on prices, it is possible that the welfare benefits of "managed competition" are allocated largely to insurers, rather than consumers, and public spending risks to increase due to higher prices (Curto et al., 2014).

An interesting comparison is the one between a subsidy program where discounts are computed via the ACA formula (as in the status quo), and a program where all the regulations are unchanged, but the value of discounts to low-income buyers is pre-set by the regulators. To adopt such policy, a key choice would regard the size of the pre-set discounts (or voucher), as the equilibrium outcomes (prices, enrollment, markups, and government spending) depend directly on this parameter.

To study this alternative without a prior on the "correct" voucher amount, I simulate equilibria for a value of the discount to a low-income buyer varying between 0-100% of the buyer's expected annual health expenditure. When the fraction is zero, this corresponds to a situation without any subsidy, which is per-se an extreme yet interesting scenario, since the legitimacy of the federal program is currently under the scrutiny of the US Supreme Court.

For each value of the voucher I simulate equilibrium prices imposing the first-order conditions (equation (8) in Section 2), and checking that the second-order conditions are satisfied. I then compute enrollment, average revenue, average per-enrollee markups, and total government outlays for the subsidy program. In Figure 3 I compare these four market outcomes to the levels obtained in the status-quo, where discounts are computed via the ACA formula and insurers' incentives internalize the extent to which higher prices imply higher discounts, leaving demand unchanged.

From this curves, one can draw a simultaneous comparison between these outcomes under the ACA, as opposed to the outcomes obtained by simulating equilibrium under a voucher program, for any given level of the voucher. This is done in Figure 4, summarizing the key result of my counterfactual.

I start by discussing the effect of a complete shut-down of the subsidy program, leaving age-rating as the only regulation affecting pricing in the Californian market; this corresponds to a voucher value equal to zero percent of the buyer's expected expenditure. Equilibrium under this scenario shows the importance of the subsidy program in guaranteeing coverage of the low-income groups. In fact, without subsidies, enrollment would fall by approximately 80% from the ACA level in 2014. While government expenditure would be null, insurers would be able to charge higher markups (+20%), since high-income buyers – becoming the majority of the enrollees – are less price-sensitive. As a result, despite richer buyers being less risky, prices would fall only slightly compared to those observed when the ACA subsidy scheme is in place.

As the voucher amount increases, enrollment and government expenditure also increase, while markups are reduced due to both, higher costs and higher price sensitivity within enrollment pools, which containin an increasing share of low-income consumers. The tradeoffs are clear: the government can cover the low-income uninsured by paying for them, but this implies a progressively larger burden on taxpayers, while markpus first decrease (due to the composition effect, with low-income being costlier to serve but also more price sensitive) but then start to increase again due to the lower and lower share of premium for which subsidized buyers are directly responsible.

The behavior of the four lines in Figure 4 deliver the main result of my analysis: adopting a voucher program with discounts equal to 70-80% of a buyer's expected health expenditure, enrollment would (weakly) increase, government spending would (weakly) decrease, while markups and average revenues would go down. This translates to a welfare transfer from insurers to consumers and/or taxpayers.

Looking closely at this region of the graph, emphasized by the two vertical lines which correspond, respectively, to a voucher equal to 70 and 80% of expected expenditure, one can also read this result in absolute magnitudes, using the graphs in Figure 3. With a voucher equal to 70% of expenditure, enrollment would be approximately equal to the one realized under ACA regulations, while government expenditure in subsidies would fall by 20%, for savings exceeding $1B. Average premiums would fall by 15% ($800), and markups by 25% ($550).

At the opposite extremum of the range of voucher values which seem to dominate the status-quo subsidy design, a voucher equal to 80% of expenditure would increase enrollment by approximately 10%, equal to 120,000 additional enrollees. In this case, government outlays would be unchanged, staying equal to the $5.5B provided in California in 2014. The change in average revenues and markups would be slightly smaller, but of similar magnitude.

These results might suffer from the use of a small dataset, which only allows me to estimate a coarse distribution of risk and preferences in the population served by the Californian exchange. However, they are indicative that a properly calibrated voucher program might dominate the subsidy scheme mandated by the ACA, transferring welfare gains away from insurers, favoring consumers and/or reducing the burden on public finances.

8 Conclusions

In this paper I provided an econometric framework to analyze competition and regulation in regulated health insurance exchanges, motivated by an empirical application to the post-ACA individual market for health insurance. I explicitly build pricing regulations in a model that allows for imperfect competition among carriers offering differentiated health insurance plans, and I show how estimates of risk and preferences across different types of buyers can be obtained even when cost data are not available. My empirical application uses 2014 data from the Californian exchange, the largest in the US, and shows how low-income buyers – who receive government support to purchase coverage – are more price sensitive and riskier. I then discuss how, although the subsidy program is critical to guarantee that these buyers purchase coverage, its design could be ameliorated, transferring welfare away from insurers, to favor consumers and/or taxpayers.

My model shows how, in a health insurance exchange, rating regulations interact with the composition of potential buyers, heterogenous in their preferences for insurance and expected health risk. An accurate representation of this heterogeneity is key to characterize insurers' incentives, and to investigate how these would react to changes in the current regulation. Leveraging on the model's predictions, and in particular assuming equilibrium pricing while imposing restrictions on the covariation between price-sensitivity and expected risk of a buyer, it is possible to estimate demand and cost curves observing prices, enrollment decisions, and some information on the composition of buyers. Although with claims data one could estimate a more flexible joint distribution of risk and preferences for insurance, my exercise overcomes this data requirement, and allows the analysis of health even exchanges when insurers or regulators are not willing to release this information.

Estimates from the model can then be used to expore a variety of policies. In this version of the paper I focus on the design of the subsidy program. In fact, as already emphasized by the literature on supplemental insurance for senior citizens, the way in which discounts to low-income buyers are determined is a key aspect of "managed competition" in health exchanges, as different designs determine the ability of insurers to extract surplus from consumers or taxpayers. My results indicate that, in ACA exchanges, the government subsidy program is critical to ensure that low-income buyers purchase health insurance; absent any support, I estimate that in equilibrium coverage would drop by eighty percent. At the same time, however, a different way to compute the amount given to the subsidized might dominate the design dictated by the ACA, where as prices increase, so do subsidies and government expenditure.

Through equilibrium simulation, I predict that a flat voucher program, where a low-income buyer receives a discount equal to 70-80% of her annual expected health expenditure, would mantain coverage levels between 100-115% of the levels obtain under the ACA scheme. At the same time, however, government outlays for subsidies would fall by 0-20%, while also average revenues and individual markups charged by insurers would drop by 12-15%, and 22-27%, respectively.

References

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

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

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

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

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

Brown, J., M. Duggan, I. Kuziemko, and W. Woolston (2014): "How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program," *American Economic Review*, 104, 3335–64.

Buchmueller, T. and J. Dinardo (2002): "Did Community Rating Induce an Adverse Selection Death Spiral? Evidence from New York, Pennsylvania, and Connecticut," *American Economic Review*, 92, 280–294.

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

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

Dafny, L., J. Gruber, and C. Ody (2014): "More Insurers Lower Premiums: Evidence from Initial Pricing in the Health Insurance Marketplaces," Working Paper 20140, National Bureau of Economic Research.

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

Dickstein, M. J., M. Duggan, J. Orsini, and P. Tebaldi (2015): "The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the ACA," Working Paper 20907, National Bureau of Economic Research.

Duggan, M. and T. Hayford (2011): "Has the Shift to Managed Care Reduced Medicaid Expenditures? Evidence from State and Local-Level Mandates," Working Paper 17236, National Bureau of Economic Research.

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

Einav, L. and A. Finkelstein (2011): "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives*, 25, 115–38.

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

Einav, L., A. Finkelstein, and P. Schrimpf (2010b): "Optimal mandates and the welfare cost of asymmetric information: Evidence from the uk annuity market," *Econometrica*, 78, 1031–1092.

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

Ericson, K. M. M. and A. Starc (2013): "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange," Working Paper 19527, National Bureau of Economic Research.

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

Hackmann, M. B., J. T. Kolstad, and A. E. Kowalski (2013): "Adverse Selection and an Individual Mandate: When Theory Meets Practice," Working Paper 19149, National Bureau of Economic Research.

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

Handel, B. R., I. Hendel, and M. D. Whinston (2015): "Equilibria in Health Exchanges: Adverse Selection vs. Reclassification Risk," *Econmetrica, forthcoming.*

Hendren, N. (2013): "Private Information and Insurance Rejections," *Econometrica*, 81, 1713–1762.

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

MaCurdy, T., J. Shafrin, M. B. Blasnek, S. Zaidi, M. Bounds, B. Lipton, D. Pham, S. Sherpa, E. Shrestha, S. Manolache, and S. Zaidi (2014): "Revised Final Report on the CY 2014 Update of the Georaphic Practice Cost Index for the Medicare Physician Fee Schedule," Tech. rep., Burlingame, CA: Acumen LLC.

Orsini, J. and P. Tebaldi (2014): "Regulated Age-Based Pricing in Health Insurance Exchanges," *Available at SSRN 2464725.*

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

Sasso, A. T. L. and I. Z. Lurie (2009): "Community rating and the market for private non-group health insurance," *Journal of Public Economics*, 93, 264 – 279.

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

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

Topkis, D. M. (1978): "Minimizing a submodular function on a lattice," *Operations research*, 26, 305–321.

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

Tables and Figures

Table 1: Standard Age-Rating Curve suggested by CMS and adopted by 48 states

Age	18-20	21	22	23	24	25	26	27	28	29	30	31
Age-factor	0.635	1	1	1	1	1.004	1.024	1.048	1.087	1.120	1.135	1.159
Age	32	33	34	35	36	37	38	39	40	41	42	43
Age-factor	1.183	1.208	1.214	1.222	1.23	1.238	1.246	1.262	1.278	1.302	1.325	1.357
Age	44	45	46	47	48	49	50	51	52	53	54	55
Age-factor	1.397	1.444	1.5	1.563	1.635	1.706	1.786	1.864	1.952	2.04	2.135	2.23
Age	56	57	58	59	60	61	62	63	64			
Age-factor	2.333	2.437	2.548	2.603	2.714	2.81	2.873	2.952	3			

Note: The table reports the standard age-rating curve suggested by CMS. For any age between 18 and 64 it reports the corresponding age-factor. This specifies the ratio between the premium of a specific product for buyers of this age, and the premium of the same product for a 21-year-old buyer in the same rating-region.

Table 2: Maximum contributions for subsidy calculation

Income as % of FPL	**up to 150%**	**150-200%**	**200-250%**	**250-400%**
Max % of income to buy benchmark plan	4%	6.3%	8.05%	9.5%
Maximum monthly premium	$58	$123	$206	$370

Note: The table shows, as a function of a buyer's income, the maximum share that can be spent on the "benchmark plan", i.e. the second cheapest Silver plan in the region. In practice, the subsidy amount is computed as the difference between the premium of this product (after the age adjustment) and the corresponding share of annual income for the buyer. The bottom row shows the maximum dollar amount that can be spent by buyers in the corresponding income "bin" for the benchmark plan.

Table 3: Premium for 21-year-old, average premium paid, and product characteristics

	(a) **Premium for 21-year-old, unsubsidized**					
Tier	Mean	St. Dev.	Min	Max	N	
All	$3,022	$900	$1,538	$5,784	380	
Catastrophic	$2,118	$317	$1,538	$2,899	79	
Bronze	$2,324	$283	$1,732	$3,023	77	
Silver	$3,038	$459	$2,045	$4,035	76	
Gold	$3,615	$597	$2,336	$4,821	74	
Platinum	$4,105	$703	$2,635	$5,784	74	
	(b) **Average premium paid, by age and income**					
	Subsidized			**Unubsidized**		
	20-29	**30-44**	**45-64**	**20-29**	**30-44**	**45-64**
All	$1,673	$1,740	$2,279	$2,995	$3,893	$8,983
Catastrophic	-	-	-	$2,115	$2,747	$6,341
Bronze	$789	$591	$103	$2,312	$3,006	$6,936
Silver	$1,441	$1,438	$1,420	$2,990	$3,888	$8,971
Gold	$2,021	$2,202	$3,160	$3,581	$4,655	$10,743
Platinum	$2,482	$2,792	$4,543	$4,061	$5,281	$12,186
	(c) **Product characteristics**					
	Annual Deductible		**Max Out-of-Pocket**			
	Mean	St. Dev.	Mean	St. Dev.	N	
All	$2,783	$2,595	$5,892	$932	380	
Catastrophic	$6,350	$0	$6,350	$0	79	
Bronze	$5,000	$0	$6,350	$0	77	
Silver*	$2,250	$0	$6,350	$0	76	
Gold	$0	$0	$6,350	$0	74	
Platinum	$0	$0	$4,000	$0	74	

Note: Summary statistics of premiums and product characteristics for the 380 products in CoCA. Panel (a) summarizes the rates for the 21-year-olds unsubsidized. Panel (b) reports average premium paid (weighted by enrollment), by type of buyer and by metal tier. Panel (c) shows the product characteristics for unsubsidized consumers, subsidized buyers receive "enhanced" Silver: deductible is reduced to $0, and maximum out-of-pocket to $2,250.

Table 4: Buyers' composition and number of carriers across regions

Variable	Age	Income	Mean	St. Dev.	Min	Max	N
Potential buyers	20-29	Unsubsidized	5527	4156	644	14158	19
		Subsidized	18817	14149	2201	48202	19
	30-44	Unsubsidized	7770	5788	868	20094	19
		Subsidized	26451	20704	2955	68410	19
	45-64	Unsubsidized	9320	6395	1059	22300	19
		Subsidized	31729	21773	3605	75920	19
Potential buyers relative to total	20-29	Unsubsidized	0.055	0.005	0.045	0.062	19
		Subsidized	0.186	0.018	0.153	0.212	19
	30-44	Unsubsidized	0.077	0.005	0.065	0.085	19
		Subsidized	0.263	0.016	0.222	0.289	19
	45-64	Unsubsidized	0.095	0.008	0.083	0.112	19
		Subsidized	0.323	0.026	0.282	0.381	19
Number of carriers			4.316	1.108	3	6	19

Note: Composition of potential buyers and number of carriers in CoCA, where an observation is a region in the exchange. Composition of potential buyers across the 20 regions is constructed from exchange's estimates (subsidized buyers) and Census data (age composition and unsubsidized buyers).

Table 5: Market shares by type of buyer, metal tier, and carrier

(a) **Decision to purchase by age and income**

	Subsidized			Unsubsidized		
Age	20-29	30-44	45-64	20-29	30-44	45-64
Enrollment as share of potential buyers	67.3 %	59.2%	78.7%	31%	32.9%	34.2%

(b) **Enrollment by metal tier and income**

	Subsidized	Unsubsidized
Minimum Coverage	-	7.6%
Bronze	24.3%	35.6%
Silver	66.2%	29.8%
Gold	5%	12.6%
Platinum	3.9%	14.4%

(c) **Enrollment by insurer and income**

	Subsidized	Unsubsidized	Number of Regions
Anthem	36.3%	40.4%	19
Blue Cross of CA	28.5%	26.9%	19
Chinese Community Health	22.0%	7.1%	2
Contra Costa Health	2.9%	1.9%	1
HealthNet	15.0%	14.8%	13
Kaiser Permanente	21.1%	20.8%	18
LA Care	9.1%	10.7%	2
Molina	2.2%	1.8%	4
Sharp	9.9%	16.4%	1
Valley	3.0%	2.3%	1
Western	3.3%	5.2%	2

Note: Summary of market shares, averages across the 20 regions in CoCA. Panel (a) shows average enrollment, as % of potential buyers. Panel (b) shows average share across metal tiers, where subsidized buyers are prohibited from purchasing Catastrophic coverage. Panel (c) averages market shares by insurer, and reports the number of regions in which each insurer operates.

Table 6: First-stage: OLS of premium for 21-years-old unsubsidized on market characteristics.

Annual premium for 21-years-old, unsubsidized		
Share over 45 (0-100)	23.11***	22.47***
	(7.159)	(7.575)
Share subsidized (0-100)	-48.81**	-44.82*
	(23.25)	(24.59)
Market Size (10,000)	-31.05***	-29.88***
	(3.145)	(3.392)
Deductible	-0.241***	-0.241***
	(0.00859)	(0.00842)
Max OOP	-0.218***	-0.218***
	(0.0304)	(0.0302)
Insurers FE	No	Yes
Observations	380	380
R-squared	0.830	0.987

Note: Robust standard errors in parentheses. An observation is a product in Covered California: that is a unique region-carrier-tier combination. The mean of the dependent variable is $3,022, as reported in Table 3.

Figure 1: Relationship between premium and market shares, varying age, income

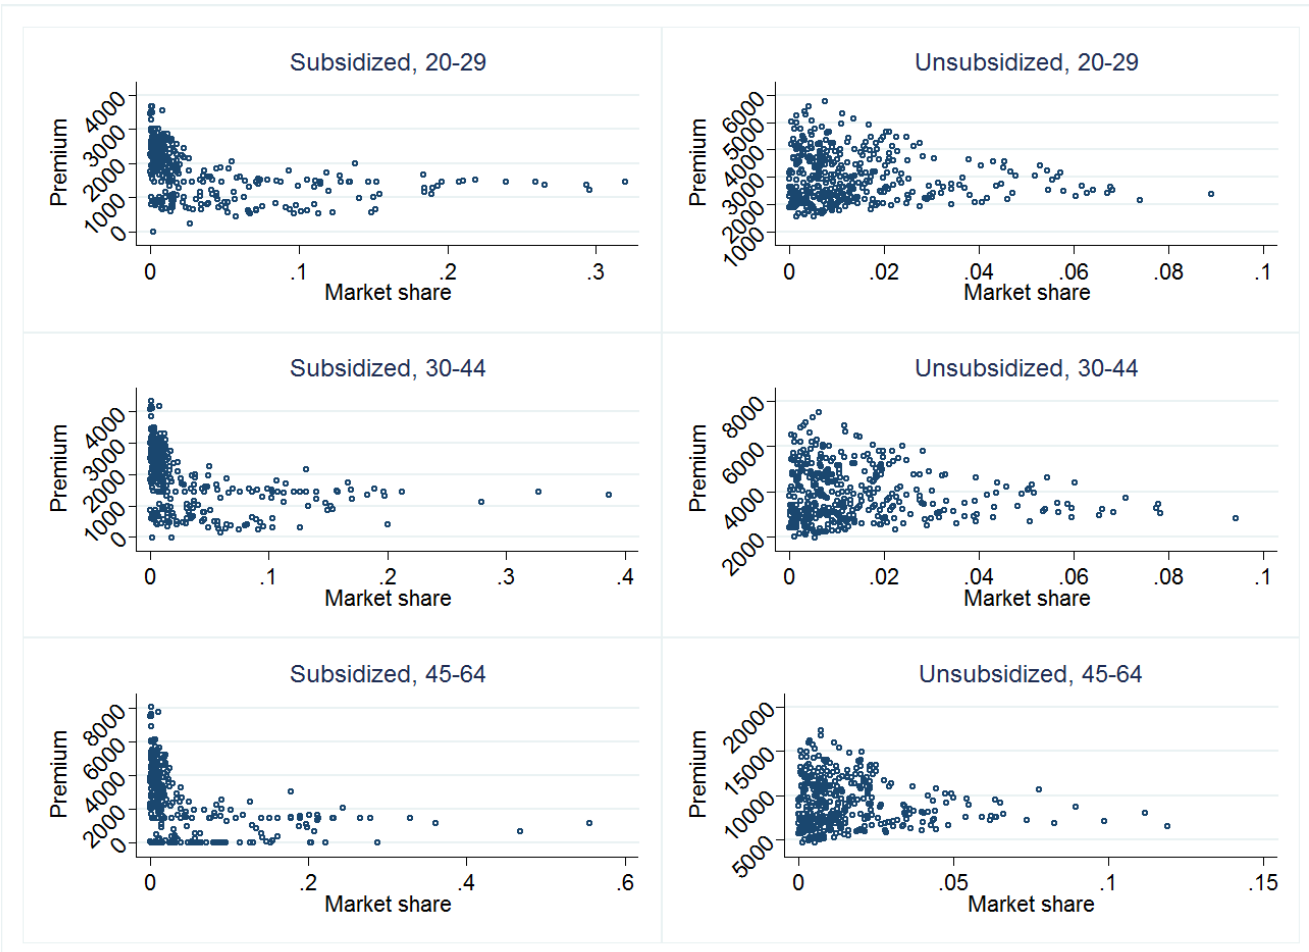

Note: Scatter plots of premiums vs. market shares, by age group and income status. Each dot is a region-product pair, adding up to 380 observations for unsubsidized buyers, and 301 for subsidized buyers for which Catastrophic coverage is not available.

Table 7: IV Logit demand estimates and implied average own-price elasticity for the six types of buyers observed in the Covered California data.

	Subsidized			**Unsubsidized**		
	20-29	**30-44**	**45-64**	**20-29**	**30-44**	**45-64**
Annual Premium	-0.00199***	-0.00153***	-0.000619***	-0.000486***	-0.000514***	-0.000113***
	(0.000140)	(9.71e-05)	(5.52e-05)	(0.000133)	(0.000114)	(4.21e-05)
Deductible	-0.000194***	-0.000189***	-6.76e-05	-0.000139***	-0.000170***	-0.000105***
	(5.46e-05)	(4.97e-05)	(5.91e-05)	(3.84e-05)	(4.36e-05)	(3.62e-05)
Max OOP	-0.000340***	-0.000345***	-0.000367***	3.94e-05	-7.50e-06	7.59e-05
	(4.07e-05)	(3.63e-05)	(5.08e-05)	(6.55e-05)	(6.86e-05)	(6.32e-05)
Insurers Fixed Effects:						
Anthem	2.874***	1.901***	1.914***	-2.212***	-1.201	-2.746***
	(0.282)	(0.232)	(0.289)	(0.733)	(0.814)	(0.706)
Blue Cross of CA	2.459***	1.478***	1.445***	-2.561***	-1.568**	-3.105***
	(0.271)	(0.213)	(0.236)	(0.718)	(0.794)	(0.690)
Chinese Community Health	2.075***	0.274	1.124***	-3.168***	-2.108**	-3.988***
	(0.374)	(0.320)	(0.293)	(0.800)	(0.946)	(0.763)
Contra Costa Health	1.172***	-0.440*	-0.450**	-4.417***	-3.724***	-5.269***
	(0.290)	(0.238)	(0.217)	(0.817)	(0.870)	(0.760)
HealthNet	1.888***	0.760***	0.642**	-3.018***	-2.048**	-3.615***
	(0.321)	(0.268)	(0.257)	(0.733)	(0.823)	(0.701)
Kaiser Permanente	2.425***	1.401***	1.327***	-2.753***	-1.773**	-3.397***
	(0.340)	(0.255)	(0.282)	(0.717)	(0.795)	(0.692)
LA Care	1.721***	0.500	-1.254***	-3.275***	-2.177***	-3.973***
	(0.375)	(0.357)	(0.291)	(0.672)	(0.755)	(0.641)
Molina	-0.560	-1.558***	-2.596***	-5.841***	-4.803***	-6.413***
	(0.389)	(0.361)	(0.559)	(0.714)	(0.776)	(0.685)
Sharp	1.330***	0.848***	-0.450*	-2.902***	-1.697**	-3.168***
	(0.283)	(0.236)	(0.257)	(0.734)	(0.804)	(0.706)
Valley	0.406	-1.181***	1.148***	-4.282***	-3.600***	-4.896***
	(0.288)	(0.234)	(0.269)	(0.784)	(0.865)	(0.761)
Western	0.978**	0.112	0.475	-3.878***	-2.789***	-4.483***
	(0.386)	(0.291)	(0.500)	(0.840)	(0.924)	(0.811)
Mean Premium Paid	$1,672.94	$1,739.59	$2,278.70	$2,994.65	$3,892.58	$8,983.16
Avg. Own-Price Elasticity	-3.370%	-2.747%	-1.495%	-1.453%	-1.995%	-1.012%
Observations	301	301	301	380	380	380
R-squared	0.924	0.949	0.840	0.952	0.944	0.952

Note: Robust standard errors in parentheses. An observation is a product in Covered California: that is a unique region-carrier-tier combination. Each column reports the parameters of the indirect utility derived by the buyer as a function of premium, product characteristics, and insurer fixed effect. The bottom panel reports average own-price elasticity for the products in the data. In the first three columns only Bronze, Silver, Gold, and Platinum plans are included since subsidized buyers are not offered Catastrophic coverage.

Table 8: Estimates of annual expected health expenditure, by income and age.

	Subsidized		Unsubsidized	
Insurer:	20-44	45-64	20-44	45-64
All	\$3,760	\$6,808	\$2,498	\$4,523
	(\$1,936 \$4,388)	(\$5,935 \$8,749)	(\$0 \$6,061)	(\$0 \$12,529)
Anthem	\$3,649	\$6,608	-	-
	(\$1,891 \$4,265)	(\$5,764 \$8,487)		
Blue Cross of CA	\$3,886	\$7,037	-	-
	(\$2,027 \$4,428)	(\$6,073 \$9,329)		
HealthNet	\$3,192	\$5,779	-	-
	(\$1,677 \$3,718)	(\$5,049 \$7,882)		
Kaiser Permanente	\$4,479	\$8,110	-	-
	(\$2,231 \$5,252)	(\$7,199 \$10,126)		
Other minor	\$3,670	\$6,645	-	-
	(\$2,259 \$4,564)	(\$5,328 \$10,930)		

Note: Estimates obtained regressing marginal revenues on marginal risk units across different groups of buyers, as in Proposition 3. 95% confidence intervals in parentheses, obtained via block-bootstrap of the entire two-step estimation procedure (a block is a rating region, level at which prices are set).

Table 9: Comparison between annual health expenditure estimated via Proposition 3 (without cost data) to estimates using the 2011 Medical Expenditure Panel Survey.

	Subsidized		Unsubsidized	
	20-44	45-64	20-44	45-64
Estimated without	\$3,760	\$6,808	\$2,498	\$4,523
cost data (Proposition 3)	(\$1,936 \$4,388)	(\$5,935 \$8,749)	(\$0 \$6,061)	(\$0 \$12,529)
Estimated from Medical	\$3,542	\$7,044	\$3,177	\$6,105
Expenditure Panel Survey	(\$2,979 \$4,106)	(\$6,500 \$7,587)	(\$2,827 \$3,528)	(\$5,214 \$6,995)

Note: The table compares, for the four groups for which I estimate costs (see Table 8), the estimates obtained without information on differential cost across age and income to estimates of health expenditure for the same groups using the 2011 Medical Expenditure Panel Survey. The first two rows are identical to the first two rows of Table 8, standard errors to construct the 95% confidence intervals for the M.E.P.S. estimates are constructed using the individual weights provided with the survey.

Figure 2: The profit-maximization problem for two plans in the Californian exchange.

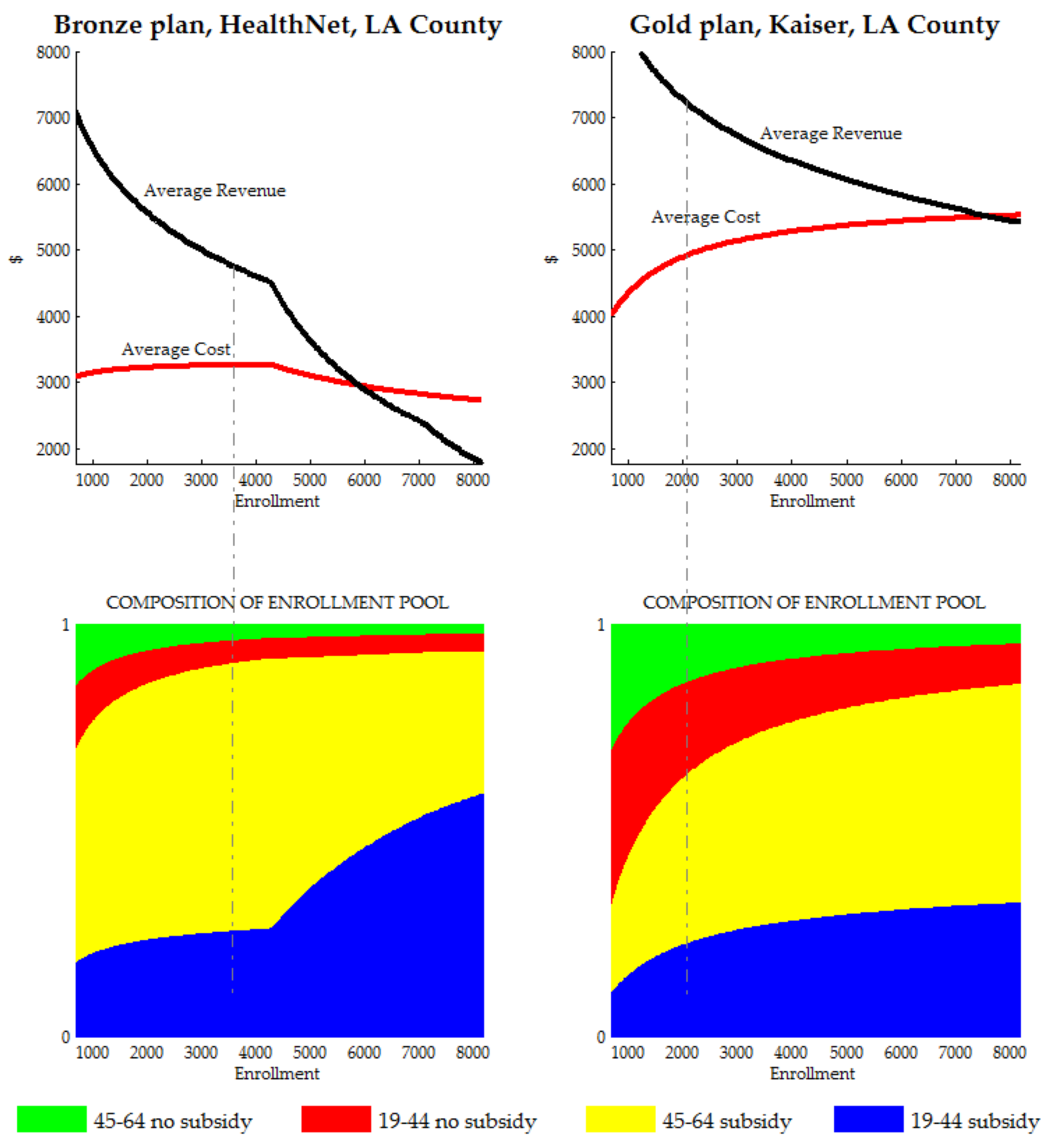

Note: Description of the profit-maximization problem for two plans in the same rating region (LA County) in CoCA. Top panels: Relationship between enrollment, average revenues, and average costs induced by changes in the premium for 21-year-olds unsubsidized. Bottom panels: Corresponding composition of enrollment. The dashed lines indicate the levels observed in the data.

Figure 3: Market outcomes under different levels of a flat voucher, compared to market outcomes under the ACA subdsidy scheme.

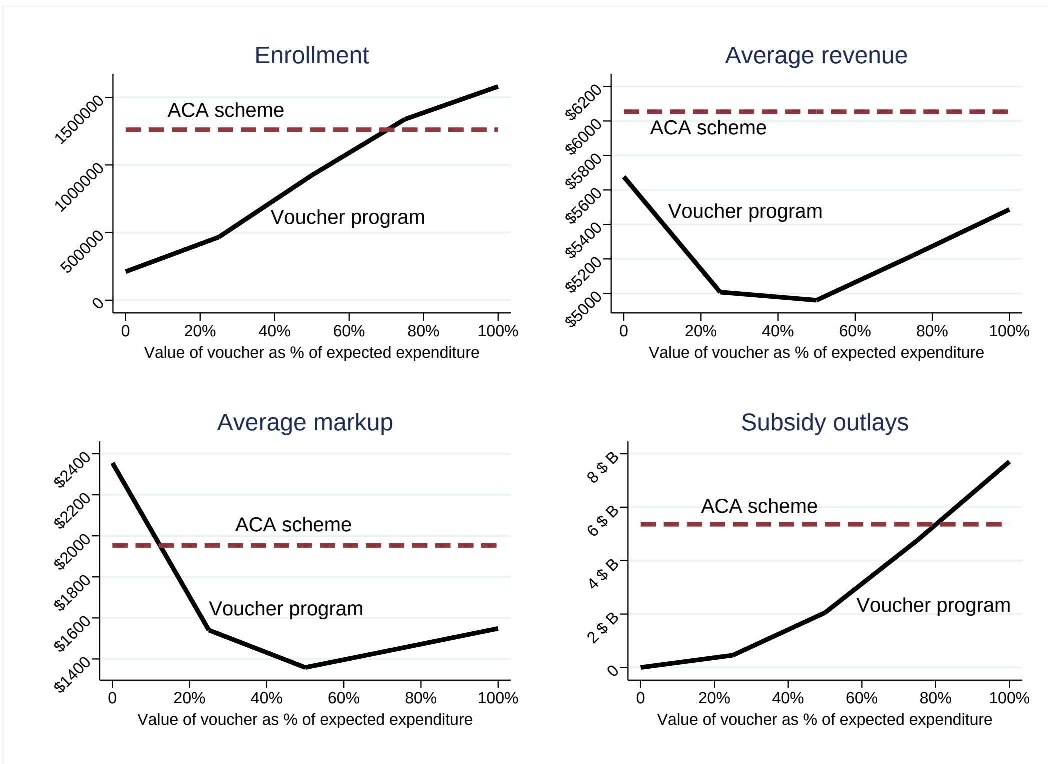

Note: For each level of the voucher program, computed as share of a buyer's annual health expenditure (horizontal axis), the figure plots enrollment, average revenue (buyers contribution plus subsidies), average individual markup, and total government outlays for the subsidy program. For each level of the voucher equilibria are simulated market-by-market, using the estimated primitives presented in Section 6.3, and leaving unchanged the age-rating restrictions. The dashed lines correspond to the levels observed under the status-quo subsidy scheme as mandated by the ACA.

Figure 4: Relative differences between equilibrium outcomes under the ACA and those obtained under a voucher program, varying the value of the voucher.

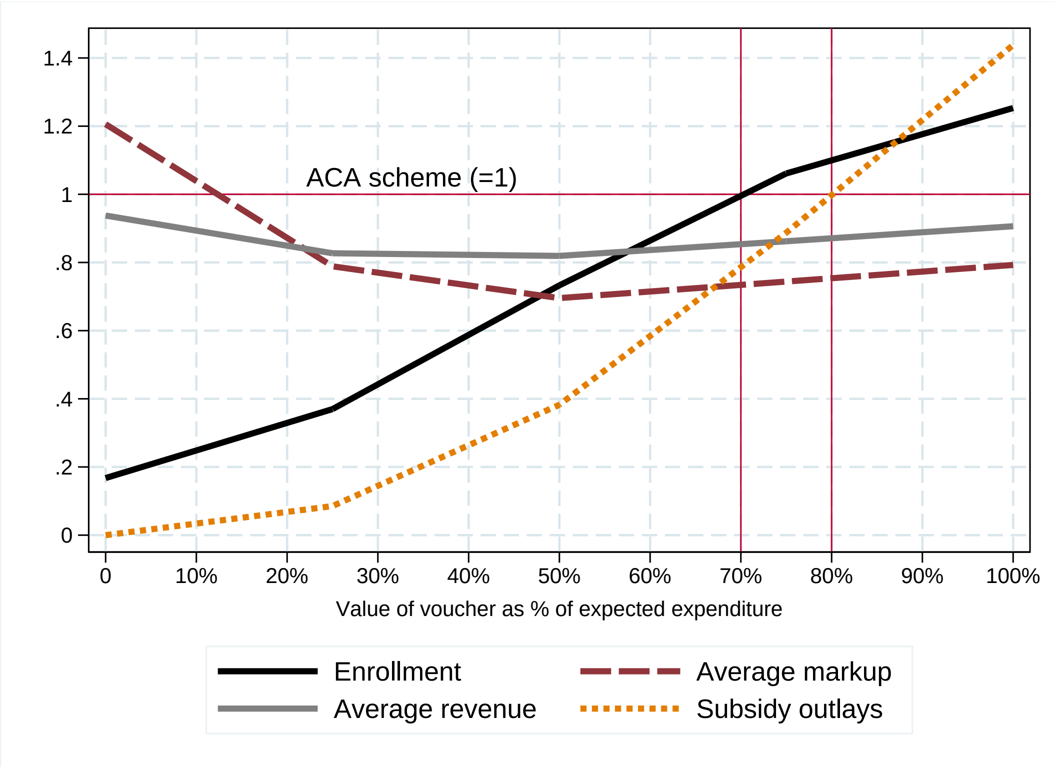

Note: Each line corresponds to a subplot in Figure 3, and it is equal to the ratio between the outcome obtained by simulating equilibrium under the voucher program and the status-quo outcome under ACA regulations.

Supplementary Appendix

Omitted Algebra

Proof of Proposition 2. To show that $\mathbb{E}\left[\varepsilon_{jm}\right]=0$, note that this can be written as

$$\begin{aligned}
& \mathbb{E}\left[\int\left(d_{jm}\left(r_m\left(\tau\right);\tau|\chi_m\right)+\sum_{j'\in J_{fm}} r_{j'm}\left(\tau\right)\frac{\partial d_{j'm}\left(r_m\left(\tau\right);\tau|\chi_m\right)}{\partial p_j}\right)\frac{\partial a\left(r_{jm}\left(\tau_0\right),\tau\right)}{\partial\left(r_j\left(\tau_0\right)\right)}dG_m\left(\tau\right)\right]- \\
& \mathbb{E}\left[\int\left(\sum_{j'\in J_{fm}}\frac{\partial c_{j'm}\left(r_m\left(\tau\right);\tau|\chi_m\right)}{\partial p_j}\right)\frac{\partial a\left(r_{jm}\left(\tau_0\right),\tau\right)}{\partial\left(r_j\left(\tau_0\right)\right)}dG_m\left(\tau\right)\right]+ \\
& \mathbb{E}\left[\int\left(d_{jm}\left(r_m\left(\tau\right);\tau|\chi_m\right)+\sum_{j'\in J_{fm}} r_{j'm}\left(\tau\right)\frac{\partial d_{j'm}\left(r_m\left(\tau\right);\tau|\chi_m\right)}{\partial p_j}\right)\frac{\partial a\left(r_{jm}\left(\tau_0\right),\tau\right)}{\partial\left(r_j\left(\tau_0\right)\right)}d\eta_m\left(\tau\right)\right]- \\
& \mathbb{E}\left[\int\left(\sum_{j'\in J_{fm}}\frac{\partial c_{j'm}\left(r_m\left(\tau\right);\tau|\chi_m\right)}{\partial p_j}\right)\frac{\partial a\left(r_{jm}\left(\tau_0\right),\tau\right)}{\partial\left(r_j\left(\tau_0\right)\right)}d\eta_m\left(\tau\right)\right] \\
= \; & 0+ \\
& \mathbb{E}_\chi\left[\int\left(d_{jm}\left(r_m\left(\tau\right);\tau|\chi_m\right)+\sum_{j'\in J_{fm}} r_{j'm}\left(\tau\right)\frac{\partial d_{j'm}\left(r_m\left(\tau\right);\tau|\chi_m\right)}{\partial p_j}\right)\frac{\partial a\left(r_{jm}\left(\tau_0\right),\tau\right)}{\partial\left(r_j\left(\tau_0\right)\right)}\mathbb{E}_\eta\left[d\eta_m\left(\tau\right)\right]\right]- \\
& \mathbb{E}_\chi\left[\int\left(\sum_{j'\in J_{fm}}\frac{\partial c_{j'm}\left(r_m\left(\tau\right);\tau|\chi_m\right)}{\partial p_j}\right)\frac{\partial a\left(r_{jm}\left(\tau_0\right),\tau\right)}{\partial\left(r_j\left(\tau_0\right)\right)}\mathbb{E}_\eta\left[d\eta_m\left(\tau\right)\right]\right] \\
= \; & 0.
\end{aligned}$$

The first equality follows because the constrained equilibrium condition (8) must hold. The second applies the law of iterated expectations, and follows by independence of η_m from other observables, and $\mathbb{E}\left[\eta_m\left(\tau\right)\right]=0$.