Managed Competition vs. Public Procurement in Individual Health Insurance*

José Ignacio Cuesta§ Pietro Tebaldi‡

March 2025
Submission to 2025 NBER SI - please do not circulate

Abstract. When sponsoring health insurance markets, governments must choose between alternative mechanisms to balance quality, cost, and, ultimately, welfare. While regulated or managed competition has been widely adopted, several forces may favor sole sourcing of coverage from insurers selected through competitive bidding. Auctions may limit inefficiencies due to adverse selection and market power, and incentivize quality provision, but may also limit variety. We set up and estimate an empirical model to investigate these trade-offs in the context of the Affordable Care Act. Our approach combines individual-level enrollment data, plan provider networks, and average plan-level claims, with a pure characteristics demand model and a model of cost as a function of risk selection into plans and plan quality. Holding plan quality fixed, procurement would generate gains from stronger competition and lower risk skimming that, in many markets, would outweigh losses from limiting variety. Allowing for endogenous plan quality, scoring rules rewarding low premiums are worse than managed competition, while those rewarding plan quality generate large welfare gains across most markets.

Keywords health insurance, competition, procurement, adverse selection, auctions
JEL Codes. H42, I13, L13

*We thank Daniel Ackerberg, Steve Berry, Vivek Bhattacharya, Liran Einav, Phil Haile, Igal Hendel, Kate Ho, Gastón Illanes, Jakub Kastl, Victoria Marone, Aviv Nevo, Ariel Pakes, Rob Porter, Mark Shepard, Paulo Somaini, Ashley Swanson, and Robert Town, and conference and seminar participants at ASHECON, Cornell University, Harvard-MIT-BU, Northwestern University, Princeton University, PUC-Chile, SoCal Structural Econometrics Conference, Università Bocconi, University of Chicago, University of Chile, University of Maryland, UNC-Duke, University of Pennsylvania, UIUC, UT Austin, University of Wisconsin-Madison and Yale University for helpful comments. Kan Xu and Chris Xue provided superb research assistance for this project. All errors are our own.§Stanford University and NBER. Email: jicuesta@stanford.edu.‡Columbia University and NBER. Email: p.tebaldi@columbia.edu.

1 Introduction

Social preferences for broad and equitable access to healthcare motivate public financing and intervention when a significant share of individuals do not have health insurance. Despite varied policy efforts, around ten percent of the U.S. population is uninsured (Baicker, Chandra and Shepard, 2023), down from about fifteen percent in 1990–2013.[1] While such uninsurance rate is exceptional among OECD economies (Baicker *et al.*, 2023), it is the norm in low- and middle-income countries (Chen *et al.*, 2022). The shortcomings of unregulated markets in achieving universal coverage are attributed to a variety of factors including low willingness-to-pay (Finkelstein, Hendren and Shepard, 2019), information asymmetries leading to adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976) and moral hazard (Pauly, 1968), behavioral biases (Chandra, Handel and Schwartzstein, 2019), and ideological opposition to government programs (Sances and Clinton, 2019).

A common approach to boost health insurance coverage is what Enthoven (1993) defined as *managed competition*. This market design combines financial incentives to support enrollees (e.g., means-tested subsidies) and insurers (e.g., risk adjustment), with regulation that constrains plan quality, pricing, profit margins, and market definition (for further discussion, see Einav and Levin 2015). Managed competition relies on household choices among multiple insurers and plans to discipline plan premiums and quality. Moreover, by offering plan variety, managed competition accommodates the differing needs of a population with heterogeneous preferences for premiums (Tebaldi, Torgovitsky and Yang, 2023), cost-sharing (Marone and Sabety, 2022), and provider networks (Shepard, 2022). However, the effectiveness of managed competition in expanding coverage may be limited by market power and adverse selection, which have been shown to be present in these markets (Decarolis, Polyakova and Ryan, 2020; Tebaldi, 2024).

In this paper, we compare managed competition to a counterfactual market design inspired by the use of public procurement for the provision of other impure public goods. Concretely, we compare managed competition to a procurement mechanism that selects a *winner-takes-all* insurer to serve an uninsured population, in line with early ideas in Diamond (1992). The appeal of procurement mechanisms is twofold. First, competition is potentially strengthened. Second, insurers need not worry about risk selection from competitors. Conversely, excluding the losing insurers from a marketplace limits plan variety, which can lower coverage and consumer surplus among those unable to enroll in their preferred option despite paying lower premiums. Procurement-like market designs

[1]The majority of the U.S. population benefits from either employer-sponsored insurance (50 percent), public insurance for the elderly (Medicare, 14 percent), or means-tested insurance (Medicaid, 20 percent).

for health insurance have received limited attention from researchers. However, there are a few examples of them in practice. For instance, many employers select their menus of employer-sponsored plans via auction-like mechanisms (Gupta, Parente and Sanyal, 2012). Moreover, Medicaid MCO, Medicare Part D, and Pre-ACA Massachusetts enrolled low-income individuals into plans via auction-like mechanisms (Forbes and Dunbar, 2022).

We start by developing a framework to compare managed competition to a class of incentive-compatible procurement mechanisms based on scoring auctions (Asker and Cantillon, 2008). We characterize managed competition as Bertrand-Nash pricing by insurers, in line with recent related research (e.g., Curto, Einav, Levin and Bhattacharya 2021; Shepard 2022; Tebaldi 2024). For procurement mechanisms, we focus on second-score auctions under which the government determines a scoring rule to select a winner, and the winner can update their bid—and potentially their plan quality—and sell insurance as long as their score is above the second-highest score resulting from the auction. These auctions are strategically simple, similar to standard second-price auctions. Importantly, with non-constant marginal and average cost curves—as one expects in insurance markets with selection—the resulting bids need not be equal to perfectly competitive premiums but rather may differ depending on risk selection across plans and into the market.

Our framework identifies four economic forces driving the relative appeal of managed competition and procurement mechanisms. First, managed competition provides more plan variety, which is valuable when consumers have heterogeneous plan preferences. Second, procurement mechanisms may place incentives for insurers to increase or decrease plan quality depending on the scoring rule. Third, selection into health insurance and across plans may change across market designs, leading to changes in premiums that are ambiguous in direction. Fourth, procurement mechanisms may be more competitive, leading to lower premiums. These economic forces are inherently context-specific. Measuring them requires combining data on consumer and plan attributes with estimates of consumer preferences and insurer costs in an equilibrium model.

The Covered California marketplace provides a compelling setting for our empirical analysis. This market is a textbook example of managed competition operating under the rules of the Affordable Care Act of 2010 (ACA), which regulates plan attributes, geographic market definition, premium-setting, premium subsidies, and cost-sharing, among others. In practice, between 2 and 7 insurers competed in each of its 19 regions between 2014 and 2017. Beyond providing a baseline for managed competition, this setting offers four more advantages. First, its regulation is well-documented and understood. Second, detailed individual-level data on enrollment choices by more than 3 million enrollees

and plan-level data on average costs are available, which we complement with detailed data on zipcode-level plan provider networks. Third, the variation across individuals and markets within California mimics well the heterogeneity of a larger, national context. Fourth, the ACA primarily serves low-income and often otherwise uninsured individuals, a relevant population for policy analysis and for which procurement mechanisms may be more appropriate.

On the demand side, we specify a pure characteristics plan demand model that allows for substantial observable and unobservable preference heterogeneity (Berry and Pakes, 2007). Whereas standard logit demand models may significantly overestimate the welfare impact of large changes in the number of available options, the pure characteristics model does not include idiosyncratic error terms and, hence, is not subject to the same concern. The combination of rich individual-level data with the regulation in our settings allows for straightforward identification arguments which are a special case of the general results in Berry and Haile (2024). In particular, we exploit pricing variation across individuals, within market, resulting from age-based risk-rating regulation and subsidies, variation in insurance generosity following discontinuities in cost-sharing reductions, and variation in the location of in-network providers and individuals' zipcode. Using this "panel-like structure" we recover preferences over plan premiums, actuarial value, and provider networks. The resulting estimates imply strong and heterogeneous responses to plan attributes. On the extensive margin, increasing premiums by \$10/month leads to a decrease in overall marketplace enrollment of 1.2 percentage points. Improving provider networks by either increasing PCP coverage within ten miles of a household residence or lowering the distance to the closest in-network hospital leads to increasing enrollment by 0.5 and 2.2 percentage points, respectively. This result, in turn, suggests that the quality of plans selected by procurement mechanisms may impact overall insurance coverage and welfare.

On the supply side, we specify an insurer cost function whereby claim payouts depend on plan generosity, plan-specific cost unobservables, but also enrollee's age and willingness-to-pay for insurance. For identification, we leverage an inversion approach based on optimal pricing conditions combined with data on average claims. Cost estimates reveal substantial selection on observables and unobservables. As expected, enrollee costs increase steeply with age. In terms of unobservables, enrollee costs are increasing on willingness-to-pay for plan generosity, generating adverse selection that is not corrected by age-rating adjustments. To account for changes in plan costs driven by changes in plan quality, we again leverage an inversion approach based on first-order conditions for plan attributes to identify adjustment costs.

Our main counterfactual results show that replacing managed competition with a

procurement mechanism would strongly impact market outcomes and welfare. Simulating equilibria under scoring auctions and comparing outcomes to those observed under managed competition, our first finding is that a procurement mechanism that selects an insurer based on premium only would mostly deliver worse outcomes than managed competition. In contrast, a scoring rule that selects an insurer based on its implied enrollment dominates managed competition. In particular, such a mechanism would increase marketplace enrollment by 1.1 percentage points and consumer surplus by $125/year, in a context where bidders were unable to adjust their plan quality. With endogenous quality provision, the gap between both polar procurement cases widens—since incentives to adjust quality operate in opposite directions—and a scoring rule rewarding enrollment leads to a positive plan quality adjustment, increasing marketplace enrollment by 1.6 percentage points and consumer surplus by $251/year relative to managed competition.

The structure of our model allows us to disentangle the roles of variety, quality, selection, and competition as underlying forces through which different mechanisms impact outcomes and welfare. We find that procurement mechanisms harm enrollment through the variety channel since consumers lose options from their menus and are left with a single plan. By lowering variety alone, the scoring rules we study reduce enrollment by 4–8 percentage points. However, the combination of other forces at least partially compensates for this negative impact. In particular, selection improves across the board, leading to lower premiums and increasing enrollment to some extent. Furthermore, more intense competition in the auction leads to lower markups, further increasing enrollment by between 1.5 and 5 percentage points. Comparing alternative designs of procurement mechanisms, we find that the variety and competition effects are stronger for a scoring rule that rewards price since such auctions tend to select low-quality and low-price insurers. Finally, we find that quality effects are quantitatively large and that they differ across scoring rules. This implies that the design of procurement mechanisms may impact plan quality in ambiguous ways depending on the incentives placed on insurers. These results suggest that the four mechanisms we identify are quantitatively important in driving the impacts of alternative mechanisms for health insurance provision.

Our paper contributes to three branches of the literature. First, to research focusing on how the design and regulation of health insurance markets to limit market failures and improve efficiency (e.g., Bundorf, Levin and Mahoney 2012; Starc 2014; Handel, Hendel and Whinston 2015a; Ho and Lee 2017, 2019; Miller, Petrin, Town and Chernew 2019; Decarolis *et al.* 2020; Jaffe and Shepard 2020; Curto *et al.* 2021; Saltzman 2021; Handel and Kolstad 2022; Marone and Sabety 2022; Dickstein, Ho and Mark 2024; Tebaldi 2024; Vatter 2024). Our paper is the first to provide a systematic comparison between managed

competition and a class of feasible alternative mechanisms akin to public procurement, that, albeit used in practice and discussed conceptually by Diamond (1992), has not received enough attention in academic and policy discussions. Our analysis provides a conceptual framework to identify the key economic forces at play and quantifies the impacts of such alternative mechanisms in a large and relevant empirical setting.

Second, we contribute to a literature focused on the value of choice. Measuring the value of new goods involves several methodological challenges. In particular, many standard demand models tend to overestimate the value of variety due to the introduction of idiosyncratic preference shocks. This issue is a common concern in this literature (Ackerberg and Rysman, 2005; Quan and Williams, 2018). We build on this literature by employing econometric methods that limit the impact of these undesirable assumptions. In particular, we adopt the pure characteristics demand model by Berry and Pakes (2007) to our large-scale empirical application with individual-level data. We are among the first empirical papers to estimate such a model (e.g., Song 2007; Pang, Su and Lee 2015). In the context of health insurance, there is an active debate around how the amount of choice impacts consumer welfare (Dafny, Ho and Varela, 2013; Abaluck and Gruber, 2023). Our application informs this discussion by applying our methodology to estimate rich distributions of consumer preferences in a compelling setting.

Third, we contribute to a literature that attempts to measure the value of plan provider networks. Provider networks have received substantial attention from policymakers and researchers in recent years. Previous research mostly relies on demand estimates that inform willingness-to-pay for hospitals (Ho, 2006; Ericson and Starc, 2015; Gowrisankaran, Nevo and Town, 2015; Ho and Lee, 2017; Prager, 2020; Shepard, 2022; Serna, 2023). Our paper complements this literature by using a novel, comprehensive dataset of provider networks with a renewed emphasis on primary care physicians and non-hospital care—which deliver as much as 89 percent of healthcare services—complementing Drake (2019).[2] Our empirical setting and strategy isolate variation in provider networks that allow for estimating rich distributions of willingness-to-pay for access to doctors and clinics, a key input in evaluating changes to health insurance institutions.

The remainder of the paper is organized as follows. Section 2 introduces a framework to compare managed competition and procurement mechanisms. Section 3 describes the institutions and data, and Section 4 develops our empirical model. Section 5 discusses our counterfactual analysis and main results. Finally, Section 6 concludes.

[2]See e.g. `https://www.ncbi.nlm.nih.gov/books/NBK201069/`; accessed March 19, 2025.

2 A Framework to Compare Managed Competition to Scoring Auctions

We study an environment in which a population of individuals do not have health insurance. Absent intervention, the market is either completely unraveled or excessively underserved relative to the planner's insurance coverage target. The policy goal—that we take as given in our analysis—is to increase coverage, and to strike the right balance between quality and cost of health insurance.

One widespread approach in this context is to set up a health insurance exchange defining rules for entry, contract design, limits to price discrimination, timing of enrollment, geographic market definition, and providing incentives to individuals (subsidies) and insurers (risk adjustment and reinsurance). This approach is referred to as managed competition, and a leading example is the creation of the ACA marketplaces in 2014. The alternative we consider is a procurement mechanism using scoring auctions to select one insurer for each market. To set up a fair comparison between market mechanisms we will hold the remaining regulations (e.g., geographic market definition, rating regulations, and subsidies) fixed. To highlight the main trade-offs between managed competition and scoring auctions, we begin by considering a simplified framework.

Environment. Every uninsured individual i is characterized by a type θ_i; the population in a market is distributed as $G(\theta)$. When individual i enrolls in a plan j with characteristics x_j—describing the insurer brand, the contract financial details such as deductibles and copays, and the set of in-network healthcare providers—they derive money-metric indirect utility $v(\theta_i, x_j)$, and they generate expected costs for the insurer selling j equal to $c(\theta_i, x_j)$. The market sponsor provides a premium subsidy $s(\theta_i)$, which the individual can use to purchase coverage. For simplicity, we omit other regulations here and introduce ACA regulatory details in our empirical analysis below.

Managed competition. Under managed competition, there are J plans in the market, so that when facing a price vector $\boldsymbol{p} = (p_1, p_2, \cdots, p_J)$ and characteristics $\boldsymbol{x} = (x_1, x_2, \cdots, x_J)$ individual i selects the plan j that maximizes $v(\theta_i, x_j) - p_j + s(\theta_i)$ as long as this leads to positive utility, and remains uninsured otherwise. We denote this event by

$$\Theta_j(\boldsymbol{p}, \boldsymbol{x}) \equiv \left\{ \theta : v(\theta, x_j) - p_j + s(\theta) \geq \max \left\{ 0, \max_{k \neq j} v(\theta, x_k) - p_k + s(\theta) \right\} \right\},$$

so that for prices and characteristics $(\boldsymbol{p}, \boldsymbol{x})$ the total enrollment and average expected cost for plan j are

$$Q_j(\boldsymbol{p}, \boldsymbol{x}) = \int \mathbf{1}\Big[\theta \in \Theta_j(\boldsymbol{p}, \boldsymbol{x})\Big] dG(\theta) \,,$$
$$AC_j(\boldsymbol{p}, \boldsymbol{x}) = Q_j^{-1}(\boldsymbol{p}, \boldsymbol{x}) \int c(\theta, x_j) \mathbf{1}\Big[\theta \in \Theta_j(\boldsymbol{p}, \boldsymbol{x})\Big] dG(\theta) \,.$$

In a (Bertrand) Nash equilibrium $(\boldsymbol{p}^\star, \boldsymbol{x}^\star)$, every plan j chooses premium and characteristics to maximize expected profits taking competitors' strategies as given. Formally,

$$p_j^\star, x_j^\star \in \arg\max_{p_j, x_j} Q_j(p_j, x_j; \boldsymbol{p}_{-j}^\star, \boldsymbol{x}_{-j}^\star) \Big(p_j - AC_j(p_j, x_j; \boldsymbol{p}_{-j}^\star, \boldsymbol{x}_{-j}^\star)\Big) \,. \tag{1}$$

The total enrollment in equilibrium is

$$Q^\star = \sum_j Q_j(\boldsymbol{p}^\star, \boldsymbol{x}^\star) = \int \underbrace{\mathbf{1}\Big[\theta : \max_j \; v(\theta, x_j^\star) - p_j^\star + s(\theta) \geq 0\Big]}_{\Theta^m(\boldsymbol{p}^\star, \boldsymbol{x}^\star) \equiv \bigcup_j \Theta_j(\boldsymbol{p}^\star, \boldsymbol{x}^\star)} dG(\theta) \,, \tag{2}$$

which highlights how removing plans from the menu of available choices will (weakly) reduce enrollment, holding premiums and characteristics fixed. To see this, note that the set $\Theta^m(\boldsymbol{p}^\star, \boldsymbol{x}^\star)$ of individuals purchasing coverage in the managed competition equilibrium shrinks when removing j-elements from the union.

While having a large number of plans in the marketplace increases variety holding premium fixed, imperfect competition between insurers can lead to two sources of inefficiencies. First, regardless of heterogeneity in risk across individuals (i.e., even if $c(\theta_i, x_j)$ does not vary with θ_i), heterogeneity in preferences and the resulting imperfect competition between insurers enable them to charge positive markups ($p_j^\star > AC_j(\boldsymbol{p}^\star, \boldsymbol{x}^\star)$). Second, in markets with adverse selection (in which $\partial AC_j(\boldsymbol{p}, \boldsymbol{x})/\partial p_j > 0$), increases in premiums from imperfect competition discourage enrollment disproportionately among low-cost individuals, further increasing average cost and premiums (Einav, Finkelstein and Cullen, 2010; Handel, Hendel and Whinston, 2015b). This can lead to lower total enrollment, possibly paired with the unraveling of high-quality, high-cost plans. Moreover, when plans also choose quality, to the extent to which higher quality attracts riskier individuals, competition between plans may lead to underprovision (Rothschild and Stiglitz, 1976).

Procurement mechanisms. We consider an alternative scenario in which the market sponsor announces a scoring rule Ψ that varies with a (scalar) bid and possibly with plan characteristics, $\Psi(b, x)$ with $\partial\Psi/\partial b < 0$. Insurers submit bids and characteristics leading to the vector of scores $\boldsymbol{\psi} = (\Psi(b_1, x_1), \Psi(b_2, x_2), \cdots, \Psi(b_J, x_J))$. The winner is j^w for which $\psi_{j^w} = \max \boldsymbol{\psi}$.

We focus on *second-score scoring auctions* (Asker and Cantillon, 2008) which—by having easy-to-compute dominant strategy solutions—are easier to implement and analyze, hence, a natural possibility for health insurance exchanges. In these auctions, the winner becomes the only plan available in the market and can optimize premium and characteristics as long as the resulting score does not fall below $\psi^{(2)}$, the second-highest element of $\boldsymbol{\psi}$. Formally, after winning the auction j^w sets (p^a, x^a) solving

$$\begin{aligned}(p^a, x^a) \in \arg\max_{p,x}\ & Q_{j^w}(p, x; +\infty, \emptyset)\Big(p - AC_{j^w}(p, x; +\infty, \emptyset)\Big) \\ \text{s.t. }\ & \Psi(p, x) \geq \psi^{(2)}\ ,\end{aligned}$$

where replacing other plans' prices with $+\infty$ and characteristics with $\emptyset$ indicates that those plans are not available.

The dominant strategy is for insurers to submit a bid and plan characteristics to achieve the highest possible score, subject to breaking even as the only option in the market. Formally, every j will submit a (b_j, x_j) such that

$$(b_j, x_j)\ \in \arg\max_{b,x} \Psi(b, x)\ \ \text{s.t.}\ \ Q_j(b, x; +\infty, \emptyset)\Big(b - AC_j(b, x; +\infty, \emptyset)\Big) = 0\ .$$

Any combination that leads to a lower score with weakly positive profits is dominated because it can "leave money on the table" when losing the auction against a score that could be beaten without incurring expected losses. Any combination that leads to strictly negative profits is dominated because winning the auction against a second-highest score that is only ε smaller would be worse than losing the auction.

Total enrollment under a procurement mechanism delivering outcomes (p^a, x^a) is

$$Q^a = \int \mathbf{1}\underbrace{\Big[\theta : v(\theta, x^a) - p^a + s(\theta) \geq 0\Big]}_{\equiv\, \Theta^a(p^a, x^a)} dG(\theta)\ . \tag{3}$$

Variety, quality, selection, and competition. Enrollment under a procurement mechanism and managed competition may differ depending on the differences between plan

premiums $\boldsymbol{p}^\star$ and p^a, and between plan characteristics $\boldsymbol{x}^\star$ and x^a. These differences in turn affect the set of types purchasing under managed competition ($\Theta^m(\boldsymbol{p}^\star, \boldsymbol{x}^\star)$) and those purchasing under the auction ($\Theta^a(p^a, x^a)$). There are several distinct channels that explain these differences across market mechanisms.

First, holding everything else constant, limiting variety harms consumers.[3] Second, quality provision may vary under different mechanisms. Auctions with scoring rules that reward quality may increase quality provision relative to managed competition, while the opposite may happen under scoring rules that primarily reward low bids. Third, in the insurance context, changing the set of available options, their premiums, and their characteristics moves different enrollees across plans, affecting their risk selection and, therefore, expected average cost. Lastly, while markups under managed competition are determined by the simultaneous interaction between competing plans, markups under the second-score auctions we consider here result from the gap between the first- and second-highest scores. These different forces co-exist and interact to generate the difference between outcomes under managed competition and scoring auctions. It may be impossible to provide theoretical predictions on their directions and magnitudes without imposing extremely stark assumptions. However, our framework can be used to quantify each channel and define a set of empirical quantities that can be measured and analyzed with an econometric model like the one we introduce in Section 4.

The *variety effect* can be measured without any changes to premium or characteristics, by simply removing all options except the winner j^w. Formally, and highlighting the dependence on the specific scoring rule:

$$\text{Variety effect}(\Psi) \equiv \int \mathbf{1}\left[\theta \in \Theta^a(p^\star_{j^w}, x^\star_{j^w})\right] - \mathbf{1}\left[\theta \in \Theta^m(\boldsymbol{p}^\star, \boldsymbol{x}^\star)\right] dG(\theta)\,, \tag{4}$$

which must be negative since $\Theta^a(p^\star_{j^w}, x^\star_{j^w}) \subseteq \Theta^m(\boldsymbol{p}^\star, \boldsymbol{x}^\star)$. The magnitude of this effect may vary across scoring rules that select different winners, affecting the difference between the two sets.

The incremental *quality effect* can be measured by holding the premium fixed at its managed competition level but considering the characteristics provided in equilibrium

[3] An alternative stand, which tilts normative recommendations toward a more paternalistic approach and limits concerns for reduced variety, is the one in which individuals make (too) many mistakes when choosing plans in insurance exchanges. See e.g. Abaluck and Gruber (2023); Brown and Jeon (2024) and references therein.

under the scoring auction:

$$\text{Quality effect}(\Psi) \equiv \int \mathbf{1}\left[\theta \in \Theta^a(p^\star_{j^w}, x^a)\right] - \mathbf{1}\left[\theta \in \Theta^a(p^\star_{j^w}, x^\star_{j^w})\right] dG(\theta)\,, \tag{5}$$

which cannot be signed without more information. In the simplest case in which x is scalar and valued by all consumers ($\partial v(\theta, x)/\partial x > 0$), the quality effect is positive if and only if $x^a > x^\star_{j^w}$. Different scoring rules affect both the sign and magnitude of the quality effect.

The incremental *selection effect* can be measured by holding markups fixed at their managed competition level but adjusting premiums by differences in expected average cost under the different market mechanisms. Concretely, let $\mu^\star$ denote the average markup under managed competition, $\mu^\star \equiv p^\star_{j^w} - AC_{j^w}(\boldsymbol{p}^\star, \boldsymbol{x}^\star)$.[4] Then, letting $AC^a \equiv AC_{j^w}(p^a, x^a; +\infty, \emptyset)$, the selection effect is

$$\text{Selection effect}(\Psi) \equiv \int \mathbf{1}\left[\theta \in \Theta^a(AC^a + \mu^\star, x^a)\right] - \mathbf{1}\left[\theta \in \Theta^a(p^\star_{j^w}, x^a)\right] dG(\theta)\,, \tag{6}$$

which captures changes in enrollment driven by changes in cost and not in markups. The selection effect is positive if and only if the winner's average cost under the auction is lower than their average cost under managed competition.[5]

Lastly, the incremental *competition effect* of the scoring auction relative to managed competition can be measured as

$$\text{Competition effect}(\Psi) \equiv \int \mathbf{1}\left[\theta \in \Theta^a(p^a, x^a)\right] - \mathbf{1}\left[\theta \in \Theta^a(AC^a + \mu^\star, x^a)\right] dG(\theta)\,, \tag{7}$$

which is positive if and only if the auction leads to lower markups than those under managed competition.

The cumulative enrollment effect of switching from managed competition to scoring auctions is the sum of these four effects. While our argument here focused on enrollment, a similar decomposition can be done for consumer surplus and other functions that aggregate outcomes of interest over subsets of types as specified above.

[4] Alternatively, one could define $\mu^\star$ as the markup of the winner, but this can sometime be very low (or near zero) and not reflect markups in the manage competition regime.

[5] These cost changes combine (i) shifts in risk composition of enrollees due to changes in variety and quality, (ii) shifts in risk composition of enrollees due to changes in premium, and (iii) changes in cost due to changes in quality.

3 Empirical Setting

3.1 Institutions

Our empirical analysis focuses on the Californian ACA marketplace, also analyzed in Saltzman, Swanson and Polsky (2021); Panhans (2019); Tebaldi, Torgovitsky and Yang (2023); Tebaldi (2024). This is a textbook example of managed competition: insurers offer a plan menu to the public in a highly regulated environment, which combines federal and state-specific rules on plan attributes, geographic market definition, and premium setting, among others. Here, we briefly summarize the regulations that are essential for our analysis.

Insurers may decide to offer plans in each of the 19 rating regions in Covered California, for which they must be able to offer adequate provider networks.[6] Plans are standardized in terms of their generosity and organized into five tiers: Bronze, Silver, Gold, Platinum, and Catastrophic. The four metal tiers cover—in expectation—60, 70, 80, and 90 percent of healthcare expenses, respectively. Catastrophic plans generally have higher cost-sharing and cannot be purchased by subsidized individuals or individuals older than 30. Individuals with income lower than 250 percent of the FPL benefit of cost-sharing reductions, which increase the generosity of Silver plans (see also DeLeire, Chappel, Finegold and Gee, 2017) up to covering 88 percent of healthcare spending.

In terms of premiums, there are a few pieces of regulation worth mentioning. First, regulation is such that insurers simply choose a scalar base premium per plan. An age-adjustment schedule then turns this base premium into an age-specific premium, limiting price discrimination across demographic groups by insurers. That is the revenue that insurers obtain from selling a plan. On the demand side, individuals may be eligible for subsidies as a function of their income level relative to the federal poverty line (FPL), where the magnitude of the subsidies is linked to the premium of the second cheapest Silver plan (see also Jaffe and Shepard, 2020). The premiums that consumers pay are net of these subsidies. In our analysis, we are careful to distinguish between insurer revenue and premiums paid by consumers.

Other regulation include risk adjustment policies (Saltzman, 2021), reinsurance programs, and a minimum insurer medical loss ratio (Cicala, Lieber and Marone, 2019). For

[6]Federal regulations (see e.g. `https://www.kff.org/affordable-care-act/issue-brief/network-adequacy-standards-and-enforcement/`) require that insurers cover essential providers and that networks are such that individuals do not need to travel too far, or wait for too long, to see physicians. Despite these rules, which are often stricter in State-based marketplaces such as Covered California, networks have been narrower than in other segments of the insurance market (Haeder, Weimer and Mukamel, 2019).

further details, we refer readers to Saltzman (2021); Tebaldi (2024) and references therein.

3.2 Data

Enrollment. Our primary data source is administrative records on the universe of individuals enrolled in plans offered in Covered California for 2014–2017. The data include individual and plan characteristics, including individual and household identifiers, age, gender, zipcode of residence, household income as a share of FPL, and plan enrollment. Enrollment information includes insurer, network type, metal tier, and premium.

For the analysis, we focus on single policyholders between 26 and 64 with income below 400 percent of FPL. The number of enrollees in the data is 3,094,738. We provide summary statistics for enrollees in the data in Table 1-A. The average enrollee is 47 years old and has an income of 229 percent of the FPL. However, there is substantial dispersion in all these attributes across the market. Regarding enrollment, 64 percent of enrollees chose a Silver plan, 27 percent chose a Bronze plan, and the remaining 9 percent chose either gold or platinum. Enrollees pay on average $1,921 in premiums per year, with substantial dispersion driven by age-based subsidies.

While our data provide a detailed characterization of enrollees and plans within Covered California, it does not provide information about consumers outside the marketplace. We exploit data from the 2009–2013 American Community Survey public use files to characterize that group. We estimate market size using a method that extends those in related work by Tebaldi *et al.* (2023) and Dickstein *et al.* (2024) by allowing for more geographic granularity to get at market size by zipcode, demographics, and year. We provide further details in Appendix A.1. Table 1-D displays summary statistics for these potential buyers. The share of potential enrollees that choose a plan in Covered California is 32 percent. This enrollment share varies substantially across zipcodes and demographics, likely reflecting heterogeneity in willingness-to-pay for the plans in the marketplace.

Provider networks. To construct measures of plan provider networks, we leverage detailed data on the universe of providers each plan in the marketplace offers. Covered California mandates that every qualified health plan (QHP) provide monthly information on all PCPs, specialists, hospitals, and other providers they cover. Our data come from quarterly snapshots of these reports. For our analysis, we use the second quarter of each year.[7] For each physician and facility, the data include the name, the national provider identification number (NPI), specialty, zipcode, and the plans it has contracted with. For

[7]There is minimal variation in provider networks reported across quarters within a year. In addition, the quality of the data for 2014 is low, and hence we impute the data from 2015 to 2014.

Table 1: Descriptive statistics

	(1)	(2)	(3)	(4)	(5)	(6)
Variable	N	Mean	SD	p10	p50	p90
A. Enrollment data						
Age	3,094,738	46.97	11.36	30.00	49.00	61.00
Income as percentage of FPL	3,094,738	229.11	62.74	160.00	215.00	325.00
Yearly premium	3,094,738	1920.56	1357.56	485.03	1596.90	3779.97
Enrolled in Bronze plan	3,094,738	0.27	0.45	0.00	0.00	1.00
Enrolled in Silver plan	3,094,738	0.64	0.48	0.00	1.00	1.00
Enrolled in Gold plan	3,094,738	0.05	0.21	0.00	0.00	0.00
Enrolled in Platinum plan	3,094,738	0.04	0.19	0.00	0.00	0.00
B. Potential buyers						
Enrolled in ACA plan	9,657,042	0.32	0.47	0.00	0.00	1.00
Under age 45	9,657,042	0.56	0.50	0.00	1.00	1.00
C. Plan networks						
HMO plan	33,416	0.53	0.50	0.00	1.00	1.00
Share of PCPs in 0-10 miles	33,416	0.25	0.22	0.00	0.20	0.56
Distance to closest hospital	33,416	8.97	11.95	0.00	5.02	21.10
D. Plan data						
HMO plan	1,429	0.51	0.50	0.00	1.00	1.00
Average revenue per enrollee	1,429	5,355.75	1,257.70	3,935.05	5,134.53	7,159.26
Average cost per enrollee	1,094	3,835.13	1,988.96	1,896.21	3,602.54	6,018.96
Average cost per enrollee | HMO	631	3,577.64	1,974.21	1,738.94	3,080.70	6,018.96
Average cost per enrollee | PPO	463	4,218.50	1,950.82	2,476.39	4,113.38	5,574.11
Average cost per enrollee | Bronze	246	2,225.33	938.56	1,084.06	2,170.79	3,533.53
Average cost per enrollee | Silver	292	4,012.69	1,249.28	2,656.11	4,113.38	5,574.11
Average cost per enrollee | Gold	280	4,998.72	1,632.77	3,215.29	5,006.43	7,254.62
Average cost per enrollee | Platinum	276	9,295.58	3,815.89	3,688.68	9,622.29	13,032.99

Notes: Panel A displays summary statistics of the enrollment data at the individual level, for enrollees in Covered California. Panel B displays summary statistics of plan market shares, revenue, and costs at the insurer-region-tier-network-year level. Revenue and cost are reported in $/year, and weighted by plan enrollment. Panel C displays summary statistics of plan networks at the insurer-region-zipcode-tier-network-year. Panel D displays summary statistics on enrollment in Covered California by potential buyers and whether a potential buyer is under age 45 at the individual level.

our analysis, we focus on the networks of PCPs and hospitals.

We leverage data on the location of both households and providers at the zipcode level to construct measures of provider networks offered by each plan in each zipcode. For physicians, we measure the quality of plan j's network in zipcode z as the share of the providers in a catchment area C_z^c around z covered by plan j. Let N_{zjt} be the number of

PCPs in zipcode z that are covered by plan j in year t. Our measure is:

$$\mathcal{N}_{jz}^{c} = \frac{\sum_{l \in C_z^c} N_{ljt}}{\sum_{l \in C_z^c} \sum_{p \in \mathcal{J}_{zt}} N_{lpt}}$$

where $\mathcal{J}_{zt}$ is the set of plans offered in zipcode z in year t. [8] We compute these measures for a catchment area c, defined as a 10-mile radius relative to the centroid of each zipcode for our main analysis.

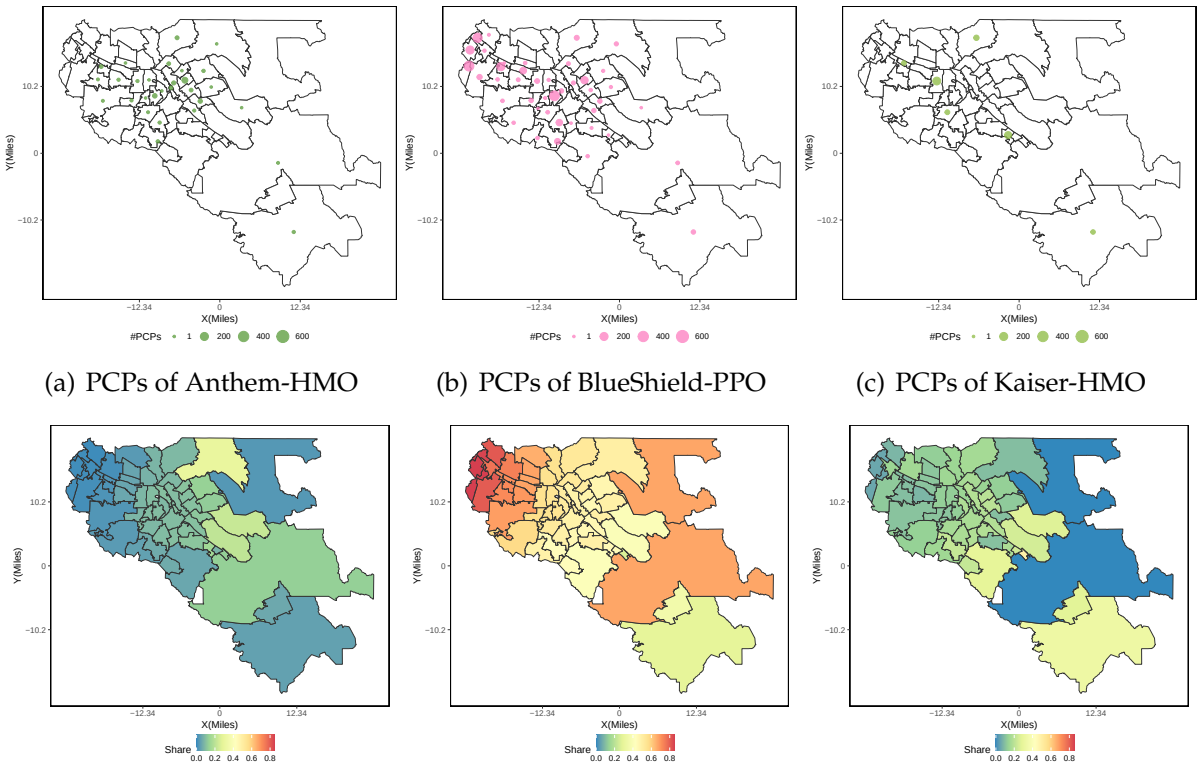
To illustrate this approach, Figure 1 displays the spatial distribution of PCPs in the Santa Clara region that Anthem-HMP, BlueShield-PPO, and Kaiser-HMO cover. Panels (a), (c), and (e) display the raw data at the zipcode level, where it is easy to note differences in coverage across plans both overall in the region and also in relative terms across zipcodes in the region. As a result, our measures of network quality for this region display substantial variation across plans, but also within a plan across zipcodes, which is what panels (b), (d), and (f) show. These patterns repeat across regions and plans. Table 1-C shows that plans, on average, cover 25 percent of available PCPs within 10 miles of a zipcode, but there is substantial variation across plans and zipcodes, with a standard deviation of 22 percent.

The spatial distribution of hospitals is more sparse than that of physicians. This motivates a different measurement of local hospital network quality. In particular, we measure hospital network quality as the distance to the closest hospital from zipcode z that is covered by plan j in year t. Table 1-C displays summary statistics for this measure. The closest hospital is around 9 miles away for the average plan and zipcode. However, there is vast variation across zipcodes, with many zipcodes having an in-network hospital for a plan but also a long tail of zipcodes that must travel more than 20 miles to the closest in-network hospital.

Cost data. To measure plan average costs, we exploit data from annual Rate Review Filings (RRF) for 2014–2017. The data report annual claims per plan member for each plan and region in Covered California. Table 1-D shows that the average cost per enrollee is around $3,800 per year, and it is 18 percent higher among HMO than PPO plans. Similarly, plan average costs are increasing in plan generosity: Silver, Gold, and Platinum plans display average costs around 2, 2.5, and 4.5 times higher than Bronze plans. These patterns likely reflect selection on provider networks and plan generosity, respectively.

[8] It is possible in the data for the same provider to have multiple practice locations within a zipcode. In the analysis, we treat these locations as one for the same provider within a zipcode. In contrast, if the same provider offers multiple locations in different zipcodes within a catchment area, we count them as separate providers.

Figure 1: PCP network data: raw data and coverage shares



(a) PCPs of Anthem-HMO (b) PCPs of BlueShield-PPO (c) PCPs of Kaiser-HMO

(d) Share of Anthem-HMO (e) Share of BlueShield-PPO (f) Share of Kaiser-HMO

Notes: Top: The figure shows the spatial distribution of centroids of zipcodes that have accessible PCPs covered by Anthem-HMO, BlueShield-PPO, and Kaiser-HMO in Region 7, Q2 2015. The size of the dots shows the total number of available PCPs of specific insurer-network in the zipcode in each panel. Bottom: The figure shows a heatmap of share within 10 miles of Anthem-HMO, BlueShield-PPO, and Kaiser-HMO in Region 7, Q2 2015. Zipcodes with no available PCPs are left empty.

3.3 Relationship between Plan Attributes and Enrollment Choices

Different dimensions of the ACA regulations generate quasi-experiments in the choice and network data that we exploit to measure the causal effects of plan attributes on plan demand. These effects are the backbone of our demand model which we introduce in Section 4.

We begin by noticing that regulation generates variation in plan premiums and actuarial value for individuals of different demographics within a zipcode, holding every other aspect of the choice environment fixed. Analogously, geographic variation in provider networks within a rating region leads to variation in network breadth across individuals residing in different zipcodes, who otherwise face identical prices and contracts and have the same age and income. It seems natural, then, to estimate fixed-effects regressions to isolate these sources of variation to document the impact of plan attributes on plan choice.

Table 2: Relationship between plan choices and plan attributes

	(1)	(2)	(3)
	1(Chosen plan) × 100		
Annual premium	-0.195***		
	(0.012)		
Actuarial value		0.129***	
		(0.003)	
Distance to closest hospital			-0.014***
			(0.003)
Share PCPs in 0-10 miles			1.859***
			(0.165)
Share PCPs in 10-20 miles			0.108
			(0.165)
Mean of dependent variable	1.412	3.325	1.412
N	213,578,845	56,307,961	213,578,845
R^2	0.143	0.155	0.090
Plan-zipcode-coarse age-income FE	Y	N	N
Plan-zipcode-age-coarse income FE	N	Y	N
Plan-age-income FE	N	N	Y
Sample	All plans	Silver plans	All plans

Notes: This table displays results from a regression of an indicator of plan choice on plan attributes and varying sets of fixed effects as indicated in the bottom panel. The regression is estimated using individual-level data. Age is collapsed to 1-year bins for fixed effects, whereas coarse age is collapsed to 4-year bins. Income is collapsed to 5 percent bins of income relative to FPL, whereas coarse income is collapsed to 10 percent bins. For column 2, the regression is estimated on a sample that only includes Silver plans. The dependent variable is scaled by 100 for clarity of the results. Clustered standard errors at the zipcode level in parentheses. ***$p < 0.01$, **$p < 0.05$, *$p < 0.1$.

For instance, to measure the effect of premiums on plan choice we estimate the following regression:

$$y_{ijrt} = \beta \text{Annual premium}_{ijrt} + \eta_{jrtz(i)w(i)} + \varepsilon_{ijrt}\,, \tag{8}$$

where y_{ijrt} is an indicator for individual i choosing plan j, Annual premium$_{ijrt}$ is the annual premium the plan j charges to an enrollee with individual i's demographics, and $\eta_{jrtz(i)w(i)}$ is a granular set of fixed effects that consists of the interaction of plan, region, year, zipcode, and demographic types $w(i)$ that interact 4-year age bins with 10-percentage point bins of income relative to FPL. Under this specification, β captures demand responses to variation in plan premiums for contracts that are otherwise identical—actuarial value does not vary within plan and narrow income groups and provider networks do not vary within plan and zipcode—among narrowly defined demographic groups. Column 1 in Table 2 shows that the relationship is negative and strong, as expected. An increase of $10/month in the plan premium is associated with a decrease of 1.7 percent in the likelihood it is chosen.

We use variants of equation (8) to measure the effect of actuarial value and provide network breadth on plan choice. For actuarial value we focus on Silver plans and isolate the variation introduced by three regulatory discontinuities at 150, 200, and 250 percent of income relative to FPL. [9] In particular, we modify the set of fixed by defining demographic types $w(i)$ as the interaction of finer 1-year age bins with coarser 20 percentage point bins of income relative to FPL. The former account for some variation in premiums across demographic groups. The latter isolate variation in actuarial value from cost-sharing reductions around the regulatory discontinuities mentioned above. Under this specification, plans only differ in actuarial value and are otherwise almost identical. Column 2 in Table 2 shows estimates of this regression on data for Silver plans. The relationship is positive and strong, as expected, with a five percentage point increase in actuarial value being associated with a 19 percent decrease in the likelihood it is chosen.

At last we focus on the effect of network breadth on plan choice. For this, we estimate yet another variant of equation (8) including network attributes as independent variables and modifying the set of fixed effects in two ways. First, by dropping the interaction with zipcode. Second, by demographic types $w(i)$ as the interaction of narrow 1-year age bins and 10-percentage point bins of income relative to FPL. This specification isolates variation in network quality of a plan while keeping plan premium, plan generosity, and other plan attributes fixed. Column 3 in Table 2 shows that higher network quality leads to higher plan demand. The relationships are quantitatively meaningful. For instance, reducing the distance to the closest in-network hospital by 5 miles leads to an increase in the likelihood of choosing a plan of 4.9 percent, whereas a 10 percentage points increase in PCP coverage within ten miles of a zipcode is associated with a 13 percent increase in the likelihood of choosing a plan. These patterns suggest that provider networks are a relevant driver of plan demand and, hence, that variety may play a role when evaluating alternatives to managed competition. Our measures of revealed preferences for provider networks complement previous work measuring network quality using the inclusive value of a hospital demand model (e.g., Capps, Dranove and Satterthwaite, 2003; Ho, 2009).

[9] The actuarial value of Silver plans falls sharply at these discontinuities: from 95 to 88 percent at an income relative to FPL of 150 percent, then further to 74 at an income relative to FPL of 200 percent, and further to 70 at an income relative to FPL of 250 percent. See also DeLeire *et al.* (2017), Lavetti, DeLeire and Ziebarth (2023), and Tebaldi (2024).

4 Econometric Model

4.1 Demand

Pure characteristics random utility model. We employ a version of the pure characteristics demand model introduced in Berry and Pakes (2007), which we tailor to accommodate the granular individual-level data available in our empirical setting.

An individual i is summarized by its type $\theta_i = (w_i, z_i, \alpha_i, \beta_i, \mu_i, \xi^{w_i})$, where w_i collects age and income; z_i indicates zipcode of residence (within a region, which we index by r throughout); and α_i, β_i, μ_i and ξ^{w_i} collect preference parameters that affect utility for a plan. In region r and year t (or simply rt), individuals pick one plan j from the set $\mathcal{J}_{rt}$ or remain uninsured.

A plan is characterized by actuarial value $AV_{jrt}^{w_i}$, varying with income due to cost-sharing reductions, a collection of network variables $\mathcal{N}_{jrt}^{z_i}$ that includes zip-specific network breadth and distance to the closest hospital, and the premium is $p_{jrt}^{w_i}$, varying with income and age as described in Section 3.1. Moreover, every plan j in rt corresponds to a unique insurer-network combination (e.g. Anthem PPO in North Los Angeles in 2015) f_{jrt}; this is the level at which plans may differ for reasons that are unobserved to the econometrician.

The indirect utility that i derives from choosing j in rt is equal to:

$$u_{jrt}(\theta_i) = -\alpha_i p_{jrt}^{w_i} + \beta_i^{av} AV_{jrt}^{w_i} + \beta_i^{n} \mathcal{N}_{jrt}^{z_i} + \mu_i + \xi_{f_{jrt}}^{w_i} \; ; \tag{9}$$

in the notation of Section 2: $v(\theta_i, x_{jrt}) - p_{jrt}^{w_i} = u_{jrt}(\theta_i)/\alpha_i$. We normalize the utility of the reinsurance outside option to $u_{0rt}(\theta_i) = 0$.

The term $\xi_{f_{jrt}}^{w_i}$ captures unobservables at the insurer-network-region-year level and is allowed to vary with age in seven bins: [26-30), [30-35), [35-40), [40-45), [45-50), [50-55) and [55-64]. Across the same seven bins, the distribution of $(\alpha_i, \beta_i, \mu_i)$ is arbitrary, while within each bin we set

$$\begin{aligned}
\alpha_i &\sim \text{LogNormal}(\overline{\alpha}^{w_i}, \sigma^{\alpha, w_i}) \, , \\
\beta_i^{av} &\sim \text{LogNormal}(\overline{\beta}^{av, w_i}, \sigma^{av, w_i}) \, , \\
\beta_i^{n} &\sim \text{Normal}(\overline{\beta}^{n, w_i}, \Sigma^{n, w_i}) \, , \\
\mu_i^{n} &\sim \text{Normal}(\overline{\mu}^{w_i}, \sigma^{\mu, w_i}) \, ;
\end{aligned} \tag{10}$$

the subscript w_i indicates that all parameters vary by age bin, the parameter $\overline{\alpha}^{w_i}$ also varies (linearly) with income measured as a percentage of the FPL, and the distribution of β_i^n is a bivariate normal with diagonal variance-covariance matrix Σ^{n, w_i}. The normally

distributed random constant μ_i generates a nesting structure insofar individuals with particularly low draws of this parameter are hardly attracted by any plan, and vice-versa individuals with particularly high draws of μ_i hardly consider the outside option.

Importantly, the model specified in (9) does not include an idiosyncratic error term, departing from traditional logit models most often employed by the literature since Berry (1994); Berry, Levinsohn and Pakes (1995). These models are quite attractive because the Gumbel distribution leads to analytical expressions for choice probabilities, simplifying estimation. However, as already noted in Ackerberg and Rysman (2005); Dafny, Ho and Varela (2013); Quan and Williams (2018), these models can significantly overestimate the impact of large changes in the number of available options on quantity and welfare, a feature that partly motivates Berry and Pakes (2007). At the cost of computational complexity in estimation and counterfactual analyses, our model in (9) does not suffer from this issue, and can be used to study counterfactual changes in choice sets without the need to correct the scale of preference shocks. Relatedly, in the absence of the idiosyncratic error term, the model in (9) requires a scale normalization. For this, we set the standard deviation of α_i equal to 0.25.[10]

Given a measure of types $G(\theta_i)$ which combines the number of individuals in each demographic and zipcode with the distribution of parameters conditional on age bin, the total enrollment in plan j in rt can be written as

$$\begin{aligned} Q_{jrt} &= \int \mathbf{1}\left[j = \arg\max_k u_{krt}(\theta_i)\right] dG(\theta_i) \\ &= \sum_{w_i}\sum_{z_i} N_{rt}(w_i, z_i) \underbrace{\int \mathbf{1}\left[j = \arg\max_k u_{krt}(\theta_i)\right] dG(\theta_i|w_i, z_i)}_{=\sigma_{jrt}(w_i, z_i)}, \end{aligned} \tag{11}$$

where we use $\sigma_{jrt}(w_i, z_i)$ to denote the conditional choice probability for an individual with demographics w_i in zipcode z_i, and $N_{rt}(w_i, z_i)$ the number of these individuals rt.

Identification. The enrollment data introduced in Section 3 collects, for every individual i, the plan characteristics included in the indirect utility in (9) as well as (Y_i, w_i, z_i). The choice outcome is $Y_i \in J_{rt}$, with $\Pr[Y_i = j w_i, z_i, rt] = \sigma_{jrt}(w_i, z_i)$. Identification of the vectors

[10]This choice is dictated by two observations. First, the ratio between average and standard deviation of the price coefficients α_i's across the seven age bins are pinned down consistently when we use different scale parameters. Their range is between 0.5 and 1.5. Second, setting the standard deviation to 1—which may seem more natural—generates a large mass of $\alpha_i \approx 0$. While this is not an issue for estimation or demand responses, it creates large numerical instability when computing consumer surplus (for which we must divide by α_i). The choice of 0.25 avoids this problem making near-zero α_i extremely rare.

of parameters $(\overline{\alpha}^{w_i}, \sigma^{\alpha,w_i}, \overline{\beta}^{av,w_i}, \sigma^{av,w_i}, \overline{\beta}^{n,w_i}, \Sigma^{n,w_i}, \overline{\mu}^{w_i}, \sigma^{\mu,w_i}, \xi^{w_i})$, one for each of the seven age bins listed above, relies on the panel structure we discussed in Section 3.3 (c.f. Table 2). This turns out to be a special case of the variation assumed in Berry and Haile (2024) to obtain nonparametric identification of demand.

In our setting, preferences over networks $(\overline{\beta}^{n,w_i}, \Sigma^{n,w_i})$ are identified from the joint variation in choices and networks conditional on region and demographics - $(Y_i, \mathcal{N}_{rt}^{z_i} | w_i, rt)$ - across zipcodes z_i. By conditioning on w_i and rt the set of plans (J_{rt}), their unobservable characteristics (ξ^{w_i}), premium and actuarial value do not vary. The key assumption to identify preferences over provider networks is that individuals who live in a specific zipcode do not have preferences for physicians and hospitals that systematically differ from those in other zipcodes. In other words, we assume that β_i^n is distributed *iid* across zipcodes conditional on age and income.

For the parameters governing the distribution of the premium coefficient $(\overline{\alpha}^{w_i}, \sigma^{\alpha,w_i})$, as in Tebaldi, Torgovitsky and Yang (2023) we rely on the fact that within coarse bins of age and income, the ACA regulation varies premiums exogenously across individuals who otherwise face the same exact menu of options. This generates price variation conditional on actuarial value, zipcode (and thus $\mathcal{N}_{rt}^{z_i}$), J_{rt}, and $\xi_{rt}^{w_i}$. In practice, one can hold fixed all characteristics other than price, and the main identification assumption is that individuals who have very similar age and income have the same distribution of α_i.[11] Similarly, preferences over insurance generosity measured by actuarial value, $(\overline{\beta}^{av,w_i}, \sigma^{av,w_i})$, are identified exploiting the discontinuous regulatory variation in the actuarial value of Silver plans at three income discontinuities (income equal to 150, 200, and 250 percent of the FPL, see also DeLeire, Chappel, Finegold and Gee, 2017; Tebaldi, 2024).

Given the distribution of $(\alpha_i, \beta_i^{av}, \beta_i^n)$, the remaining parameters of the model are the distribution of the constant term μ_i and the unobservables terms collected in $\xi_{rt}^{w_i}$, one for every combination of insurer-network-region-year. These must match aggregate shares of every product, in every market, after accounting for variation in prices. As usual, a plan j that generates - on average across individuals - a relatively low value of utility net of ξ^{w_i} and yet (within the age bin) registers a large share of enrollment in rt relative to its competitors is associated to a larger value of $\xi_{f_{jrt}}^{w_i}$, and vice-versa. Similarly, the mean of μ_i is pinned down by the share of individuals choosing the outside option. As for a nesting parameter in a nested logit model, σ^{μ,w_i} matches the differences in substitution patterns between plans within the market and between the inside plans and the outside option.

[11] We also considered an alternative approach using an instrumental variable strategy as in Tebaldi (2024); this led to similar results and we preferred to use the rich set of "fixed effects" $\xi_{rt}^{w_i}$ and estimate the model via maximum likelihood without control function.

Estimation. We estimate all the parameters of the model jointly via simulated maximum likelihood within each of the seven age groups. For this, we construct choice probabilities for each individual as a function of parameters and then maximize the log-likelihood

$$\mathcal{L} = \sum_i \log \left\{ \prod_j \sigma_{jrt}(w_i, z_i)^{\mathbf{1}[Y_i = j]} \right\} . \tag{12}$$

Unlike the widely used logit and probit discrete choice models with idiosyncratic error terms with full support, the pure characteristic model with microdata we employ here can predict $\sigma_{jrt}(w_i, z_i)$ for a given parameter value (or sometimes any value).[12] We currently handle this instances by replacing $\sigma_{jrt}(w_i, z_i) = 0$ with the lowest possible number ϱ for which - at machine precision - $\log(\varrho) > -\infty$. We then verify ex-post that the share of individuals in our dataset for which this verifies is small. With the current specification, this is between 4-9 percent across the seven estimation datasets, one for every age bin.

To compute $\sigma_{jrt}(w_i, z_i)$ we we follow the steps outlined in Berry and Pakes (2007). First, when i selects j this implies that for all $k \neq j$

$$\alpha_i(p_{krt}^{w_i} - p_{jrt}^{w_i}) \geq \beta_i^{av}(AV_{krt}^{w_i} - AV_{jrt}^{w_i}) + \beta_i^{n}(\mathcal{N}_{krt}^{z_i} - \mathcal{N}_{jrt}^{z_i}) + \mu_i(\mathbf{1}[k \neq 0] - \mathbf{1}[j \neq 0]) + \xi_{f_{krt}}^{w_i} - \xi_{f_{jrt}}^{w_i} \; ;$$

or equivalently that

$$\alpha_i \geq \underline{\Delta}_i \equiv \max_{k: p_{krt}^{w_i} > p_{jrt}^{w_i}} \frac{\beta_i^{av}(AV_{krt}^{w_i} - AV_{jrt}^{w_i}) + \beta_i^{n}(\mathcal{N}_{krt}^{z_i} - \mathcal{N}_{jrt}^{z_i}) + \mu_i(\mathbf{1}[k \neq 0] - \mathbf{1}[j \neq 0]) + \xi_{f_{krt}}^{w_i} - \xi_{f_{jrt}}^{w_i}}{\alpha_i},$$

$$\alpha_i \leq \overline{\Delta}_i \equiv \min_{k: p_{krt}^{w_i} < p_{jrt}^{w_i}} \frac{\beta_i^{av}(AV_{krt}^{w_i} - AV_{jrt}^{w_i}) + \beta_i^{n}(\mathcal{N}_{krt}^{z_i} - \mathcal{N}_{jrt}^{z_i}) + \mu_i(\mathbf{1}[k \neq 0] - \mathbf{1}[j \neq 0]) + \xi_{f_{krt}}^{w_i} - \xi_{f_{jrt}}^{w_i}}{\alpha_i}.$$

The quantities $\underline{\Delta}_i$ and $\overline{\Delta}_i$ are a function of data and parameters. For any value of $\boldsymbol{\xi}$ and a set of M Halton draws from the distribution of $\boldsymbol{\beta}$, one obtains a collection $\{\underline{\Delta}^m, \overline{\Delta}^m\}_{m=1}^M$, and can approximate the choice probabilities as

$$\sigma_{jrt}(w_i, z_i) \approx M^{-1} \sum_m \int_{\underline{\Delta}_i^m}^{\overline{\Delta}_i^m} \phi\left(\frac{\log(\alpha) - \overline{\alpha}^{w_i}}{\sigma^{\alpha, w_i}} \right) d\alpha \; , \tag{13}$$

where $\phi(\cdot)$ is the density of a standard Gaussian distribution.

[12] For a given specification, one can solve a series of linear programs to pre-determine the number of observations in a dataset for which the observed choice $Y_i = j$ will have zero probability for any value of demand parameters.

The right-hand side of (13) only depends on data and parameters, and can be used to evaluate the likelihood of $Y_i = j$ in (12). We maximize $\mathcal{L}$ employing Artelys' Knitro multi-algorithm nonlinear programming procedure, for which we input simple logit estimates as a starting point.[13]

4.2 Cost and Insurer Conduct

Expected cost, equilibrium pricing, and selection. We specify a cost model that is similar to the one employed in Tebaldi (2024). An individual of type θ_i who enrolls in plan j in rt generates expected claims for the insurer in the amount of

$$c_{jrt}(\theta_i) = av_j \underbrace{\exp[K_{jrt} + \zeta \text{age}_{w_i} + \lambda \beta_i^{av}/\alpha_i]}_{\text{Expected medical spending}}, \tag{14}$$

where the term av_j indicates the approximate actuarial value of plan j that in Covered California is fixed to 60 for Bronze plans, 70 for Silver plans, 80 for gold plans, and 90 for platinum plans.[14] Log-annual expected medical spending is a linear function of a plan-specific term K_{jrt}, enrollee age, and their willingness-to-pay for actuarial value, as measured by β_i^{av}/α_i. The coefficient λ captures selection on unobservable determinants of preferences: if positive, individuals who are more willing to pay for generous insurance carry higher risk, and vice-versa.

At the observed premiums and characteristics, our model-predicted expected average cost for plan j in rt is therefore

$$AC_{jrt} = Q_{jrt}^{-1} \int c_{jrt}(\theta_i) \mathbf{1}\left[j = \arg\max_k u_{krt}(\theta_i) \right] dG(\theta_i),$$

which we can compute for any combination of $(K_{jrt}, \zeta, \lambda)$ using demand parameters.

To model insurers' pricing behavior under the ACA we must enrich the framework of Section 2 to incorporate federal rating rules. Insurers are not setting premiums directly. Instead, they must pick one scalar-valued "base rate" τ_{jrt} for every plan j in every region-year rt, and this is then mapped into premiums as a function of age and income after applying age-adjustments and subsidies: $p_{jrt}^{w_i} = \rho(\tau_{jrt}, w_i)$.[15]

---

[13] These are MLE estimates from a model in which we add an idiosyncratic type 1 extreme value error term to each alternative in the choice set and set to zero the standard deviations for all random coefficients.

[14] This differs from the random variable $AV_j^{w_i}$ entering the utility specification in (9) because of cost-sharing subsidies for low-income enrollees.

[15] The rating rules in $\rho(\cdot)$ have been studied extensively in Jaffe and Shepard (2020); Tebaldi (2024), here

With these rating rules operating in the background, if $\boldsymbol{\tau}^{\star}_{rt}$ is the vector of rates set in a Nash equilibrium in rt, imposing (1) amounts to equalize marginal revenue and marginal cost for each plan j:

$$\begin{aligned}\frac{\partial}{\partial \tau_{jrt}} \sum_{k \in J_{f_j}} R_{krt}(\boldsymbol{\tau}^{\star}_{rt}) &= \frac{\partial}{\partial \tau_{jrt}} \sum_{k \in J_{f_j}} Q_{krt}(\boldsymbol{\tau}^{\star}_{rt}) AC_{krt}(\boldsymbol{\tau}^{\star}_{rt}) \\ &= \sum_{k \in J_{f_j}} AC_{krt}(\boldsymbol{\tau}^{\star}_{rt}) \frac{\partial Q_{krt}(\boldsymbol{\tau}^{\star}_{rt})}{\partial \tau_{jrt}} + Q_{krt}(\boldsymbol{\tau}^{\star}_{rt}) \frac{\partial AC_{krt}(\boldsymbol{\tau}^{\star}_{rt})}{\partial \tau_{jrt}},\end{aligned} \tag{15}$$

where—emphasizing the dependence on base rates—total revenue on the left-hand side is computed as $R_{krt}(\boldsymbol{\tau}_{rt}) = \sum_i \sigma_{krt}(w_i, z_i)(\rho(\tau_{jrt}, w_i) + \text{subsidy}_i)$, and J_{f_j} collects the plans offered by the insurer selling j. The right-hand side in (15) is the sum of the marginal cost across these plans. If either ζ or λ in (14) are not zero, so that expected costs vary across individual types, in this model $AC_{krt}(\boldsymbol{\tau}^{\star}_{rt}) \neq \frac{\partial AC_{krt}(\boldsymbol{\tau}^{\star}_{rt})}{\partial \tau_{jrt}}$, as it is typically the case in a selection market.

Identification. The model features one unobservable for every supply-side observation (K_{jrt}) in addition to the parameters ζ and λ which govern cost heterogeneity across individuals within plan. Rather than imposing additional cross-plan restrictions on K_{jrt},[16] to identify cost parameters we combine observed average claims data with the implications of supply-side optimal pricing in (15). The argument is constructive and adapts the widely used inversion of oligopoly pricing conditions (Rosse, 1970; Bresnahan, 1987; Berry, Levinsohn and Pakes, 1995) to an insurance market, where the discrepancy between average and marginal cost is a key measure of the sign and magnitude of selection (Einav, Finkelstein and Cullen, 2010).

First, consider any given value of (ζ, λ), and—simplifying notation since all terms are evaluated at the observed rates $\boldsymbol{\tau}^{\star}_{rt}$—one can rewrite (15) as

$$\sum_{k \in J_{f_j}} \frac{\partial R_{krt}}{\partial \tau_{jrt}} = \sum_{k \in J_{f_j}} e^{K_{jrt}} \left(H^{(\zeta,\lambda)}_{krt} \frac{\partial Q_{krt}}{\partial \tau_{jrt}} + Q_{krt} \frac{\partial H^{(\zeta,\lambda)}_{krt}}{\partial \tau_{jrt}} \right), \text{ for all } j,$$

where $H^{(\zeta,\lambda)}_{jrt} \equiv e^{-K_{jrt}} AC_{jrt}$ only captures heterogeneity in cost due to the characteristics of individuals choosing plan j. In matrix notation this becomes $\boldsymbol{MR}_{rt} = (\Xi(\zeta, \lambda) \cdot \Omega)\, e^{\boldsymbol{K}_{rt}}$, where

we hold it fixed throughout and we focus instead on the market mechanism.

[16] Since average claims are observed, one alternative route is to impose cross-plan restrictions on K_{jrt} and project average claims on a set of observable characteristics and the composition of enrollment. This approach has been followed by Bundorf, Levin and Mahoney (2012) and Tebaldi (2024). The latter provides evidence in support of the oligopoly pricing assumption that we adopt in our analysis throughout.

Ω is an ownership matrix with generic element $\omega(j,k) = \mathbf{1}[f_j = f_k]$, $e^{\boldsymbol{K}_{rt}} = \left(e^{K_{1rt}}, \cdots, e^{K_{Jrt}}\right)$, $\boldsymbol{MR}_{rt}^{\intercal} \equiv (\sum_{k \in J_{f_1}} \frac{\partial R_{krt}}{\partial \tau_{1rt}}, \ldots, \sum_{k \in J_{f_J}} \frac{\partial R_{Jrt}}{\partial \tau_{Jrt}})$, and

$$\Xi(\zeta,\lambda) \equiv \begin{bmatrix} H_{1rt}^{(\zeta,\lambda)} \frac{\partial Q_{1rt}}{\partial \tau_{1rt}} + Q_{1rt} \frac{\partial H_{1rt}^{(\zeta,\lambda)}}{\partial \tau_{1rt}} & \cdots & H_{Jrt}^{(\zeta,\lambda)} \frac{\partial Q_{Jrt}}{\partial \tau_{1rt}} + Q_{Jrt} \frac{\partial H_{Jrt}^{(\zeta,\lambda)}}{\partial \tau_{1rt}} \\ \vdots & \ddots & \\ H_{1rt}^{(\zeta,\lambda)} \frac{\partial Q_{1rt}}{\partial \tau_{Jrt}} + Q_{1rt} \frac{\partial H_{1rt}^{(\zeta,\lambda)}}{\partial \tau_{Jrt}} & & H_{Jrt}^{(\zeta,\lambda)} \frac{\partial Q_{Jrt}}{\partial \tau_{Jrt}} + Q_{Jrt} \frac{\partial H_{Jrt}^{(\zeta,\lambda)}}{\partial \tau_{Jrt}} \end{bmatrix}.$$

From this expression, for a given (ζ, λ), assuming equilibrium pricing one obtains:

$$\widehat{\boldsymbol{K}}_{rt}(\zeta,\lambda) \equiv \log\left\{\left(\Xi(\zeta,\lambda)\cdot\Omega\right)^{-1}\right\}\boldsymbol{MR}_{rt}\,,$$

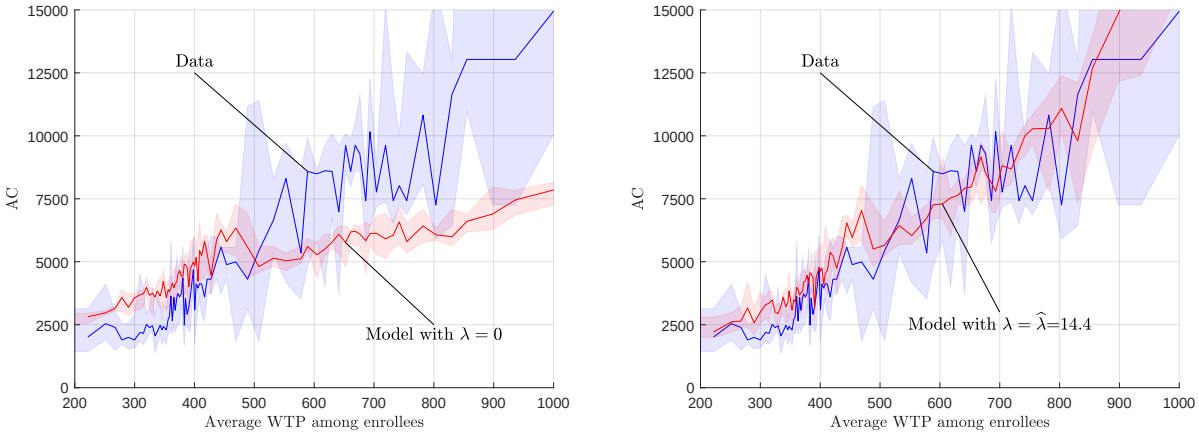
and a corresponding value of model-predicted average cost: $\widehat{\boldsymbol{AC}}_{rt}(\zeta,\lambda) = e^{\widehat{\boldsymbol{K}}_{rt}(\zeta,\lambda)}\boldsymbol{H}_{rt}^{(\zeta,\lambda)}$. When setting $\zeta = \lambda = 0$, the above procedure is *identical* to the case with constant marginal cost since $\boldsymbol{H}_{jrt}^{(0,0)} \equiv av_j$ for all j, and $\widehat{\boldsymbol{AC}}_{rt}(0,0)$ would be simply equal to the vector of constant marginal costs that rationalizes observed rates as optimal for every insurer.

Following these steps returns a measures of average and marginal cost for every plan, and every value of selection parameters ζ and λ. Following the intuition of Einav, Finkelstein and Cullen (2010), one can then use observed average claims data to find the values $(\widehat{\zeta},\widehat{\lambda})$ that satisfy $AC_{jrt}^{\text{obs}} = \widehat{AC}_{jrt}(\widehat{\zeta},\widehat{\lambda})$. These will be unique as long as the econometrician observes sufficient variation across jrt in the within-plan difference between the distribution of age and β_i^{av}/α_i among enrollees and the distribution of the same variables among marginal buyers. Intuitively: when $\lambda > 0$ plans enrolling many individuals with high willingness-to-pay but with marginal buyers who have, on average, lower β_i^{av}/α_i are predicted to have a difference between marginal cost (inferred from optimal pricing) and average cost (observed in the data) when compared to plans for which the composition of marginal and average buyers in terms of β_i^{av}/α_i is similar.

Estimation. Since the identification argument above is constructive it lends itself to an analogous estimation strategy. To retain more statistical power for the estimation of λ, we calibrate ζ to 0.0379 to fit the relationship between log cost and age among insured individuals in the Medical Expenditure Panel Survey (see Tebaldi, 2024).

We then search over a grid of λ parameters for the value that minimizes the enrollment-weighted average measure of $\left|AC_{jrt}^{\text{obs}} = \widehat{AC}_{jrt}(\widehat{\zeta},\widehat{\lambda})\right|$ across jrt observations. Figure 2 illustrates this procedure by plotting the behavior of plan-level average costs obtained assuming optimal rating for different values of λ against the average willingness-to-pay for actuarial value among the plan's enrollees. In both panels the blue line corresponds to

Figure 2: Identification of relationship between preferences and cost



(a) Fit of AC^{obs} setting $\lambda = 0$

(b) Fit of AC^{obs} at estimated λ

Notes: This figure displays the relationship between the willingness-to-pay for plan actuarial value and enrollee average costs for different values of the selection parameter λ. In both figures the blue line approximates this relationship using observed average costs and the average willingness for actuarial value among plan's enrollees obtained from our demand model. The red line is the same relationship but replacing observed average cost with the predicted $\widehat{AC}_{jrt}(\widehat{\zeta}, \widehat{\lambda})$ obtained assuming optimal pricing. In both panels $\zeta = 0.0379$ matching the age-cost relationship in the Medical Expenditure Panel Survey. Panel (a) sets the selection parameter $\lambda = 0$ hence assuming no selection on unobservable preferences conditional on age. Panel (b) sets the selection parameter to the estimated value $\lambda = 14.4$.

the relationship between plan-level average costs as observed in our data and the composition of enrollment in terms of willingness-to-pay for actuarial value. The figure shows the significance of allowing for heterogeneity in costs across individuals of the same age. Setting $\lambda = 0$ model-predicted and observed average cost would track each other very poorly, while this is not the case for $\lambda > 0$, indicating adverse selection.

4.3 Adjustments to Plan Characteristics

Our analysis maintains the assumption that what we observe under the status-quo (managed competition) regime is an equilibrium in rates *and* characteristics. As discussed in Section 2, when considering procurement auctions as a counterfactual market mechanism, it is natural to consider whether and how insurers would change the characteristics of their plans, both upon bidding and after winning the auction. While demand estimates—paired with rates and rating rules—can be used to measure the (possibly negative) benefits of changing characteristics such as network and brand quality in terms of plan revenues, we must enrich the cost model to capture how these changes would impact cost.

We follow an approach that is similar to the one in Wollmann (2018); Barahona, Otero

and Otero (2023); Cuesta, Noton and Vatter (2024).[17] First, we assume that

$$K_{jrt}(\mathcal{N}_{f_{jrt}}, \xi_{f_{jrt}}) = \overline{\kappa}_{jrt} + (\eta^n + \epsilon^n_{jrt}) \underbrace{(\mathcal{N}_{f_{jrt}} - \mathcal{N}^\star_{f_{jrt}})}_{\substack{\text{Change from} \\ \text{observed network}}} + (\eta^\xi + \epsilon^\xi_{jrt}) \underbrace{(\xi_{f_{jrt}} - \xi^\star_{f_{jrt}})}_{\substack{\text{Change from} \\ \text{observed brand quality}}}, \tag{16}$$

where $(\epsilon^n_{jrt}, \epsilon^\xi_{jrt})$ are *iid* mean zero error terms. These capture unmodeled determinants of network and brand adjustment costs. This specification has the desirable property that, at the observed networks $\mathcal{N}^\star$ and (estimated) brand quality terms $\xi^\star$, $K_{jrt}(\mathcal{N}^\star_{f_{jrt}}, \xi^\star_{f_{jrt}}) = \overline{\kappa}_{jrt}$, which is identified as discussed above.

To recover the cost adjustment parameters η^n and η^ξ we once again impose the equilibrium assumption (1) and take the necessary first-order conditions with respect to changes in primary care physician networks, hospital networks, and brand quality. Letting insurers adjust these characteristics at the region-year-network level, this leads to

$$\eta^n + \sum_{j \in J_f} \epsilon_{jrt} = \left(\sum_{j \in J_f} \frac{\partial R_{jrt}}{\partial \mathcal{N}_{f_{jrt}}} - \frac{\partial Q_{jrt}}{\partial \mathcal{N}_{f_{jrt}}} AC_{jrt} - Q_{jrt} \frac{\partial H^{(\widehat{\zeta},\widehat{\lambda})}_{jrt}}{\partial \mathcal{N}_{f_{jrt}}} e^{\overline{\kappa}_{jrt}} \right) \left(\sum_{j \in J_f} Q_{jrt} H^{(\widehat{\zeta},\widehat{\lambda})}_{jrt} e^{\overline{\kappa}_{jrt}} \right)^{-1}, \tag{17}$$

from which we can obtain an estimate of η^n by taking expectations of the right-hand side across insurers, regions, and years since $\mathbb{E}[\sum_{j \in J_f} \epsilon_{jrt}] = 0$ conditional on other observables. The same approach can be repeated for all network characteristics collected in $\mathcal{N}_{frt}$ as well as the brand quality ξ_{frt}.

4.4 Model Estimates

Extensive margin demand responses. Table 3-A displays extensive margin demand responses to changes in plan attributes. We report extensive margin sensitivity as the change in the probability of picking a plan in the marketplace in response to when all characteristics for the plans in the marketplace vary. There are a few patterns in our results that are worth discussing. First, not surprisingly, individuals have a preference for lower premiums. Increasing all plans' premiums by \$10/month leads to a decrease in marketplace enrollment of 1.22 percent, with younger individuals displaying stronger responses. Second, there are strong preferences for provider networks. Individuals prefer plans with higher coverage of physicians and with closer in-network hospitals. Increasing

[17] As discussed in Section 5.5, while we acknowledge that a fully-fledged model of endogenous cost determination microfounding how hospital and physician prices are determined (see e.g. Gowrisankaran, Nevo and Town, 2015; Ho and Lee, 2017, 2019) would be desirable, it goes beyond the scope of our analysis.

Table 3: Extensive margin demand sensitivity and willingness-to-pay for plan attributes

	(1)	(2)	(3)	(4)	(5)
		Age below 45		Age above 45	
	All	Low income	High income	Low income	High income
A. Δ% in market enrollment					
Δ^+\$10/month in premium	-1.22	-1.37	-1.02	-1.25	-1.20
Δ^+0.05 in PCP share in 0-10 miles	0.50	0.59	0.72	0.32	0.51
Δ^-2 miles in distance	2.17	2.51	2.66	1.81	1.99
B. Willingness-to-pay (\$/year)					
Δ^+5 p.p in actuarial value	614.93	427.05	1,113.83	413.67	522.4
Δ^+0.05 in PCP share in 0-10 miles	53.50	36.65	100.68	33.73	44.36
Δ^-2 miles in distance	489.19	418.26	1,066.69	198.25	253.10

Notes: Panel A displays percentage changes in marketplace enrollment in response to changes in all plans' attributes. Panel B displays the willingness-to-pay for a change in a plan attribute. Columns in the top row indicate subpopulation of potential buyers according to age around 45 years, and income around 250 percent of FPL. Each row indicates a change to plan attributes.

the share of PCPs covered within 0-10 miles of an individual's zipcode by five percentage points, for all plans, leads to an increase in marketplace uptake of 0.5 percent. Decreasing the distance to the closest in-network hospital by two miles—a more abstract exercise but meaningful to measure preferences for hospital distance—would increase enrollment by 2.17 percent. Overall, individuals with higher income display stronger demand responses to plans' provider networks and weaker responses to premiums.

Willingness-to-pay for plan attributes. At the estimated demand parameters we obtain the distribution of willingness-to-pay for plan attribute k, in the notation above measured by β_i^k/α_i. Table 3-B summarizes the average willingness-to-pay for actuarial value and provider networks across all potential buyers and for particular subpopulations. There are several patterns worth noting in these results. First, consumers in this market feature a high willingness-to-pay for plan generosity, as the average consumer is willing to pay \$615/year for an increase in actuarial value of five percentage points, which is around a third of the average annual paid premium. Willingness-to-pay for plan generosity is stronger among higher-income consumers. Second, consumers display a substantial willingness-to-pay for provider networks, as the average consumer is willing to pay \$54/year to increase the share of PCPs within 10 miles that a plan covers by five percentage points and \$489/year to have the closest in-network hospital 2 miles closer to its zipcode, or almost a fourth of the average annual paid premium. The willingness-to-pay for provider networks is higher for relatively younger and higher-income consumers. Taken together, these patterns suggest that changing plan characteristics (both horizontal and vertical) in

Table 4: Substitution patterns in response to changes in plan attributes

		Diversion of plan enrollment upon change in a plan's attributes (%)											
		Anthem			Blue Shield			Kaiser			Other		
Insurer	Tier	Br.	Sil.	G./Pl.	Br.	Sil.	G./Pl.	Br.	Sil.	G./Pl.	Br.	Sil.	G./Pl.
		(A) Δ^+ \$10/month in premium											
Anthem	Bronze	-23.93	5.28	1.03	5.17	3.09	0.46	2.90	0.05	0.11	1.91	1.26	0.80
	Silver	10.61	-30.52	1.07	2.52	14.94	1.21	2.81	2.74	1.22	7.77	10.54	3.45
Blue Shield	Bronze	8.65	0.13	0.00	-57.92	40.80	0.61	2.03	0.05	0.00	2.30	1.06	0.61
	Silver	2.61	5.64	0.03	13.53	-34.80	3.52	0.37	1.14	0.11	0.21	4.86	1.37
Kaiser	Bronze	4.72	0.29	0.00	1.47	0.65	0.03	-29.33	5.46	3.86	5.95	2.88	1.88
	Silver	0.51	3.27	0.00	0.10	4.10	0.00	9.35	-27.81	1.00	0.23	7.71	0.01
		(B) Δ^+ 0.05 in PCP share in 0-10 miles											
Anthem	Bronze	26.00	-7.04	-3.21	-4.97	-5.58	-1.10	0.04	0.03	-0.03	-0.36	-0.62	-0.79
	Silver	-35.56	60.14	-8.74	1.89	-8.90	-0.61	2.66	1.48	1.11	7.52	8.09	2.95
Blue Shield	Bronze	-19.09	-0.65	-0.27	105.43	-58.30	-13.57	-2.11	-0.09	-0.13	-2.31	-1.98	-2.21
	Silver	-6.89	-5.40	-0.52	-20.75	53.55	-10.63	-0.47	-0.29	-0.18	-0.32	-2.97	-2.03
Kaiser	Bronze	0.12	0.24	0.00	-1.28	-0.55	-0.10	-16.68	7.16	3.02	3.99	1.66	1.12
	Silver	0.00	1.26	-0.01	-0.14	-1.52	-0.08	4.30	-14.68	3.26	0.04	6.59	-0.12
		(C) Δ^- 2 miles in distance to closest hospital											
Anthem	Bronze	33.08	-3.22	-2.97	-4.12	-7.82	-2.05	-1.04	-0.26	-0.32	-1.95	-2.01	-2.85
	Silver	-39.14	104.94	-9.63	0.06	-16.51	-3.70	0.97	-0.46	0.21	5.42	-2.11	-3.11
Blue Shield	Bronze	-29.4	-3.32	-1.29	142.14	-54.04	-18.13	-3.94	-0.39	-0.70	-6.24	-5.36	-9.32
	Silver	-10.13	-6.07	-1.27	-14.24	64.15	-10.48	-1.22	-0.62	-0.52	-1.66	-6.30	-5.38
Kaiser	Bronze	-1.60	-0.92	-0.09	-1.38	-1.60	-0.28	2.53	8.26	0.56	-1.65	-0.99	-1.51
	Silver	-1.98	-2.64	-0.35	-0.71	-4.99	-0.62	-6.55	31.54	-4.66	-1.38	-1.90	-3.17

Notes: This table displays demand responses to changes in plan attributes. The left panel displays the plan that is changing its attributes. Columns 1–12 display demand responses for 12 (groups of) plans as indicated in the top rows, where B, S, and O indicate Bronze, Silver, and Other plans, respectively. Panels A, B, and C display percentage changes in plan enrollment in response to changes in premiums, covered PCP share, and distance to in-network hospital, as indicated. Each cell reports $\Delta q_k/q_j$, where k is the plan in the top row and j is the plan in the left panel, which is the one changing its attributes.

the marketplace induces substantial demand responses which are the main outcome of our interest in our analysis of counterfactual market mechanism.

Substitution patterns. Estimates of substitution patterns across plans upon changes in characteristics are a key input to measure how equilibrium outcomes vary across alternative market mechanisms. Table 4 displays a summary of the substitution patterns implied by our demand model, and a few of these are worth highlighting.

First, demand for a specific plan is highly responsive: changing premiums or provider networks leads to large changes in own plan demand across plans. For example, if any of the three main insurers in the market (Anthem, Blue Shield, and Kaiser in Table 4) increased its monthly rate by \$10 for any of their Bronze or Silver plans, we estimate that total demand for the plan would fall by around a third. Second, while the same insurer would

Table 5: Cost estimates

Change	$\Delta \log c(\theta_i, x_j)$
Δ^+\$100 WTP for 5% AV	0.72
Δ^+10 years of age	0.38
Δ^+0.1 pp share PCP in 0-10 miles	0.022
Δ^-0.5 miles to nearest hospital	0.081
Δ^+1% brand quality ξ	0.039

Notes: Changes in log cost as a function of individual characteristics or in response to changes in plan attributes and individual characteristics, as indicated in the left panel.

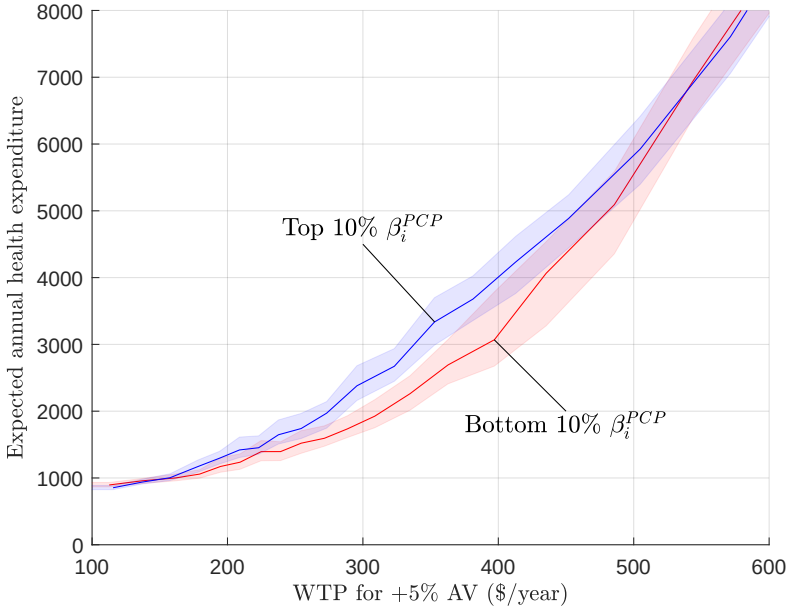
recapture some of that demand in their other metal tiers—particularly in Bronze plans when increasing premiums for Silver plans—many of their enrollees would substitute towards competing insurers. When substituting across insurers, consumers mostly choose plans in the same tier as the one changing its attributes. These rich substitution patterns are an important feature of the (pure characteristics) demand model with observed and unobserved heterogeneity, which can be estimated by leveraging the granularity of the data provided by the marketplace.

Cost estimates. The key cost parameters are summarized in Table 5. We estimate that 10 years of age increases cost by approximately 38%, while an additional \$100 in willingness to pay for insurance generosity as measured by actuarial value increases cost by almost two-thirds.

Imposing the additional assumption that observed network breadth and brand quality result from Nash equilibrium behavior under managed competition we find that increasing these characteristics would increase insurer cost. Quantitatively, our estimates of η^n and η^ξ indicate that insurers' costs would increase by slightly more than 2% if network breadth for primary-care physicians were to increase by 1 percentage point. Adding hospitals, which in our simple model of amounts to reducing hospital distance, is costlier. We measure that reducing hospital distance by 0.5 miles would increase insurers' costs by 8%. Lastly, increasing brand quality as captured by ξ_f in our demand model by 1% increases cost by almost 4%.

A key quantity of interest in our analysis is the relationship between preference for insurance generosity, preferences for provider networks, and expected costs. Figure 3 illustrated this relationship by combining our demand and cost estimates. We find a significant degree of adverse selection on willingness-to-pay for actuarial value (β_i^{av}/α_i above). Individuals for which this is lower than \$300 (for a 5% increase in AV) have

Figure 3: Relationship between cost and willingness-to-pay



Notes: This figure displays the relationship between the willingness-to-pay for plan actuarial value on the x-axis and enrollee average costs on the y-axis, separately for individuals at the top (blue) and bottom (red) deciles of the distribution of preferences over PCP coverage.

expected annual health expenditure lower than $2,000, while these grow to $4,000 on average among those with β_i^{av}/α_i around $400, and higher than $6,000 if willingness-to-pay is higher than $500. Moreover, as in Shepard (2022); Serna (2023, 2024), we find that individuals who have stronger preferences for provider network breadth (corresponding to larger values of β_i^n) have higher cost. Quantitatively, our findings imply that for average values of other preference parameters buyers with with network preferences in the top ten percent of the population expect medical spending that can be $1,000 (or 25% of the average plan-level claims in the data) greater than those in the bottom ten percent.

5 Managed Competition vs. Public Procurement

5.1 Counterfactual Setup

We are now ready to use our empirical model to compare market outcomes and welfare under managed competition and second-score scoring auctions. We proceed by keeping the body of regulation from our setting roughly in place and then compare equilibrium outcomes under the following environments:

Managed competition under imperfect competition: Insurers offer differentiated plans and

compete by setting premiums in a Nash-Bertrand game (see e.g., Jaffe and Shepard 2020 and Tebaldi 2024, among others). State and federal regulations impose a division of the market in rating regions, rating rules across age and subsidies as a function of income and age, and standardize coverage generosity across metal tiers.

Procurement mechanisms: The regulator announces a scoring rule Ψ and a budget, which we set at the region-year level as equal to the observed public spending under the ACA in the same market. Insurers bid knowing their own demand and cost curves by submitting rates and plan characteristics roughly as described in Section 2. The winner is the insurer with the highest score and offers the four metal tiers setting premiums and characteristics freely as long as the corresponding score remains above the second-highest. Consumers then choose whether or not to enroll in one of the winner's plans. Importantly, to facilitate the comparison between mechanisms, as we do for the government budget we also keep rating regions, rating rules, subsidies, and standardization of metal tiers fixed and identical to the ACA design.

A key input for this comparison is the scoring rule Ψ. At the moment we consider two simple scoring rules that capture the essential trade-offs between prioritizing low premiums versus adjusting scores for plans' quality. While in future analyses we plan to discuss more specific (and realistic) implementations, we now consider

$$\begin{aligned} \textit{Min bid:} \quad & \Psi(b_j, x_j) = -b_j, \\ \textit{Max } Q\textit{:} \quad & \Psi(b_j, x_j) = Q^a(p_j, x_j)\,; \end{aligned}$$

Min bid is a scoring rule that ranks insurers rewarding low rates (subject to spending the budget in subsidies), and *Max Q* is a scoring rule that ranks insurers by the total marketplace enrollment they would generate as the sole supplier of insurance at the rate submitted in their bid. This second score captures the possibility that the regulator uses information on how individuals may prefer one insurer over another when evaluating competing bids

Simulation details. Since we inverted equilibrium first-order conditions to estimate cost parameters, we observe outcomes under our baseline scenario of managed competition between multi-plan insurers. Our demand estimates allow us to augment observed enrollment and average costs with measures of consumer surplus.

To analyze scoring auctions we rely on the fact that dominant strategies are easy to calculate. We consider both auctions in which we hold networks and brand characteristics fixed, and auctions in which we let insurers adjust these upon bidding and after winning

the auction. Given a score and a budget in a region-year, for every insurer observed in the market we find the bid (and, in the more general case, network and brand quality changes) that maximize the Ψ conditional on breaking even and using the sponsor's budget. We then determine the winner, breaking ties randomly. The final outcomes of the auction (enrollment, cost, consumer surplus, and profits) result from solving a constrained monopolist problem in which the insurer can freely set rates (and network and brand quality) subject to the corresponding score remaining above the second-highest.

We repeat these steps separately for each of the 76 region-year markets we observe in our data, for both scores, and either holding characteristics fixed and letting them vary within a certain range (+/- 20%) around their status quo level.

5.2 Implications for Market Outcomes and Welfare

A summary of the main results from our counterfactual analysis is displayed in Table 6, which shows plan offerings, costs, premiums, enrollment, and welfare outcomes across equilibria in the different market mechanisms we consider. Column (1) describes the baseline case of managed competition observed in the data: the average market has around five insurers and 19 plans, and—on average across markets *without* weighting by population—37.6 percent of potential buyers enroll in a marketplace plan. The average cost (again, averaged across markets) is slightly lower than $4,200, and the markup per person per year is $516. Across other columns, we find that the scoring auctions we consider would have a significant impact on market outcomes.

Auctions holding fixed plan characteristics. For illustration, we focus first on the setting with fixed networks. Regardless of the scoring rule, a procurement mechanism leads to a monopoly insurer serving the market, hence bringing the supply of plans down to between four and five, depending on the insurer each mechanism selects (some offer both an HMO and a PPO plan within a metal tier, we leave these within-insurer plan offerings fixed throughout). In terms of market outcomes, a scoring rule that only rewards low rates like *Min bid* in our simulations would select a low-cost insurer and thus lead to average costs 5 percent lower than under managed competition. Moreover, average paid premiums would fall to almost half the level of managed competition—note that this fall combines lower premiums charged due to stronger competition with changes in selection towards younger and healthier enrollees. Nevertheless, a *Min bid* scoring rule would most likely select a low-quality insurer. As a result, despite the lower premiums we find that enrollment in the marketplace would fall by almost two percentage points (5.3%) relative to managed competition. This result suggests an important role for quality, preference heterogeneity, and variety.

Table 6: Managed competition vs. second-score scoring auctions in Covered California

	(1)	(2)	(3)	(4)	(5)
		Δ relative to managed competition			
		Fixed networks		Endogenous networks	
Outcome	Managed competition	*Min bid*	*Max Q*	*Min bid*	*Max Q*
Number of plans in the menu	18.8	-14.8	-14.3	-14.5	-14.2
Average cost ($/year)	4,196.5	-221.9	-18.0	-1,028.7	-24.9
Average post-subsidy premium ($/year)	1,713.8	-877.7	-483.1	-1,599.7	-366.7
Network and quality adjustment (%)	–	–	–	-20.0	3.0
Enrollment share (%)	37.62	-1.98	1.09	-0.77	1.57
CS per potential buyer ($/year)	–	-40.8	125.3	1.8	251.1
Profit per buyer ($/year)	515.9	-394.2	-320.3	-503.5	-204.5

Notes: This table shows equilibrium market and welfare outcomes across market designs. Column (1) shows results for managed competition under imperfect competition. Columns (2) and (3) show results for procurement mechanisms under fixed networks. Columns (4) and (5) show results for procurement mechanisms under endogenous networks. All results are simple averages across region-years.

Indeed, when we move on to considering a scoring rule that places more weight on quality, like *Max Q* in our current analysis, we find that even though premiums decrease by about half as much as under *Min bid*, enrollment in the marketplace increases. More precisely, a procurement mechanism with a *Max Q* scoring rule would lead to average costs only slightly lower than under managed competition, average paid premiums 28 percent lower, and a 1.1 percentage point increase in the enrollment share in the marketplace. These results suggest that—even holding fixed the characteristics of plans we observe in under the status quo managed competition regime—the losses in variety associated with procurement mechanisms may be more than compensated by lower premiums as long as enough quality is provided.

In terms of consumer welfare, the impacts of procurement mechanisms range from losses of $394 per-person per-year relative to managed competition for the *Min bid* scoring rule to gains of $125/year for *Max Q*, consistently with the enrollment patterns. On the supply side, the decreases in average costs that we find for both scoring rules relative to managed competition are coupled with decreases in premiums, which is large enough that average profits decrease under both scoring rules.

Auctions with endogenous plan characteristics. Accounting for potential responses in quality by insurers has meaningful effects on the equilibrium resulting under scoring auctions. By focusing on price, a *Min bid* procurement mechanism would place incen-

tives for insurers to lower their quality, so as to lower their cost and be able to offer a lower premium. Our results indeed show that average costs and average paid premiums would fall strongly in this scenario relative to managed competition. The decrease in premiums is large enough to lower the adverse impact on enrollment relative to managed competition to 0.77 percentage point—less than half the adverse impact of *Min bid* with fixed networks—and make the change in market design essentially welfare-neutral for consumers. This is essentially a case of a race to the bottom among insurers.

In contrast, a *Max Q* procurement mechanism would place incentives to possibly increase quality provision. Indeed, we find that the winning insurer would increase its quality, which implies that average costs and average paid premiums fall less relative to managed competition than under fixed networks. However, consumers value networks enough that this higher quality provision leads to a stronger positive impact of a shift to procurement on overall enrollment of 1.6 percentage points and, similarly, a stronger impact on average consumer surplus at $251/year.

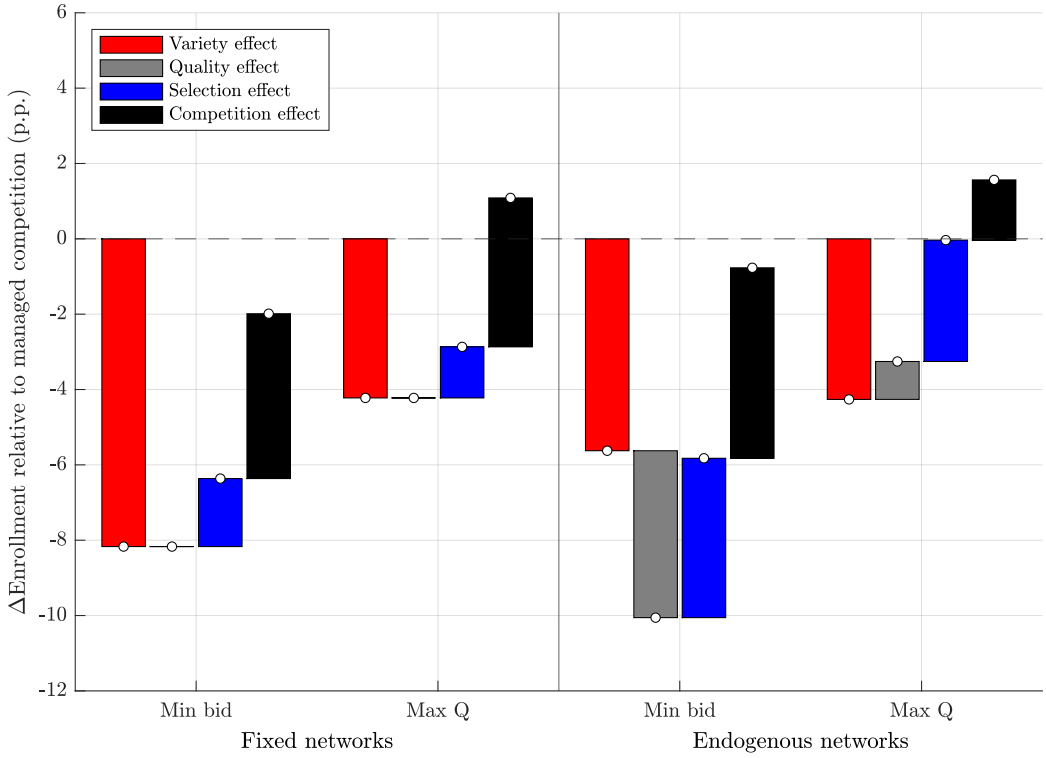
5.3 Mechanisms: Variety, Quality, Selection and Competition Effects

Our framework in Section 2 identifies four economic forces determining the overall impacts of moving from managed competition to procurement mechanisms for health insurance provision. To disentangle such economic forces, we consider a decomposition for the roles of changing variety, changing quality, changing selection, and changing competition for each scoring rule relative to a baseline of managed competition. This decomposition builds on equations (4) through (7) in Section 2.

The main results from this decomposition for market enrollment are displayed in Figure 4.[18] The *variety effect* is always negative since consumers lose options from their choice sets and are left with a single plan when moving from managed competition to a procurement mechanism. However, the magnitude of the variety effect varies substantially across scoring rules. In particular, a scoring rule that places more emphasis on premiums like *Min bid* has a variety effect leading to stronger decreases in enrollment than a scoring rule more likely to select a higher-quality plan, like *Max Q*. Moreover, the magnitude of the variety effect is sizable. Absent network adjustments, moving towards a monopoly plan decreases overall enrollment by eight and four percentage points for *Min bid* and *Max Q*, respectively, consistent with the vast preference heterogeneity documented earlier in the paper. The *quality effect* is only present in contexts in which insurers may adjust their quality under the procurement mechanisms. Our results show that this effect is

[18]Figure A.1 shows the same decomposition for consumer surplus, which displays similar patterns than the results for enrollment.

Figure 4: Mechanisms contributing to changes in equilibrium outcomes



Notes: This figure shows a decomposition of the impacts of procurement mechanisms relative to managed competition under imperfect competition on marketplace enrollment in percentage points. The decomposition has four components: variety effect, quality effect, selection effect, and markup effects. The effects are reported additively, such that for each bar, the white dot indicates the total effect up to that point in the decomposition, and the difference between subsequent white dots indicates the marginal effect of each component of the decomposition. All results are simple averages across region-years.

quantitatively meaningful and changes signs across scoring rules. As discussed in the previous section, *Min bid* and *Max Q* scoring rules place opposite incentives for quality provision. Indeed, we find a negative quality effect of 4.1 percentage points of enrollment for the former and a positive one for the latter, albeit weaker.

In principle, the *selection effect* could also vary across scoring rules. However, our findings suggest that it is always positive for the cases we study, yet it varies substantially. Finally, we find that the *competition effect* is always positive, reflecting that procurement mechanisms enhance competition across insurers and induce them to offer lower premiums. This effect is particularly stronger for the *Min bid* scoring rule because it emphasize price more when selecting an insurer. In terms of magnitude, the competition effect leads to increases in enrollment ranging from 1.8 to 5.5. percentage points.

Taken together, these results suggest there is a natural trade-off when designing procurement mechanisms for health insurance. Scoring rules with more focus on price will

harm consumers more in terms of variety, but also benefit them more in terms of leading to lower premiums. Changes in quality and selection are harder to sign. In our setting, this trade-off is such that managed competition would mostly dominate procurement mechanisms with a *Min bid* scoring rule in the aggregate, but the opposite holds for procurement mechanisms with a *Max Q* scoring rule. These results suggest that the four mechanisms we identify are potentially quantitatively important in driving the impacts of the market designs we plan to study.

5.4 Heterogeneity across Markets

There is substantial heterogeneity across markets in the appeal of a procurement approach to health insurance relative to managed competition. Indeed, our results show that, for the simple context of fixed networks, a *Min bid* scoring rule delivers better outcomes than managed competition in only 52 and 56 percent of the markets in terms of enrollment and consumer surplus, respectively. Those shares are much larger for *Max Q*, at 67 percent and 68 percent of the markets in terms of enrollment and consumer surplus, respectively. Moreover, while for 42 and 44 percent of the markets we find that *both* scoring rules dominate managed competition in terms of enrollment and consumer surplus, for several regions, only one of them does so. Concretely, in around a fourth of the regions, *Max Q* dominates managed competition, while managed competition dominates *Min bid*. The opposite is true in roughly ten percent of regions. These patterns are consistent with the extent of preference heterogeneity discussed earlier in the paper, highlighting that the relative appeal of these mechanisms depends on market specifics.

A natural dimension to consider when trying to explain heterogeneity across markets is baseline market structure. The government defines markets using rating regions that differ in market structure, preferences, and potential buyers. These variables, in turn, affect premium-setting incentives both in the managed competition baseline environment and under the procurement mechanism counterfactual environment. To inform the extent of heterogeneity in how procurement mechanisms would impact different markets, we implement our counterfactual analysis separately for regions with 4 or fewer and 5 or more insurers under the managed competition baseline. Note that we do not interpret this exercise as isolating the impact of competition since these regions may differ in other ways than the number of firms in the market.

The relative performance of managed competition and procurement mechanisms differs substantially across markets with different baseline market structure. Table 7 shows the results. Indeed, the impact of a shift from managed competition to a procurement mechanism is mostly higher in markets with fewer insurers at baseline. Our results show

Table 7: Heterogeneity across markets in impacts of procurement market designs

		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
		Market		Consumer		Decomposition of enrollment impacts (p.p.)			
Scoring	Market	enrollment (p.p.)		surplus ($/year)		Variety	Quality	Selection	Competition
rule	structure	Mean	% ≥ 0	Mean	% ≥ 0	effect	effect	effect	effect
A. Fixed networks									
Min bid	4- insurers	-2.37	52.38	-30.33	61.90	-8.63	0.00	1.05	5.21
	5+ insurers	-1.52	53.94	-53.84	50.00	-7.60	0.00	2.74	3.35
Max Q	4- insurers	1.41	71.43	181.4	66.67	-3.78	0.00	0.62	4.57
	5+ insurers	0.69	61.76	56.04	70.59	-4.77	0.00	2.27	3.19
B. Endogenous networks									
Min bid	4- insurers	0.07	61.90	113.90	80.95	-4.84	-4.47	2.84	6.53
	5+ insurers	-1.80	35.29	-136.73	47.06	-6.60	-4.39	5.94	3.24
Max Q	4- insurers	1.74	85.71	278.28	90.48	-3.83	0.66	3.76	1.15
	5+ insurers	1.35	82.35	217.49	91.18	-4.79	1.43	2.55	2.16

Notes: This table shows equilibrium market and welfare outcomes across market designs. Column (1) shows results for managed competition under imperfect competition. Columns (2) and (3) show results for procurement mechanisms under fixed networks. Columns (4) and (5) show results for procurement mechanisms under endogenous networks. All results are simple averages across region-years for each scoring rule and category of market structure as noted in the left panel.

that the average impact of such a policy on consumer surplus would be higher in markets with fewer insurers at baseline, and that holds across scoring rules and whether insurers adjust their quality provision.

To understand what drives these patterns, we compute the decomposition from the previous section separately for regions with 4 or fewer and 5 or more insurers under the managed competition baseline. Table 7 shows the results. We find that the *variety effect* tends to be more negative in markets with more insurers at baseline, essentially capturing the fact that these markets naturally lose more variety from a shift to a procurement mechanism that selects a monopoly insurer. In addition, we find that the *competition effect* is mostly higher in markets with fewer insurers at baseline, likely driven by larger markups at baseline in markets in which insurers face fewer competitors and hence is further from competitive outcomes.

5.5 Discussion

Our results shed light on the potential impacts of moving away from managed competition towards a procurement-like approach using scoring auctions in the context of government-sponsored individual health insurance. In this section, we have implemented an empirical comparison between these mechanisms. Our analysis has some limitations

worth discussing. While we allow for plan premiums and quality to adjust in our counterfactual analysis, there are other elements in the environment that remain unmodeled. As we discuss below, many of these concerns are likely to be quantitatively minor in our setting since the ACA marketplace is a relatively small and fringe segment of the market—particularly for the large insurers accounting for most of the enrollment. However, these issues may be relevant in other settings or if a broader use of procurement mechanisms was to be considered. Extensions of this project will aim at providing robustness exercises to stress test our results against these concerns.

First, we focus on procurement auctions where the set of bidders is comprised of the insurers that were active in the local market under managed competition. By not modeling entry, we may be underestimating the competitive effects of procurement mechanisms if bidding costs are low and there are other potential entrants, hence the extent to which it may lower premiums relative to managed competition. The opposite would hold if bidding costs were high enough to deter entry, possibly leading to overestimating the competitive effects of procurement mechanisms. In our context, given insurers already have provider networks in local markets set up (either for this or other insurance segments), the former scenario is more likely. Furthermore, this marketplace may provide opportunities for small innovative insurers to the market and expand.

Second, while we model endogenous plan quality choices, we do so in a stylized and parsimonious way. A natural concern may be that providers—being aware of their relevance for insurers to make the best bid in a procurement mechanism—could attempt to extract more surplus from insurers, leading to premium increases (Ho and Lee, 2019). This is a more likely scenario in concentrated provider markets and in settings where insurers must offer complete networks and cannot exclude providers. On the other hand, by ensuring that—upon winning—the insurer would be the sole supplier in the marketplace this could shift bargaining power in favor of lower costs. Again, given the ACA marketplace is a small share of most insurers' operations and that these firms already have networks in place for their other business segments (i.e., Medicare Advantage and employer-sponsored coverage), changes in the design of the ACA marketplace are unlikely to impact those.

Third, having a single insurer in place may lead to economies of scale that could potentially translate to lower premiums.[19] These forces would make our results a lower bound for the potential gains from a procurement mechanism, although by the arguments

[19]The presence of scale effects and their impacts on prices have already been documented in other procurement settings (Allende *et al.*, 2024). Moreover, there is also some evidence of scale economies in insurance Serna (2023).

above regarding this being a limited segment of the market one may not expect this to be quantitatively meaningful.

There are two dynamic concerns regarding procurement mechanisms. On the demand side, consumers may be forced to switch insurers upon changes in the sole provider in their local market, incurring switching costs that could lead to welfare losses. Two natural ways to deal with this concern in practice would be to either auction the market only for new enrollees in every period or for several years at once. On the supply side, by being excluded from the marketplaces, insurers may decide to invest less in their provider networks or to leave the market altogether, weakening competition in subsequent procurement auctions. While this is a possibility, by the argument that ACA marketplaces are a relatively small segment of the market, we would not expect this to be an issue. Alternatively, policymakers could explore multi-award auctions as a way to keep insurers in the market and protect future competition.

Finally, throughout our analysis, we focused on the market definitions of Covered California. Market definition is a key choice variable for a government and likely a relevant one when designing a procurement mechanism. For example, a government could define markets by age groups in light of preference heterogeneity along that dimension. In future work, we plan to study alternative market definitions to better understand how to tailor procurement mechanisms to heterogeneity in preferences, market size, and baseline market structure.

6 Conclusion

We consider and empirically evaluate procurement mechanisms as alternatives to managed competition for government-sponsored health insurance markets. Rather than promoting oligopoly competition between insurers or directly providing a public option, the government could set up a procurement mechanism where the winner becomes the sole provider for a given market. On the one hand, consumers would face limited choices; on the other, inefficiencies due to risk selection and imperfect competition would be reduced, and quality may be less distorted.

We use data from the Covered California ACA marketplace to measure these trade-offs and the overall impacts of adopting a procurement mechanism relative to the managed competition baseline. In this context, we observe individual-level enrollment choices, plan attributes, including the universe of provider networks, as well as plan average costs. These data enable us to measure the distribution of preferences over plans as well as a rich insurer cost function. With these primitives, we simulate the impacts of procurement mechanisms on both market outcomes and welfare, isolate the main economic forces

driving these impacts, and discuss how the design of procurement auctions determines these impacts.

Our results suggest that the potential gains from increased competition and reduced risk selection could outweigh the costs of limiting choice on average across local markets in our setting. However, the design of the scoring rule in procurement auctions is pivotal in achieving those impacts.

References

Abaluck, J. and Gruber, J. (2023). When Less is More: Improving Choices in Health Insurance Markets. *The Review of Economic Studies*, **90** (3), 1011–1040.

Ackerberg, D. A. and Rysman, M. (2005). Unobserved Product Differentiation in Discrete-Choice Models: Estimating Price Elasticities and Welfare Effects. *The RAND Journal of Economics*, **36** (4), 771–788.

Akerlof, G. A. (1970). The Market for "Lemons": Quality Uncertainty and the Market Mechanism. *The Quarterly Journal of Economics*, **84** (3), 488–500.

Allende, C., Atal, J. P., Carril, R., Cuesta, J. I. and Gonzalez-Lira, A. (2024). Centralizing Procurement: The Roles of Scale, Selection, and Variety, manuscript.

Asker, J. and Cantillon, E. (2008). Properties of Scoring Auctions. *The RAND Journal of Economics*, **39** (1), 69–85.

Baicker, K., Chandra, A. and Shepard, M. (2023). Achieving Universal Health Insurance Coverage in the United States: Addressing Market Failures or Providing a Social Floor? *Journal of Economic Perspectives*, **37** (2), 99–122.

Barahona, N., Otero, C. and Otero, S. (2023). Equilibrium effects of food labeling policies. *Econometrica*, **91** (3), 839–868.

Berry, S. (1994). Estimating Discrete-Choice Models of Product Differentiation. *RAND Journal of Economics*, **25** (2), 242–262.

—, Levinsohn, J. and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, **63** (4), 841–890.

— and Pakes, A. (2007). The Pure Characteristics Demand Model. *International Economic Review*, **48** (4), 1193–1225.

Berry, S. T. and Haile, P. A. (2024). Nonparametric identification of differentiated products demand using micro data. *Econometrica*, **92** (4), 1135–1162.

Bresnahan, T. F. (1987). Competition and collusion in the american automobile industry: The 1955 price war. *The Journal of Industrial Economics*, pp. 457–482.

Brown, Z. Y. and Jeon, J. (2024). Endogenous information and simplifying insurance choice. *Econometrica*, **92** (3), 881–911.

Bundorf, M. K., Levin, J. and Mahoney, N. (2012). Pricing and Welfare in Health Plan Choice. *American Economic Review*, **102** (7), 3214–48.

Capps, C., Dranove, D. and Satterthwaite, M. (2003). Competition and Market Power in Option Demand Markets. *The Rand Journal of Economics*, **34** (4), 737–763.

Chandra, A., Handel, B. and Schwartzstein, J. (2019). Chapter 6 - behavioral economics and health-care markets. In B. D. Bernheim, S. DellaVigna and D. Laibson (eds.), *Handbook of Behavioral Economics - Foundations and Applications 2*, vol. 2, North-Holland, pp. 459–502.

Chen, S., Geldsetzer, P., Chen, Q., Moshabela, M., Jiao, L., Ogbuoji, O., Sie, A., Atun, R. and Bärnighausen, T. (2022). Health Insurance Coverage in Low- and Middle-Income Countries Remains Far from the Goal of Universal Coverage. *Health Affairs*, **41** (8), 1142–1152.

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

Cuesta, J. I., Noton, C. and Vatter, B. (2024). *Vertical Integration and Plan Design in Healthcare Markets*. Working Paper 32833, National Bureau of Economic Research.

Curto, V., Einav, L., Levin, J. and Bhattacharya, J. (2021). Can Health Insurance Competition Work? Evidence from Medicare Advantage. *Journal of Political Economy*, **129** (2), 570–606.

Dafny, L., Ho, K. and Varela, M. (2013). Let Them Have Choice: Gains from Shifting Away from Employer-Sponsored Health Insurance and toward an Individual Exchange. *American Economic Journal: Economic Policy*, **5** (1), 32–58.

Decarolis, F., Polyakova, M. and Ryan, S. P. (2020). Subsidy Design in Privately Provided Social Insurance: Lessons from Medicare Part D. *Journal of Political Economy*, **128** (5), 1712–1752.

DeLeire, T., Chappel, A., Finegold, K. and Gee, E. (2017). Do Individuals Respond to Cost-sharing Subsidies in their Selections of Marketplace Health Insurance Plans? *Journal of health economics*, **56**, 71–86.

Diamond, P. (1992). Organizing the Health Insurance Market. *Econometrica*, **60** (6), 1233–1254.

Dickstein, M. J., Ho, K. and Mark, N. (2024). Market Segmentation and Competition in Health Insurance. *Journal of Political Economy*, **132** (1), 96–148.

Drake, C. (2019). What are Consumers Willing to Pay for a Broad Network Health Plan?: Evidence from Covered California. *Journal of Health Economics*, **65**, 63 – 77.

Einav, L., Finkelstein, A. and Cullen, M. R. (2010). Estimating welfare in insurance markets using variation in prices. *The quarterly journal of economics*, **125** (3), 877–921.

— and Levin, J. (2015). Managed Competition in Health Insurance. *Journal of the European Economic Association*, **13** (6), 998–1021.

Enthoven, A. C. (1993). The History and Principles of Managed Competition. *Health affairs*, **12** (suppl 1), 24–48.

Ericson, K. M. and Starc, A. (2015). Measuring Consumer Valuation of Limited Provider Networks. *The American Economic Review*, **105** (5), 115–119.

Finkelstein, A., Hendren, N. and Shepard, M. (2019). Subsidizing Health Insurance for Low-income Adults: Evidence from Massachusetts. *American Economic Review*, **109** (4), 1530–1567.

Forbes, M. and Dunbar, S. (2022). *Understanding Medicaid Managed Care Procurement Practices Across States*. Tech. rep., MACPAC.

Gowrisankaran, G., Nevo, A. and Town, R. J. (2015). Mergers When Prices Are Negotiated: Evidence from the Hospital Industry. *American Economic Review*, **105** (1), 172–203.

Gupta, A., Parente, S. T. and Sanyal, P. (2012). Competitive Bidding for Health Insurance Contracts: Lessons from the Online HMO Auctions. *International Journal of Health Care Finance and Economics*, **12**, 303–322.

Haeder, S. F., Weimer, D. and Mukamel, D. B. (2019). A consumer-centric approach to network adequacy: access to four specialties in california's marketplace. *Health Affairs*, **38** (11), 1918–1926.

Handel, B., Hendel, I. and Whinston, M. D. (2015a). Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk. *Econometrica*, **83** (4), 1261–1313.

—, — and — (2015b). Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk. *Econometrica*, **83** (4), 1261–1313.

— and Kolstad, J. (2022). The Affordable Care Act After a Decade: Industrial Organization of the Insurance Exchanges. *Annual Review of Economics*, **14** (1), 287–312.

Ho, K. (2006). The Welfare Effects of Restricted Hospital Choice in the US Medical Care Market. *Journal of Applied Econometrics*, **21** (7), 1039–1079.

— (2009). Insurer-Provider Networks in the Medical Care Market. *American Economic Review*, **99** (1), 393–430.

— and Lee, R. (2017). Insurer Competition in Health Care Markets. *Econometrica*, **85** (2), 379–417.

— and Lee, R. S. (2019). Equilibrium Provider Networks: Bargaining and Exclusion in Health Care Markets. *American Economic Review*, **109** (2), 473–522.

Jaffe, S. and Shepard, M. (2020). Price-Linked Subsidies and Imperfect Competition in Health Insurance. *American Economic Journal: Economic Policy*, **12** (3), 279–311.

Lavetti, K., DeLeire, T. and Ziebarth, N. R. (2023). How do Low-income Enrollees in the Affordable Care Act Marketplaces Respond to Cost-sharing? *Journal of Risk and Insurance*, **90** (1), 155–183.

Marone, V. R. and Sabety, A. (2022). When Should There Be Vertical Choice in Health Insurance Markets? *American Economic Review*, **112** (1), 304–42.

Miller, K. S., Petrin, A., Town, R. and Chernew, M. (2019). *Optimal Managed Competition Subsidies*. Working Paper 25616, National Bureau of Economic Research.

Pang, J.-S., Su, C.-L. and Lee, Y.-C. (2015). A Constructive Approach to Estimating Pure Characteristics Demand Models with Pricing. *Operations Research*, **63** (3), 639–659.

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

Pauly, M. V. (1968). The Economics of Moral Hazard: Comment. *The American Economic Review*, **58** (3), 531–537.

Prager, E. (2020). Healthcare Demand under Simple Prices: Evidence from Tiered Hospital Networks. *American Economic Journal: Applied Economics*, **12** (4), 196–223.

Quan, T. W. and Williams, K. R. (2018). Product variety, across-market demand heterogeneity, and the value of online retail. *The RAND Journal of Economics*, **49** (4), 877–913.

Rosse, J. N. (1970). Estimating cost function parameters without using cost data: Illustrated methodology. *Econometrica: Journal of the Econometric Society*, pp. 256–275.

Rothschild, M. and Stiglitz, J. (1976). Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information. *The Quarterly Journal of Economics*, **90** (4), 629–649.

Saltzman, E. (2021). Managing Adverse Selection: Underinsurance versus Underenrollment. *The RAND Journal of Economics*, **52** (2), 359–381.

—, Swanson, A. and Polsky, D. (2021). *Inertia, Market Power, and Adverse Selection in Health Insurance: Evidence from the ACA Exchanges*. Working Paper 29097, National Bureau of Economic Research.

Sances, M. W. and Clinton, J. D. (2019). Who Participated in the ACA? Gains in Insurance Coverage by Political Partisanship. *Journal of Health Politics, Policy, and Law*, **44** (3), 349–379.

Serna, N. (2023). Non-Price Competition and Risk Selection Through Hospital Networks.

— (2024). What Can Market Structure Tell Us About Equilibrium Contracts in Healthcare?

Shepard, M. (2022). Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange. *American Economic Review*, **112** (2), 578–615.

Song, M. (2007). Measuring Consumer Welfare in the CPU market: An Application of the Pure-characteristics Demand Model. *The RAND Journal of Economics*, **38** (2), 429–446.

Starc, A. (2014). Insurer Pricing and Consumer Welfare: Evidence from Medigap. *The RAND Journal of Economics*, **45** (1), 198–220.

Tebaldi, P. (2024). Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA. *The Review of Economic Studies*, **92** (1), 586–620.

—, Torgovitsky, A. and Yang, H. (2023). Nonparametric Estimates of Demand in the California Health Insurance Exchange. *Econometrica*, **91** (1), 107–146.

Vatter, B. (2024). Quality disclosure and regulation: Scoring design in medicare advantage. *Econometrica*, **forthcoming**.

Wollmann, T. G. (2018). Trucks without bailouts: Equilibrium product characteristics for commercial vehicles. *American Economic Review*, **108** (6), 1364–1406.

Online Appendix

Managed Competition vs. Public Procurement in Individual Health Insurance

José Ignacio Cuesta and Pietro Tebaldi

A Details on Data Construction

A.1 Measuring the Number of Potential Buyers

The enrollment data describe choices among plans in Covered California, but does not provide information on the outside option for this population. Hence, we develop a procedure to measure the set of potential buyers and hence outside share. This procedure is similar to those adopted in recent work by Tebaldi *et al.* (2023) and Dickstein *et al.* (2024).

We start by estimating the likelihood of a person being a potential buyer using data for California from the American Community Survey (ACS). We define potential buyers in Covered California as individuals that are either individually insured or uninsured. We then estimate the following linear regression:

$$\text{Potential buyer}_i = f(X_i) + \varepsilon_i$$

where $f(X_i)$ is a flexible function of individual characteristics. In particular, we first define 8 coarse age groups given by [20-26), [26-31), [31-36), [36-41), [41-46), [46-51), [51-56), [56-61) and [61-64]; and 7 coarse FPL groups given by [100-140), [140-150), [150-200), [200-250), [250-300), [300-350) and [350-400]. The function $f(X_i)$ then includes the interactions between indicators for these two sets of variables with continuous age and income as a share of FPL; interactions for indicators for regions and gender with continuous age and income as a share of FPL; interactions between indicators for regions and gender; and interactions between indicators for age and income bins and gender. This regression is weighted using ACS person weights. Let the fitted likelihood of individual i being a potential buyer be τ_i, which takes on the same value for any individual with the same age, FPL, sex, region and year.

The next step is to use the estimates $\hat{\tau}_i$ to obtain the number of potential buyers among consumers of a particular type. Let d index types and d_i encode the type of an individual as defined by gender and the bins of age and income as a share of FPL defined above. Define $\mathcal{D}_{drt} = \{i : d_i = d, r_i = r, t_i = t\}$ denote the set of individuals of type d. We compute

the number of potential buyers of type d in region r and year t as:

$$NPB_{drt} = \sum_{i \in \mathcal{D}_{drt}} \hat{\tau}_i \cdot w_i$$

where w_i are person weights in the ACS.

The final step in the procedure is to move from the number of potential buyers of a given type at the region level, to the zipcode level. This is not feasible to implement using ACS data only, using population counts are not available from such data at the zipcode level. We complement that data source with population counts at the zipcode level from `www.greatdata.com`. Using these data, we calculate the zipcode population weights w_z, such that $\sum_{z \in r} w_z = 1 \quad \forall r$, and then calculate the number of potential buyers of a given type in a zipcode as:

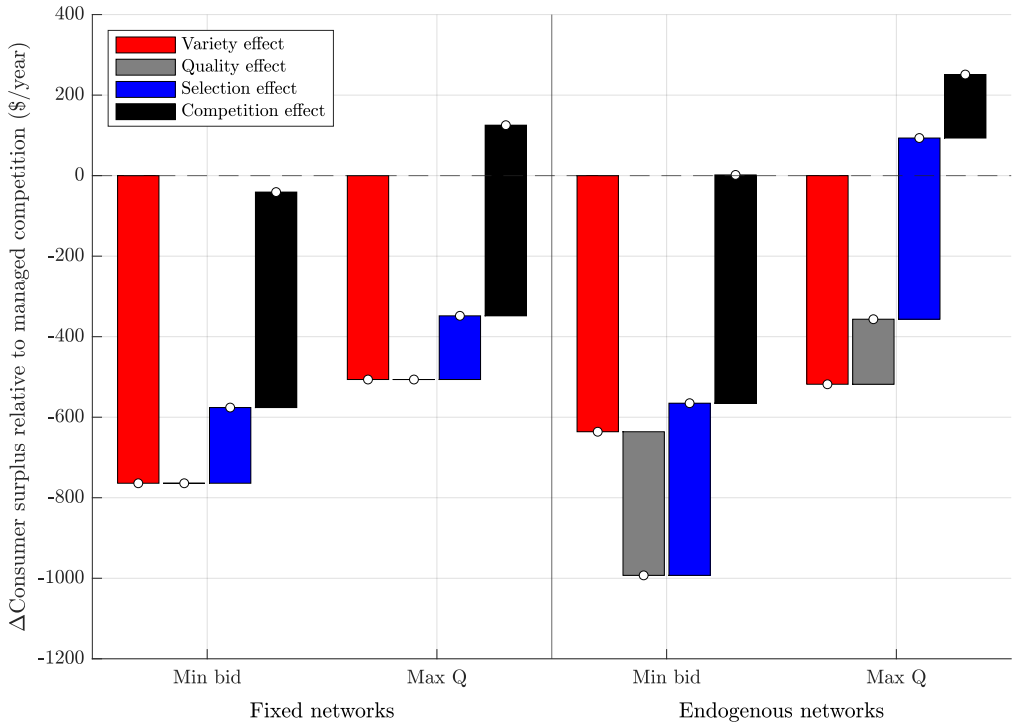
$$NPB_{dzt} = NPB_{drt} \cdot w_z$$

which requires the assumption that the distribution of population across zipcodes within a region is unrelated to individual types.

A.2 Augmenting the Enrollment Data to Account for the Outside Share

We estimate demand using the individual-level data from the marketplace and a maximum likelihood estimator. In the context of that strategy, we incorporate our measures of market size by augmenting the data to match the outside shares implied by our estimates of the number of potential buyers. We start by calculating the number of potential buyers of particular type θ_i in a given region and year that choose the outside option by taking the difference between $NPB_{\theta zt}$ and the number of enrollees in Covered California in the enrollment data. Let $N^0_{\theta zt}$ denote this number. For each type θ in a zipcode and year, we then randomly draw $N^0_{\theta zt}$ individuals from the enrollment data, and set their choice to be the outside option. We then append these synthetic enrollees to our dataset and proceed with estimation as described in Section 4.1.

Figure A.1: Decomposition of impacts of procurement market design on consumer surplus



Notes: This figure shows a decomposition of the impacts of procurement mechanisms relative to managed competition under imperfect competition on consumer surplus per potential buyer in percentage points. The decomposition has four components: variety effect, quality effect, selection effect, and markup effects. The effects are reported additively, such that for each bar, the white dot indicates the total effect up to that point in the decomposition, and the difference between subsequent white dots indicates the marginal effect of each component of the decomposition. All results are simple averages across region-years.