Public Procurement vs. Regulated Competition in Selection Markets*

José Ignacio Cuesta[§] Pietro Tebaldi[‡]

August 2025

Abstract. A common approach to markets with adverse selection is to regulate competition to minimize inefficiencies, while preserving consumer choice among firms. We study the role of procurement auctions—leading to sole provision by the winning firm—as an alternative market design. Relative to regulated competition, auctions affect product variety, quality, markups, and remove cream-skimming incentives. We develop a framework to study this comparison and apply it to individual health insurance in the US. We find that procurement auctions would increase consumer welfare in most markets, mainly by limiting inefficiencies from adverse selection and market power, and by increasing quality.

Keywords: procurement, adverse selection, managed competition, health insurance

JEL Codes: H42, I13, L13

*We thank Daniel Ackerberg, John Asker, Juan Pablo Atal, Nano Barahona, Steve Berry, Vivek Bhattacharya, Liran Einav, Phil Haile, Igal Hendel, Kate Ho, Gastón Illanes, Jakub Kastl, Victoria Marone, Aviv Nevo, Ariel Pakes, Rob Porter, Katja Seim, Mark Shepard, Paulo Somaini, Ashley Swanson, Robert Town, Ben Vatter, and conference and seminar participants at ASHECON, Cornell University, Harvard-MIT-BU, Imperial College, Northwestern University, Princeton University, PUC-Chile, SoCal Structural Econometrics Conference, UCLA, 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

Adverse selection threatens market efficiency, often prompting government intervention via regulation and public funding (see e.g., Einav, Finkelstein, and Mahoney, 2021, and references therein). When designing these regulations, policymakers typically hesitate to limit consumer choice among firms, aiming to preserve product variety and competition. Consequently, many insurance (e.g., health, home, auto, or crop) and financial markets (e.g., pensions and annuities, mortgages, or credit cards) operate under *regulated competition*, with multiple firms competing under a combination of policies restricting prices, quality, or markups, alongside mandates and subsidies. However, preserving choice between competitors can lead to inefficiencies due to enduring cream-skimming incentives (Rothschild and Stiglitz, 1976), and designing regulation that balances quality and costs in selection markets remains an active area of research and policy debate.

Instead of regulating competition *within the market*, policymakers could make firms compete *for the market* (Demsetz, 1968) by relying on public procurement. Indeed, adverse selection implies that average cost slopes downward, hinting at a parallel with natural monopolies (c.f. Kong, Layton, and Shepard, 2024), where procurement is a common approach. In addition to affecting product variety, markups, and quality, under adverse selection procurement also removes the threat of rivals poaching low-risk, profitable customers, thereby reducing inefficiencies from cream-skimming incentives. Adopting procurement in selection markets is a policy option already discussed by Diamond (1992, 1994) for health insurance, and currently employed in a few relevant contexts.[1] Nevertheless, it is yet to be formally compared to regulated competition.

In this paper, we compare public procurement to regulated competition in markets with adverse selection. We develop a model of price and quality competition that outlines the trade-offs these mechanisms pose in terms of product variety, quality, costs, and markups. Our analysis focuses on procurement mechanisms using scoring auctions—a common approach to balance costs and quality (Asker and Cantillon, 2010). We then discuss how data from a regulated oligopoly can be used to estimate the demand and cost primitives needed to evaluate counterfactual scoring auctions empirically. We apply this framework to US individual health insurance, currently regulated under the Affordable Care Act (ACA). Our findings suggest that procurement mechanisms would benefit most consumers through increased quality, lower costs, and lower markups. These effects are stronger in markets with more severe adverse selection or with fewer insurers.

[1] E.g., in employer-sponsored and low-income health insurance in the US (Gupta, Parente, and Sanyal, 2012; Shepard and Forsgren, 2023; Zewde and Perez, 2024), and pension funds in Chile (Harrison, Parada-Contzen, and Villena, 2023).

In Section 2, we develop a stylized framework to compare procurement auctions and regulated competition. Firms are vertically and horizontally differentiated, whereas consumers have heterogeneous preferences correlated with costs, leading to adverse selection. The baseline is regulated competition, under which firms compete on price and quality. The counterfactual is a second-score scoring auction (SSSA henceforth; Asker and Cantillon, 2008), in which firms submit bids consisting of price and quality. A scoring rule then maps bids into scores; the highest-scoring firm wins monopoly rights and must supply a product with a score at least as high as the second-highest score. SSSAs are a compelling procurement design for our exercise, as they can flexibly balance price and quality to accommodate government objectives. Moreover, like second-price auctions, their solution and analysis do not require assuming a high degree of firms' strategic sophistication.

Our framework isolates four channels through which a shift from regulated competition to procurement impacts consumers. First, holding prices and quality fixed, SSSAs worsen outcomes for consumers by reducing product variety (Dafny, Ho, and Varela, 2013). Second, SSSAs can lead to either higher or lower quality, depending on how strongly the scoring rule favors quality relative to price. Third, by changing the selection of consumers in the winning firm, average cost may fall, leading to lower prices. Finally, through competition at the bidding stage, SSSAs can reduce markups and further lower prices. The directions and magnitudes of these effects are ambiguous, rendering the comparison between regulated competition and procurement an empirical question.

For this reason, in Section 3 we outline our approach to recovering the demand and cost primitives needed to compare an observed equilibrium under regulated competition to counterfactual SSSAs. For demand, we consider a pure characteristics random utility model that accommodates observable and unobservable preference heterogeneity (Berry and Pakes, 2007). While logit models tend to overestimate the welfare impacts of large changes in choice sets, the pure characteristics model avoids this by omitting idiosyncratic error terms. With—increasingly common—individual-level data, these models can be estimated via maximum likelihood, avoiding the challenges associated with aggregate data (Pang, Su, and Lee, 2015). For cost, we combine data on average costs with optimal price- and quality-setting assumptions. This identifies cost heterogeneity across products, how costs depend on individual preferences and demographics, and how quality adjustments affect costs. With demand and cost, we can compute optimal bids, determine the winner, and solve for equilibrium outcomes under any SSSA scoring rule, enabling a direct comparison with the regulated competition status quo.

We apply this framework to evaluate a procurement approach in the US individual health insurance market. Health insurance is a compelling setting, since several countries organize under regulated competition—often labeled as *managed competition* in healthcare (Enthoven, 1993; Einav and Levin, 2015). Designing regulation for this market is the focus of extensive research and reform, yet, as Baicker, Chandra, and Shepard (2023) notes, the resulting "patchwork of policies" has fallen short of stated policy goals.[2] In this context, there has been discussion suggesting that procurement may enhance market efficiency (Diamond, 1992, 1994; Montanera, Mishra, and Raghu, 2022). Importantly, our goal is not to propose a specific alternative to the ACA, but to assess the potential for procurement to improve outcomes relative to the status quo in a relevant setting with adverse selection, and to examine its drivers.

Our setting is the Covered California ACA marketplace during 2014–2017, which we describe in Section 4. This is a textbook example of regulated competition, where consumers choose among plans offered by up to seven highly regulated insurers. We augment the individual-level data on enrollment and plan characteristics used in Tebaldi, Torgovitsky, and Yang (2023) and Tebaldi (2025) with granular data on plan provider networks as well as provider and individual locations. These data are crucial for measuring horizontal differentiation across plans: Whereas networks are common to all the enrollees of a plan, they are differentially attractive across enrollees based on their location. We combine this observation with regulatory variation in prices and actuarial value to identify heterogeneous preferences over plan characteristics, including provider networks.

We estimate our model for this setting in Section 5. We begin by estimating a pure characteristics demand model with random coefficients on observed plan attributes (premium, actuarial value, and network breadth), and insurer-network-region-year-specific unobservables. We find that the average individual is willing to pay $70 per year—or 3.6 percent of the average annual premium of $1,920—for a five percentage point increase in primary care physician coverage within 10 miles, and $40 per year for the same increase within 10-20 miles. Having the closest in-network hospital two miles closer is valued at $45 per year. Valuations are highly heterogeneous: Individuals in the 90*th* percentile value provider access more than five times the population average. In terms of plan generosity, individuals are indifferent between a five percentage point increase in actuarial value and a $290 decrease in annual premium, with substantial heterogeneity across the population.

The empirical cost model is identified and estimated by combining demand estimates with

[2] For example, subsidies to consumers can distort competition between insurers (Decarolis, 2015; Decarolis, Polyakova, and Ryan, 2020; Jaffe and Shepard, 2020; Tebaldi, 2025); risk adjustment does not fully remove cream-skimming incentives (Brown, Duggan, Kuziemko, and Woolston, 2014), can lead to upcoding (Geruso and Layton, 2020), and can increase markups (Starc, 2014; Mahoney and Weyl, 2017; Einav, Finkelstein, and Tebaldi, forthcoming).

observed average costs and the assumption that plan prices and characteristics reflect an equilibrium. We find that expected costs increase with age by an average of 73 percent every ten years. Moreover, costs vary with preferences. A $100 higher willingness to pay for five percentage points in actuarial value predicts 3.5 percent higher expected costs, and stronger preferences for provider networks also predict higher expected costs—in line with Shepard (2022). In terms of changes in provider networks or brand quality, we find that expected costs increase by 3–3.5 percent when either: the density of in-network primary care physicians rises by ten percentage points, the closest in-network hospital is five miles closer, or the valuation of brand quality is $100 higher.

Having estimated demand and cost primitives, we can compare regulated competition and procurement. In Section 6, we simulate counterfactual SSSAs for a range of scoring rules with varying relative weights on a plan's quality and premium. Our findings suggest that SSSAs that reward plan quality can significantly improve average consumer surplus relative to regulated competition, at $250 per person-year. In terms of mechanisms, the gains from lower costs ($370), lower markups ($310), and higher quality ($70) outweigh the losses from reduced variety ($500). Net of changes in government spending, consumer gains exceed $200 per person per year, or about $2.4 billion over our four-year period. These impacts occur when scoring rules properly weight plan quality. If the weight is too low, insurers reduce quality relative to regulated competition, making consumers worse off. If the weight is too high, quality and costs rise to inefficient levels. Importantly, absent the gains from addressing adverse selection through procurement, the result would flip, and consumers would be better off on average under ACA-regulated competition.

Beyond aggregate impacts, we find that SSSAs outperform the status quo in 87 percent of region-year markets. Inspecting this heterogeneity across markets reveals that gains from SSSAs are larger in markets with fewer insurers, higher market concentration, or more severe adverse selection—consistent with the comparative statics of our stylized framework. These results provide guidance on when procurement can be advantageous relative to regulated competition.

Related Literature. Since the work of Demsetz (1968), Spence (1975), and Shleifer (1985), a rich theoretical literature has examined the trade-offs between quality, cost, and imperfect competition. Notably, McGuire and Riordan (1995) and Wolinsky (1997) compare sole-sourcing arrangements to regulated competition. Here, we revisit this question with a focus on markets with adverse selection—absent in earlier models—and emphasize the product variety and preference heterogeneity omitted by Diamond (1992, 1994) in his discussion of procurement for health insurance.

We also speak to the growing body of empirical work on market design in selection markets.

In the health insurance context, this includes—only recently and among many others—Decarolis (2015), Decarolis, Polyakova, and Ryan (2020), Jaffe and Shepard (2020), Marone and Sabety (2022), Dickstein, Ho, and Mark (2024), Vatter (2025). Handel and Ho (2021) provides a broad review, expanded to focus on ACA-specific studies by Handel and Kolstad (2022) and Baicker, Chandra, and Shepard (2023). Beyond health insurance, Einav, Finkelstein, and Mahoney (2021) review several studies on policy design for selection markets; for financial markets, these include Crawford, Pavanini, and Schivardi (2018), Illanes and Padi (2019), Benetton (2021), and Cuesta and Sepúlveda (2021). Our contribution is to consider procurement as an alternative to regulated competition. Relatedly, Montanera, Mishra, and Raghu (2022) compares second-price procurement auctions to risk adjustment in the context of low-income US health insurance (Medicaid MCOs). This is relevant, since risk adjustment (despite some shortcomings, c.f. footnote 2 above) is widely used to address cream-skimming inefficiencies. Unlike us, Montanera, Mishra, and Raghu (2022) do not model plan variety, endogenous quality, or enrollment decisions, yet they also find that procurement could improve welfare relative to a combination of competition and risk adjustment.

A key input to our analysis is a measure of the value of variety. Previous research has emphasized that relying on demand models with idiosyncratic shocks to random utility when evaluating changes in choice sets is problematic (Ackerberg and Rysman, 2005; Dafny, Ho, and Varela, 2013; Quan and Williams, 2018). Although pure characteristics models (Bajari and Benkard, 2005; Berry and Pakes, 2007; Song, 2007) provide a solution, computational issues make estimation challenging (Pang, Su, and Lee, 2015). Here, we show how to use individual-level data to estimate the model by maximum likelihood. The advantages of this approach are that it eliminates assumptions on the distribution of idiosyncratic preference shocks often dictated by computational feasibility, and aligns the demand model with recent identification results by Berry and Haile (2024).

Finally, our cost identification approach applies the inversion of equilibrium conditions (Rosse, 1970; Bresnahan, 1987; Berry, Levinsohn, and Pakes, 1995) to a selection market, where the difference between marginal and average cost informs cost heterogeneity across consumers (Einav, Finkelstein, and Cullen, 2010). This allows us to relax parametric assumptions that restrict costs across products, as in Bundorf, Levin, and Mahoney (2012) and Tebaldi (2025), among others. Moreover, by allowing costs to adjust with endogenous quality, we contribute to recent research on competition with endogenous product attributes, including Fan (2013), Wollmann (2018), Crawford, Shcherbakov, and Shum (2019), Barahona, Otero, and Otero (2023) and, in healthcare, Ho and Lee (2019), Ghili (2022), Serna (2023), Cuesta, Noton, and Vatter (2025), and Vatter (2025).

2 Framework

We begin with a stylized framework to illustrate the essential trade-offs between a differentiated products oligopoly and a procurement auction in a market with adverse selection. Our analysis takes the perspective of a planner who regulates an oligopoly, observes its equilibrium, and has the institutional power to change the market mechanism. For ease of exposition, we avoid incorporating specific regulations in this section, which will vary on a case-by-case basis; we do account for ACA regulations in our empirical application below.

Oligopoly and Regulated Competition Environment. There is a population of individuals, indexed by i, and J firms. Each firm j offers a contract (e.g., an insurance policy or a loan) with scalar quality $\delta_j \geq 0$ and price p_j. When purchasing contract j, individual i derives indirect utility $U_{ij} = -\alpha_i p_j + \beta_i \delta_j + \epsilon_{ij}$, with $\alpha_i, \beta_i > 0$. The expected cost for the firm is $C_{ij} = \kappa_j \delta_j (1 + \rho \beta_i)$, where $\kappa_j, \rho \geq 0$. Individuals can also choose an outside option $j = 0$ with $U_{i0} = 0$ for all i.

The β_i component of an individual's type describes their preferences for vertical contract quality, while the vector $\boldsymbol{\epsilon}_i = (\epsilon_{i1}, \epsilon_{i2}, \ldots, \epsilon_{iJ})$ measures horizontal tastes for contracts. The larger the dispersion in $\boldsymbol{\epsilon}_i$, the more horizontal tastes drive contract choice. The costs of supplying contract j increase in quality at rate κ_j, and adverse selection is present if $\rho > 0$, since this implies that individuals with stronger preferences for quality also have higher costs.

The expected profit for firm j when setting quality and price (δ_j, p_j) under oligopoly is

$$\Pi_j^o(\delta_j, p_j; \boldsymbol{\delta}_{-j}, \boldsymbol{p}_{-j}) = \int \sigma_j(\delta_j, p_j; \boldsymbol{\delta}_{-j}, \boldsymbol{p}_{-j}, \beta_i)(p_j - C_{ij}) dF(\beta_i) , \tag{1}$$

where the probability that a consumer with preference β_i chooses j is $\sigma_j(\delta_j, p_j; \boldsymbol{\delta}_{-j}, \boldsymbol{p}_{-j}, \beta_i) \equiv \int \mathbf{1}\{U_{ij} \geq U_{ik} \ \ \forall k\} dF(\alpha_i, \boldsymbol{\epsilon}_i | \beta_i)$, or $\sigma_j(\beta_i)$ for notational simplicity. In a Nash equilibrium $(\boldsymbol{\delta}^\star, \boldsymbol{p}^\star)$, firms choose quality and prices optimally, taking as given their competitors' choices. Without stronger assumptions, heterogeneity across individuals and $\rho > 0$ prevent a general characterization of equilibrium, which is beyond our scope. Nevertheless, it is worth noting that more severe adverse selection—as implied by a larger value of ρ—means that, ceteris paribus, a firm has a weaker incentive to increase quality, since

$$\frac{\partial^2 \Pi_j^o}{\partial \rho \partial \delta_j} = -\kappa_j \int \beta_i \left(\sigma_j(\beta_i) + \delta_j \frac{\partial \sigma_j(\beta_i)}{\partial \delta_j} \right) dF(\beta_i) < 0 .$$

As ρ increases, the marginal return from a higher δ_j decreases, since marginal revenue does not change while marginal cost increases.

This environment relates to many prior studies on selection markets focused on understanding and designing regulation for contracts, pricing conduct, mandates, or subsidies.[3] These regulations are pervasive due to the inefficiencies that arise with adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976). Regulated (or managed) competition refers to the resulting, often heavily regulated, market environment (see e.g., Enthoven, 1993 and Einav and Levin, 2015, for the case of healthcare, or Wolinsky, 1997 for a discussion of regulating oligopolistic firms). In this framework, consumer choice between firms is a key principle for disciplining equilibrium outcomes. However, this principle may fail to achieve efficiency when adverse selection is present, since firms have incentives to attract low-cost individuals from their competitors. As highlighted by Kong, Layton, and Shepard (2024), this creates a symmetry between adverse selection and natural monopolies.

Procurement Via Second-Score Scoring Auctions. As an alternative, the planner could make firms compete *for the market* rather than *within the market* (Demsetz, 1968). For this purpose, one possibility is to set up a second-score scoring auction (SSSA henceforth).

In an SSSA, the planner selects and publicly announces a scoring rule. We focus on quasi-linear scoring rules (see e.g., Asker and Cantillon, 2008), which are functions of quality and a monetary bid b: $\Psi(\delta, b) = \psi(\delta) - b$, where $\psi(\cdot)$ is continuous and differentiable. Every firm j submits (δ_j, b_j), and the resulting scores are calculated and ranked. The firm with the highest score wins the rights to supply the market at any combination (δ^a, p^a) such that the implied score does not fall below the second-highest score in the auction. Formally, the winner is j^w such that $\psi(\delta_{j^w}) - b_{j^w} \geq \psi(\delta_k) - b_k$ for all k, breaking ties randomly without loss of generality. After winning, j^w sets (δ^a, p^a) by solving

$$\max_{\delta,p} \Pi^a_{j^w}(\delta, p) \quad \text{s.t.} \quad \psi(\delta) - p \geq \max_{k \neq j^w} \psi(\delta_k) - b_k \,, \tag{2}$$

where with a slight abuse of notation $\Pi^a_{j^w}(\delta, p) \equiv \Pi^o_{j^w}(\delta, p; \emptyset, \infty)$, indicating that individuals can only choose to purchase j^w or nothing, since other contracts are not available.

Akin to second-price auctions, bidder behavior in SSSAs is strategically simple. Specifically, it is weakly dominant for every firm j to submit a bid (δ_j, b_j) that maximizes the resulting score

[3] Studies of insurance with similar demand and cost structure—or that can be mapped into it—include Cohen and Einav (2007), Einav, Finkelstein, and Cullen (2010), Handel, Hendel, and Whinston (2015), Jaffe and Shepard (2020), Geruso, Layton, McCormack, and Shepard (2023), and Tebaldi (2025), among others. For credit markets and selection markets more broadly, this framework nests Einav, Finkelstein, and Schrimpf (2010), Mahoney and Weyl (2017), Azevedo and Gottlieb (2017), Crawford, Pavanini, and Schivardi (2018), and Cuesta and Sepúlveda (2021), among others.

subject to breaking even.[4] That is, every bid (δ_j, b_j) maximizes $\psi(\delta) - b$ subject to $\Pi_j^a(\delta, b) = 0$. Denoting by $\overline{b}_j(\delta)$ the break-even price when offering quality δ as the sole contract in the market, and by $\overline{\delta}_j(b)$ its inverse, for any scoring rule, it is weakly dominant to submit a bid on this curve. Formally, b_j maximizes $\psi(\overline{\delta}_j(b)) - b$, and $\delta_j = \overline{\delta}_j(b_j)$. These functions are therefore key to studying equilibrium in SSSAs, and can be constructed with demand and cost estimates.[5]

Graphical Example. We provide an example of a switch from an oligopoly to procurement via SSSAs. We focus on a case with $\alpha_i = 1$, β_i taking values β^{low} or β^{high}, and $j \in \{A, B\}$. For clarity, we ignore ϵ_i. Figure 1-a shows iso-utility curves for the two individual types, and the price-quality combination of the oligopoly equilibrium, $A^\star$ and $B^\star$. High-β types select $j = A$, with higher quality and price, whereas low-β types select $j = B$. From the planner's viewpoint, the points $A^\star$ and $B^\star$ are observed, and the indifference curves can be obtained from estimating demand. In Figure 1-b, we overlay the sets of break-even price-quality pairs for each firm, $\overline{\delta}_A(b)$ and $\overline{\delta}_B(b)$.

Moving to outcomes under SSSAs, Figure 1-c illustrates an SSSA with a scoring rule with a high enough $d\psi/d\delta$ as to increase both contracts' quality relative to oligopoly. Each firm $j \in \{A, B\}$ submits a bid (δ_j, b_j) such that $d\overline{\delta}_j/db = (d\psi/d\delta)^{-1}$. In this example, A obtains the highest score and the second-highest score—achieved by B—defines the constraint A must satisfy when choosing (δ^a, p^a) to maximize profits: $\psi(\delta^a) - p^a \geq \psi(\delta_B) - b_B$, or $\delta^a \geq \psi^{-1}(\psi(\delta_B) - b_B + p^a)$. Combining this with A's profit-maximizing behavior—imposing $\Pi_A^a(\delta^a, p^a) \geq 0$, or, equivalently, $\delta^a \leq \overline{\delta}_A(p^a)$—defines the shaded region. The outcome of the SSSA must be inside this set. The winner selects the exact point by solving a constrained monopoly problem based on demand and cost curves. All individuals are better off within this set relative to oligopoly.

In contrast, Figure 1-d shows a case in which ψ does not incentivize quality provision and B wins the auction. As before, the shaded region is the set of possible outcomes from the SSSA. For any point in this set, the losses for consumers of type β^{high}—who choose A under oligopoly—cannot be compensated by quality or price changes due to changes in average cost or markups.

[4]For any (δ_j, b_j) leading to positive profits, the firm could increase its score (and winning probability) by lowering b_j or increasing δ_j. For any (δ_j, b_j) leading to negative profits, the firm could win and incur losses. Among break-even (δ_j, b_j) combinations, the firm maximizes its winning probability by choosing the one leading to the highest score.

[5]As a technical note, $\overline{b}_j(\delta)$ needs to exist, requiring the demand and average cost curves with quality δ to cross at least once. A sufficient condition for this is that $\partial AC_j^a(\delta, p)/\partial p < 1$, where $AC_j^a(\delta, p)$ is the average cost when j is the only contract available to consumers. This is a "no unraveling" condition, tightly linked to what Kong, Layton, and Shepard (2024) identifies as necessary for the existence of a Nash equilibrium.

Figure 1: From oligopoly equilibrium to SSSA: Illustration with $\beta_i \in \{\beta^{\text{low}}, \beta^{\text{high}}\}$, $J = 2$

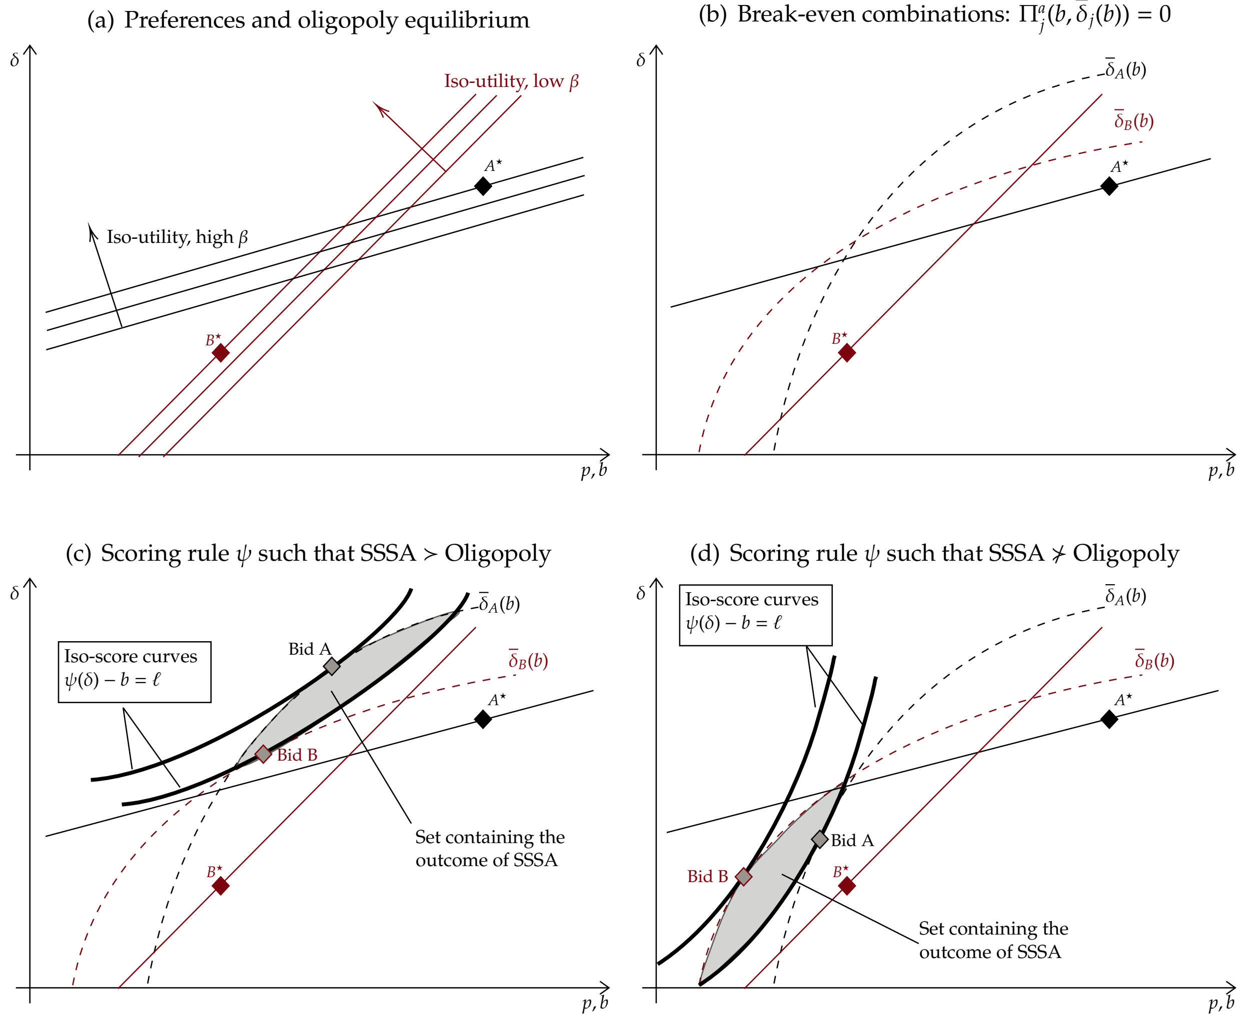

Notes: The figure illustrates the comparison between an oligopoly equilibrium and SSSAs for a simple case with two firms $j = A, B$, $\alpha_i = 1$, and two types $\beta_i \in \{\beta^{\text{low}}, \beta^{\text{high}}\}$. The top-left panel shows the price-quality combination set in equilibrium, $A^\star$ and $B^\star$, as well as iso-utility curves for the two types of consumers. High-β selects A, while low-β selects B. The top-right panel overlays the graph with the two sets of break-even price-quality combinations, one for every firm. The bottom-left panel shows the outcomes of a SSSA in which $d\psi/d\delta$ is set sufficiently high to induce an equilibrium in which both types are better off. In the bottom-right panel, instead, $d\psi/d\delta$ is lower, such that quality drops. The SSSA in this case makes high-β types worse off, while low-β types would prefer this scoring rule to the one in the bottom-left panel. The main text provides further details.

Trading Off Variety, Quality, Selection, and Markups. The impact of a shift from regulated competition to procurement through SSSAs on consumer surplus would be

$$\Delta \text{CS} = \int \left(\max\{0, U^a_{ij^w}(p^a)\} - \max_{j \in \{0,\dots,J\}} \underbrace{-\alpha_i p_j^\star + \beta_i \delta_j^\star + \epsilon_{ij}}_{U^\star_{ij}} \right) di\,, \tag{3}$$

where $U^a_{ij^w}(p) \equiv -\alpha_i p + \beta_i \delta^a + \epsilon_{ij^w}$. This can be decomposed into four effects capturing the key economic forces at play: changes in variety, quality, costs, and markups.

A clear disadvantage of SSSAs is that consumer choice is limited to the winner. Ceteris paribus, if prices and quality were not affected by a switch from regulated competition to procurement auctions, consumers would be worse off due to the lower variety. For example, even in Figure 1-c, low-β individuals are worse off when only A is available prior to price and quality adjustments. Formally, we denote as the *variety effect* the change in consumer surplus due to a shift from an oligopoly equilibrium $(\boldsymbol{\delta}^\star, \boldsymbol{p}^\star)$ to an auction with winner j^w and outcomes (δ^a, p^a), namely

$$\text{Variety Effect} = \int \left(\max\{0, U^\star_{ij^w}\} - \max_{j \in \{0,\dots,J\}} U^\star_{ij} \right) di \leq 0 \,,$$

which is always negative for rational consumers. An exception arises if consumers are prone to mistakes or misinformed, leading to suboptimal choices (Handel, 2013; Abaluck and Gruber, 2023; Brown and Jeon, 2024; Vatter, 2025). This could justify limiting variety, but it is beyond our scope.

While limiting variety harms consumers, quality provision in SSSAs depends on how the scoring rule rewards higher quality. When bidding, firm j has an incentive to increase δ as long as $d\bar{\delta}/db \geq (d\psi/d\delta)^{-1}$. Therefore, with sufficiently high $d\psi/d\delta$, SSSAs can induce higher quality than under regulated competition, while the opposite may occur under scoring rules that primarily reward low prices—the bottom panels in Figure 1 illustrate these cases. To capture the impact of changes in quality, we measure the incremental *quality effect* by holding the price fixed at what the winning contract j^w would set under regulated competition, $p^\star_{j^w}$, but changing its quality to δ^a, as

$$\text{Quality Effect} = \int \left(\max\{0, U^a_{ij^w}(p^\star_{j^w})\} - \max\{0, U^\star_{ij^w}\} \right) di \,,$$

which is positive if and only if $\delta^a > \delta^\star_{j^w}$.

Cost changes are another channel through which a shift toward SSSAs could impact consumers. In particular, the average cost for the winning contract also changes, due to (i) shifts in the composition of buyers due to changes in variety and quality, (ii) shifts in the composition of buyers due to price changes, and (iii) changes in cost due to changes in quality. We measure the incremental *cost effect* (or *selection effect*) by holding markups fixed at their level under regulated competition, $\mu^{\star}_{jw}$, but adjusting prices to reflect changes in average cost from shifting to auctions:

$$\text{Cost/Selection Effect} = \int \Big(\max\{0, U^{a}_{ijw}(AC^{a}_{jw} + \mu^{\star}_{jw})\} - \max\{0, U^{a}_{ijw}(p^{\star}_{jw})\} \Big)\, di\,,$$

where AC^{a}_{jw} is the average cost of the winning contract in the SSSA. This effect is positive if the winner's average cost under the SSSA is lower than its average cost under regulated competition.

Lastly, while under regulated competition markups are determined by the strategic interaction between competing contracts, markups in SSSAs result from the gap between the first- and second-highest scores. We measure this incremental *markup effect* as

$$\text{Markup Effect} = \int \Big(\max\{0, U^{a}_{ijw}(p^{a})\} - \max\{0, U^{a}_{ijw}(AC^{a}_{jw} + \mu^{\star}_{jw})\} \Big)\, di\,,$$

which sign depends on the degree of product differentiation under oligopoly, and demand and cost elasticities to price and quality.

The impact of a shift from regulated competition to procurement auctions is inherently context-specific. However, restricted versions of this framework can deliver some results for this comparison. Assuming logit demand is one way to gain tractability. In particular, under logit demand—which implies large losses from limiting variety—there are conditions under which SSSAs deliver higher welfare than regulated competition, as the following remark shows:

Remark 1. *With $J \geq 2$, if $\epsilon_{ij} = \zeta \varepsilon_{ij}$ with ε_{ij} iid following a Gumbel distribution, and $\kappa_j = \kappa$ for all j, given an oligopoly equilibrium $(\boldsymbol{\delta}^{\star}, \boldsymbol{p}^{\star})$ with average costs and markups equal to $AC^{\star}_j$ and $\mu^{\star}_j$, for $j = 1, \ldots, J$, letting $j^{\star}(\alpha_i, \beta_i) = \arg\max_j -\alpha_i p^{\star}_j + \beta_i \delta^{\star}_j$, there exists an SSSA such that $\delta^a \geq \max_j \delta^{\star}_j$, and, if*

$$\underbrace{-\zeta \log(J)}_{\textit{Variety}} + \underbrace{\beta_i(\delta^a - \delta^{\star}_{j^{\star}(\alpha_i,\beta_i)})}_{\textit{Quality}} + \alpha_i\big(\underbrace{AC^{\star}_{j^{\star}(\alpha_i,\beta_i)} - \overline{b}(\delta^a)}_{\textit{Cost/Selection}} + \underbrace{\mu^{\star}_{j^{\star}(\alpha_i,\beta_i)}}_{\textit{Markup}}\big) > 0 \quad \textit{for all } \alpha_i, \beta_i\,,$$

then consumer surplus under the SSSA is strictly greater than under the oligopoly equilibrium.

We prove this in Appendix A, where we also discuss the comparison between SSSAs and oligopoly for simple logit cases with fixed quality and without adverse selection.

The sufficient condition in Remark 1 points to four empirical objects that can predict the potential of SSSAs to improve outcomes relative to oligopoly. There are two non-price channels. First, variety losses likely increase in the number of products J, and with the variance of idiosyncratic taste shocks. Second, the gains from quality are bounded (below) by the difference between the quality resulting in the SSSA—as determined by the scoring rule—and the quality selected by every type under oligopoly. To these, one must add price effects, which are more relevant when individuals are more price-sensitive—as captured by a larger α_i. Price changes combine changes in average cost and markups. Importantly, even ignoring quality changes, SSSAs can outperform oligopoly by combining gains from the selection and markup channels. Even absent selection, while one may think that more horizontal dispersion in consumer preferences implies a disadvantage for SSSAs, this is not generally true since such dispersion also increases markups.

3 Estimation

To compare an observed oligopoly with counterfactual procurement auctions, one needs estimates of the distribution of preferences and costs, as well as the cost of increasing quality. These inputs allow for building $\overline{\delta}_j(b)$, computing bids by maximizing $\psi(\overline{\delta}_j(b)) - b$, solving the winner's problem in equation (2), and calculating counterfactual outcomes.

Demand. A shift from an oligopoly to a procurement mechanism entails a drastic reduction in the number of choices available to consumers, making the modeling of preference heterogeneity crucial. Demand models with idiosyncratic and fully supported preference shocks ϵ_{ij} (e.g., logit, mixed logit, and probit) have the undesirable property that any change in the choice set $\mathcal{J}$ changes the space of unobservable characteristics (Ackerberg and Rysman, 2005; Berry and Pakes, 2007).

To circumvent this issue, we note that the recent diffusion of individual-level datasets, paired with growing computational capabilities, enables using the pure characteristics demand model (Bajari and Benkard, 2005; Berry and Pakes, 2007). In this model, the indirect utility function is

$$U_{ij} = -\alpha_i p_j + \delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i) \,, \tag{4}$$

where $\delta(\cdot)$ is known up to the parameter vector $\boldsymbol{\beta}_i$ (which could include a product-level unobservable ξ_j), $\boldsymbol{x}_j$ is a vector of observable characteristics of contract j, and $\boldsymbol{z}_i$ is a vector of observable characteristics of individual i. The demand system is characterized by the distribution of α_i and $\boldsymbol{\beta}_i$, and variation in these parameters along with $\boldsymbol{z}_i$ captures vertical and horizontal differentiation.[6] Berry and Haile (2024) provides sufficient conditions for identification of this model, but, for any specific application, institutional details must point to the necessary sources of exogenous

[6]To draw a parallel to the model of Section 2, one could set $\epsilon_{ij} = \delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i) - \mathbb{E}_{\boldsymbol{z}_i, \boldsymbol{\beta}_i}[\delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i)]$ and notice that the scope of horizontal preferences—by which different individuals disagree in ranking products when all prices are zero—depends on the joint variation in $\boldsymbol{z}_i$ and $\boldsymbol{\beta}_i$.

variation in prices and product characteristics.

For estimation, individual-level data on choices and characteristics prove extremely useful insofar as the parameters of (4) can be estimated via maximum likelihood. Following the steps in Berry and Pakes (2007), we define

$$\underline{\Delta}_{ij}(\boldsymbol{\beta}_i) \equiv \max_{k:p_k>p_j} \frac{\delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i) - \delta(\boldsymbol{x}_k, \boldsymbol{z}_i; \boldsymbol{\beta}_i)}{p_k - p_j}, \text{ and } \overline{\Delta}_{ij}(\boldsymbol{\beta}_i) \equiv \min_{k:p_k<p_j} \frac{\delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i) - \delta(\boldsymbol{x}_k, \boldsymbol{z}_i; \boldsymbol{\beta}_i)}{p_k - p_j}, \tag{5}$$

such that the likelihood of observing individual i choosing j is

$$\ell_{ij} = \int \left(\int_{\underline{\Delta}_i(\boldsymbol{\beta}_i)}^{\overline{\Delta}_i(\boldsymbol{\beta}_i)} dF(\alpha_i | \boldsymbol{\beta}_i; \boldsymbol{z}_i) \right) dF(\boldsymbol{\beta}_i; \boldsymbol{z}_i) \approx M^{-1} \sum_{m=1}^{M} \int_{\underline{\Delta}_i(\overline{\boldsymbol{\beta}}+\Sigma^{\beta}\boldsymbol{\nu}_i^m)}^{\overline{\Delta}_i(\overline{\boldsymbol{\beta}}+\Sigma^{\beta}\boldsymbol{\nu}_i^m)} dF\big(\alpha_i | \underbrace{\overline{\boldsymbol{\beta}} + \Sigma^{\beta}\boldsymbol{\nu}_i^m}_{\substack{m^{\text{th}} \text{ simulated} \\ \text{draw of } \boldsymbol{\beta}_i}}; \boldsymbol{z}_i\big), \tag{6}$$

where the term on the right is the numerical integral with respect to the distribution $F(\boldsymbol{\beta}_i; \boldsymbol{z}_i)$. For this, the researcher draws M vectors $\boldsymbol{\nu}_i^m$ iid from a known mean-zero distribution (typically normal or lognormal). The parameters of the model are the mean vector $\overline{\boldsymbol{\beta}}$, the variance-covariance matrix Σ^{β}, and the parameters entering the distribution $F(\alpha_i | \boldsymbol{\beta}_i; \boldsymbol{z}_i)$. As long as this admits a density $f(\alpha_i | \boldsymbol{\beta}_i; \boldsymbol{z}_i)$ that is differentiable with respect to a parameter vector, the individual and total likelihoods are also differentiable with respect to the same parameters.

Computational problems can arise due to how ℓ_{ij} responds to changes in $\overline{\boldsymbol{\beta}}$ or Σ^{β}. Absent variation in $\boldsymbol{z}_i$—as in Berry and Pakes (2007), Song (2007), and Pang, Su, and Lee (2015)—the simulated likelihood and the predicted choice shares are non-smooth. In particular, the derivatives jump discontinuously at points of the parameter space, implying that the solution of the problems in equation (5) changes for a non-negligible mass of draws of $\boldsymbol{\beta}_i$. The availability of individual-level data on choices and characteristics mitigates this issue. Intuitively, individual characteristics play a role similar to that of idiosyncratic shocks to random utility, and ensure smoother choice shares and likelihoods. Specifically, variation in $\boldsymbol{z}_i$ reduces the mass of simultaneous ties in the solutions of (5), which define the boundaries of integration in (6). See Appendix B for details.

Cost. The next step is to recover the cost function, taking demand estimates as given. We consider scenarios in which, in addition to observing prices, characteristics, and choices, and having exogenous variation to identify demand, one also observes realized average costs AC_j for every contract j, which can be computed from individual-level costs when available. Identification combines such cost data with the implications of optimal pricing, adapting the inversion of oligopoly pricing conditions (Rosse, 1970; Bresnahan, 1987; Berry, Levinsohn, and Pakes, 1995) to a selection market, where discrepancies between average and marginal cost inform the sign and magnitude of selection (Einav, Finkelstein, and Cullen, 2010).

Generalizing the cost model from above to include $\boldsymbol{z}_i$ and $\boldsymbol{x}_j$, we write

$$C_{ij} = \kappa_j(\boldsymbol{x}_j)(1 + \rho(\boldsymbol{z}_i, \boldsymbol{\beta}_i)) \,, \tag{7}$$

which includes three sets of parameters: parameters that describe cost heterogeneity across products ($\kappa_j(\cdot)$ varies across j's), parameters that describe how changes in characteristics impact cost ($\partial \kappa_j / \partial \boldsymbol{x}_j$), and parameters that govern cost heterogeneity across types (the function $\rho(\cdot)$). By not restricting how $\kappa_j(\cdot)$ may differ across j's, the model imposes weaker assumptions on the cost function structure than previous work by Bundorf, Levin, and Mahoney (2012) or Tebaldi (2025).[7]

Under the maintained assumption that the observed $(\boldsymbol{p}^\star, \boldsymbol{x}^\star)$ consists of an equilibrium of the differentiated products oligopoly, one can write the necessary first-order condition for prices as

$$MR_j = MC_j = \kappa_j(\boldsymbol{x}_j^\star) \left(\frac{\partial Q_j}{\partial p_j} + \int \rho(\boldsymbol{z}_i, \boldsymbol{\beta}_i) \frac{\partial \sigma_j(\boldsymbol{z}_i, \boldsymbol{\beta}_i)}{\partial p_j} dF(\boldsymbol{z}_i, \boldsymbol{\beta}_i) \right) . \tag{8}$$

In this expression, marginal revenue with respect to price, MR_j, the derivative of total sales, $\partial Q_j / \partial p_j$, and the derivative of sales conditional on $\boldsymbol{z}_i$ and $\boldsymbol{\beta}_i$, $\partial \sigma_j(\boldsymbol{z}_i, \boldsymbol{\beta}_i) / \partial p_j$, as well as the distribution $F(\boldsymbol{z}_i, \boldsymbol{\beta}_i)$, depend only on demand parameters. Therefore, for any given function $\rho(\cdot)$, the value of $\kappa_j(\boldsymbol{x}_j^\star)$ at the observed $\boldsymbol{x}_j^\star$ can be obtained from the FOC in equation (8).[8] Let this value be $\kappa_j^\rho(\boldsymbol{x}_j^\star)$, of which there is one for every j and every possible $\rho(\cdot)$. Since AC_j is observed, $\rho(\cdot)$ must be such that

$$AC_j Q_j = \kappa_j^\rho(\boldsymbol{x}_j^\star) \left(Q_j + \int \rho(\boldsymbol{z}_i, \boldsymbol{\beta}_i) \sigma_j(\boldsymbol{z}_i, \boldsymbol{\beta}_i) dF(\boldsymbol{z}_i, \boldsymbol{\beta}_i) \right) . \tag{9}$$

Appendix C shows that there is a unique function $\rho(\cdot)$ that satisfies equations (8) and (9) as long as the variation in prices and characteristics identifying demand also varies the composition of average ($\sigma_j(\boldsymbol{z}_i, \boldsymbol{\beta}_i) dF(\boldsymbol{z}_i, \boldsymbol{\beta}_i)$) and marginal ($\partial \sigma_j(\boldsymbol{z}_i, \boldsymbol{\beta}_i) / \partial p_j dF(\boldsymbol{z}_i, \boldsymbol{\beta}_i)$) buyers across contracts. Intuitively,

[7]Since average claims are observed, these articles impose functional form assumptions that restrict the differences between $\kappa_j(\cdot)$ and $\kappa_k(\cdot)$ when $k \neq j$. Under these assumptions, costs are identified without the need to impose conduct assumptions such as equilibrium pricing. Alternatively, under the same cross-plan restrictions, Tebaldi (2025) discusses identification of costs assuming equilibrium pricing but without observing average claims. Here, we combine the observability of average claims with conduct assumptions, which allows us to leave $\kappa_j(\cdot)$ fully unrestricted.

[8]Here we discuss single-product firms. Our application considers the case of multi-product firms.

absent cost heterogeneity across types, marginal cost is equal to average cost, which is observed in the data. Therefore, $\rho(\cdot)$ is identified by the joint distribution of differences between marginal revenue and average cost and the $(z_i, \boldsymbol{\beta}_i)$-composition of marginal and average buyers. If contracts with larger $\boldsymbol{\beta}_i$'s among their average buyers relative to their marginal buyers have $MR_j < AC_j$, ρ must increase in $\boldsymbol{\beta}_i$, and vice-versa; similarly for z_i.

As a last step, imposing necessary FOCs with respect to $\boldsymbol{x}_j$ identifies how costs vary with changes in product characteristics. Formally:

$$\frac{\partial R_j}{\partial \boldsymbol{x}_j} = \kappa_j(\boldsymbol{x}_j^\star)\left(\frac{\partial Q_j}{\partial \boldsymbol{x}_j} + \int \rho(z_i, \boldsymbol{\beta}_i)\frac{\partial \sigma_j(z_i, \boldsymbol{\beta}_i)}{\partial \boldsymbol{x}_j} dF(z_i, \boldsymbol{\beta}_i)\right) + \frac{\partial \kappa_j(\boldsymbol{x}_j^\star)}{\partial \boldsymbol{x}_j}\int \rho(z_i, \boldsymbol{\beta}_i)\sigma_j(z_i, \boldsymbol{\beta}_i) dF(z_i, \boldsymbol{\beta}_i)\,, \quad (10)$$

from which one can obtain $\partial \kappa_j(\boldsymbol{x}_j^\star)/\partial \boldsymbol{x}_j$. One can then use equation (10) to identify and estimate a parametric model of how changes in contract characteristics impact cost.

4 Empirical Application: California ACA Marketplace

4.1 Institutional Setting

Our empirical analysis focuses on the California ACA marketplace (see also Panhans, 2019; Saltzman, 2021; Tebaldi, Torgovitsky, and Yang, 2023; Tebaldi, 2025). In this market, insurers offer a menu of plans in a highly regulated environment, which combines federal and state rules on plan attributes, market definition, and premium setting, among others. We briefly summarize the regulations that are essential for our analysis.

Insurers may decide to offer plans in each of 19 geographic rating regions, for which they must offer adequate provider networks.[9] Plans are standardized into five tiers based on their generosity: 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 have higher cost-sharing and cannot be purchased by individuals who are subsidized or older than 30. Individuals with income below 250 percent of the federal poverty line (FPL) benefit from cost-sharing reductions, which increase Silver plans' generosity (DeLeire, Chappel, Finegold, and Gee, 2017).

A few pieces of regulation on premiums are worth mentioning. Insurers choose a scalar base

[9] Federal regulations (see e.g., `https://www.kff.org/affordable-care-act/issue-brief/network-adequacy-standards-and-enforcement/`; accessed on July 7, 2025) require coverage of essential providers and 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 offered by ACA plans have been narrower than in other segments of the insurance market (Haeder, Weimer, and Mukamel, 2019).

premium per plan. An age-adjustment schedule then turns this base premium into age-specific premiums, limiting price discrimination across demographic groups. These premiums are the revenue insurers obtain from selling a plan. Individuals may receive subsidies based on their income relative to the FPL, with the subsidy amount being linked to the premium of the second-cheapest Silver plan (Jaffe and Shepard, 2020). Consumers pay premiums net of these subsidies.

Other regulations include budget-neutral risk adjustment (Saltzman, 2021), reinsurance (Kim and Li, 2025), and minimum loss ratio (Cicala, Lieber, and Marone, 2019). For further details, we refer readers to Saltzman (2021) and Tebaldi (2025) and references therein.

4.2 Data

Our primary data source includes administrative records for 2014–2017 obtained directly from Covered California; versions of these data have also been used by Tebaldi, Torgovitsky, and Yang (2023) and Tebaldi (2025). To measure plan horizontal differentiation and then estimate preferences, we augment these data with granular information on plan provider networks.

Provider Networks. Covered California mandates that every qualified health plan provide information on all primary care physicians (PCPs), specialists, hospitals, and other providers they cover. Our data comes from quarterly snapshots of these reports, and we focus on the networks of PCPs and hospitals. See Supplemental Appendix S.1.1 for more details.

We measure plan j's network breadth for individuals in zipcode z as the share of the providers in a catchment area C_z^c within distance $c \in \{0-10 \text{ miles}, 10-20 \text{ miles}\}$ around z covered by plan j. Let PCP_{zjt} be the number of PCPs in zipcode z covered by plan j in year t, and PCP_{zt}^{any} the number of PCPs in z covered by at least one plan in Covered California. Our measure is

$$\mathcal{N}_{jz}^c = \frac{\sum_{z' \in C_z^c} PCP_{z'jt}}{\sum_{z' \in C_z^c} PCP_{z't}^{\text{any}}} .$$

Figure 2 displays the spatial distribution of PCPs in the Santa Clara region that Anthem-HMP, BlueShield-PPO, and Kaiser-HMO cover. Figure 2-a to 2-c 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, $\mathcal{N}_{jz}^{0-10}$ displays substantial variation across plans and within a plan across zipcodes, which is what Figures 2-d to 2-f show. For our purposes, the figure highlights that, ceteris paribus, consumers would prefer BlueShield as the winner of an SSSA in this region, yet this preference is likely stronger in the northwest and

Figure 2: PCP network data: Raw data and network breadth

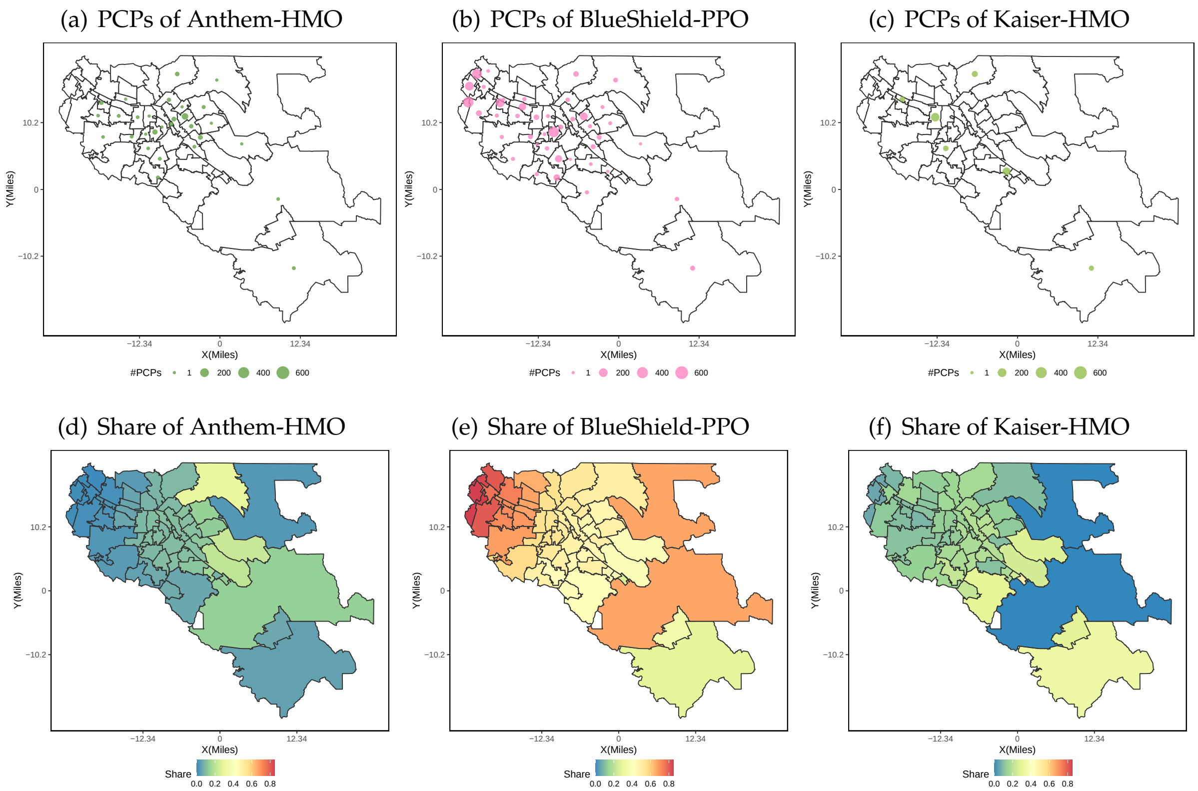

Notes: The top panel shows the spatial distribution of in-network PCPs for Anthem-HMO, BlueShield-PPO, and Kaiser-HMO in Santa Clara in 2015. Dot sizes indicate the number of in-network PCPs in the zipcode. The bottom panel shows a heatmap of the share of in-network PCPs within 10 miles for the same plans. Zipcodes without in-network PCPs are left empty.

less so in the southeast. Table 1-A summarizes our provider network data, showing that plans, on average, cover 25 percent of available PCPs within 10 miles of a zipcode, and that 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 hospital network quality, which we measure as the distance to the closest hospital from zipcode z covered by plan j in year t. Table 1-A 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. Many zipcodes have an in-network hospital for a plan, but a long tail of zipcodes require traveling more than 20 miles to the closest in-network hospital.

Enrollment and Cost. The data include individual and plan characteristics, including individual and household identifiers, age, gender, zipcode of residence, household income as a share of FPL, premiums, and plan enrollment. 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.

Table 1: Descriptive statistics

Variable	(1) N	(2) Mean	(3) SD	(4) p10	(5) p50	(6) p90
A. 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
Share of PCPs in 10-20 miles	33,416	0.27	0.21	0.03	0.22	0.55
Distance to closest hospital	33,416	8.97	11.95	0.00	5.02	21.10
B. 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
Annual 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
C. 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
D. Potential buyers						
Enrolled in ACA plan	9,657,042	0.32	0.47	0.00	0.00	1.00

Notes: Panel A displays summary statistics of plan networks at the insurer-region-zipcode-tier-network-year. Panel B displays summary statistics of the enrollment data at the individual level. Panel C displays summary statistics of plan revenues and costs at the insurer-region-tier-network-year level. Revenue and cost are reported in $/year and weighted by plan enrollment. Panel D provides summary statistics on enrollment by potential buyers.

We provide summary statistics for enrollees in the data in Table 1-B. 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 9 percent chose either Gold or Platinum. Yearly premiums average $1,921, with large variation driven by age-based subsidies.

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

Lastly, while our data provide a detailed characterization of enrollees and plans within Covered California, they do not provide information about consumers outside the marketplace. We use data from the American Community Survey to estimate market size by five-digit zipcode, demographics, and year. We provide further details in the Supplemental Appendix S.1.2. Table 1-D displays summary statistics for these potential buyers. The share of potential enrollees who 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 insurance.

4.3 Evidence on the Main Drivers of Consumer Choice

The regulation of the ACA generates quasi-experiments that we exploit to estimate the causal effects of plan attributes on plan demand. We begin by noting that regulation induces variation in plan premiums and actuarial value for individuals of different demographics within a zipcode, holding all other aspects of the choice environment fixed. Similarly, 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, contracts, and have the same age and income. It is then natural to estimate fixed-effects regressions to isolate these sources of variation to document the impact of plan attributes on plan choice.

To measure the effect of premiums on plan choice, we estimate

$$y_{ijrt} = \beta \text{Annual premium}_{ijrt} + \eta_{jrtz(i)w(i)} + \varepsilon_{ijrt}\ , \tag{11}$$

where y_{ijrt} is an indicator for individual i choosing plan j, Annual premium$_{ijrt}$ is the premium the plan j charges to an enrollee with i's demographics, and $\eta_{jrtz(i)w(i)}$ is a granular set of fixed effects for the interaction of plan, region, year, zipcode, and demographic types $w(i)$ that combine 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 plans that are otherwise identical among narrowly defined demographic groups—actuarial value does not vary within plan and narrow income groups, and provider networks do not vary within plan and zipcode. 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.

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

We use variants of equation (11) to measure the effect of actuarial value and provider network breadth on plan choice. For actuarial value, we focus on Silver plans and isolate the variation from three regulatory discontinuities at 150, 200, and 250 percent of income relative to FPL.[10] 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 accounts for some variation in premiums across demographic groups, while the latter isolates variation in actuarial value from cost-sharing reductions at the regulatory thresholds. Under this specification, plans differ only in actuarial value and are otherwise almost identical. Column 2 in Table 2 shows estimates of this regression. The relationship is positive and strong: a five percentage point increase in actuarial value is associated with a 19 percent decrease in the likelihood it is chosen.

Finally, we focus on the effect of network breadth on plan choice. We estimate another variant of

[10]The actuarial value of Silver plans falls sharply at these discontinuities: from 95 to 88 percent, then to 74 percent, and then further to 70 percent, at income levels relative to FPL of 150, 200, and 250 percent, respectively. See also DeLeire, Chappel, Finegold, and Gee (2017), Lavetti, DeLeire, and Ziebarth (2023), and Tebaldi (2025).

equation (11), using network attributes as independent variables and modifying the fixed effects in two ways. First, by removing the interaction with zipcode. Second, by defining 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 a plan's network while keeping premium, generosity, and other attributes fixed. Column 3 in Table 2 shows that higher network quality significantly increases plan demand. For instance, reducing the distance to the closest in-network hospital by 5 miles increases the likelihood of choosing a plan by 4.9 percent, whereas a 10 percentage point increase in PCP coverage within ten miles of a zipcode does so by 13 percent. 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 regulated competition.

5 Econometric Specification and Estimates

5.1 Demand

Specification. We adopt the pure characteristics demand model (Berry and Pakes, 2007) and tailor it to our granular individual-level data. An individual i is summarized by its type $\theta_i = (w_i, z_i, \alpha_i, \beta_i, \mu_i, \xi_i)$, where w_i collects age and income, z_i indicates zipcode of residence (within a region, which we index by r), and α_i, β_i, μ_i and ξ_i collect preference parameters that affect utility for a plan. In region r and year t, individuals pick one plan j from the set $\mathcal{J}_{rt}$ or remain uninsured.

A plan is characterized by actuarial value $AV_{jrt}^{w_i}$, which varies with income due to cost-sharing reductions, and the collection of network variables $\mathcal{N}_{jrt}^{z_i}$. This includes zipcode-specific network breadth at 0-10 and 10-20 miles radii, and distance to the closest hospital. The premium $p_{jrt}^{w_i}$ varies with income and age as described in Section 4. Moreover, every plan j in rt corresponds to a unique insurer-network combination f_{jrt} (e.g., Anthem PPO in North Los Angeles in 2015); this is the level at which plans may differ for reasons that are unobserved to the econometrician.

The indirect utility that consumer i derives from choosing j in rt is

$$u_{jrt}(\theta_i) = -\alpha_i p_{jrt}^{w_i} + \beta_i^{av} AV_{jrt}^{w_i} + \beta_i^{n} \mathcal{N}_{jrt}^{z_i} + \mu_i + \xi_{i f_{jrt}} \; ; \tag{12}$$

and we specify preference heterogeneity as

$$\alpha_i \sim \text{LogNormal}\left(\overline{\alpha}^{w_i}, \sigma^{\alpha,w_i}\right) \qquad \beta_i^{av} \sim \text{LogNormal}\left(\overline{\beta}^{av,w_i}, \sigma^{av,w_i}\right)$$

$$\beta_i^{n} \sim \text{LogNormal}\left(\overline{\beta}^{n,w_i}, \Sigma^{n,w_i}\right) \qquad (\mu_i, \xi_{if_{jrt}}) \sim \text{Normal}\left(\begin{bmatrix} \overline{\mu}^{w_i} \\ \overline{\xi}^{w_i}_{f_{jrt}} \end{bmatrix}, \begin{bmatrix} \sigma^{\mu,w_i} & 0 \\ 0 & \sigma^{\xi,w_i} \end{bmatrix}\right).$$

We normalize the utility of the outside option of not purchasing insurance to $u_{0rt}(\theta_i) = 0$.

All parameters vary flexibly with age in seven bins: [26-30), [30-35), [35-40), [40-45), [45-50), [50-55), and [55-64]. The term $\xi_{if_{jrt}}$, with mean and variance $\overline{\xi}^{w_i}_{f_{jrt}}$ and σ^{ξ,w_i}, captures unobservables at the insurer-network-region-year level, while the term μ_i is a shock to the preference for insurance. This generates the familiar nesting structure with inside goods in one nest—all the plans in the marketplace—and the outside option in a separate nest. Individuals with particularly low draws of this parameter are hardly attracted by any plan, whereas individuals with particularly high draws of μ_i hardly consider the outside option. 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 multivariate lognormal with diagonal variance-covariance matrix Σ^{n,w_i}.

Given a measure of types $G(\theta_i)$ combining 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

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

where $\sigma_{jrt}(w_i, z_i)$ denotes 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 4 collects, for every individual i, the plan characteristics included in the indirect utility in (12) as well as (Y_i, w_i, z_i). The choice outcome is $Y_i \in \mathcal{J}_{rt}$, with $\Pr[Y_i = j | w_i, z_i, rt] = \sigma_{jrt}(w_i, z_i)$. Identification of the vectors of parameters $(\overline{\alpha}^{w_i}, \sigma^{\alpha,w_i}, \overline{\beta}^{av,w_i}, \sigma^{av,w_i}, \overline{\beta}^{n,w_i}, \Sigma^{n,w_i}, \overline{\mu}^{w_i}, \sigma^{\mu,w_i}, \overline{\xi}^{w_i}, \sigma^{\xi,w_i})$, one for each of the seven age bins listed above, relies on the panel structure discussed in Section 4.3 (c.f. Table 2). This is a special case with sufficient variation to guarantee identification of demand (Berry and Haile, 2024).

The parameters governing the distribution of the premium coefficient $(\overline{\alpha}^{w_i}, \sigma^{\alpha,w_i})$, are identified by 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, as in Tebaldi, Torgovitsky, and Yang (2023). This generates premium variation conditional on actuarial value, zipcode (and thus networks $\mathcal{N}_{rt}^{z_i}$), and $\mathcal{J}_{rt}$. In practice, one can hold fixed all characteristics other than premium, and the main identification assumption is that individuals who have very similar age and income have the same distribution of α_i. 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, 2025).

Preferences over provider 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 $\mathcal{J}_{rt}$ (and thus their unobservable characteristics), premiums, and actuarial value do not vary. The key assumption to identify preferences over 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.

Given the distribution of $(\alpha_i, \beta_i^{av}, \beta_i^n)$, the remaining parameters of the model are the distributions of the constant μ_i and of the unobservable terms collected in ξ_i, with one parameter $\overline{\xi}_{fjrt}^{w_i}$ for each combination of insurer-network-region-year and age bin. These must match plan aggregate shares in every market after accounting for variation in prices. As usual, an insurer f that generates—on average across individuals—a relatively low utility net of ξ_i and yet—within the age bin—registers a large enrollment share in rt relative to its rivals is associated to a larger $\xi_{fjrt}^{w_i}$. Similarly, the mean of μ_i is pinned down by the share of individuals choosing the outside option.

Demand Estimates. We estimate the demand parameters by maximizing the log-likelihood

$$\mathcal{L} = \sum_i \log \left\{ \prod_j \sigma_{jrt}(w_i, z_i)^{1[Y_i = j]} \right\} . \tag{13}$$

Appendix B complements our previous discussion in Section 3 and presents our estimation steps in detail. Importantly, the variation in w_i and z_i, combined with the extent to which—due to ACA regulations—premiums vary more granularly than the unobservable term $\overline{\xi}^{w_i}$, motivate us to maximize $\mathcal{L}$ directly rather than estimating the model via method of moments.

Table 3-A displays extensive margin demand responses to changes in plan attributes, calculated

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

	All	Age below 45		Age above 45	
		Low income	High income	Low income	High income
A. Δ% in market enrollment					
Δ^{+}\$10/month in premium	-2.23	-2.57	-2.67	-1.77	-2.30
Δ^{+}0.05 in PCP share in 0-10 miles	1.29	1.11	1.33	1.18	1.64
Δ^{+}0.05 in PCP share in 10-20 miles	0.70	0.81	0.94	0.54	0.66
Δ^{-}2 miles to closest hospital	0.79	0.99	1.15	0.56	0.69
B. Willingness to pay (\$/year)					
Δ^{+}5 p.p in actuarial value	287.7	193.89	218.46	380.88	386.85
Δ^{+}0.05 in PCP share in 0-10 miles	69.85	53.08	59.92	85.21	86.22
Δ^{+}0.05 in PCP share in 10-20 miles	38.42	38.08	42.57	36.57	36.26
Δ^{-}2 miles to closest hospital	45.43	48.20	54.23	38.98	39.00

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 subpopulations of potential buyers according to age around 45 years, and income around 250 percent of FPL. Each row indicates a change to plan attributes.

as the change in the probability of enrolling in the marketplace in response to a change in the characteristics of all plans. A few patterns that are worth noting. First, unsurprisingly, individuals prefer lower premiums: Increasing all plans' premiums by \$10/month decreases marketplace enrollment by 2.2 percent, with stronger responses from younger individuals. Second, there are strong preferences for provider networks: Individuals prefer plans with broader physician networks and closer in-network hospitals. Increasing all plans' share of PCPs covered within 0-10 miles of an individual's zipcode by five percentage points raises enrollment by 1.3 percent. Decreasing the distance to the closest in-network hospital by two miles—an abstract exercise but meaningful to measure preferences for hospital distance—increases enrollment by 0.8 percent.

Using these estimates, we obtain the distribution of willingness to pay for plan attribute k, which is β_i^k/α_i. Table 3-B summarizes the results, and there are several patterns worth noting. First, consumers feature a high willingness to pay for plan generosity, as the average consumer is willing to pay \$287/year for an increase in actuarial value of five percentage points, or around 15 percent of the average annual paid premium. Willingness to pay for plan generosity is stronger among higher-income and older consumers. Second, consumers display a substantial willingness to pay for provider networks, as the average consumer is willing to pay \$70/year to increase the share of PCPs within 10 miles that a plan covers by five percentage points and \$38/year to have the closest in-network hospital two miles closer to their zipcode.

Figure 3: Distribution of willingness to pay conditional on age

(a) Willingness to pay for actuarial value

(b) Willingness to pay for network breadth

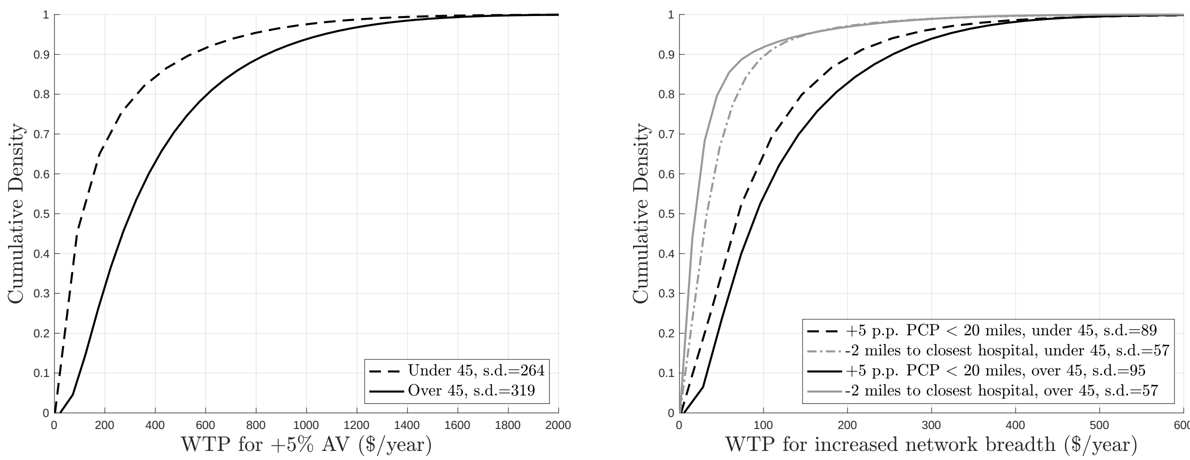

Notes: These figures display CDFs of willingness to pay for a change in a plan attribute based on our demand estimates, separately for individuals below age 45 (dashed) and above age 45 (solid).

Beyond these population averages, Figure 3 shows that there is substantial willingness to pay for plan attributes. For instance, the willingness to pay for a five percentage point increase in actuarial value is more than twice as much at the 75*th* than the 25*th* percentile. Similarly, the willingness to pay for broader networks is around three times higher at the 75*th* than the 25*th* percentile. This vast preference heterogeneity—particularly along horizontal attributes—makes variety a more relevant feature to consider when designing a market, which will play a key role in our comparison between regulated competition and procurement auctions below.

Table 4 displays a summary of the substitution patterns implied by our demand model. Our estimates suggest that demand for a specific plan is highly responsive: changing premiums or provider networks leads to large changes in own plan demand. For example, if any of the three main insurers (Anthem, Blue Shield, and Kaiser) increased the monthly rate by $10 for either of their Bronze or Silver plans, we estimate that demand for the plan would fall by between 16 and 33 percent. While the insurer would recapture some of that demand in their other metal tiers—particularly in Bronze plans when increasing Silver plans' premiums—many of their enrollees would switch to 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 a feature of the demand model with observed and unobserved heterogeneity that we estimate.

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	-16.28	3.98	0.96	1.55	0.76	0.10	1.37	0.02	0.12	1.82	0.64	0.36
	Silver	7.41	-21.66	2.22	0.15	3.45	0.07	0.24	1.14	0.13	0.42	4.05	0.48
Blue Shield	Bronze	2.29	0.14	0.00	-33.02	21.02	0.18	1.55	0.04	0.00	1.46	0.73	0.10
	Silver	0.61	2.43	0.04	12.84	-25.98	3.76	0.32	1.09	0.26	0.15	2.22	0.63
Kaiser	Bronze	1.63	0.20	0.00	1.25	0.41	0.02	-15.77	3.26	2.15	1.12	0.78	0.27
	Silver	0.16	2.66	0.00	0.10	2.73	0.00	8.38	-19.36	0.97	0.05	2.89	0.00
		B. Δ^{+} 0.05 in PCP share in 0-10 miles											
Anthem	Bronze	8.31	-2.35	-0.47	-0.91	-0.43	-0.05	-0.60	-0.03	-0.05	-0.60	-0.34	-0.19
	Silver	-4.04	9.43	-1.05	-0.09	-1.41	-0.05	-0.10	-0.39	-0.07	-0.24	-0.91	-0.24
Blue Shield	Bronze	-1.40	-0.09	-0.00	14.53	-7.76	-0.60	-0.74	-0.03	-0.01	-0.58	-0.40	-0.10
	Silver	-0.43	-0.89	-0.03	-7.89	13.96	-1.86	-0.20	-0.36	-0.12	-0.10	-0.87	-0.33
Kaiser	Bronze	-0.72	-0.10	-0.00	-0.60	-0.20	-0.02	6.32	-1.23	-0.91	-0.38	-0.29	-0.11
	Silver	-0.09	-0.95	-0.00	-0.05	-1.09	-0.00	-4.33	9.49	-1.45	-0.04	-0.93	-0.02
		C. Δ^{-} 2 miles in distance to closest hospital											
Anthem	Bronze	4.26	-1.18	-0.23	-0.38	-0.22	-0.02	-0.29	-0.01	-0.03	-0.44	-0.14	-0.09
	Silver	-1.90	5.10	-0.60	-0.03	-0.86	-0.03	-0.05	-0.22	-0.03	-0.13	-0.58	-0.13
Blue Shield	Bronze	-0.60	-0.04	-0.00	6.86	-3.64	-0.26	-0.34	-0.01	-0.00	-0.34	-0.16	-0.03
	Silver	-0.19	-0.50	-0.02	-3.31	6.51	-0.94	-0.09	-0.18	-0.07	-0.06	-0.45	-0.15
Kaiser	Bronze	-0.35	-0.05	-0.00	-0.27	-0.11	-0.01	3.62	-0.71	-0.56	-0.21	-0.14	-0.05
	Silver	-0.04	-0.56	-0.00	-0.01	-0.61	-0.00	-2.62	5.75	-0.92	-0.01	-0.58	-0.01

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 Br, Sil, and G/Pl indicate Bronze, Silver, and Gold/Platinum 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.

5.2 Cost

Cost Functions and Insurer FOCs. When an individual of type θ_i enrolls in plan j in region-year rt the insurer expects costs equal to

$$c_{jrt}(\theta_i, \mathcal{N}_{jrt}, \overline{\xi}_{f_{jrt}}) = \underbrace{\kappa_{jrt}(\mathcal{N}_{jrt}, \overline{\xi}_{f_{jrt}}) \cdot AV_j}_{\text{Common across individuals}} \cdot \underbrace{\exp[\rho^{\text{age}}\text{age}_i + \rho^{AV} WTP^{AV}(\theta_i) + \rho^{\text{net}} WTP^{\text{net}}(\theta_i)]}_{\text{Expected medical spending, varying across individuals}}, \quad (14)$$

where, to capture selection, the log-annual expected medical spending is a linear function of age, willingness to pay for actuarial value, and willingness to pay for network breadth. The parameters in ρ capture the strength of selection. The term AV_j indicates the actuarial value of plan j, which in Covered California is fixed at 60 for Bronze plans, 70 for Silver plans, 80 for Gold plans, and 90 for Platinum plans.[11] For given $\mathcal{N}^\star$ and $\overline{\boldsymbol{\xi}}^\star$, we model the common component of costs as

$$\kappa_{jrt}(\mathcal{N}_{jrt}, \overline{\xi}_{f_{jrt}}) = \overline{\kappa}_{jrt} \exp\left[(\boldsymbol{\eta}^n + \lambda^n_{jrt})(\mathcal{N}_{jrt} - \mathcal{N}^\star_{jrt}) + (\boldsymbol{\eta}^\xi + \lambda^\xi_{jrt})(\overline{\xi}_{f_{jrt}} - \overline{\xi}^\star_{f_{jrt}}) \right] ,$$

where $\overline{\boldsymbol{\kappa}}$ and $\boldsymbol{\eta} = (\boldsymbol{\eta}^n, \boldsymbol{\eta}^\xi)$ are the parameter vectors to be estimated, while $(\lambda^n_{jrt}, \lambda^\xi_{jrt})$ are iid mean zero error terms capturing unmodeled determinants of network and brand adjustment costs. The expected average cost, of which we observe a noisy realization, is

$$AC_{jrt} = Q^{-1}_{jrt} \int c_{jrt}(\theta_i, \mathcal{N}_{jrt}, \overline{\xi}_{f_{jrt}})\, \mathbf{1}\left[j = \arg\max_k u_{krt}(\theta_i) \right] dG(\theta_i) . \tag{15}$$

Assuming that the observed data come from a Nash equilibrium adds to equation (15) two sets of moments corresponding to first-order conditions with respect to rates and characteristics. In terms of pricing, according to federal rating rules, insurers must pick a scalar-valued base rate τ_{jrt} for every plan j in every region-year rt. This is then mapped into premiums by age and income by applying age adjustments and subsidies. With these rating rules operating in the background, at the equilibrium rates $\boldsymbol{\tau}^\star_{rt}$, equalizing marginal revenue and marginal cost gives

$$\frac{\partial}{\partial \tau_{jrt}} \sum_{k \in \mathcal{J}_{f_j}} R_{krt}(\boldsymbol{\tau}^\star_{rt}) = \sum_{k \in \mathcal{J}_{f_j}} AC_{krt}(\boldsymbol{\tau}^\star_{rt}) \frac{\partial Q_{krt}(\boldsymbol{\tau}^\star_{rt})}{\partial \tau_{jrt}} + Q_{krt}(\boldsymbol{\tau}^\star_{rt}) \frac{\partial AC_{krt}(\boldsymbol{\tau}^\star_{rt})}{\partial \tau_{jrt}} , \tag{16}$$

where—emphasizing the dependence on base rates—total revenue is denoted by $R_{krt}(\boldsymbol{\tau}_{rt})$, and $\mathcal{J}_{f_j}$ collects the plans offered by the insurer selling j. The right-hand side in (16) is the sum of the marginal cost across these plans. If $\boldsymbol{\rho} \neq 0$, so that expected costs vary across individual types, $\frac{\partial AC_{krt}(\boldsymbol{\tau}^\star_{rt})}{\partial \tau_{jrt}} \neq 0$, as it is typically the case in a selection market.

The first-order conditions with respect to network breadth and brand quality imply

$$\eta^n + \sum_{j \in \mathcal{J}_f} \lambda^n_{jrt} = \left(\sum_{j \in \mathcal{J}_f} \frac{\partial R^\star_{jrt}}{\partial \mathcal{N}_{f_{jrt}}} - \frac{\partial Q^\star_{jrt}}{\partial \mathcal{N}_{f_{jrt}}} AC^\star_{jrt} - Q^\star_{jrt} \frac{\partial (AC^\star_{jrt}/\overline{\kappa}_{jrt})}{\partial \mathcal{N}_{f_{jrt}}} \right) \left(\sum_{j \in \mathcal{J}_f} Q^\star_{jrt} AC^\star_{jrt} \right)^{-1} , \tag{17}$$

where we use the star superscript to highlight that all functions are evaluated at the observed rates, networks, and estimated brand quality. The first-order conditions with respect to ξ are analogous.

[11] This differs from the variable $AV_j^{w_i}$ entering the utility specification in (12) because of cost-sharing subsidies.

Figure 4: Identification of selection parameters ρ

(a) Fit of AC setting $\rho = 0$ (b) Fit of AC at estimated $\widehat{\rho}$

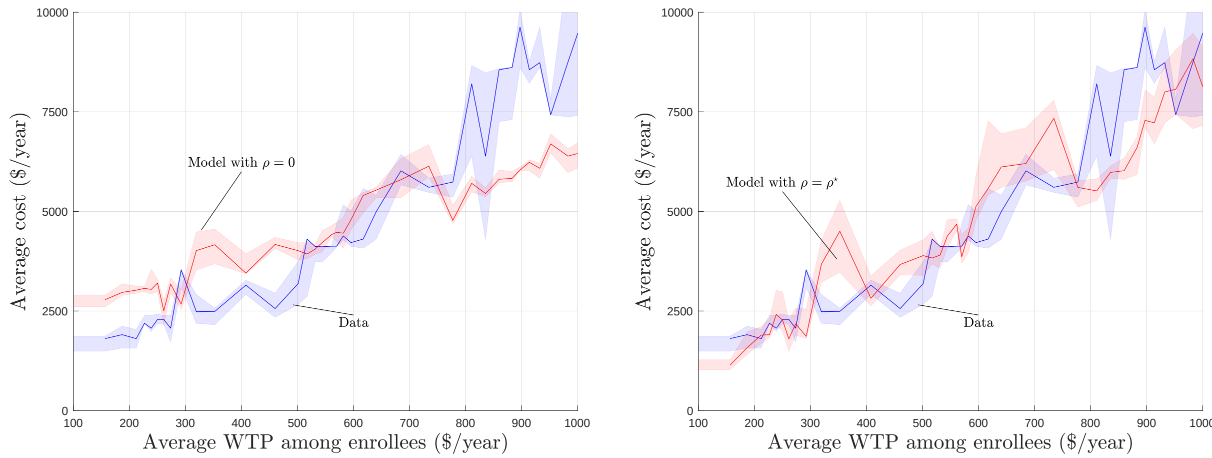

Notes: This figure displays the relationship between average WTP for plan actuarial value (obtained from the demand model) and enrollee average costs for different values of ρ. In both figures, the blue line approximates this relationship using observed average costs, whereas for the red line, we replace the observed average cost with the predicted value obtained assuming equilibrium pricing. Panel (a) sets $\rho = 0$, hence assuming no cost differences across individuals. Panel (b) sets ρ to its estimated value (see Table 5).

Identification. The parameters of the cost function in (14) are identified as detailed in the constructive argument in Section 3 and Appendix C. To summarize, from the demand estimates, we obtain marginal revenue for each plan, which equals marginal cost under the assumption of equilibrium pricing. Since we observe average cost, differences between marginal and average cost can be projected on the composition of average and marginal buyers across plans to identify the parameters that describe cost heterogeneity across individuals. The key identifying variation amounts to differential selection of types across plans, induced by differences in the composition of potential buyers across markets, and by the number and characteristics of insurers in each market. The argument is illustrated in Figure 4. For $\rho = 0$, Figure 4-a shows that the model would systematically overpredict (underpredict) average cost for plans whose enrollees have relatively low (high) willingness to pay for actuarial value. Instead, at the estimated $\hat{\rho}$, Figure 4-b shows that the relationship between the enrollees' composition and average cost predicted by the model tracks closely the one obtained by combining the data and the demand estimates.

Identification of the cost of adjusting networks and brand quality is more straightforward. From the demand model, we compute the revenue changes these adjustments would generate, which in equilibrium must equal the corresponding cost changes. Netting out cost changes from shifts in the types selecting plans, residual cost changes identify η as the average of the right-hand side of equation (17).

Figure 5: Relationship between cost and WTP

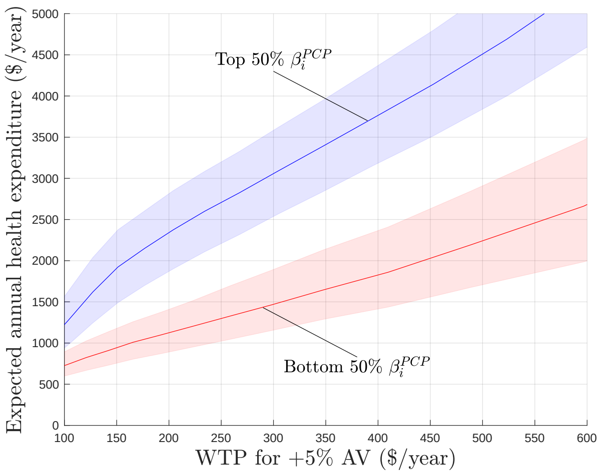

Table 5: Cost parameters

	Δ% in expected cost
Selection parameters (ρ)	
+10 years of age	+73.30%
+ $100 WTP^{AV}	+3.76%
+ $100 WTP^{net}	+22.22%
Cost of quality provision (η)	
+10 p.p. in PCP share in 0-10 miles	+2.98%
-5 miles to closest hospital	+3.33%
$100/year brand quality $\overline{\xi}$	+3.67%

Notes: This figure displays the relationship between willingness to pay for plan actuarial value and enrollee average costs, separately for individuals above and below the median of preferences over PCP coverage. The table displays expected responses to changes to changes in consumer age and willingness to pay for plan attributes (driven by ρ), as well as to changes in plan attributes (driven by η).

Cost Estimates. Figure 5 and Table 5 illustrate the resulting cost estimates, which result from the estimation steps that follow the identification argument. See Appendix C for details.

Our estimates of ρ provide evidence for adverse selection. Table 5 shows how costs vary with individual characteristics. A 10-year increase in age leads to a 73 percent increase in expected medical spending, while a $100 increase in annual willingness to pay for five percentage points in actuarial value raises cost by 3.8 percent. Lastly, a $100 increase in willingness to pay for networks—measured as the willingness to pay to increase PCP coverage within 20 miles by five percentage points and reduce the distance to the closest hospital by two miles—increases expected cost by 22 percent. For η, we estimate that increasing PCP coverage by 10 percentage points or decreasing the distance to the closest hospital by two miles increases expected cost by approximately three percent. Moreover, providing an additional $100 per year in average brand quality increases expected cost by 3.7 percent.

Figure 5 illustrates the relationship between expected cost, willingness to pay for actuarial value, and willingness to pay for network breadth. Higher willingness to pay for actuarial value leads to higher costs, leading to adverse selection. Moreover, individuals who value broader provider networks have higher expected costs and a steeper relationship between cost and willingness to pay for actuarial value, as in Shepard (2022) and Serna (2023, 2024). These properties of the cost functions suggest that adverse selection is a quantitatively relevant factor in comparing regulated competition and procurement auctions in our context.

6 Procurement Auctions vs. ACA Marketplaces

6.1 Counterfactual Setup

We compare the observed equilibrium in the Covered California marketplace to SSSAs. Since our estimator of the cost function involves inverting first-order conditions, our baseline of regulated competition is equivalent to the outcomes observed in the Covered California data.

Simulating Second-Score Scoring Auctions. We consider auctions in which the winner offers four plans—one per metal tier—and individuals can choose to purchase a plan or remain uninsured. To focus on the comparison between oligopoly and SSSAs, we keep the age-adjustment regulation and subsidies fixed. Since government spending critically determines outcomes, we want to avoid results to reflect large swings in spending rather than changes in market design. For this, we add a transfer of T dollars per enrollee paid directly to the winning insurer, calibrated to match observed spending to spending under auctions with fixed plan quality. To simplify the insurer's pricing problem, we fix the ratio of premiums across metal tiers to reflect their actuarial value; for example, the ratio between pre-subsidy Bronze and Silver plan premiums is equal to 0.6/0.7. Therefore, every insurer bid consists of a scalar premium bid and deviations from the status quo in terms of network breadth and brand quality.

To compute outcomes under SSSAs, we begin by using the demand and cost estimates to compute the break-even combination of premiums and quality for every insurer in every region-year market. Formally, for every insurer f in region-year rt, we consider a grid of changes in network breadth and brand quality. For every $(\mathcal{N}, \overline{\xi})$ in this grid, we compute

$$\Pi^a_{frt}(b, \mathcal{N}, \overline{\xi}) = \sum_{j \in \mathcal{J}_f} \int \underbrace{\sigma^a_{jrt}(b, \mathcal{N}, \overline{\xi}, \theta_i)}_{\text{Probability } i \text{ chooses } j} \Big(\underbrace{b \times \text{Age factor}(w_i) \times AV_j}_{\text{Revenue when } i \text{ chooses } j} - \underbrace{c_{jrt}(\theta_i, \mathcal{N}, \overline{\xi})}_{\text{Cost when } i \text{ chooses } j} + T \Big) \, dG_{rt}(\theta_i) \, ,$$

varying with the scalar bid b.[12] Then, we compute $\overline{b}_f(\mathcal{N}, \overline{\xi})$ by solving $\Pi^a_f(\overline{b}_f(\mathcal{N}, \overline{\xi}), \mathcal{N}, \overline{\xi}) = 0$. For any scoring rule, insurer f picks a bid on $\overline{b}_f(\mathcal{N}, \overline{\xi})$.

We consider the class of SSSAs with scoring rules defined by

$$\Psi(b, \mathcal{N}, \overline{\xi}) = \psi \times CS^a(0, \mathcal{N}, \overline{\xi}) - b \, ,$$

[12] The function $\sigma^a_{jrt}(b, \mathcal{N}, \overline{\xi}, \theta_i)$ is the share of individuals of type θ_i purchasing j when only plans offered by insurer f are available, pre-subsidy premiums are $b \times$ Age Factor$(w_i) \times AV_j$, subsidies are held fixed to the ACA value for the same individual, and networks and mean brand quality are set to $\mathcal{N}$ and $\overline{\xi}$.

where $\psi \geq 0$, and $CS^a(b, \mathcal{N}, \overline{\xi})$—a scalar measure of plan quality—is the consumer surplus generated when the plans available are offered by a winner setting rate b, network $\mathcal{N}$, and average brand quality $\overline{\xi}$. Given a value of ψ, insurer f's optimal bid is obtained by maximizing $\Psi(\overline{b}_f(\mathcal{N}, \overline{\xi}), \mathcal{N}, \overline{\xi})$ over the grid of network and brand quality values, and setting $b = \overline{b}_f(\mathcal{N}, \overline{\xi})$. The winner w is the firm f with the highest score. The final solution $(b^a, \mathcal{N}^a, \overline{\xi}^a)$ maximizes $\Pi^a_{wrt}(b, \mathcal{N}, \overline{\xi})$ subject to $\Psi(b, \mathcal{N}, \overline{\xi})$ being weakly larger than the second-highest score, as stated in equation (2).

To study how the design of scoring rules in SSSAs impacts quality provision incentives and market outcomes, we simulate SSSAs both under endogenous quality provision and with fixed quality, whereby insurers are constrained to offer the same quality as under regulated competition.

An Example: San Francisco 2016. To illustrate how a shift from regulated competition to procurement auctions works in our framework, we focus on the San Francisco market in 2016 as an example, indexing firms by $f \in \{A, B, C, D, E\}$. Figure 6 shows the results, displaying the monthly rate—prior to age adjustments and subsidies, averaged across metal tiers—on the horizontal axis, and average consumer surplus if plans were free—$CS^a(0, \mathcal{N}, \overline{\xi})$—-on the vertical axis of each plot.

We begin by comparing the outcomes under regulated competition to insurers' bids with fixed networks and brand quality. These bids b_f solve $\Pi^a_f(b_f, \mathcal{N}^\star_f, \overline{\xi}^\star_f) = 0$ for each f. Figure 6-a shows that all bids are lower than the average rates under regulated competition, despite offering the same quality. For endogenous quality, Figure 6-b shows the break-even price-quality combinations $\overline{\delta}_f(b)$ for each insurer f—the empirical analog of Figure 1. For any scoring rule, each insurer's bid is the price-quality pair on $\overline{\delta}_f(b)$ that yields the highest score, which we can compute.

Figures 6-c and 6-d show the outcomes of an SSSA with $\psi = 0.02$, for fixed and endogenous quality, respectively. With near-zero ψ, this scoring rule places almost no weight on quality, making iso-score curves almost vertical. The winner is the insurer with the lowest rate, $w = C$, and the runner-up is insurer E. Insurer C can modify its plans to maximize profits as long as the outcome is to the left of the runner-up's iso-score curve. With fixed quality, C's final premium is roughly equal to E's bid, and consumer surplus is lower than under regulated competition. With endogenous quality, the performance of this SSSA worsens: Insurers are incentivized to offer low quality, and average consumer surplus is \$40 per person-year lower than under regulated competition.

This pattern changes when the scoring rule rewards quality. Figures 6-e and 6-f show results for an SSSA with $\psi = 2$, under which iso-score curves are upward-sloping. With fixed quality, insurer E wins the auction, and C is the runner-up, resulting in average consumer surplus nearly identical to the status quo. With endogenous quality, all insurers increase quality under this higher

Figure 6: Counterfactual SSSAs: San Francisco 2016 as an example

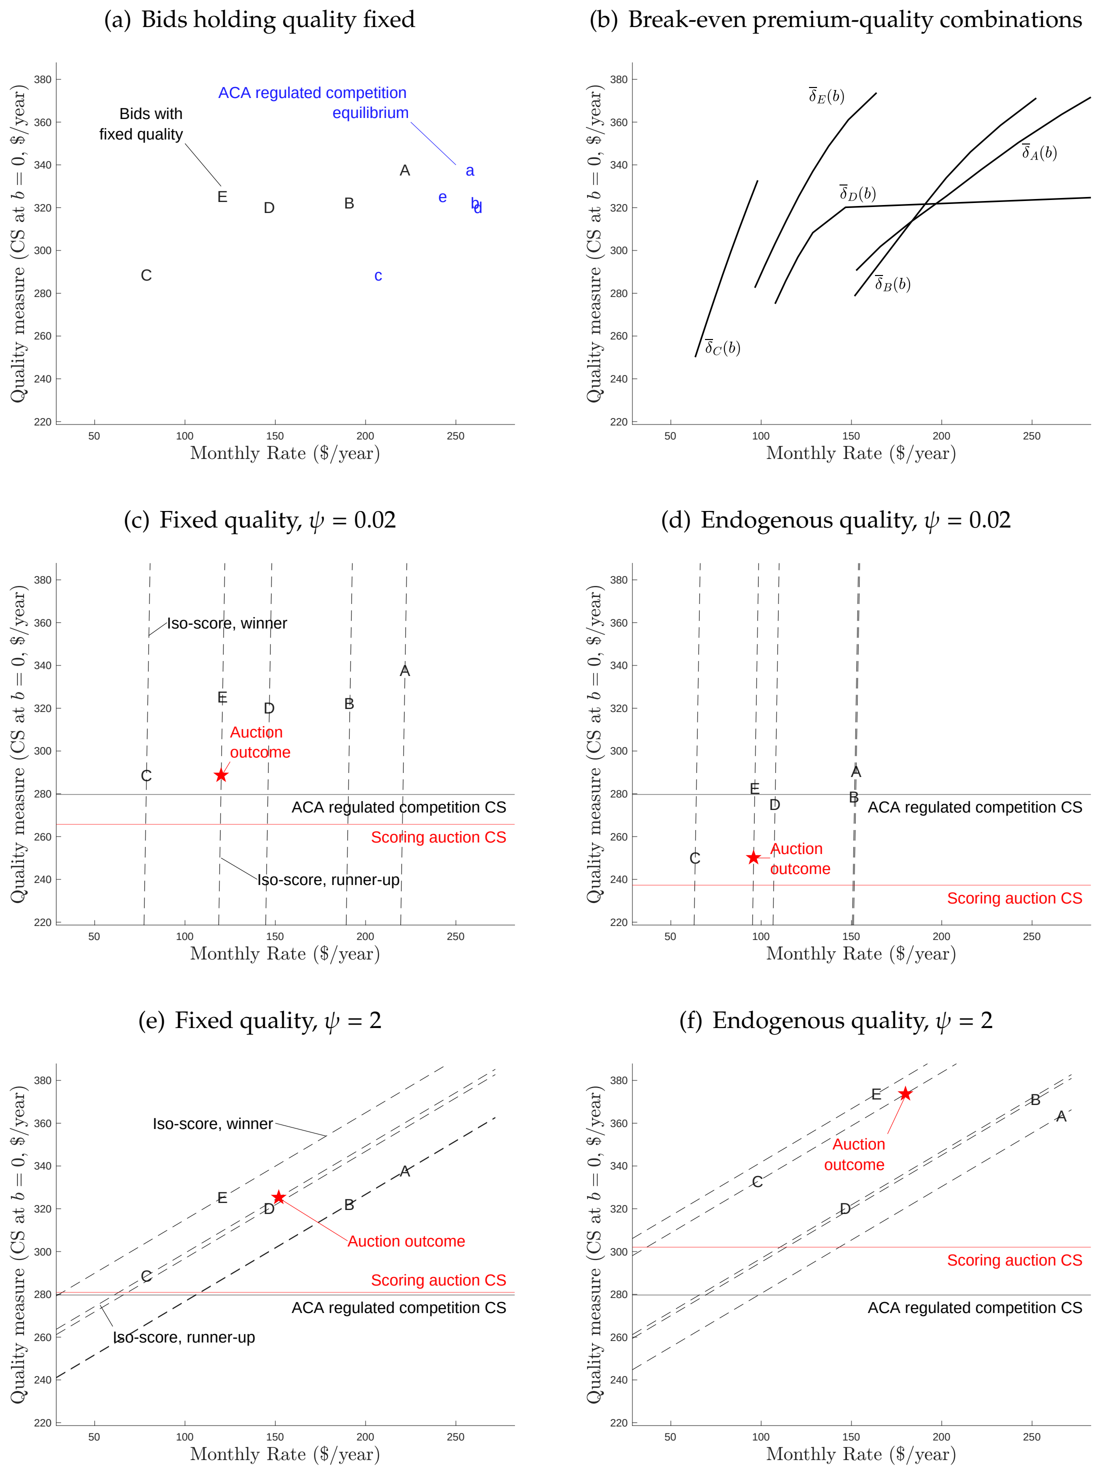

Notes: The figure illustrates a shift from regulated competition to procurement auctions in the San Francisco region in 2016. The x-axis shows the monthly rate—prior to age adjustments and subsidies, averaged across metal tiers—and the y-axis shows average consumer surplus evaluated at a monthly rate of zero. Panel (a) shows regulated competition outcomes (lowercase, blue) and insurers' bids under fixed quality (uppercase, black). Panel (b) shows the break-even price-quality combinations for each insurer. Panels (c) and (d) show results for a scoring rule with a low weight on consumer surplus (ψ = 0.02). Panels (e) and (f) show results for a scoring rule with a higher weight on consumer surplus (ψ = 0.2). Finally, panels (c) and (e) show results for fixed quality, whereas panels (d) and (f) show results for endogenous quality. In panels (c)–(f), dashed lines indicate iso-score curves, the red star indicates the SSSA outcome, and the solid lines indicate consumer surplus under regulated competition (black) and SSSA (red).

ψ, raising average consumer surplus by $22 per person-year relative to regulated competition.

6.2 Aggregate Results

We compute counterfactual outcomes for every region-year for SSSAs with a range of scoring rules $\psi \in [0, 2]$. Figure 7 shows how consumer surplus, enrollment, average cost, markups, and government spending vary with ψ for the cases of fixed and endogenous quality.[13] We focus the discussion of results on the more relevant case of endogenous quality.

SSSAs generally increase average consumer surplus relative to regulated competition, as shown by Figures 7-a and 7-b. For scoring rules rewarding low premiums ($\psi = 0$), adopting SSSAs would have a modest impact of $50 per person-year, with average consumer surplus decreasing in nearly half of the markets. However, as scoring rules place more weight on quality, consumer surplus gains increase, reaching nearly $250 per person-year at $\psi = 0.95$, or 13 percent of the annual premium under regulated competition (cf. Table 1). Moreover, average consumer surplus would increase in 87 percent of markets. Further increasing the weight the scoring rule places on quality worsens outcomes, reflecting premium-sensitivity and auctions becoming less competitive.

Most changes in consumer surplus stem from how changes in premiums and quality affect enrollees, whereas the participation margin—and overall insurance rate—do not vary much between regulated competition and procurement. Figure 7-c shows that SSSAs with a scoring rule placing no weight on quality would lower enrollment by three percentage points relative to baseline, whereas a scoring rule with a weight on quality of $\psi = 0.95$ would increase overall enrollment by slightly less than half a percentage point. The pattern of impacts on enrollment mirrors that on average consumer surplus. The lack of impact on participation is due to the relatively low enrollment in the marketplace (32 percent), despite the large premium subsidies offered by the ACA. As a result, our demand system features a high degree of nesting along the participation margin, leading to small effects on total enrollment in our simulations.

Changes in average cost are an important driver of the difference between SSSAs and regulated competition, driven by changes in quality and the selection of individuals across plans. As Figure 7-d illustrates, these changes are sizable. An SSSA with $\psi = 0$ leads to a decrease in average cost of $1,200—around 30 percent lower than under the regulated competition baseline (cf. Table 1). Increasing ψ leads to higher quality and smaller cost reductions, but, at $\psi = 0.95$, average cost

[13] We calibrate the per-enrollee transfer amount T to ensure that the change in government spending between regulated competition and SSSAs with fixed quality is approximately zero; this is evident from Figure 7-f.

Figure 7: Average effects across markets and scoring rules

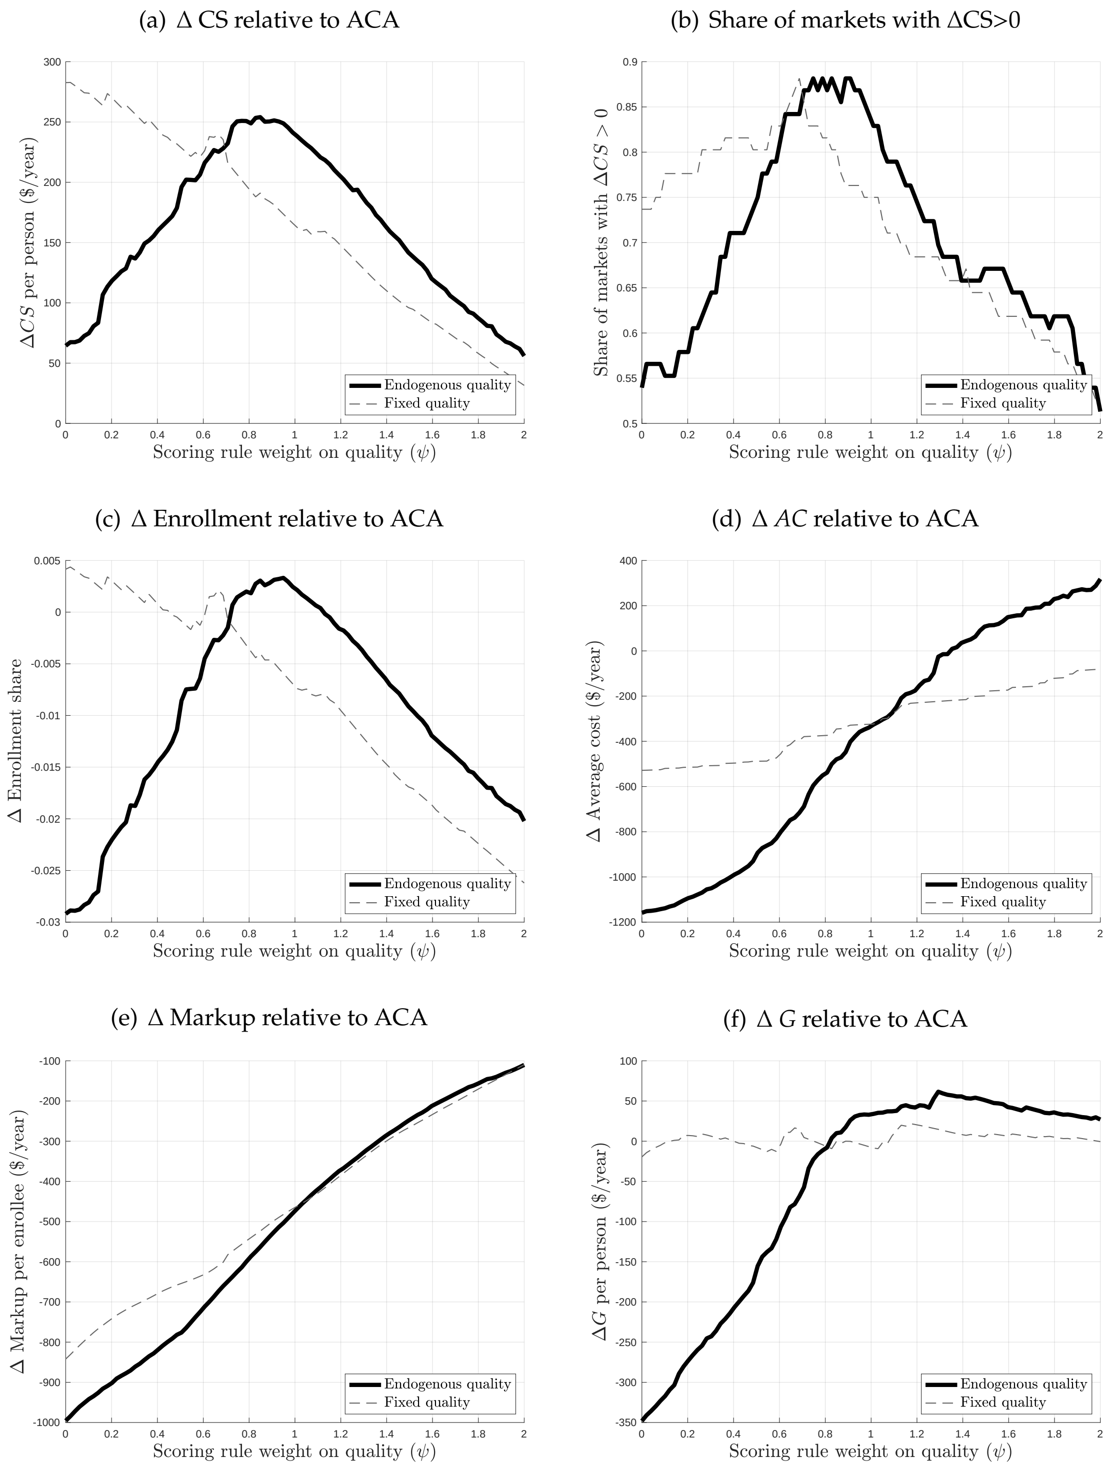

Notes: This figure shows the impacts of a shift from regulated competition to SSSAs. Each panel reports results for a specific outcome, as highlighted in the corresponding labels. The x-axis indicates the weight the scoring rule places on consumer surplus, $\psi \in [0, 2]$. Each figure reports results for the cases with endogenous quality and fixed quality. Averages across the 76 region-year markets.

Figure 8: Quality provision and auction competitiveness

(a) Δ Quality: bid and SSSA outcome

(b) Monthly rate: bid and SSSA outcome

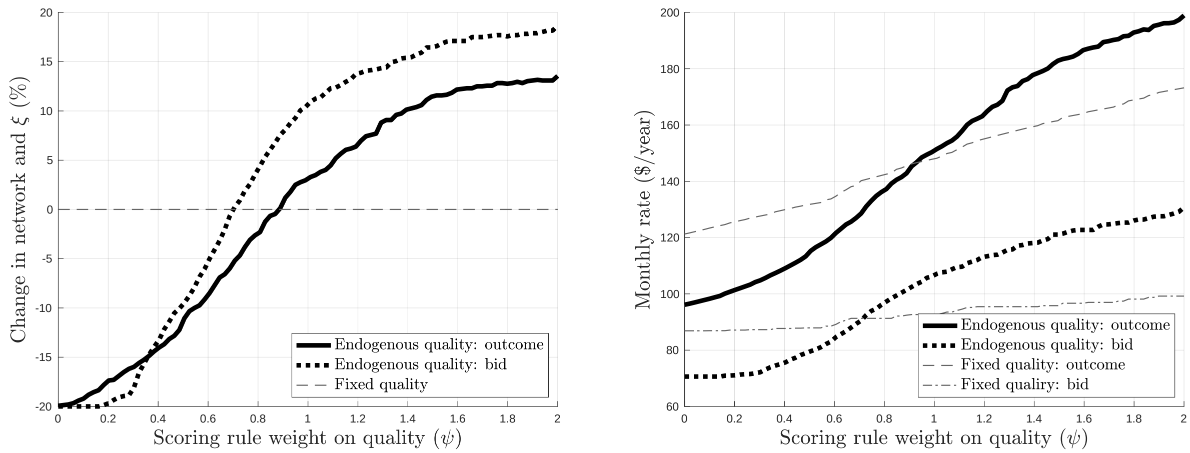

Notes: This figure shows the impacts of a shift from regulated competition to SSSAs on quality and monthly rates. The x-axis indicates the weight the scoring rule places on consumer surplus, $\psi \in [0, 2]$. Each figure reports results for the bid and outcome of the winning insurer, separately for the cases with endogenous quality and fixed quality. Averages across the 76 region-year markets.

reductions are nevertheless still as large as 10 percent ($400) of the average cost under regulated competition. By enhancing competition for the market, SSSAs also decrease markups, as shown by Figure 7-e. The winner's average markup is lower than the average markup under the ACA for all values of ψ. At the consumer-optimal value of $\psi = 0.95$, markup reductions are larger than average cost reductions at $500, or 26 percent of the average premium paid under the ACA.

Lastly, Figure 7-f plots changes in government spending. At low values of ψ, SSSAs reduce spending, while at $\psi = 0.95$ SSSAs increase spending, although by less than $50 per person-year—the consumer surplus gains in Figure 7-a are four times larger. Net of spending changes, these average gains are $200 per person-year, adding up to nearly $2.4 billion during our study period.

Quality Provision and Auction Competitiveness. Two key channels through which an SSSA's scoring rule determines outcomes are its implied incentives for quality provision and its influence on competition in the auction. The gap between the first- and second-highest scores in an SSSA provides a measure of competitiveness, and is a key driver of the adjustments in price and quality from the winning bid to the final auction outcome. To explore this, Figure 8 plots the changes in quality relative to regulated competition and the monthly rates—which translate into premiums—at the winning bid and at the final auction outcome, for a range of scoring rules $\psi \in [0, 2]$.

For low values of ψ, insurers bid by lowering quality relative to the ACA status quo. For scoring rules $\psi < 0.35$, insurers lower quality even more than what they would if they were a monopolist—

after winning, the insurer increases quality relative to its bid. To win the auction, they must lower rates as much as possible, which triggers a race to the bottom in quality provision. Lower quality implies lower costs, which in the bidding stage are passed through to lower premiums, in turn making the market more attractive for low-cost types who value quality less. These auctions are very competitive, as reflected by the small differences in quality between winning bids and final outcomes. Moreover, in terms of monthly rates, the spread between winning bids and final outcome is around \$30 per month. Such low-ψ SSSAs increase consumer surplus because the resulting reductions in costs and markups are larger than the losses implied by lower quality.

Focusing on scoring rules near the consumer-optimal value ($\psi = 0.95$), we find that SSSAs differ substantially from the status quo. Insurer bids offer higher quality, while greater emphasis on both vertical and horizontal plan differentiation weakens competition in the auctions, as reflected by a wider gap between the solid and dotted lines in Figure 8-b. On average, the winning bid implies increases in network breadth and brand quality of approximately 10 percent relative to regulated competition. After winning, insurers ultimately deliver more modest quality improvements, around 3 percent above the status quo. The gap between bid and final rates also widens from \$30 to about \$45 per month. Higher quality leads to higher costs, and weaker auction competition translates into higher markups. Nevertheless, consumers value the quality improvements enough that auctions with ψ near 0.95 yield the highest consumer surplus.

For scoring rules with higher weights on quality ($\psi > 1$), outcomes worsen: while insurers are incentivized to further improve quality—and indeed do so—costs escalate and eventually exceed those under the status quo. These increases reflect both the cost of higher quality and the compounding effects of adverse selection: higher costs lead to higher premiums, worsening risk selection and further raising costs. Since willingness to pay for coverage and quality is strongly correlated with cost, scoring rules placing excessive weight on quality can trigger adverse selection spirals—even under SSSAs. This tension between quality provision, competition, and adverse selection is reflected in the inverted-U shape of consumer surplus and enrollment in Figure 7.

6.3 Mechanisms: Variety, Quality, Cost and Markup Effects

In Section 2, we introduced a decomposition of the full effect of switching from regulated competition to SSSAs into four key economic forces: variety, quality, cost (or selection), and markup effects. In Figure 9, we present results for such decomposition for SSSAs with endogenous quality provision under two benchmark scoring rules: one that does not reward quality ($\psi = 0$), and the

Figure 9: Channels through which procurement auctions and regulated competition differ

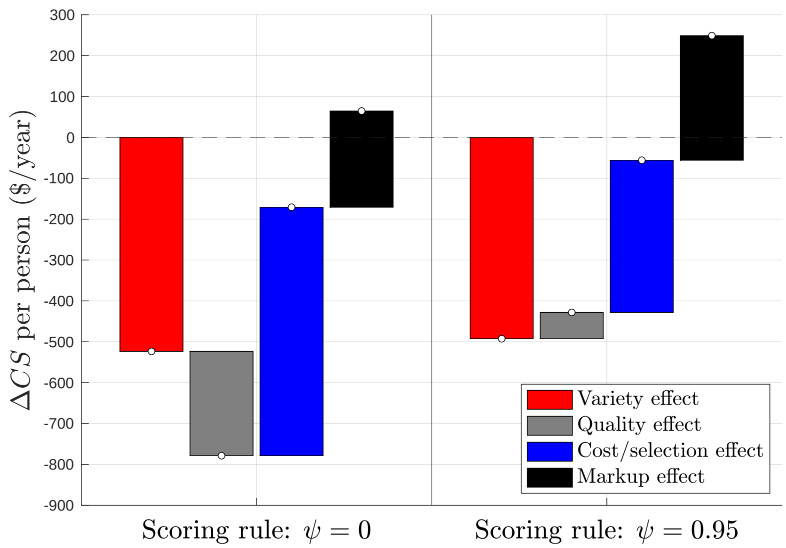

Notes: This figure shows a decomposition of the impacts of SSSAs relative to regulated competition on consumer surplus per person-year. The decomposition has four components: variety effect, quality effect, cost effect, and markup effect. The effects are reported additively, such that for each bar, the 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. The left panel shows results for a scoring rule with $\psi = 0$, whereas the right panel shows results for $\psi = 0.95$. Averages across the 76 region-year markets.

one that achieves the highest impacts on consumer surplus and enrollment ($\psi = 0.95$).

The *variety effect* is always negative, as consumers lose choice and are limited to a single insurer under SSSAs. The magnitude of this effect varies modestly across scoring rules. In particular, switching to SSSAs reduces average consumer surplus by $520 per person-year when $\psi = 0$, and by $490 when $\psi = 0.95$. These substantial losses are consistent with the significant preference heterogeneity we estimate. It is therefore unsurprising that, holding premiums and quality fixed and ignoring equilibrium effects, restricting consumers to the winning insurer in an SSSA is so harmful. However, quality, cost, and markup effects must also be taken into account.

Unlike variety, the impact of SSSAs on quality provision depends strongly on the scoring rule. Indeed, the *quality effect* changes signs from $\psi = 0$ to $\psi = 0.95$. Under $\psi = 0$, quality falls, and consumer surplus decreases by $260, adding to the losses from lower variety, totaling nearly $800 per person-year, or 50 percent of annual premiums. Instead, under $\psi = 0.95$, quality increases, raising consumer surplus by $70 and mitigating 14 percent of the losses from lower variety.

The *cost effect* and *markup effect* are also key drivers of consumer welfare gains. This aligns with pervasive adverse selection and imperfect competition under regulated competition, and with SSSAs affecting these margins, as illustrated by Figures 7-d and 7-e. For SSSAs with $\psi = 0$, a large cost effect of \$610 per person-year results from addressing adverse selection and lowering quality. Because losses in terms of variety and quality are severe, this positive cost effect is not enough to make procurement overall beneficial. However, the reduction in markups adds \$250 in consumer surplus, which flips the total to a gain of nearly \$70 per person-year. For $\psi = 0.95$, the cost and markup effects are more balanced. Consumers gain \$370 and \$310 from cost and markup reductions, respectively. In both cases, the cost effect is critical for SSSAs to outperform regulated competition. Even at $\psi = 0.95$, where higher quality rules out cost savings from quality reductions, consumers would be worse off under SSSAs by \$120 per person-year without a positive cost effect. This highlights that addressing adverse selection is key to enabling gains from procurement.

6.4 Heterogeneity Across Markets and Predictors of Procurement's Potential

Our findings provide some support for procurement auctions, but mask substantial heterogeneity across markets. As noted in Section 2, general results comparing SSSAs and regulated competition are difficult to prove. However, under the restricted environment in Remark 1, we can provide some guidance by linking our results to the attributes of the markets in our sample.

Remark 1 identifies three factors that predict the success of SSSAs: (i) low baseline variety, (ii) high baseline markups, and (iii) strong adverse selection. We explore how the market-level impacts of shifting to SSSAs on consumer surplus vary along these margins. First, Figure 10-a shows that the gains from SSSAs are driven by markets with less baseline variety, as measured by the number of insurers in the ACA, decreasing by a factor of five in markets with six or more insurers. Second, to study the relationship between the gains from SSSAs and baseline markups, we proxy the latter with the HHI under the ACA. Figure 10-b shows that gains from SSSAs are higher in markets with higher HHI and much smaller in less concentrated markets with an HHI below 2,500. Lastly, to measure adverse selection in each market at baseline, we compute the difference between the observed average cost under the ACA and the average cost of the SSSA monopolist—a high value of this variable suggests stronger baseline adverse selection. Figure 10-c shows that the gains from SSSAs are concentrated among markets with more severe adverse selection under regulated competition. All these patterns align with Remark 1.

Altogether, the combination of theoretical results (from a simplified setting) and empirical findings from our econometric model suggests that SSSAs are more likely to outperform regulated competition in markets with lower variety, higher markups, or more severe adverse selection.

Figure 10: Heterogeneity in the effects of SSSAs on consumer surplus

(a) ΔCS by number of insurers (b) ΔCS by HHI (c) ΔCS by change in AC

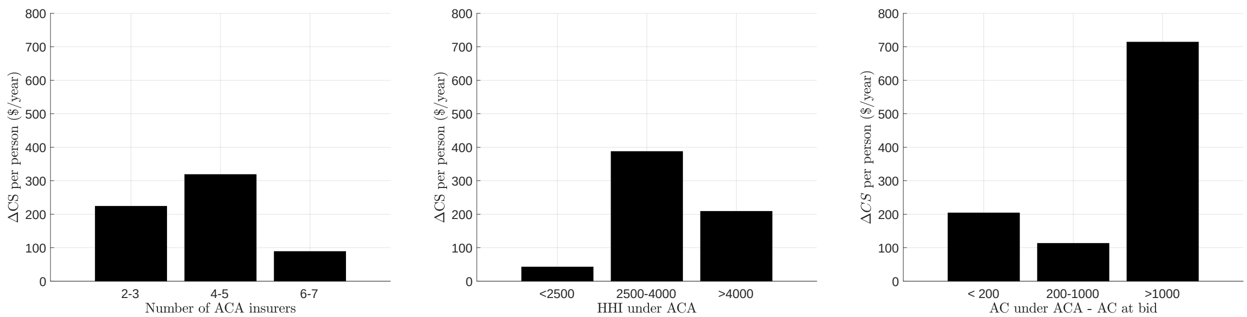

Notes: This figure illustrates the relationship between the impacts of a shift from regulated competition to SSSAs on consumer surplus and key market characteristics. The y-axis displays the change in consumer surplus per person-year. The x-axis in panels (a), (b), and (c) displays bins of number of insurers, HHI, and average cost under oligopoly relative to SSSA-regulated monopoly, respectively. Results for SSSAs with a scoring rule with $\psi = 0.95$.

7 Discussion

Our analysis comes with considerations and caveats that are worth discussing.

Entry. We focus on auctions where the bidders are the firms observed under regulated competition. By not modeling entry, we may underestimate the competitive effects of SSSAs if other potential entrants are willing to bid. The opposite would occur if bidding costs were high enough to deter entry, leading us to overestimate competitive effects. In our empirical context, most insurers already have provider networks in place for other insurance segments (e.g., employer-sponsored or Medicare Advantage), making the former scenario more plausible. Furthermore, auctions may provide opportunities for small, innovative insurers to enter the market and expand.

Input Markets. While we model endogenous quality, we do so in a stylized and parsimonious way. A concern is that input providers (e.g., hospitals in our application), aware of their relevance for firms to win a procurement auction, could try to extract more surplus from bidders (insurers), leading to cost increases not captured by our model (Ho and Lee, 2019). This is more likely in more concentrated input markets and when regulators enforce strict minimum quality requirements. Conversely, by ensuring the winner is the sole supplier, an SSSA could shift bargaining power toward bidders and imply lower costs. In Table 6, we assess the robustness of our main results to large changes in input prices, and hence costs. Even for significant increases (up to 15 percent), average consumer surplus increases under SSSAs relative to the status quo. Since the individual segment is a small segment of the overall health insurance market, cost increases capable of changing our main findings qualitatively seem unlikely.

Table 6: Robustness to Changes in Input Prices and Weak Enforcement of Auction Outcomes

	Baseline	A. Change in input prices						B. Weak enforcement		
		−20%	−10%	−5%	+5%	+10%	+20%	−10%	−20%	−50%
ΔCS per person ($/year)	+249	+573	+416	+330	+160	+81	−84	+182	+115	−86
Share of markets with ΔCS>0	0.87	1.00	0.97	0.92	0.72	0.62	0.30	0.78	0.75	0.39

Notes: This table displays results from a sensitivity analysis of our counterfactual results for an SSSA with $\psi = 0.95$. The first column displays baseline results from Section 6.2. Panel A displays results from our counterfactual analysis in scenarios in which insurer costs change by between −20 and +20 percent upon winning the SSSA. Panel B displays results from scenarios in which the winner can violate the minimum score constraint in equation (2). In practice, we solve the winner's problem after lowering the second-highest score in the auction by a range of values between 10 and 50 percent.

Economies of Scale. Having a single insurer in place may lead to economies of scale that could potentially translate to lower costs (Williamson, 1976).[14] These forces would make our results a lower bound for the potential gains from a procurement mechanism. Indeed, Table 6 shows how, in our application, further cost reductions could increase consumer surplus gains, and that eventually all markets would prefer SSSAs to regulated competition.

Weak Enforcement of Auction Outcomes. Regulators may face obstacles (e.g., time inconsistency, contract incompleteness, political economy, costly monitoring) to enforce that auction outcomes are delivered as contracted (Bajari, Houghton, and Tadelis, 2014; Decarolis, 2014). While the severity of this is context-specific, the fact that similar rules on product quality are imposed under regulated competition mitigates the concern. In addition, in Table 6 we verify that our results are robust to scenarios in which the winner can supply worse contracts than under perfect enforcement. As we relax the constraint in the winner's profit-maximization problem in equation (2)—by up to 37 percent of the second-highest score—we continue to find that SSSAs lead to higher average consumer surplus than regulated competition.

Switching Costs. In terms of demand-side dynamics, consumers may be forced to switch contracts upon changes in the sole provider in their local market across subsequent auctions, incurring switching costs that could lead to welfare losses. Two natural ways to mitigate this concern in practice would be to either auction the market only for new enrollees in every period, or to do so for several years at once. Investigating the welfare relevance of switching costs would be critical for these important market design considerations.

[14] Scale effects and their price effects in procurement have already been documented (Allende, Atal, Carril, Cuesta, and Gonzalez-Lira, 2025). Moreover, there is also some evidence of scale economies in health insurance (Serna, 2023).

Mistakes in Consumer Choice. A body of work documents behavioral biases that lead to suboptimal consumer choices, such as when faced with too many options (Abaluck and Gruber, 2023), or when choice requires costly attention (Brown and Jeon, 2024). These findings support paternalistic regulations that limit available products. These frictions would suggest that our estimated welfare losses from lower variety are biased upward, potentially understating the gains from procurement.

Dynamic Competition. In terms of supply-side dynamics, a natural concern is that by selecting a single winner, SSSAs may induce other firms to invest less in their quality or to leave the market altogether, weakening competition in subsequent auctions. This concern deserves a careful, case-by-case consideration. For our application to the ACA, since marketplaces remain a small share of the national insurance market, this does not seem to be a major concern. Moreover, different insurers can be winners in different regional markets, further limiting the possibility of exit.

Market Definition. 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. In our setting, regulators could define markets by different geographic regions or across demographic groups in light of preference heterogeneity. We leave it for future work to study alternative market definitions that tailor procurement mechanisms to heterogeneity in preferences, market size, and baseline market structure.

Auction Format. One might ask why we focused on SSSAs, and what counterfactuals would look like if one were to consider first-score scoring auctions—in which the winning firm would not be allowed to adjust prices or quality relative to the submitted bid—or auctions allowing more than one firm to offer products in the market. A compelling feature of SSSAs is that they make supply decisions strategically simple despite adverse selection, since firms need not worry about their rivals attempting to cream skim low-cost types by offering cheaper, lower-quality products. This property would not be present in first-score auctions or multi-award auctions.

Private vs. Public Provision. Since we consider counterfactuals in which each market is supplied exclusively by one firm, one might ask how this compares to direct provision by the government through a public option. To this point, we note that auctions continue to leverage the profit-maximizing motive of private firms to achieve higher quality at the lowest cost; it is not clear whether and how direct public supply could (efficiently) pursue the same objectives. Studying a public option would require additional assumptions about its objective function and cost structure.

8 Conclusion

We discuss public procurement—in the form of second-score scoring auctions—as an alternative to regulated competition in markets with adverse selection. Rather than promoting regulated competition between numerous firms (e.g., insurers in the context of healthcare), a regulator could determine a scoring rule that weights quality and prices, and solicit firms to submit bids and compete in the auction. In this scenario, consumers would face fewer choices, but inefficiencies from adverse selection and imperfect competition would decrease, and—when incentivized by the scoring rule—quality would increase.

To determine whether this alternative design has meaningful potential, we outline the steps to recover the demand and cost primitives needed to simulate counterfactual equilibria in second-score scoring auctions. We apply this framework to the Covered California ACA health insurance marketplace for 2014–2017. In counterfactual simulations, we find that auctions with a scoring rule placing sufficient weight on quality would lead to higher consumer welfare (net of government spending) than the status quo. The combination of reduced adverse selection, lower markups, and higher quality would more than offset the harm to consumers from limiting their options when choosing health insurance. These patterns are stronger in markets with fewer than five insurers, higher insurer concentration, or more severe adverse selection. While we do not propose a specific alternative to the ACA, our findings suggest that procurement mechanisms may improve outcomes in markets where adverse selection and imperfect competition undermine efficiency, warranting further research into their design.

A Second-Score Scoring Auctions vs. Oligopoly in Logit Models

A.1 Proof of Remark 1

With $J \geq 2$ and symmetric costs, bids are symmetric, and, for sufficiently high $d\psi/d\delta$, bid quality is greater than $\max_j \delta^\star_j$. Moreover, since bids are symmetric, the winner has no further room to optimize, so that $p^a = \bar{b}(\delta^a)$. Therefore, if ε_{ij} is iid Gumbel, consumer surplus under the SSSA is $CS^a = \int \log\left(1 + e^{\frac{-\alpha_i \bar{b}(\delta^a) + \beta_i \delta^a}{\zeta}}\right) di$. Given the inequality in Remark 1, for any α_i and β_i:

$\dfrac{-\alpha_i \bar{b}(\delta^a) + \beta_i \delta^a}{\zeta} > \log(J) + \dfrac{-\alpha_i p^\star_{j^\star(\alpha_i,\beta_i)} + \beta_i \delta^\star_{j^\star(\alpha_i,\beta_i)}}{\zeta}$, or equivalently $e^{\frac{-\alpha_i \bar{b}(\delta^a) + \beta_i \delta^a}{\zeta}} > J e^{\frac{-\alpha_i p^\star_{j^\star(\alpha_i,\beta_i)} + \beta_i \delta^\star_{j^\star(\alpha_i,\beta_i)}}{\zeta}}$.

Given the definition of $j^\star(\alpha_i, \beta_i)$, $1 + e^{\frac{-\alpha_i \bar{b}(\delta^a) + \beta_i \delta^a}{\zeta}} > 1 + \sum_j e^{\frac{-\alpha_i p^\star_j + \beta_i \delta^\star_j}{\zeta}}$ for all α_i, β_i, and one obtains $CS^a = \int \log\left(1 + e^{\frac{-\alpha_i \bar{b}(\delta^a) + \beta_i \delta^a}{\zeta}}\right) di > \int \log\left(1 + \sum_j e^{\frac{-\alpha_i p^\star_j + \beta_i \delta^\star_j}{\zeta}}\right) di = CS^\star$.■

A.2 Implications of Remark 1

We consider the comparison between SSSAs and oligopoly in the following simple cases.

Simple logit, exogenous quality, symmetric contracts, no adverse selection. If $\alpha_i \equiv \alpha$ and $\beta_i \equiv \beta$ for all i, contracts are symmetric, and firms cannot adjust quality, the condition in Remark 1 is both necessary and sufficient for SSSAs to dominate oligopoly in terms of consumer surplus, and it simplifies to $\mu^\star > \zeta \log(J)/\alpha$, where $\mu^\star$ is the markup in the symmetric oligopoly equilibrium. In the simple logit case, this is $\mu^\star = \frac{\zeta}{\alpha}\frac{J}{J-1+\sigma_0}$, where σ_0 denotes the share of individuals choosing the outside option. Therefore, SSSAs dominate oligopoly as long as $\frac{J}{J-1+\sigma_0} > \log(J)$, or $\sigma_0 < \frac{J}{\log(J)} - J + 1$. This is always true for $J = 2$, never true for $J \geq 4$, and depends on the market share of the outside good—which is inversely proportional to markups—for $J = 3$.

Simple logit, exogenous quality, asymmetric contracts, no adverse selection. Consider now the case of J^ℓ low-quality plans with quality δ^ℓ and J^h high-quality plans with $\delta^h > \delta^\ell$, with $c^h > c^\ell$. We keep focusing on the simple case of $\alpha_i \equiv \alpha$ and $\beta_i \equiv \beta$ for all i and there is no adverse selection. As long as the scoring rule ψ rewards quality sufficiently, and $J^h \geq 2$, the SSSA's winner is a high-quality plan and the final price p^a is c^h. Letting $\sigma^{\star,h}$ and $\mu^{\star,h}$ denote the equilibrium market share and markup of high-quality plans under oligopoly, one has $\alpha\mu^{\star,h}/\zeta = (1-\sigma^{\star,h})^{-1}$, and similarly for low-quality plans. Remark 1 then implies:

a. if $\sigma^{\star,h} > \max\left\{\sigma^{\star,\ell}, 1-\left(\log(J)\right)^{-1}\right\}$, then $CS^a > CS^\star$;

b. if $\sigma^{\star,\ell} > \max\left\{\sigma^{\star,h}, 1-\left(\log(J) + \frac{\alpha}{\zeta}(c^h - c^\ell) - \frac{\beta}{\zeta}(\delta^h - \delta^\ell)\right)^{-1}\right\}$, then $CS^a > CS^\star$.

Both conditions are intuitive and relate to the analysis in Anderson and De Palma (2001). A larger market share of a product type indicates a higher average utility, implying that other products charge relatively high markups. If high-quality products h have the highest market share in the oligopoly equilibrium, the only downside of the SSSA comes from the decrease in the number of options, from J to one. If, instead, low-quality products have the highest market share in the oligopoly equilibrium, SSSAs have the additional downside of increasing average cost from c^ℓ to c^h, but also the additional benefit of increasing quality from δ^ℓ to δ^h.

B Pure Characteristics Demand Model with Individual-Level Data

B.1 The Role of Individual-Level Data

We consider the pure characteristics demand model (Bajari and Benkard, 2005; Berry and Pakes, 2007) in a situation in which, in addition to product-level market shares and characteristics $(s_j, \boldsymbol{x}_j)$, one observes individual-level choices and characteristics: $(Y_i, \boldsymbol{z}_i)$. Utility is specified in equation (4) of Section 3 as $U_{ij} = -\alpha_i p_j + \delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i)$. The numerical approximation of the likelihood is:

$$\Pr[Y_i = j] = \ell_{ij} \approx M^{-1} \sum_{m=1}^{M} \int_{\underline{\Delta}_{ij}(\boldsymbol{\beta}_i^m)}^{\overline{\Delta}_{ij}(\boldsymbol{\beta}_i^m)} dF\big(\alpha_i \big| \underbrace{\overline{\boldsymbol{\beta}} + \Sigma^{\beta} \boldsymbol{\nu}_i^m}_{\boldsymbol{\beta}_i^m}; \boldsymbol{z}_i\big) \, ; \qquad \begin{aligned} \underline{\Delta}_{ij}(\boldsymbol{\beta}_i) &\equiv \max\nolimits_{k: p_k > p_j} \tfrac{\delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i) - \delta(\boldsymbol{x}_k, \boldsymbol{z}_i; \boldsymbol{\beta}_i)}{p_k - p_j} \, ; \\ \overline{\Delta}_{ij}(\boldsymbol{\beta}_i) &\equiv \min\nolimits_{k: p_k > p_j} \tfrac{\delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i) - \delta(\boldsymbol{x}_k, \boldsymbol{z}_i; \boldsymbol{\beta}_i)}{p_k - p_j} \, . \end{aligned}$$

The set of draws $\{\boldsymbol{\nu}_{\boldsymbol{i}}^{\boldsymbol{m}}\}_{m=1}^M$ is treated as fixed and, combined with $\overline{\boldsymbol{\beta}}$ and Σ^β, leads to the draws $\{\boldsymbol{\beta}_{\boldsymbol{i}}^{\boldsymbol{m}}\}_{m=1}^M$ for numerical integration. A larger value of M reduces the simulation error in the calculation of ℓ_{ij}. Importantly, if the data only contain product-level information, $\boldsymbol{x}_j$ and $s_j = \mathbb{E}_i[Y_i = j]$, the above reduces to the market share expression in Berry and Pakes (2007) and Pang, Su, and Lee (2015), after omitting $\boldsymbol{z}_i$ and replacing $\boldsymbol{\nu}_i^m$ with $\boldsymbol{\nu}^m$. As Berry and Pakes (2007) notes, if the conditional distribution of α_i has a density (which we denote f) which is a differentiable function of a parameter vector, then s_j, ℓ_{ij}, and $\mathcal{L} = \sum_i \log(\Pr[Y_i])$ are all differentiable functions of this parameter.

Problems can arise due to how ℓ_{ij} responds to changes in $\overline{\boldsymbol{\beta}}$ or Σ^β. Considering $\overline{\boldsymbol{\beta}}$—the case of Σ^β is analogous—by applying the Leibniz integral rule and the Envelope Theorem, one obtains:

$$\begin{aligned} \frac{\partial \ell_{ij}}{\partial \overline{\boldsymbol{\beta}}} &= M^{-1} \sum_{m=1}^{M} f(\overline{\Delta}_{ij}(\boldsymbol{\beta}_i^m) | \boldsymbol{\beta}_i^m; \boldsymbol{z}_i) \frac{\partial}{\partial \overline{\boldsymbol{\beta}}} \left(\frac{\delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i^m) - \delta(\boldsymbol{x}_{\overline{k}_i^m}, \boldsymbol{z}_i; \boldsymbol{\beta}_i^m)}{p_{\overline{k}_i^m} - p_j} \right) \\ &- M^{-1} \sum_{m=1}^{M} f(\underline{\Delta}_{ij}(\boldsymbol{\beta}_i^m) | \boldsymbol{\beta}_i^m; \boldsymbol{z}_i) \frac{\partial}{\partial \overline{\boldsymbol{\beta}}} \left(\frac{\delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \boldsymbol{\beta}_i^m) - \delta(\boldsymbol{x}_{\underline{k}_i^m}, \boldsymbol{z}_i; \boldsymbol{\beta}_i^m)}{p_{\underline{k}_i^m} - p_j} \right) \\ &+ M^{-1} \sum_{m=1}^{M} \int_{\underline{\Delta}_{ij}(\boldsymbol{\beta}_i^m)}^{\overline{\Delta}_{ij}(\boldsymbol{\beta}_i^m)} \frac{\partial}{\partial \overline{\boldsymbol{\beta}}} dF\big(\alpha_i \big| \boldsymbol{\beta}_i^m; \boldsymbol{z}_i\big) \, , \end{aligned} \tag{B.1}$$

where we use $\overline{k}_i^m \in \mathcal{J}$ to denote the solution of the minimization problem defining $\overline{\Delta}_i(\boldsymbol{\beta}_i^m)$, and similarly for $\underline{k}_i^m$. The last term in equation (B.1) is continuous as long as $\delta(\boldsymbol{x}_j, \boldsymbol{z}_i; \cdot)$ and $f(\alpha_i | \cdot; \boldsymbol{z}_i)$ are continuous functions of $\boldsymbol{\beta}_i$. However, the first two terms are not, since there are discontinuities for every $\overline{\boldsymbol{\beta}}$ at which one of the two discrete-valued functions $\overline{k}_i^m$ and $\underline{k}_i^m$ changes its value.

The incidence of these discontinuities is more severe if either $\boldsymbol{z}_i$ does not vary (or varies too little), or M is not sufficiently large relative to the dimension of $\boldsymbol{\beta}_i$. Taking the derivative of the total likelihood, one has

$$\frac{\partial \mathcal{L}}{\partial \overline{\boldsymbol{\beta}}} = M^{-1} \sum_{i=1}^{I} \frac{1}{\Pr\left[Y_i\right]} \sum_{m=1}^{M} \frac{\partial \ell_{iY_i}}{\partial \overline{\boldsymbol{\beta}}},$$

and, as the number of individuals I and the corresponding variation in $\boldsymbol{z}_i$ grow, the number of points at which this derivative is discontinuous (and the size of the "jumps") decay very quickly. Relatedly, the same is true if one considers the approximated share of product j: $s_j = M^{-1} I^{-1} \sum_i \sum_m \ell_{ij}$, whose derivative with respect to utility parameters also becomes increasingly smooth as I increases. Intuitively, individual-level data play a role similar to that of idiosyncratic error terms in common logit or probit models, making the event of any one consumer being exactly indifferent between the same two products a null event. In our tests, the log-likelihood derived from individual-level data becomes smooth at moderate values of M.

B.2 Estimation Details of Pure Characteristics Demand Model for Covered California

The model specification in Section 5 requires estimating

$$(\overline{\alpha}^{w_i}, \sigma^{\alpha,w_i}, \overline{\beta}^{av,w_i}, \sigma^{av,w_i}, \overline{\boldsymbol{\beta}}^{n,w_i}, \Sigma^{n,w_i}, \overline{\mu}^{w_i}, \sigma^{\mu,w_i}, \overline{\boldsymbol{\xi}}_{rt}^{w_i}, \sigma^{\xi,w_i}),$$

with one such collection by age bin [26-30), [30-35), [35-40), [40-45), [45-50), [50-55), and [55-64], and letting $\overline{\alpha}^{w_i}$ also vary linearly with income as a share of the FPL.

In the absence of idiosyncratic shocks to random utility, the model requires a scale normalization: We set the standard deviation of α_i equal to 0.25.[15] We then estimate all the parameters 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 in equation (13). To compute $\sigma_{jrt}(w_i, z_i)$, we follow the steps outlined above. When

[15]This choice is dictated by two observations. First, the ratio between the average and standard deviation of the price coefficients α_i across the seven age bins is pinned down consistently when we use different scale normalizations. 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 by making near-zero α_i extremely rare.

i selects j, this implies that for all $k \neq j$

$$\alpha_i \geq \underline{\Delta}_{ij} \equiv \max_{k: p^{w_i}_{krt} > p^{w_i}_{jrt}} \frac{\beta^{av}_i (AV^{w_i}_{krt} - AV^{w_i}_{jrt}) + \beta^n_i (\mathcal{N}^{z_i}_{krt} - \mathcal{N}^{z_i}_{jrt}) + \mu_i (\mathbf{1}[k \neq 0] - \mathbf{1}[j \neq 0]) + \xi_{i f_{krt}} - \xi_{i f_{jrt}}}{p^{w_i}_{krt} - p^{w_i}_{jrt}},$$
$$\alpha_i \leq \overline{\Delta}_{ij} \equiv \min_{k: p^{w_i}_{krt} < p^{w_i}_{jrt}} \frac{\beta^{av}_i (AV^{w_i}_{krt} - AV^{w_i}_{jrt}) + \beta^n_i (\mathcal{N}^{z_i}_{krt} - \mathcal{N}^{z_i}_{jrt}) + \mu_i (\mathbf{1}[k \neq 0] - \mathbf{1}[j \neq 0]) + \xi_{i f_{krt}} - \xi_{i f_{jrt}}}{p^{w_i}_{krt} - p^{w_i}_{jrt}}.$$

The quantities $\underline{\Delta}_{ij}$ and $\overline{\Delta}_{ij}$ are a function of data and parameters.

For any value of $(\overline{\beta}^{av,w_i}, \sigma^{av,w_i}, \overline{\boldsymbol{\beta}}^{n,w_i}, \Sigma^{n,w_i}, \overline{\mu}^{w_i}, \sigma^{\mu,w_i}, \overline{\boldsymbol{\xi}}^{w_i}_{rt}, \sigma^{\xi,w_i})$, and a fixed set of $M = 200$ Halton draws for every sampled individual i, we obtain a collection $\{\underline{\Delta}^m_{ij}, \overline{\Delta}^m_{ij}\}^M_{m=1}$, and approximate

$$\sigma_{jrt}(w_i, z_i) \approx M^{-1} \sum_{m=1}^{M} \int_{\underline{\Delta}^m_{ij}}^{\overline{\Delta}^m_{ij}} \phi \left(\frac{\log(\alpha) - \overline{\alpha}^{w_i}}{\sigma^{\alpha,w_i}} \right) d\alpha \,, \tag{B.2}$$

where $\phi(\cdot)$ is the density of a standard Gaussian distribution. For estimation, we sample I^{w_i} individuals, varying between 82,826 and 106,265 across age groups.

Unlike the widely used logit and probit discrete choice models with idiosyncratic error terms with full support, the pure characteristics model with individual-level data we employ here can predict $\sigma_{jrt}(w_i, z_i) = 0$ for a given parameter value (or sometimes any value).[16] We handle these instances by replacing $\sigma_{jrt}(w_i, z_i) = 0$ with the lowest possible number ϱ for which $\log(\varrho) > -\infty$ at machine precision. We then verify ex-post that the share of individuals in our dataset for which this occurs is small, and that it decreases with M. With the current specification, this is 4–9 percent across the seven estimation datasets, one for every age bin.

The right-hand side of (B.2) only depends on data and parameters, and can be used to evaluate the likelihood of $Y_i = j$ in (13). We maximize $\mathcal{L}$ employing Artelys' Knitro multi-algorithm nonlinear programming procedure, for which we provide simple logit estimates as starting values.[17]

[16] 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.

[17] 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.

C Details of Cost Identification and Estimation

C.1 Cost Identification: A Formal Argument

To formalize the constructive argument outlined in Section 3, we begin by combining equations (8) and (9) to obtain

$$MR_j \left(\frac{\partial Q_j}{\partial p_j} + \int \rho(z_i, \boldsymbol{\beta}_i) \frac{\partial \sigma_j(z_i, \boldsymbol{\beta}_i)}{\partial p_j} dF(z_i, \boldsymbol{\beta}_i) \right)^{-1} = AC_j Q_j \left(Q_j + \int \rho(z_i, \boldsymbol{\beta}_i) \sigma_j(z_i, \boldsymbol{\beta}_i) dF(z_i, \boldsymbol{\beta}_i) \right)^{-1} ,$$

or, after rearranging,

$$\frac{MR_j}{\partial Q_j / \partial p_j} \left(1 + \int \rho(z_i, \boldsymbol{\beta}_i) dF_j^{\text{avg}}(z_i, \boldsymbol{\beta}_i) \right) = AC_j \left(1 + \int \rho(z_i, \boldsymbol{\beta}_i) dF_j^{\text{mg}}(z_i, \boldsymbol{\beta}_i) \right),$$

where $F_j^{\text{avg}}(z_i, \boldsymbol{\beta}_i) \equiv \frac{\sigma_j(z_i, \boldsymbol{\beta}_i)}{Q_j} F(z_i, \boldsymbol{\beta}_i)$ is a (probability) measure describing the $(z_i, \boldsymbol{\beta}_i)$-composition of average buyers of product j and, similarly, $F_j^{\text{mg}}(z_i, \boldsymbol{\beta}_i) \equiv \frac{\partial \sigma_j(z_i, \boldsymbol{\beta}_i)/\partial p_j}{\partial Q_j / \partial p_j} F(z_i, \boldsymbol{\beta}_i)$ describes the composition of its marginal buyers. This expression ultimately leads to

$$\frac{MR_j}{\partial Q_j / \partial p_j} - AC_j = \int \rho(z_i, \boldsymbol{\beta}_i) \left[AC_j dF_j^{\text{mg}} - \frac{MR_j}{\partial Q_j / \partial p_j} dF_j^{\text{avg}} \right] , \tag{C.1}$$

which admits a unique solution $\rho(\cdot)$ for all j as long as $\mathbb{K}_j(z_i, \boldsymbol{\beta}_i) \equiv AC_j F_j^{\text{mg}}(z_i, \boldsymbol{\beta}_i) - \frac{MR_j}{\partial Q_j / \partial p_j} F_j^{\text{avg}}(z_i, \boldsymbol{\beta}_i)$ is an invertible kernel. Formally: if $\int a(z_i, \boldsymbol{\beta}_i) d\mathbb{K}_j(z_i, \boldsymbol{\beta}_i) = 0$, then $a(\cdot) = 0$ (almost surely).

Informally, this requires rich variation across products in how average and marginal buyers differ. To obtain some intuition, notice that the left-hand side of (C.1) is the difference between average marginal revenue and average cost. If these are equal for all j, and the kernel $\mathbb{K}_j$ is invertible, ρ must be constant and equal to zero, ruling out adverse (or advantageous) selection. When the left-hand side of (C.1) is nonzero, how it varies across j's jointly with the variation in differences between marginal and average buyers (on the right-hand side) identifies how cost varies with z_i and $\boldsymbol{\beta}_i$. Once ρ is identified, $\kappa_j(\boldsymbol{x}_j^\star)$ is the value that satisfies equations (8) and (9).

C.2 Details of Cost Estimation

The above provides a constructive argument that can be generalized to the case of multiproduct oligopoly. This is what we do in our application in Section 5, in which we estimate costs as follows.

First, consider any given value of the parameter vector $\boldsymbol{\rho} = (\rho^{\text{age}}, \rho^{AV}, \rho^{\text{net}})$, and—simplifying notation since all terms are evaluated at the observed rates $\boldsymbol{\tau}^{\star}_{rt}$—one can rewrite (16) as

$$\sum_{k\in\mathcal{J}_{f_j}} \frac{\partial R_{krt}}{\partial \tau_{jrt}} = \sum_{k\in\mathcal{J}_{f_j}} \overline{\kappa}_{jrt}\left(H^{\rho}_{krt}\frac{\partial Q_{krt}}{\partial \tau_{jrt}} + Q_{krt}\frac{\partial H^{\rho}_{krt}}{\partial \tau_{jrt}}\right), \text{ for all } j,$$

where $H^{\rho}_{jrt} \equiv AC_{jrt}/\overline{\kappa}_{jrt}$ describes heterogeneity in cost due to the characteristics of individuals choosing plan j, or observed actuarial value differences. In matrix notation this becomes $\boldsymbol{MR}_{rt} = \left(\Xi(\boldsymbol{\rho})\cdot\Omega\right)\overline{\boldsymbol{K}}_{rt}$, where Ω is an ownership matrix with generic element $\omega(j,k) = \mathbf{1}[f_j = f_k]$, $\overline{\boldsymbol{K}}_{rt} = \left(\overline{\kappa}_{1rt},\ldots,\overline{\kappa}_{Jrt}\right)$, $\boldsymbol{MR}^{\intercal}_{rt} \equiv \left(\sum_{k\in J_{f_1}}\frac{\partial R_{krt}}{\partial \tau_{1rt}},\ldots,\sum_{k\in J_{f_J}}\frac{\partial R_{Jrt}}{\partial \tau_{Jrt}}\right)$, and

$$\Xi(\boldsymbol{\rho}) \equiv \begin{bmatrix} H^{(\rho)}_{1rt}\frac{\partial Q_{1rt}}{\partial \tau_{1rt}} + Q_{1rt}\frac{\partial H^{(\rho)}_{1rt}}{\partial \tau_{1rt}} & \cdots & H^{(\rho)}_{Jrt}\frac{\partial Q_{Jrt}}{\partial \tau_{1rt}} + Q_{Jrt}\frac{\partial H^{(\rho)}_{Jrt}}{\partial \tau_{1rt}} \\ \vdots & \ddots & \\ H^{(\rho)}_{1rt}\frac{\partial Q_{1rt}}{\partial \tau_{Jrt}} + Q_{1rt}\frac{\partial H^{(\rho)}_{1rt}}{\partial \tau_{Jrt}} & & H^{(\rho)}_{Jrt}\frac{\partial Q_{Jrt}}{\partial \tau_{Jrt}} + Q_{Jrt}\frac{\partial H^{(\rho)}_{Jrt}}{\partial \tau_{Jrt}} \end{bmatrix}.$$

From this expression, for a given $\boldsymbol{\rho}$, by assuming equilibrium pricing one obtains

$$\widehat{\boldsymbol{K}}_{rt}(\boldsymbol{\rho}) \equiv \left(\Xi(\boldsymbol{\rho})\cdot\Omega\right)^{-1}\boldsymbol{MR}_{rt}\text{ , and },\ \widehat{\boldsymbol{AC}}_{rt}(\boldsymbol{\rho}) = \widehat{\boldsymbol{K}}_{rt}(\boldsymbol{\rho})\boldsymbol{H}^{(\rho)}_{rt},$$

where $\widehat{\boldsymbol{AC}}_{rt}(\boldsymbol{\rho})$ is the model-predicted value of average cost given $\boldsymbol{\rho}$. When setting $\boldsymbol{\rho} = \mathbf{0}$, the above procedure is *identical* to the case with constant marginal cost since $\boldsymbol{H}^{\mathbf{0}}_{jrt} \equiv AV_j$ for all j, and $\widehat{\boldsymbol{AC}}_{rt}(\mathbf{0})$ is equal to the vector of constant marginal costs that rationalizes observed rates as optimal.

These steps return a measure of the average and marginal cost for every plan-market combination, varying with the selection parameters in $\boldsymbol{\rho}$. Following the intuition of Einav, Finkelstein, and Cullen (2010), and the formal discussion above, we use observed average claims data to find the values $\widehat{\boldsymbol{\rho}}$ that satisfy $AC^{\text{obs}}_{jrt} = \widehat{AC}_{jrt}(\widehat{\boldsymbol{\rho}})$. These will be unique as long as there is sufficient variation across jrt in the within-plan difference between the distribution of age, $WTP^{\text{AV}}(\theta_i)$, and $WTP^{\text{net}}(\theta_i)$ among enrollees and the distribution of the same variables among marginal buyers. Intuitively, when $\rho^{\text{age}} > 0$, plans enrolling many older individuals but with marginal buyers who are, on average, younger 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 age composition of marginal and average buyers is similar. Similar arguments apply to ρ^{AV} and ρ^{net}.

For estimation, we search over a grid of $\boldsymbol{\rho}$ parameters for the vector $\widehat{\boldsymbol{\rho}}$ that minimizes the enrollment-weighted average measure of $|AC^{\text{obs}}_{jrt} - \widehat{AC}_{jrt}(\widehat{\rho})|$ across *jrt* observations. We then recover $\overline{\boldsymbol{K}}_{rt} = \widehat{\boldsymbol{K}}_{rt}(\widehat{\rho})$. The values of the parameters $\boldsymbol{\eta}$ governing the costs of quality adjustments are estimated in a second step, by taking the average of the right-hand side of equation (17).

References

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

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

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

Allende, C., J. P. Atal, R. Carril, J. I. Cuesta, and A. Gonzalez-Lira (2025): "Centralizing Procurement: The Roles of Scale, Selection, and Variety," Manuscript.

Anderson, S. P. and A. De Palma (2001): "Product Diversity in Asymmetric Oligopoly: Is the Quality of Consumer Goods Too Low?" *The Journal of Industrial Economics*, 49, 113–135.

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

——— (2010): "Procurement when Price and Quality Matter," *The Rand Journal of Economics*, 41, 1–34.

Azevedo, E. M. and D. Gottlieb (2017): "Perfect Competition in Markets With Adverse Selection," *Econometrica*, 85, 67–105.

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

Bajari, P. and C. L. Benkard (2005): "Demand Estimation with Heterogeneous Consumers and Unobserved Product Characteristics: A Hedonic Approach," *Journal of Political Economy*, 113, 1239–1276.

Bajari, P., S. Houghton, and S. Tadelis (2014): "Bidding for Incomplete Contracts: An Empirical Analysis of Adaptation Costs," *American Economic Review*, 104, 1288–1319.

Barahona, N., C. Otero, and S. Otero (2023): "Equilibrium Effects of Food Labeling Policies," *Econometrica*, 91, 839–868.

Benetton, M. (2021): "Leverage Regulation and Market Structure: A Structural Model of the UK Mortgage Market," *The Journal of Finance*, 76, 2997–3053.

Berry, S., J. Levinsohn, and A. Pakes (1995): "Automobile Prices in Market Equilibrium," *Econometrica*, 63, 841–890.

Berry, S. and A. Pakes (2007): "The Pure Characteristics Demand Model," *International Economic Review*, 48, 1193–1225.

Berry, S. T. and P. A. Haile (2024): "Nonparametric Identification of Differentiated Products Demand using Micro Data," *Econometrica*, 92, 1135–1162.

Bresnahan, T. F. (1987): "Competition and Collusion in the American Automobile Industry: The 1955 Price War," *The Journal of Industrial Economics*, 457–482.

Brown, J., M. Duggan, I. Kuziemko, and W. Woolston (2014): "How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program," *American Economic Review*, 104, 3335–3364.

Brown, Z. Y. and J. Jeon (2024): "Endogenous Information and Simplifying Insurance Choice," *Econometrica*, 92, 881–911.

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

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

Cohen, A. and L. Einav (2007): "Estimating Risk Preferences from Deductible Choice," *American Economic Review*, 97, 745–788.

Crawford, G. S., N. Pavanini, and F. Schivardi (2018): "Asymmetric Information and Imperfect Competition in Lending Markets," *American Economic Review*, 108, 1659–1701.

Crawford, G. S., O. Shcherbakov, and M. Shum (2019): "Quality Overprovision in Cable Television Markets," *American Economic Review*, 109, 956–995.

Cuesta, J. I., C. Noton, and B. Vatter (2025): "Vertical Integration and Plan Design in Healthcare Markets," Working Paper 32833, National Bureau of Economic Research.

Cuesta, J. I. and A. Sepúlveda (2021): "Price Regulation in Credit Markets: A Trade-off between Consumer Protection and Credit Access," *Available at SSRN 3282910*.

Dafny, L., K. Ho, and M. Varela (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, 32–58.

Decarolis, F. (2014): "Awarding Price, Contract Performance, and Bids Screening: Evidence from Procurement Auctions," *American Economic Journal: Applied Economics*, 6, 108–32.

——— (2015): "Medicare part d: Are insurers gaming the low income subsidy design?" *American Economic Review*, 105, 1547–1580.

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

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

Demsetz, H. (1968): "Why Regulate Utilities?" *Journal of Law and Economics*, 11, 55–65.

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

Diamond, P. A. (1994): "Two Improvements on the Clinton Framework," *Journal of Economic Perspectives*, 8, 61–66.

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

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

Einav, L., A. Finkelstein, and N. Mahoney (2021): "The IO of Selection Markets," in *Handbook of Industrial Organization*, Elsevier, vol. 5, 389–426.

Einav, L., A. Finkelstein, and P. Schrimpf (2010): "Optimal Mandates and the Welfare Cost of Asymmetric Information: Evidence from the UK Annuity Market," *Econometrica*, 78, 1031–1092.

Einav, L., A. Finkelstein, and P. Tebaldi (forthcoming): "Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies," *Journal of Political Economy Microeconomics*.

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

Enthoven, A. C. (1993): "The History and Principles of Managed Competition," *Health Affairs*, 12, 24–48.

Fan, Y. (2013): "Ownership consolidation and product characteristics: A study of the US daily newspaper market," *American Economic Review*, 103, 1598–1628.

Geruso, M. and T. Layton (2020): "Upcoding: Evidence from Medicare on Squishy Risk Adjustment," *Journal of Political Economy*, 128, 984–1026.

Geruso, M., T. J. Layton, G. McCormack, and M. Shepard (2023): "The Two-Margin Problem in Insurance Markets," *Review of Economics and Statistics*, 105, 237–257.

Ghili, S. (2022): "Network Formation and Bargaining in Vertical Markets: The Case of Narrow Networks in Health Insurance," *Marketing Science*, 41, 501–527.

Gupta, A., S. T. Parente, and P. Sanyal (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., D. Weimer, and D. B. Mukamel (2019): "A Consumer-centric Approach to Network Adequacy: Access to Four Specialties in California's Marketplace," *Health Affairs*, 38, 1918–1926.

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

Handel, B. and K. Ho (2021): "The Industrial Organization of Health Care Markets," in *Handbook of Industrial Organization*, Elsevier, vol. 5, 521–614.

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

Handel, B. R. (2013): "Adverse selection and inertia in health insurance markets: When nudging hurts," *American Economic Review*, 103, 2643–2682.

Harrison, R., M. Parada-Contzen, and M. Villena (2023): "Can Auctions Increase Competition in the Pension Funds Market? The Chilean Experience," *Journal of Policy Modeling*, 45, 975–993.

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

Illanes, G. and M. Padi (2019): "Retirement Policy and Annuity Market Equilibria: Evidence from Chile," Working Paper 26285, National Bureau of Economic Research.

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

Kim, P. and A. Li (2025): "Reinsurance, Financial Frictions, and Subsidy Efficiency: Evidence from Health Insurance," Working paper.

Kong, E., T. Layton, and M. Shepard (2024): "Adverse Selection and (un)Natural Monopoly in Insurance Markets," Working Paper 33187, National Bureau of Economic Research.

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

Mahoney, N. and G. Weyl (2017): "Imperfect Competition in Selection Markets," *Review of Economics and Statistics*.

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

McGuire, T. G. and M. H. Riordan (1995): "Incomplete Information and Optimal Market Structure Public Purchases from Private Providers," *Journal of Public Economics*, 56, 125–141.

Montanera, D., A. N. Mishra, and T. Raghu (2022): "Mitigating risk selection in healthcare entitlement programs: A beneficiary-level competitive bidding approach," *Information Systems Research*, 33, 1221–1247.

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

Panhans, M. (2019): "Adverse Selection in ACA Exchange Markets: Evidence from Colorado," *American Economic Journal: Applied Economics*, 11, 1–36.

Quan, T. W. and K. R. Williams (2018): "Product Variety, across-market Demand Heterogeneity, and the Value of Online Retail," *The RAND Journal of Economics*, 49, 877–913.

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

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

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

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, 578–615.

Shepard, M. and E. Forsgren (2023): "Do Insurers Respond to Active Purchasing? Evidence from the Massachusetts Health Insurance Exchange," *Journal of Risk and Insurance*, 90, 9–31.

Shleifer, A. (1985): "A Theory of Yardstick Competition," *The RAND Journal of Economics*, 319–327.

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

Spence, A. M. (1975): "Monopoly, Quality, and Regulation," *The Bell Journal of Economics*, 417–429.

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

Tebaldi, P. (2025): "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," *Review of Economic Studies*, 92, 586–620.

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

Vatter, B. (2025): "Quality Disclosure and Regulation: Scoring Design in Medicare Advantage," *Econometrica*, 93, 959–1001.

Williamson, O. E. (1976): "Franchise Bidding for Natural Monopolies-in General and with Respect to CATV," *The Bell Journal of Economics*, 7, 73–104.

Wolinsky, A. (1997): "Regulation of Duopoly: Managed Competition vs Regulated Monopolies," *Journal of Economics & Management Strategy*, 6, 821–847.

Wollmann, T. G. (2018): "Trucks without Bailouts: Equilibrium Product Characteristics for Commercial Vehicles," *American Economic Review*, 108, 1364–1406.

Zewde, N. and V. Perez (2024): "Incentives and Operations of Medicaid Managed Care Plans in New York State: Implications for Procurement Design and Market Evolution," *INQUIRY: The Journal of Health Care Organization, Provision, and Financing*, 61, 00469580241258653.

Supplemental Appendix for Online Publication

Procurement Auctions vs. Regulated Competition in Selection Markets

José Ignacio Cuesta and Pietro Tebaldi

S.1 Details on Data Construction

S.1.1 Provider Networks

To construct measures of provider networks offered by each plan, we leverage detailed data on the universe of providers that each plan offers. Covered California mandates every qualified health plan to provide monthly information on all PCPs, specialists, hospitals, and other individual and facility providers covered by them. 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 our analysis, we use the second quarter of each year. 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. In addition, we focus on the networks of PCPs and hospitals. We leverage the fact that we observe the location of both households and providers at the zipcode level to construct measures of provider networks offered by each plan in each zipcode. It is possible in the data for the same provider to have multiple practice locations within a zip code. In such cases, we treat these locations as one for the same provider within a zip code. In contrast, if the same provider offers multiple locations in different zip codes within a catchment area, we count them as separate providers.

S.1.2 Measuring the Number of Potential Buyers

The enrollment data describe choices among plans in Covered California, but do not provide information on the outside option for this population. Hence, we develop a procedure to measure the set of potential buyers and hence the 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 who are either individually insured or uninsured. We then estimate the 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 $f(X_i)$ 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 $\mathcal{D}_{drt} = \{i : d_i = d, r_i = r, t_i = t\}$ as 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.

S.1.3 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 type θ_i in a given region and year who 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 5.1.