Review of Economic Studies (2025) **92**, 586–620 https://doi.org/10.1093/restud/rdae020

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

Pietro Tebaldi
Department of Economics, Columbia University and NBER, USA

First version received February 2023*; Editorial decision July* 2023*; Accepted February* 2024 *(Eds.)*

Regulations to design private yet publicly sponsored health insurance markets are increasingly adopted in many OECD countries. Here I combine data and economic theory to analyse the interaction between insurers' competition and the design of premium subsidies in determining equilibrium outcomes. My empirical model includes adverse selection, rich heterogeneity in preferences for vertically and horizontally differentiated plans and accommodates alternative assumptions on pricing conduct. In the context of the Affordable Care Act in the U.S., I estimate the joint distribution of preferences and expected cost using Californian administrative records on 3.4 million plan choices between 2014 and 2017, combined with plan and survey data on medical claims. An empirical horse race between conduct assumptions favours oligopoly pricing over perfect competition. Considering alternative subsidy designs shows that, in equilibrium, shifting subsidy generosity toward the "young invincibles" would lower premiums for all enrolees while increasing enrolment and profits.

Key words: ACA, Subsidies, Health insurance, Healthcare reform, Health exchanges

JEL codes: I11, I13, L5

1. INTRODUCTION

Welfare losses from adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976; Einav *et al.*, 2010), consumption externalities (Pauly, 1970; Summers, 1989; Mahoney, 2015), and affordability concerns (Wagstaff and van Doorslaer, 2000; Bundorf and Pauly, 2006) justify the growing role of governments in regulating and supporting premium payments in private health insurance markets (Einav and Levin, 2015). These regulations are increasingly relevant across many OECD countries (Colombo and Tapay, 2004), including the U.S. (as reviewed in Handel and Ho, 2021; Handel and Kolstad, 2022), Germany (Atal *et al.*, 2022), the Netherlands (de Ven and Schut, 2008; Roos and Schut, 2012), Switzerland (Holly *et al.*, 1998), Israel (Brammli-Greenberg *et al.*, 2018), Chile (Atal, 2019; Cuesta *et al.*, 2019), and Uruguay (Fleitas, 2020).

The strategic response of imperfectly competitive insurers to subsidy design was already highlighted for the case of prescription drugs by Decarolis (2015a, 2015b), and further

---

The editor in charge of this paper was Thomas Chaney.

analysed in Decarolis *et al.* (2020). For insurance covering medical care more broadly, Finkelstein *et al.* (2019), Jaffe and Shepard (2020), and Shepard (2022) analyse the premium subsidy program for low-income adults that played a key role in increasing healthcare access in Massachusetts since 2006. These studies consistently draw attention to how individuals and insurers are responsive to regulatory details. Therefore, understanding the ways in which subsidy design impacts market outcomes remains critical for the delineation of future policy.

For this purpose, economic theory provides useful equilibrium comparative static predictions that can be explored empirically. Given theoretical insights, quantifications in a specific context require estimates of the joint distribution of preferences and costs (Einav *et al.*, 2010) and careful modelling of *how* and *how much* insurers' compete.

In this article, I begin by noticing that, as seen also in the Netherlands and Switzerland, in the marketplaces introduced under the *Patient Protection and Affordable Care Act* (ACA) in the U.S., individuals pay subsidized premiums that may vary with income but not with age. Since expected claims and market-based pre-subsidy premiums increase with age, such subsidy design is more generous toward older individuals. Older age is also a strong predictor of willingness-to-pay for insurance. Therefore, as also noted in Graetz *et al.* (2018), this design might conflict with the goal of achieving higher levels of insurance coverage while limiting costs. Providing more generous enrolment incentives to costlier individuals increases average cost and puts upward pressure on equilibrium premiums.

Without carefully considering equilibrium responses, one might conclude that a change in subsidy design in favour of younger individuals would penalize older ones. This is the case in Tebaldi *et al.* (2023), in which demand is estimated non-parametrically but supply and equilibrium adjustments are not modelled, and differences between insurers are ignored. Considering only demand responses, lowering subsidies for relatively older individuals, and increasing subsidies for younger ones would penalize the former, albeit *average* consumer surplus would increase. Additional political economy and equity considerations pose further obstacles to a design under which older individuals would experience premium increases.

Instead, allowing premiums to re-equilibrate leads to different conclusions. Depending on the joint distribution of preferences and costs and on the intensity of competition between insurers, it might be possible to lower subsidies for older individuals, increase subsidies for younger ones, while ensuring that in the new equilibrium *all buyers face lower premiums* while total profits also increase. The intuition is simple: the changes to the subsidy design are such that the composition of enrolees becomes younger, therefore average cost is lower and elasticity of demand is higher. Both forces put downward pressure on premiums, and the resulting reduction can be sufficiently large to compensate older individuals by more than the amount by which their subsidy was lowered to begin with. This argument holds whether insurers exercise market power, although magnitudes depend on pricing conduct.

Other theoretical implications of alternative designs depend more critically on insurers' conduct. For example, market power leads to inefficiently higher prices under "price-linked" subsidies (c.f. Jaffe and Shepard, 2020), a design also adopted under the ACA and in Switzerland (Kreier and Zweifel, 2010). If insurers were perfectly competitive, as modelled theoretically in Azevedo and Gottlieb (2017), and empirically in Einav *et al.* (2010), Handel *et al.* (2015), and Dickstein *et al.* (2024), price-linked subsidies would not generate the strategic responses initially observed by Decarolis (2015a).

Taking stock of these observations, the goal of the remainder of the article is to empirically study competition and quantify alternative subsidy designs in the context of the ACA marketplaces. For this, I combine data from the first four years of the Californian marketplace, *Covered California*, with a model of equilibrium pricing that encompasses the ACA regulatory details. Importantly, the model is flexible in terms of the joint distribution of preferences and costs conditional on age, and in terms of insurers' pricing conduct, which are key determinants of the equilibrium effects of adjustments to subsidies by age and of the impact of price-linked subsidies.

The estimates of demand are obtained using individual-level premiums and enrolment data for 3.4 million plan choices observed during the 2014–2017 period, which—similarly to Finkelstein *et al*. (2019) and Tebaldi *et al*. (2023)—I combine with survey measures of uninsurance and subsidy eligibility by age, income, and geographic region. Leveraging the richness of individual-level enrolment records, I estimate a mixed-logit discrete choice model of insurance demand to obtain measures of preferences and demand heterogeneity by age.

The raw data highlight that subsidized premiums are approximately constant in age, while older individuals are significantly more likely to enrol. This is per-se suggestive of age heterogeneity in preferences. To identify demand parameters, I rely on two aspects of ACA regulations. First, discrete variation in cost-sharing reductions induces sharp discontinuities in the actuarial value of the so-called Silver plans at three income thresholds (see also Hinde, 2017; Lavetti *et al*., 2023). Second, community rating restrictions lead to a "Waldfogel instrument" identification strategy (c.f. Berry and Waldfogel, 1999; Waldfogel, 2003).[1] Indeed, age-composition is a strong predictor of regional variation in prices. Assuming that—conditional on age and income—preferences are independent from market demographics, I use a control function to correct for premium endogeneity when estimating demand.

As expected, consistently with the literature focusing on demand in health insurance marketplaces (see also Chan and Gruber, 2010; Panhans, 2019; Saltzman, 2019), I find that younger individuals are less willing to pay for insurance and more responsive to premium increases. On average, those younger than 44 value a ten percentage points increase in actuarial value less than \$350 per-year. Older individuals value this more than \$400, and more than \$700 when older than 55. If monthly premiums increase by \$10, enrolment of individuals younger than 44 would drop by more than 6%, while enrolment among those older than 55 would be 3–3.8% lower. In terms of scope for market power, I estimate an average elasticity between 1.4 and 2.5 for the "Silver plans" chosen by 68% of enrolees.

To estimate expected insurance costs incorporating adverse selection, the model employs plan-level average claims data (as in Bundorf *et al*., 2012) and individual-level healthcare spending information from the Medical Expenditure Panel Survey. Expected annual medical spending can vary across individuals and plans. Not having access to individual claims data, the baseline model rules out moral hazard,[2] while it captures selection by letting expected medical spending for an individual vary observably with age, and unobservably with willingness-to-pay for insurance generosity. Variation in the composition of buyers across plans inducing variation in average claims identifies heterogeneity in costs across individuals with different willingness-to-pay.[3]

Cost estimates indicate adverse selection, due to the strong correlation between preferences for coverage and expected costs. An age increase of ten years implies 38% higher medical

---

[1] The intuition is that the ACA allows insurers to set only one baseline premium for every plan in each geographic region. Then, pre-determined pricing schedules are used to transform baseline premiums to the premiums faced by buyers of different age. Because this regulation links profits across heterogeneous buyers to the same univariate decision, when setting base prices insurers must consider the composition of buyers (see also Orsini and Tebaldi, 2017).

[2] Appendix B shows that my results on the effect of subsidy design are robust to allowing for a degree of moral hazard significantly more severe than what it is assumed in the ACA risk adjustment model (Pope *et al*., 2014), or estimated in Lavetti *et al*. (2023).

[3] Supplemental Appendix S1 (available online) introduces a novel result providing sufficient conditions for identification of cost curves in selection markets from supply-side assumptions; this adapts to selection markets results dating back to Rosse (1970) and Bresnahan (1987).

spending. An increase in willingness-to-pay (for ten percentage points in actuarial value) of \$500 per year implies 40% higher medical spending.

Prior to considering alternative subsidy designs I use the estimates of the model to set up a horse race between alternative conduct assumptions. Although I fall short of providing a formal statistical test, empirical support for alternative supply models is desirable because, as noted above, conduct impacts the effect size of counterfactual designs. For this exercise to be conceptually sound, it is important to highlight that demand and cost estimates are obtained without imposing any conduct assumption.

Combining demand estimates with the details of rating regulations, subsidy design, and risk adjustment, I can compute average cost, average revenue, marginal cost, and marginal revenue for each plan. I find that risk-adjusted marginal revenues are, on average, 4.9% [3.4%, 6.4%] larger than marginal costs. In comparison, average revenues are, on average, 26% [23.7%, 28.3%] larger than average cost. This shows that—relative to the perfect competition benchmark—oligopoly pricing appears more consistent with observed market outcomes.

When calculating equilibrium under counterfactual subsidy designs, I begin by corroborating the findings of Decarolis (2015a, 2015b) and Jaffe and Shepard (2020) in the ACA context, I find that under oligopoly pricing price-linked subsidies increase premiums, and lower enrolment and consumer surplus. This effect is more pronounced in concentrated markets (with 2 or 3 insurers), where I find that premiums would drop, on average, by 27% and enrolment would increase by 12.5% (from 0.32 to 0.36) if subsidies were not linked to premiums. In markets with a larger number of insurers, I find a more moderate effect: premiums would drop by 5.5% and enrolment increase by 3.5%. Price-linked subsidies are non distortionary when I simulate competitive equilibria.

Shifting subsidy generosity away from older buyers and toward younger ones leads to equilibria in which buyers face lower premiums, while total profits and consumer surplus increase. This result is robust to different subsidy designs and conducts assumptions. Using vouchers, where I increase the under-35 by \$600 and lower the over-35 by \$100, I find that—in equilibrium—premiums for the under-35 would decrease by \$553, while premiums for the over-35 would decrease by \$76 (despite the lower subsidy). Consumer surplus would increase by \$104 per-person per-year, and per-enrolee subsidy would be \$68 lower. Despite the lower per-person subsidy, because total enrolment is higher (under-35 enrolment increases from 0.28 to 0.39, and over-35 enrolment from 0.32 to 0.33) total government spending increases.

Ultimately, the trade-off between lower uninsurance and total government spending is a matter of political economy debate that is far beyond the scope of this article. What I want to highlight here is that, in a market with adverse selection, it is possible to shift generosity of subsidies avoiding that any group of market participants is worse off.[4] It is also important to note that, although subsidized prices are a key driver of enrolment, there can be other effective ways to incentivize specific groups to participate in the market and affect equilibrium (Cox *et al.*, 2015; Domurat *et al.*, 2021).

There are other reasons to be wary of making (utilitarian) welfare conclusions. First, I do not model competing reasons for individuals to select plans. A growing literature (see *e.g.* Domurat *et al.*, 2021; Drake *et al.*, 2022; Saltzman *et al.*, 2021) documents behavioural biases (*e.g.* lack of information or inertia) that affect enrolment decisions and might need to be accounted for when calculating consumer surplus. Second, since I discuss counterfactual subsidy designs that would

---

[4] I avoid aggregate welfare considerations that would require to put a weight on public spending in this market. For a discussion about why regulators might want to subsidize health insurance although individuals value it less than its costs, I refer the reader to Finkelstein *et al.* (2019) and references therein.

increase total profits, this might trigger entry, a decision that I treat as exogenous and fixed. Common wisdom might suggest that "the more the merrier", as a larger number of participating insurers would lead to stiffer competition and benefit government and consumers. However, in markets with asymmetric information the effect of entry can be ambiguous, and the interaction between public policy, entry, and market outcomes remains an important object for future research (Ryan, 2023).[5] Lastly, here I focus only on equilibrium adjustments to premiums. This is first-order in the Californian marketplace in which the regulator has determined tight requirements for non-premium plan characteristics (see Section 2). However, if insurers could easily adjust other contract characteristics (Miller *et al.*, 2022), increasing the share of young invincibles in a market could also trigger quality adjustments and therefore affect consumer surplus in ways that might differ across age groups. For all these reasons, here I focus primarily on enrolment and spending while further work is needed to accurately measure changes in welfare under alternative policies.

Other related literature. In addition to the aforementioned articles, here I speak directly to a growing body of work analysing the effect of different regulations in government-sponsored health insurance markets, reviewed in Handel and Ho (2021). For the U.S., the analysis of competition and market design in Medicare Advantage and Medicare Part D is more mature, with focus on subsidies in Decarolis (2015a), Decarolis *et al*. (2020), Curto *et al*. (2021) and Miller *et al*. (2022). General studies on equilibria in health insurance exchanges are pioneered by Handel *et al*. (2015), and theoretical implications of alternative policies are the focus of Mahoney and Weyl (2017) and Veiga (2023).

The US health insurance market for those younger than 65 has been analysed primarily through the lenses of the Massachusetts healthcare reform, which served as a blueprint for the ACA (Gruber, 2010). Graves and Gruber (2012) shows the effects on premiums, and Hackmann *et al*. (2015) the effect on enrolment and costs. The role of mandates is considered in Chandra *et al*. (2011), Sommers *et al*. (2018), and then studied under the ACA by Saltzman (2019). Risk adjustment is the focus of Geruso *et al*. (2019), McGuire *et al*. (2020) (which extend the analysis to Germany and the Netherlands), and Saltzman (2021). Panhans (2019) measures adverse selection in Colorado. Fang and Krueger (2022) focus on the ACA impact on labour markets. Geruso *et al.* (2023) study the interaction policies and the two margins of enrolment and coverage choice. Dickstein *et al.* (2024) analyse the relationship between individual marketplaces and small-group insurance. Marone and Sabety (2022) consider the choice of whether or not to provide vertically differentiated plans, Polyakova and Ryan (2019) measure the incidence of subsidies across demographic groups, and Cicala *et al*. (2019) the regulations limiting insurers' markups. Reviews of the ACA and related literature are provided in Blumenthal *et al.* (2020) and Handel and Kolstad (2022).

2. ACA REGULATIONS AND DATA

2.1. *Institutional background and regulations*

As of 2013, 17% of U.S. citizens younger than 65 did not have health insurance coverage (Smith and Medalia, 2014). To address this, in 2014 the ACA instituted health insurance marketplaces in each of the 50 states. ACA marketplaces operate separately across states, but they all follow

---

[5] Indeed, the Californian regulator has been monitoring and sometime limiting entry by large insurers. See *e.g.* https://www.latimes.com/business/la-fi-obamacare-unitedhealth-20150116-story.html.

TABLE 1

Standardized plan characteristics in 2015 covered California

Panel (a): Characteristics by metal tier before cost-sharing reductions

Tier	Annual deductible	Annual max out-of-pocket	Primary visit	E.R. visit	Specialist visit	Preferred drugs	Advertised AV
Bronze	\$5,000	\$6,250	\$60	\$300	\$70	\$50	60%
Silver	\$2,250	\$6,250	\$45	\$250	\$65	\$50	70%
Gold	\$0	\$6,250	\$30	\$250	\$50	\$50	80%
Platinum	\$0	\$4,000	\$20	\$150	\$40	\$15	90%

Panel (b): Silver plan characteristics after cost-sharing reductions

Income (%FPL)	Annual deductible	Annual max out-of-pocket	Primary visit	E.R. visit	Specialist visit	Preferred drugs	Advertised AV
200–250% FPL	\$1,850	\$5,200	\$40	\$250	\$50	\$35	74%
150–200% FPL	\$550	\$2,250	\$15	\$75	\$20	\$15	88%
100–150% FPL	\$0	\$2,250	\$3	\$25	\$5	\$5	95%

Source: Section 6,460 of title 10 of the California Code of Regulations; 21 May 2014.

similar institutions and regulations as mandated by the federal reform.[6] Following the ACA, the Tax Cuts and Jobs Act of 2017, the American Rescue Plan Act of 2021, and the Inflation Reduction Act of 2022 modified certain provisions of the federal law. These changes took place following my study period in this article, yet it will be important for future work to tailor empirical models to the most current regulatory design.

Rating regions. A state is divided into geographic rating regions—groups of counties or zip codes—defining the level at which decisions by buyers and insurers take place (Dickstein *et al.*, 2015). Insurers can decide whether to offer plans and cover individuals in any given region, as long as they can offer an adequate network of healthcare providers. Different plans are classified into five coverage levels: Catastrophic, Bronze, Silver, Gold, and Platinum.

Metal tiers. The four metal tiers represent increasing generosity of insurance, measured (and advertised) as "actuarial value", an estimate of the share of healthcare spending covered by the plan: 60% for Bronze, 70% for Silver, 80% for Gold, and 90% or more for Platinum. Catastrophic plans have higher cost sharing, and generally cannot be purchased by subsidized buyers, nor by buyers older than 30, with few exceptions.[7]

In some states, including California, regulators have determined that, within each metal tier, cost-sharing characteristics are fully standardized across insurers. Deductible, coinsurance, and copayments are fixed. Plans still differ in terms of brand, hospital networks, and possibly Rx formularies. Table 1 summarizes a number of plan characteristics for each metal tier, as mandated by Covered California.

Adjusted community rating. One important provision of the ACA is that insurers are not allowed to freely adjust premiums as a function of a buyer's observable characteristics. Characteristics that can affect annual premiums are the buyer's age (see also Ericson and Starc, 2015; Orsini and Tebaldi, 2017) and, in some states, tobacco use, but even these adjustments are done

---

[6] States can choose between instituting their own marketplace, relying on the federal platform, or adopting a state-federal partnership model.

[7] Source: https://www.kff.org/health-reform/issue-brief/explaining-health-care-reform-questions-about-health-insurance-subsidies/; last accessed on 26 January 2022.

in a pre-specified way. California does not allow tobacco-based premium adjustments; therefore, here I focus on age-adjustments, which are central to my analysis.

Considering a rating region, each plan j is associated with a single "base" premium, say b_j. This is translated to age-adjusted (pre-subsidy) premium using given age adjustment factors, equal for all products. As shown in (1) below, when covering a buyer i under plan j, the insurer receives a revenue R^i_j equal to the product of b_j and the corresponding age adjustment, an increasing function of Age_i.[8]

$$\begin{aligned}
\textbf{Insurer decision:}\quad & \text{base premium } b_j \\
\textbf{Insurer revenue:}\quad & R^i_j = b_j \times \text{Adjustment}(\text{Age}_i) \\
\textbf{ACA subsidy:}\quad & S^i = \max\{0, R^i_{j^{2S}} - \overline{P}(\text{Income}_i)\},\ j^{2S} = \text{2nd-cheapest Silver} \\
\textbf{ACA premium:}\quad & P^i_j = \max\{0, R^i_j - S^i\}.
\end{aligned} \tag{1}$$

Premium subsidies. Although R^i_j is the amount collected by the insurer, enroled individuals who are eligible for premium tax credits—or simply subsidies henceforth—pay less than this amount. Eligibility and subsidy generosity are determined by the individual household's annual income: if this is less than four times the federal poverty level (FPL), the individual premium for the second cheapest Silver plan in the region is capped at a federally mandated maximum affordable amount. The resulting subsidy applies to any plan available in the region. This subsidy design is described formally in (1) above. For individual i, the premium of the second cheapest Silver plan in the region is capped at the maximum affordable amount equal to $\overline{P}(\text{Income}_i)$, and the individual-specific subsidy amount S^i is calculated to match this constraint. The premiums for all plans are lowered by S^i; subsidized premiums must be positive.

Under this subsidy design, for a given income level, individuals of different age can enrol in a Silver plan for exactly the same premium. Differences in subsidized premium across insurers and plans are instead increasing in age, while not varying with income. As a result, all plans with base premiums lower than the second cheapest Silver—which generally include all Bronze plans—are cheaper for older buyers, holding income fixed. Conversely, plans with base premiums higher than the second cheapest Silver—which generally include all Gold and Platinum plans—are more expensive for older buyers.

Cost-sharing reductions. Another ACA regulation relevant during my study period is the provision of cost-sharing reductions, available for individuals who enroled in a Silver plan with income lower than 2.5 times the FPL. For this group, the federal government covers part of their out-of-pocket spending, de facto increasing the actuarial value of Silver plans from 70% to 95% for income levels between 1 and 1.5 times the FPL, 88% for income levels between 1.5 and 2 times the FPL, and 74% for income levels between 2 and 2.5 times the FPL. Covered California achieved these changes in actuarial value in a standardized way, by altering deductible and copayments as summarized in Table 1.[9]

Risk adjustment. To limit concerns of cream skimming by insurers, the ACA introduced a budget-neutral scheme of risk-adjustment transfers. Simply put, insurers covering enrolment

---

[8] The age adjustment is equal to 1 for 21-year-old buyers, and increases smoothly to 1.4 at age 45, and finally reaches 3 at age 64. Details for all ages are shown in figure 2b.

[9] Following my study period, the Tax Cuts and Jobs Act of 2017 interrupted the funding of cost-sharing reductions, after a legal dispute over the appropriation of federal funds: c.f. *House versus Burwell*, *House versus Price*.

pools that end up being riskier than the market average receive transfers from their competitors; these transfers, by construction, add up to zero within the state. As described formally in Pope *et al.* (2014), the transfer applying to each plan is calculated by multiplying the state-level average revenue by a plan-level risk score, which can be positive or negative. The score is positive if the enrolees selecting the plan are riskier than the state average, after adjusting for the factors that are already priced in (*e.g.* age, geography, and metal tier), and it is negative otherwise. Saltzman (2021) studies the implications of ACA risk adjustment for equilibrium outcomes; here I model it and then hold it fixed throughout my analysis.[10]

Other regulations. Other ACA regulations included two temporary market stabilization programs, reinsurance and risk corridors, income-based tax penalties for individuals not purchasing coverage, and a minimum medical loss ratio of 80%.[11] I do not model these explicitly, a simplification partly dictated by data limitations. Incorporating these policies in a tractable empirical model is left to future work.

Coverage options and premiums are set and made public before the beginning of open enrolment, which takes place during the late months of each calendar year. Eligible individuals compare and purchase plans offered in their region of residence; coverage lasts for the following calendar year, as long as premium payments are honoured. Diamond *et al.* (2018) recently discuss the relationship between medical spending and interruptions of premium payments.

2.2. *Data sources and summary statistics*

2.2.1. Enrolment files. Covered California provided me with individual-level enrolment files covering the 2014–2017 period, in response to four Public Records Acts requests. For every purchase event, I observe individual and household identifiers, along with age, zip code, county, rating region, plan identifier, premium paid, and income group. Income is reported in discrete bins, but one can use the pricing regulations in (1) to determine income with higher precision, I use 5% FPL bins.

As in Finkelstein *et al.* (2019), I narrow my focus to adults aged 26–64, without dependent children, and beneficiaries of premium subsidies. This group accounts for 78% of plan selections in Covered California during my observation period, for a total of 3.38 million individuals.

---

[10] Risk adjustment in ACA marketplaces does not feature any payments from the government. This is radically different from non-budget-neutral risk adjustment schemes in which the government provides risk-based transfers to each insurer, as it is the case in other federally sponsored markets such as Medicare Advantage (Brown *et al.*, 2014; Geruso and Layton, 2020), or Medicare Part D (Decarolis, 2015a; Decarolis *et al.*, 2020).

[11] Federal reinsurance was mandatory between 2014 and 2016, collecting a fixed amount for every policy sold by any issuer ($63, $44, and $27 in 2014, 2015, and 2016, respectively), and compensating a share (100%, 50%, 50%) of claims between an attachment point ($45,000, $45,000, $90,000), and a cap ($250,000, equal for all three years). Risk corridors were intended to facilitate a target variable profit margin of 20% between 2014 and 2016. Insurers not spending at least 77% of premiums in claims would pay into the system, and insurers spending more than 83% would be eligible for funds. The program was not guaranteed to pay out, since dues could be larger than revenues. For example, in 2014 insurers were due a total of $2.8 billion, while only owing $362 million; the program paid only 12 cents for every dollar owed to insurers. An "individual mandate" tax penalty (see *e.g.* Saltzman, 2019) was charged to individuals choosing to remaining uninsured, and not qualifying for exemptions. These included "affordability exemptions". As a result, the individual mandate was only weakly enforced, particularly in the subsidy-eligible population I study in this article. Penalty revenues did not exceed 20% of hypothetical penalty payments (Miller, 2017), and the mandate penalty was ultimately lifted (starting in 2019) by the Tax Cuts and Jobs Act of 2017. Medical-loss-ratio adjusted for quality improvements is a measure of the share of an insurer's collected premiums spent in medical claims and quality improvements. Under the ACA, this ratio must not be less than 0.8. Other studies (*e.g.* Starc, 2014) have leveraged these limits explicitly to estimate empirical models of insurance supply. In my application, I do not impose medical loss ratio regulations; I estimate an average medical loss ratio of 0.85, and this remains above 0.8 across all my counterfactuals.

TABLE 2

Summary statistics

	Individual-level data (person-year)					
	Enroled (Covered CA) $N = 3, 381, 971$		Eligible (ACS draws) $N = 12, 433, 055$		Surveyed (MEPS) $N = 20, 171$	
Age	45.8	(11.7)	44	(11.4)	43.8	(11)
Income (FPL %)	214.2	(63.7)	233.6	(75.5)	257.2	(81.1)
Annual Premium	1, 477	(1, 265)	–	(–)	–	(–)
Annual Subsidy	3, 928	(2, 625)	–	(–)	–	(–)
Medical Spending	–	(–)	–	(–)	4, 111	(12, 900)
Choose Bronze (0/1)	0.241	(0.428)	–	(–)	–	(–)
Choose Silver (0/1)	0.682	(0.466)	–	(–)	–	(–)
Choose Gold (0/1)	0.041	(0.199)	–	(–)	–	(–)
Choose Platinum (0/1)	0.035	(0.185)	–	(–)	–	(–)
	Plan-level data (region-year-insurer-tier)					
	Market share within region-year (Covered CA) $N = 1{,}104$		Base prem. quantity-weighted (Covered CA) $N = 1{,}104$		Avg. claims quantity-weighted (RRF) $N = 1{,}099$	
By insurer:						
Anthem (76 region-years)	0.058	(0.122)	3, 019	(519)	3, 804	(759)
Blue Shield (76 region-years)	0.09	(0.16)	3, 190	(631)	4, 133	(1, 851)
Health Net (33 region-years)	0.033	(0.084)	2, 685	(350)	3, 317	(1, 663)
Kaiser (69 region-years)	0.102	(0.129)	3, 228	(648)	4, 210	(2, 015)
Other nine insurers	0.035	(0.07)	2, 554	(583)	2, 296	(1, 727)
By metal tier:						
Bronze	0.072	(0.079)	2, 445	(378)	2, 199	(935)
Silver	0.178	(0.189)	3, 107	(525)	3, 908	(1, 233)
Gold	0.012	(0.026)	3, 679	(664)	4, 834	(1, 658)
Platinum	0.01	(0.012)	4, 177	(720)	9, 089	(3, 890)

Note: The table summarizes data sources. In the Enroled panel, each observation is an individual in the Covered California enrolment sample, covering all purchases that took place during the 2014–2017 period, restricted to subsidized adults without dependent children. The Eligible panel corresponds to the sample of individuals constructed from the American Community Survey, consisting of subsidy-eligible adults who are either uninsured or privately insured, covering the 2013–2016 period. The Surveyed panel corresponds to the 2014–2017 Medical Expenditure Panel Survey, restricted to individuals who are privately insured and with income between 100 and 400% FPL. The panels of Market shares and Base premiums report summary statistics from the Covered California enrolment sample. The Average claims panel summarized the 2016–2019 rate review filings matched to the Covered California sample. Standard deviations in parentheses.

Excluding dependents, who under the ACA can be as old as 25, the coverage decisions for this group are simpler, and easier to analyse. Moreover, since off-exchange plans are not eligible for subsidies, excluding the unsubsidized population mitigates concerns that enrolment files may miss many individuals purchasing coverage outside the marketplace.

The top-left panel of Table 2 summarizes the enrolment data. Average age among subsidized adults in Covered California is 45.8 (with standard deviation 11.7), while average income is 214.2 (63.7)% of the FPL. Individuals pay, on average, \$1,477 (\$1,265) per-person, per-year, receiving subsidies that are, on average, more than 2.5 times as large. In terms of metal tier, 24% of enroled individuals choose a Bronze plan, while 68% choose a Silver plan. Gold and Platinum plans are selected more rarely.

Figure 1a plots how insurer revenue, subsidized premium, and the difference between Bronze and Silver premium vary across enrolees of different age. The average amount collected by the issuers increases in age, from \$3,000 per-year on average at 26 to over \$8,000 for buyers older

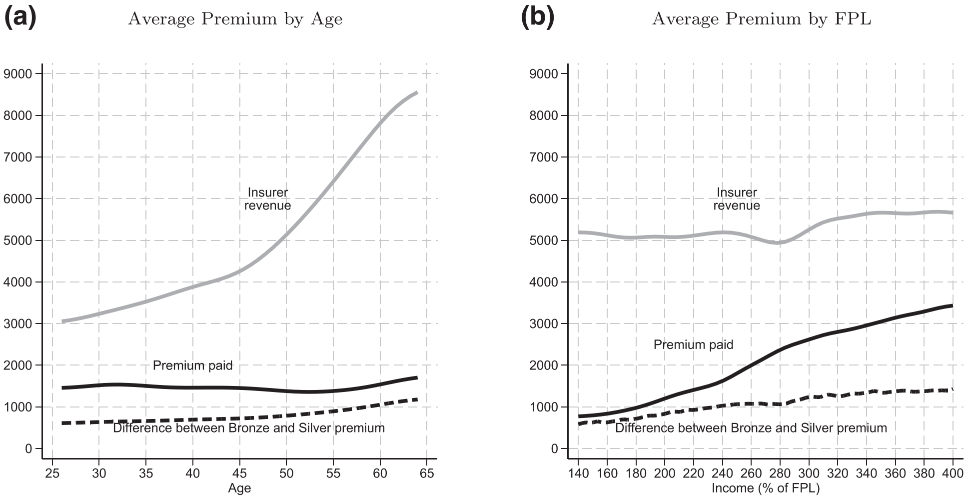

FIGURE 1

Premiums by age and income

Note: The figure illustrates the relationship between average revenue collected by the insurer (gray line), average subsidized premium paid by the individual (black line), and average difference between Bronze and Silver premiums for the individual (dashed line), as a function of age (left panel) and FPL (right panel). For revenue and premium, each observation is one individual in the Enrolment sample, for the difference between Bronze and Silver premium, each observation is one individual in the Eligible sample. Local polynomial with Gaussian kernel; bandwidth = 2 for panel 1a, bandwidth = 10 for panel 1b.

than 60. According to the ACA subsidy design, subsidized buyers do not face these increases. Premium paid is approximately constant in age, with very small variations around its average value due to differences in plan selection. At the same time, the average difference between the subsidized premiums of Bronze and Silver plans is increasing in age, from approximately \$800 to \$1,200 per-year; older individuals have to pay a higher amount to obtain more generous coverage. The relationship between income and premium is illustrated in Figure 1b. Average insurer revenues do not differ too much across individuals with different income, while premium paid is increasing, since subsidies become lower.

The bottom-left of Table 2 summarizes market shares at the plan level (insurer-year-region-metal-network; $N = 1,104$), there are between 3 and 7 insurers active in every region-year combination. Four players—Anthem, Blue Shield, Health Net, and Kaiser—are present across a large number of markets, while the nine remaining insurers are only available in a small number of regional markets, or for a limited number of years. Market shares of Anthem, Blue Shield, Health Net, and Kaiser are, on average, between 3.3 and 10%, but they vary widely across regions and years, reflecting differences in premiums, set of competitors, provider network or brand attractiveness. In terms of metal tier, a single Silver plan covers, on average, 17.8% of enrolees in a region-year pair, more than twice as large as the average share of Bronze plans. A Gold or Platinum plan covers, on average, 1% of the market.

2.2.2. Rate review filings. I use realized claims information as reported in the annual Rate Review Filings (RRF); these are released by the Center of Medicare & Medicaid Services, and publicly available.[12] As in Bundorf *et al.* (2012) and Saltzman (2021), while I observe enrolment at a granular, individual-level data, my cost measures are aggregated to a coarser level, and noisier. Enriching my analysis to incorporate individual-level claims information would be an

---

[12] Source: https://www.cms.gov/CCIIO/Resources/Data-Resources/ratereview.

important extension of my work, which would be particularly relevant to obtain more precise, externally valid measures of the effect of counterfactual policies.

In the RRF, insurers must declare average experienced claims per-member month. For rate review taking place in 2016, the experience period is 2014; for 2017 rate reviews, the experience period is 2015; and so on and so forth. My analysis uses 2016–2019 RRF. I link RRF to Covered California enrolment files using HIOS-14 (a plan-insurer identifier), enrolment year, and metal tier information. The resulting sample of plans for which I observe a measure of realized average claims consists of 1,099 unique insurer-region-year-tier-network combinations, which covers 99.5% of the 1,104 plans I observe in the enrolment data and use in my analysis.[13] In terms of enrolment, the sample of plans for which I observe RRF information covers 76% of the 3.4 million individuals included in my enrolment sample.

The bottom-right of Table 2 reports the summary statistics of realized average claims, by insurer, and by metal tier. Differences across insurers reflect a combination of plan selection, risk composition of enrolment pools, regional heterogeneity, and differences in firms' cost functions. Costs vary widely across metal tiers. A Bronze plan records, on average, claims amounting to $2,199 per-enrolee, per-year (with standard deviation $935). This compares to Silver plans, with average claims for $3,908 ($1,233) per-enrolee-year, and Gold plans, with average claims for $4,834 ($1,658). Platinum plans register much higher claims, with an average of $9,089 per-enrolee per-year.

2.2.3. Survey data.

American community survey. I construct measures of potential buyers by age, income, rating region, and enrolment year using the American Community Survey (ACS) public use file, downloaded from IPUMS (Ruggles *et al*., 2015). The procedure is similar to the one adopted by Finkelstein *et al*. (2019) and Tebaldi *et al*. (2023).[14]

As shown in Table 2, eligible buyers are, on average, two years younger and higher income (+20% FPL) relative to marketplace enrolees. Figure 2a shows more details of the relationship between age and the share of potential buyers choosing to purchase marketplace coverage, measured after combining enrolment files with the ACS. The monotone relationship between age and enrolment is evident: the average enrolment probability among under-40 individuals is between 0.22 and 0.25, this then increases with age until 0.38 for individuals aged between

---

[13] Some plans change HIOS-14 code over time or leave the marketplace. When this is the case, I cannot match enrolment to RRF. Sometimes groups of plans offered by the same insurer in the same year report the same measure of average claims, pooling across metal tiers, or pooling across rating regions. This adds noise to my measures of realized costs.

[14] For every year between 2013 and 2016, I use the corresponding 5-year ACS sample to measure potential marketplace enrolees for the following enrolment year. Each individual is a potential buyer in the marketplace if they report being either uninsured or privately insured. For every buyer, I observe age, household income, a person weight, and the public use micro data (PUMA) area of residence. Using a PUMA-to-county crosswalk, I assign individuals to the Covered California rating regions. An adjustment to this procedure is needed to account for the fact that the PUMA identifiers can be split across multiple counties, and so in some cases also multiple ACA rating regions. I allocate individuals to each rating region it overlaps using the population of the zip codes in the PUMA as weights. Finally, I merge enrolees and potential buyers for every year, rating region, age, and income cell (in 5% FPL bins). Using person weights, this leaves me with 13,265,960 (synthetic) potential buyers for the 2014–2017 enrolment years, which I then match to the enrolment file. For example, if in the 2013 ACS there are three individuals who are either uninsured or privately insurer, live in Region 5, are aged 50, and have income between 150 and 155% FPL, and the sum of their person weights is 20, the dataset of potential buyers contains 20 individuals in 2014, Region 5, age 50, and FPL cell 150–155. If there are five enrolees in the same year-region-demographic combination, I measure a total marketplace share conditional on these observables equal to 0.25.

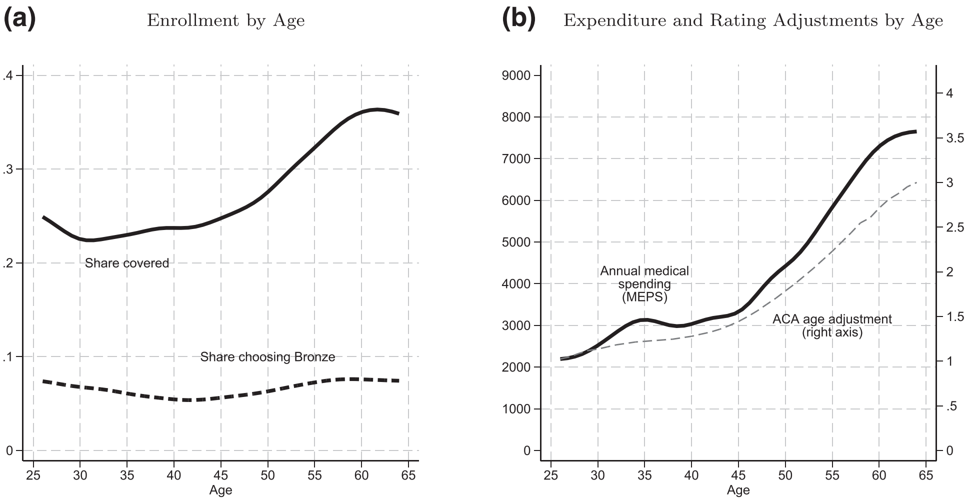

FIGURE 2

Enrolment, medical spending, and rating adjustments by age

Note: The solid (dashed) black line in the left panel illustrates the relationship between age and the probability of choosing a marketplace (Bronze) plan, measured in the Eligible sample. Local polynomial with Gaussian kernel; bandwidth = 2. The solid black line in the right panel illustrates the relationship between age and annual medical expenditure in the Medical Expenditure Panel Survey; Gaussian kernel with bandwidth = 2. The dashed grey line in the right panel indicates for every age the corresponding ACA age rating adjustment—Adjustment(Age_i) in (1)—measured on the right vertical axis.

60 and 64. Relating this pattern to the fact that average premium paid does not increase in age (Figure 1a) suggests that older individuals are more willing to pay for marketplace coverage. This is supported further by the extent to which the share of individuals choosing a Bronze plan is approximately constant in age, despite the increasing difference in premium relative to other tiers.

Medical expenditure panel survey. The last dataset employed in my analysis consists of the 2014–2017 public use files of the Medical Expenditure Panel Survey (MEPS; https://meps.ahrq.gov/), measuring medical spending for a representative sample of the U.S. population. I focus on individuals who are privately insured, with age and household income in the same range as the observations in the enrolment data. The resulting sample of 20,171 individuals is summarized in Table 2. Average annual medical spending is equal to $4,111, with standard deviation $12,900. In the next section, these data are used to estimate a parameter describing the relationship between age and total medical spending conditional on being insured, controlling for differences across years and macro areas observed in the survey.

Figure 2b plots the relationship between average annual medical spending as a function of age. The graph also shows—measured on the right axis—the ACA age adjustments to pre-subsidy premium. The ratio of a plan revenue from a 64-year-old to revenue from a 26-year-old is 3, while in the Medical Expenditure Panel Survey the ratio of medical spending between the two age groups is higher than 3.5. Average medical spending is slightly higher than $2,000 at 26, approximately $4,000 at 47 and higher than $7,500 after 60.

3. EMPIRICAL MODEL

3.1. *Demand*

A potential buyer i is defined by a pair $(\mathbf{z}_i, \theta_i)$, where $\mathbf{z}_i$ is a vector of observed characteristics (age, income, and rating region: $\mathbf{z}_i = (z_i^{\text{Age}}, z_i^{\text{Inc}}, z_i^{\text{Reg}})$), while θ_i is a scalar unobservable which may affect preferences for insurance and expected costs. If the base premium for plan j in region m and year t is b_{jmt}, with $\mathbf{b}_{mt} = \{b_{1mt}, \ldots, b_{Jmt}\}$, the premium paid by i when choosing j is $p_{ijmt} = P_j(\mathbf{b}_{mt}, \mathbf{z}_i)$; the function P captures age adjustments and subsidies, as defined by the regulations in (1).

The random indirect utility of i when purchasing j in region m, year t, is defined by $u_{ijmt} = -\alpha_t(\mathbf{z}_i)p_{ijmt} + \delta_{jmt}(\mathbf{z}_i, \theta_i) + \varepsilon_{ijmt}$, where

$$\delta_{jmt}(\mathbf{z}_i, \theta_i) \equiv \beta_t\,(\mathbf{z}_i, \theta_i)\,AV_{ij}^{D} + \boldsymbol{\mu}_t(\mathbf{z}_i)\mathbf{x}_{jmt} + \gamma_t(\xi_{jmt}; \mathbf{z}_i);$$

for $j = 0$, corresponding to not purchasing marketplace coverage, $p_{i0mt} = \delta_{i0mt} = 0$. This is a normalization; the premium for each plan can be interpreted as net of the expected tax penalty. The error terms ε_{ijmt} are drawn iid from the type one extreme value distribution. The premium coefficient $\alpha_t(\mathbf{z}_i)$ varies across years, and across observable characteristics $\mathbf{z}_i$. The same applies to the coefficient on actuarial value AV_{ij}^{D} (as observed by individuals upon selecting plans, reflecting cost-sharing reductions), but this coefficient can also vary along the unobservable dimension θ_i. The vector $\mathbf{x}_{jmt}$ collects a constant term, and indicators for insurers, and HMO provider networks, with coefficients collected in $\boldsymbol{\mu}_t(\mathbf{z}_i)$ varying across $\mathbf{z}_i$ and t.

Importantly, the scalar-valued term ξ_{jmt} represents unobservable characteristics specific to a *jmt* triplet (*e.g.* quality and breadth of provider networks, drug formularies, or brand preferences), which affect utility through the function γ_t. Being known to insurers, these characteristics can affect pricing decisions, and must be accounted for to avoid endogeneity concerns when estimating demand.

Following McFadden (1973), the probability of purchasing j in region m, year t, for individuals with characteristics $(\mathbf{z}_i, \theta_i) = (\mathbf{z}, \theta)$ is

$$q_{jmt}(\mathbf{z}, \theta) = \frac{e^{-\alpha_t(\mathbf{z}_i)P_j(\mathbf{b}_{mt}, \mathbf{z}_i) + \delta_{jmt}(\mathbf{z}, \theta)}}{1 + \sum_{k=1}^{J} e^{-\alpha_t(\mathbf{z}_i)P_k(\mathbf{b}_{mt}, \mathbf{z}_i) + \delta_{kmt}(\mathbf{z}, \theta)}}. \tag{2}$$

Given the distribution G_{mt} of $(\mathbf{z}, \theta)$ in region m, year t, total enrolment in plan j is

$$Q_{jmt} = \int q_{jmt}(\mathbf{z}, \theta)dG_{mt}(\mathbf{z}, \theta). \tag{3}$$

The difference between the demand model in (3) and standard discrete choice models with heterogeneous consumers (*e.g.* Nevo, 2001) lies in how rating regulations and subsidies determine enrolment responses to insurers' pricing decisions.

Taking the partial derivative of enrolment of plan j with respect to the base premium of plan k one obtains

$$\begin{aligned}\frac{\partial Q_{jmt}}{\partial b_{kmt}} &= \int \frac{\partial q_{jmt}(\mathbf{z}, \theta)}{\partial b_{kmt}} dG_{mt}(\mathbf{z}, \theta)\\ &= \sum_{\ell=1}^{J} \int \frac{\partial P_\ell(\mathbf{b}_{mt}, \mathbf{z})}{\partial b_{kmt}} \left(\alpha_t(\mathbf{z})q_{jmt}(\mathbf{z}, \theta)q_{\ell mt}(\mathbf{z}, \theta)\right) dG_{mt}(\mathbf{z}, \theta).\end{aligned} \tag{4}$$

Equation (4) highlights how changes in base premiums do not affect enrolment directly, since the effect on premiums paid by consumers is mediated by the term $\frac{\partial P_\ell(\mathbf{b}_{mt}, \mathbf{z})}{\partial b_{kmt}}$. This captures the change in premium of plan ℓ charged to buyers with characteristics $\mathbf{z}$ in response to an infinitesimal change in the base premium of plan k. Under the ACA, the regulations in (1) imply that, if k is the second cheapest Silver plan in the region, $\frac{\partial P_k(\mathbf{b}_{mt},\mathbf{z})}{\partial b_{kmt}} = 0$, while, for all $\ell \neq k$, $\frac{\partial P_\ell(\mathbf{b}_{mt},\mathbf{z})}{\partial b_{kmt}} < 0$. For other plans, $\frac{\partial P_k(\mathbf{b}_{mt},\mathbf{z})}{\partial b_{kmt}} = \text{Adjustment}(z_i^{\text{Age}})$, while for all $\ell \neq k$, $\frac{\partial P_\ell(\mathbf{b}_{mt},\mathbf{z})}{\partial b_{kmt}} = 0$.

3.2. *Cost*

The insurer expected claims from covering an individual i with characteristics $(\mathbf{z}_i, \theta_i)$ under plan j, in region m, year t are equal to

$$\kappa_{jmt}(\mathbf{z}_i, \theta_i) = AV_j^S L_{jmt}(\mathbf{z}_i, \theta_i), \qquad \text{where} \qquad L_{jmt}(\mathbf{z}_i, \theta_i) = e^{\phi_{jmt} + \eta(\mathbf{z}_i, \theta_i)}. \tag{5}$$

Claims are the product of the actuarial value of a plan (for some plans $AV_j^S \neq AV_{ij}^D$ due to cost-sharing reductions) and the expected total health expenditure of the individual, $L_{jmt}(\mathbf{z}_i, \theta_i)$, which may vary with individual and plan characteristics. Differences in claims across individuals define the main feature of a selection market: buyers with different preferences have different risk and expected insurable costs. Differences in claims across insurers, regions, and years, reflect differences in provider networks, negotiated prices, and insurers' strategies to manage their members' access to healthcare.

Importantly, the cost model specified in (5) does not allow expected medical spending to vary with coverage generosity, ruling out "moral hazard" (c.f. Einav and Finkelstein, 2018). In Appendix B, I relax this assumption, estimating cost functions and reproducing my main results for a range of moral hazard parameters.

At the plan level, expected average cost is equal to

$$AC_{jmt} = \frac{1}{Q_{jmt}} \int \kappa_{jmt}(\mathbf{z}, \theta) q_{jmt}(\mathbf{z}, \theta) dG_{mt}(\mathbf{z}, \theta), \tag{6}$$

and I assume that the observed average claims are equal to νAC_{jmt}, where the shock $\nu \geq 0$ is iid across jmt, and such that $\mathbb{E}[\ln(\nu) | G(\mathbf{z}, \theta), \mathbf{x}, \boldsymbol{\xi}, \mathbf{b}] = 0$.

3.3. *Identification*

3.3.1. Parametric and functional form assumptions. The parametric assumptions on $\alpha_t(\mathbf{z})$ and $\delta_{jmt}(\mathbf{z}, \theta)$ are detailed in Appendix A; all parameters are allowed to vary flexibly by year, and across seven six-years-wide age bins: $A^1 = \{26, \ldots, 31\}$, $A^2 = \{32, \ldots, 37\}, \ldots, A^6 = \{56, \ldots, 61\}$, $A^7 = \{62, 63, 64\}$. The result is a set of 644 parameters. The definitions of $\beta_t(\mathbf{z}, \theta)$ and $G(\theta|\mathbf{z})$ imply that the coefficient on actuarial value is log-normally distributed with year-age-bin-specific parameters. Unobserved heterogeneity and observed demographics are independent: $G_{mt}(\mathbf{z}, \theta) = G_{mt}(\mathbf{z}) G(\theta)$, where $G_{mt}(\mathbf{z})$ is observed.

On the cost side,

$$\eta(\mathbf{z}, \theta) = \eta^{\text{Age}} z^{\text{Age}} + \eta^{\text{WTP}} \frac{\beta_t(\mathbf{z}, \theta)}{\alpha_t(\mathbf{z})}, \qquad \text{and} \qquad \phi_{jmt} = \phi_t^1 + \phi_m^2 + \boldsymbol{\phi}^3 \text{Insurer}_{jmt}. \tag{7}$$

This allows individual medical spending to vary with age, and—to model adverse selection—with the willingness-to-pay for generosity of coverage. The remaining cost parameters are a combination of a constant, year, region, and insurer indicators.

3.3.2. Control function and actuarial value discontinuities. Identification of demand relies on regional variation in premiums conditional on age-bin and year, on discontinuous variation in actuarial value of Silver plans across buyers with different income, and on variation in the set of insurers and plans across markets.

To obtain instruments for premium, the ACA marketplaces are a setting in which the presence of rating restrictions across demographic groups leads to an intuitive Waldfogel IV (c.f. Berry and Waldfogel, 1999; Waldfogel, 2003). Insurers set base premiums responding to the distribution of demographic characteristics in a rating region, $G_{mt}(\mathbf{z})$, since this affects the shape of Q_{jmt} and AC_{jmt} as shown in (3) and (6). Identification assumes that, *conditional* on a buyer's age and income, preference do not depend on the distribution of demographics in the same geographic area, yet this affects base premiums, which should be higher in relatively older regions, and vice-versa (see also Orsini and Tebaldi, 2017; Polyakova and Ryan, 2019). Formally,

$$\mathbb{E}\left[\xi_{jmt}|G_{mt},\mathbf{z},\mathbf{x}\right]=0, \qquad \text{while} \qquad \mathbb{E}\left[b_{jmt}G_{mt}|\mathbf{z},\mathbf{x}\right]\neq 0,$$

implying $\mathbb{E}[P_j(\mathbf{b}_{mt},\mathbf{z})G_{mt}|\mathbf{z},\mathbf{x}]\neq 0$.

To obtain a control function one can use the residual $\widehat{\xi}_{jmt}$ of a regression of base premium projected on product characteristics and share of potential buyers in the region-year who are aged under-35 (the excluded IV):

$$b_{jmt}=\lambda^{35}\int \mathbf{1}\left[z^{\text{Age}}\leq 35\right]dG_{mt}(\mathbf{z})+\lambda^{\text{Tier}}+\lambda^{\text{Year}}+\lambda^{\text{Insurer}}+\xi_{jmt}. \tag{8}$$

Regression results and F-statistics are reported in Supplemental Appendix Table S1, the variation in the instrument and the corresponding variation in b_{jmt} are illustrated in Figure 3. The first stage OLS estimate of the effect of age-composition of potential buyers on base premium is $\widehat{\lambda}^{35}=-5,208$, with robust standard error 896. This implies that a 0.1 increase in the share of potential buyers aged under-35 corresponds to a \$521 reduction in base premium.

To identify the effect of actuarial value on indirect utility, as governed by $\beta_t(\mathbf{z},\theta)$, the ACA marketplaces feature discontinuities in AV_{ij}^{D} across the cost-sharing reduction thresholds (see Table 1). This institutional feature, which has also been used in Lavetti *et al.* (2023) to identify demand and cost responses to coverage generosity, implies that at three income thresholds Silver plans become suddenly less attractive, and that the choice to enrol in the marketplace is either costlier or it leads to lower coverage.

The three discontinuities correspond to $z_i^{\text{Inc}}=150,200,250$; the actuarial value of Silver plans drops from 95 to 88, then from 88 to 74, and finally from 74 to 70. As shown in Figure 3, the strongest effect is observed at $z_i^{\text{Inc}}=200$, when Silver plans become suddenly worse than Gold and Platinum plans. The 16% drop in actuarial value induces a 9.8% reduction in the probability of choosing a Silver plan.

3.3.3. Cost identification. To identify cost parameters my approach is similar to the one in Bundorf *et al.* (2012), since I observe demand at the individual level while costs are measured at the plan level. Identification of plan-level determinants of costs is standard, yet here I need to allow costs to vary also within plan across individuals who differ in age and unobservable willingness-to-pay for coverage.

Using a simplified notation, the intuition is as follows. Let u_i denote the utility for coverage of individual i. After controlling for a rich set of plan characteristics, and for the age composition

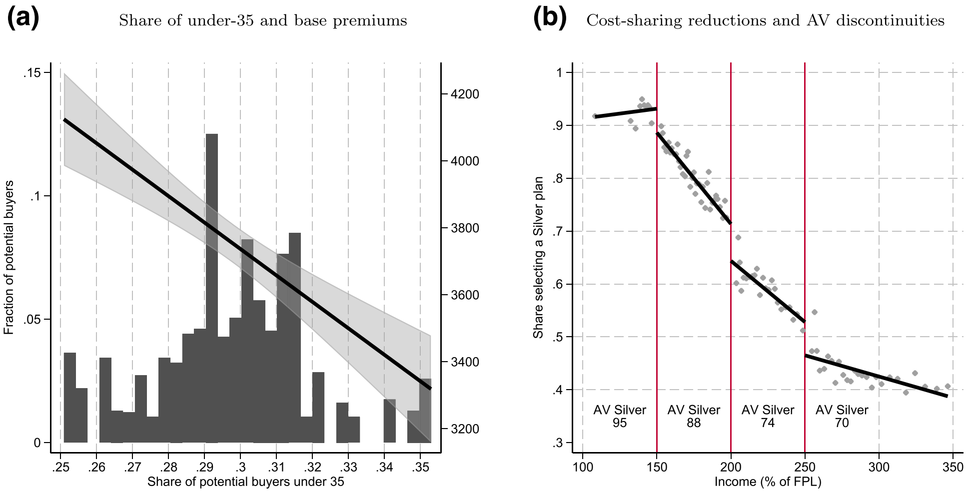

FIGURE 3

Demand identification: control function and actuarial value discontinuities

Note: The figure illustrates the variation underlying identification of premium and actuarial value coefficients. The left panel shows the histogram of the share of potential buyers younger than 35 for each *jmt* combination in the data. The figure also plots the linear relationship between b_{jmt} (measured on the right vertical axis) and the instrument, $\int \mathbf{1}[z^{\text{Age}} \le 35]dG_{mt}(\mathbf{z})$, with confidence intervals. See also Supplementary Appendix Table S1. The right panel is a binned scatter plot of the share of enrolees selecting a Silver plan as a function of income (as % of FPL). The linear relationship between the two variables is allowed to vary discontinuously at the three cut-off values corresponding to the discontinuity in actuarial value of Silver plans due to cost-sharing reductions (c.f. Section 2).

of buyers, the "residual average cost" for plan j can be written as

$$C_j = \int c(u_i)dF\left(u_i|i \text{ chooses } j\right), \tag{9}$$

where the function $c(u_i)$ is our object of interest, and $F(u_i|i \text{ chooses } j)$ describes the composition of buyers of j in terms of preferences for insurance. In its simplicity, (9) highlights the key requirement for identification of $c(u_i)$: *variation in the composition of buyers—$F(u_i|i$ chooses $j)$—across plans for which the researcher can assume that differences in the function $c(u_i)$ are known, or controlled for using observables*. In other words, one needs shifters of buyers' composition excluded from cost functions.[15],[16],[17]

Considering now the details of my application, I need to identify the cost parameters in equations (5) and (7). The Medical Expenditure Panel Survey (see Figure 2b) allows me to calibrate the parameter η^{Age}, which governs the age evolution of average annual medical spending

[15] Examples of these shifters include variation in the set of competing plans (a version of "BLP instruments", c.f. Berry *et al.*, 1995), or variation in the composition of potential buyers in terms of sociodemographic characteristics affecting willingness-to-pay for insurance (a version of "Waldfogel instruments", c.f. Waldfogel, 2003).

[16] Although a nonparametric inversion of equation (9) with respect to $c(\cdot)$ is theoretically feasible, it would require (arguably never observed) full support variation in $F(\cdot|i \text{ chooses } j)$ across j. Therefore, a parametrization similar to the one I adopted in equations (5) and (7) is necessary.

[17] Following a similar intuition, Supplemental Appendix S1 discusses how data on average claims could be replaced by supply-side equilibrium assumptions to obtain sufficient conditions under which costs functions are identified. This extends the well-known inversion of the first-order conditions which dates back to Rosse (1970) and Bresnahan (1987) to markets with asymmetric information. When claims are available one can then compare observed costs to those estimated under different assumptions on supply, and thus make an informed choice between alternative models of insurer behaviour (see Section 4.2).

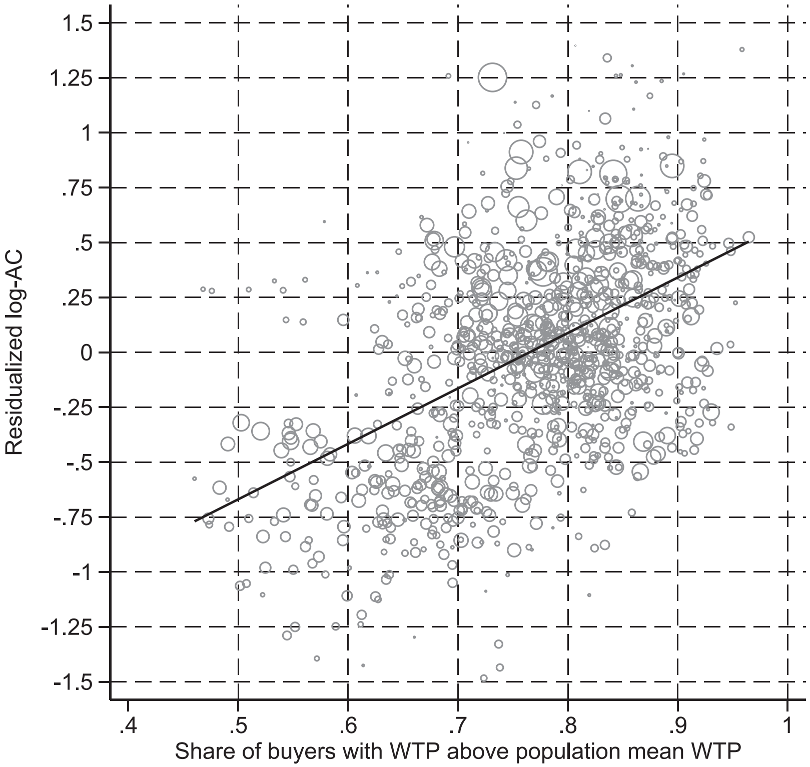

FIGURE 4

Identifying variation for η^{WTP}

Note: The figure shows a scatter plot and OLS fit of the residual of log-average claims adjusted for actuarial value and projected on region, year, insurer, and age-composition of the plan's enrolees (Y-axis) on the share of buyers with WTP for AV ($\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$) above the population average (X-axis). Observations are weighted by enrolment, each representing a *jmt* combination for which average claims are available. This variation identifies η^{WTP}, a positive relationship indicates $\eta^{\text{WTP}} > 0$.

when insured.[18] Then, as discussed above and illustrated in Figure 4, identification of η^{WTP} relies on the empirical relationship between average claims and composition of enrolment in terms of $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$ (derived from the demand estimates), after controlling for actuarial value, region, insurer, and age composition. Intuitively, if residualized claims are higher for plans covering a larger share of individuals with high $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$, as it is the case in Figure 4, $\eta^{\text{WTP}} > 0$, and vice versa.[19]

---

[18] For this purpose, I minimize

$$\frac{1}{N_{\text{MEPS}}} \sum_{\ell \in \text{MEPS}} \left\| Y_\ell - e^{\eta^{\text{Age}} \text{Age}_\ell + \text{Year}_\ell + \text{Region}_\ell} \right\|, \tag{10}$$

where Y_ℓ is the annual medical spending of individual ℓ observed in the survey, and Region_ℓ is a Medical Expenditure Panel Survey macro area. The parameter η^{Age} is very robust across specifications and estimated precisely; see Supplemental Appendix Table S5.

[19] More formally, note that equation (7) restricts the way in which insurer, year, and region affect medical spending. Given these restrictions, and since I assume that $\mathbb{E}[\ln(\nu)|G(\mathbf{z},\theta),\mathbf{x},\boldsymbol{\xi},\mathbf{b}] = 0$, the residual correlation between AC_{jmt}/AV_j^S and the density of $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$ within a given *jmt* combination pins down η^{WTP}.

TABLE 3
Summary of demand estimates by age group

	Age 26–31	Age 32–37	Age 38–43	Age 44–49	Age 50–55	Age 56–61	Age 62–64
Mean WTP	249.6	293.8	333.5	395.8	507.5	684.8	853.5
for 10% AV increase	(9.3)	(10.2)	(12.7)	(10.9)	(14.4)	(16.4)	(20.7)
St. Dev. of WTP	202.6	231.3	250.1	304.4	373.3	495.5	609.3
for 10% AV increase	(5.7)	(6)	(6.7)	(6.1)	(7.2)	(9.2)	(11.4)
% Change in enrolment if	−7.434	−6.822	−6.552	−5.69	−4.86	−3.832	−3.137
+$120/year in all Premium	(0.203)	(0.224)	(0.215)	(0.136)	(0.108)	(0.097)	(0.078)
% Change in Silver Enrolment	−2.356	−2.478	−2.113	−2.272	−1.887	−1.732	−1.492
if +1% in all Silver Premiums	(0.074)	(0.081)	(0.059)	(0.06)	(0.047)	(0.033)	(0.026)
Control Function	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Year-Specific Parameters	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Insurer-Year Fixed-Effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
N. Individuals	2, 335, 251	2, 050, 631	1, 814, 069	1, 764, 925	1, 822, 717	1, 841, 849	803, 613

Note: The table summarizes the estimates of preferences for insurance and sensitivity to premiums conditioning on different age groups. The reported parameters are functions of the demand parameters in Supplemental Appendix Tables S2 and S3. Standard errors in parentheses, obtained as the empirical standard deviation across 100 independent random draws of the demand parameters using the estimated variance-covariance matrix. The WTP for a 10% AV increase is equal to the ratio $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$, this varies across individuals both unobservably with θ, and observably with age, year, and income.

3.4. *Estimation results*

Estimation follows the steps detailed in Supplemental Appendix S2.

3.4.1. Demand estimates. The full set of demand parameters is reported in Supplemental Appendix Tables S2 and S3. Supplemental Appendix Table S4 shows the impact of the control function on demand estimates. Omitting $\widehat{\xi}_{jmt}$ would lead to estimates of premium coefficients approximately 4% lower, and to lower estimates of willingness-to-pay for coverage for most individuals.

Table 3 illustrates how demand for ACA-sponsored insurance varies with buyer's age. For each of the seven age bins used for estimation, the table summarizes the distribution of willingness-to-pay for actuarial value. The table also reports extensive margin semi-elasticity of demand—measured as the percentage drop in the probability of purchasing marketplace coverage if all annual premiums increase by $120—and average own-price elasticity of demand for Silver plans, equal to the percentage drop in the share of buyers selecting a Silver plan if the plan's premium increases by 1%. The extent to which "older buyers demand more" is consistent with intuition and with patterns in the raw data.

Average willingness-to-pay for a 10% increase in actuarial value increases steadily with age, from $250 among those aged between 26 and 31, to $334 between 38 and 43, $508 between 50 and 55, reaching the average value of $854 among those aged between 62 and 64. This average increase is accompanied by a larger variance: the standard deviation at 26–31 (32–37) is $203 ($231), while at 56–61 (62–64) it is more than twice as large, equal to $496 ($609).

Increasing all annual premiums by $120 (third row of Table 3) is equivalent to lowering subsidies by $10 per-person, per-month, while holding fixed insurers' decisions. I find that this would lower participation of buyers younger than 31 by 7.4%, compared to 6.8% among those aged between 32 and 37, and 6.5% among those aged between 38 and 43. The extensive margin response to a change in all premiums is much smaller for older buyers. Conditional on age being between 56 and 61, if all premiums increase by $120 enrolment drops by 3.8%. For the oldest age bin, 62–64, I estimate that average extensive margin semi-elasticity is equal to 3.1%.

Supplemental Appendix Figure S1 shows the entire distribution of willingness-to-pay and extensive margin response to premium across individuals. These estimates of how marketplace demand responds to subsidies complement (and align with) the estimates of closely related parameters obtained in other studies.[20]

The fourth row of Table 3 shows the estimates of the elasticity of Silver enrolment to Silver premiums. This is calculated as the percent change in enrolment in Silver plans if the premium of all Silver plans (which varies by age-income-region-year) increases by 1%. The elasticity of under-50 individuals is between 2.3 and 2.5%, while for older individuals this is between 1.5 and 1.9%.

Interpreting these estimates it is important to highlight that my model of plan choice is static, and does not consider plan switching and consumers' inertia or state dependence (see *e.g.* Drake *et al.*, 2022; Saltzman *et al.*, 2021). The extent to which older individuals are less sensitive to premium changes might denote that they have higher inertia (or switching costs), or that the exogenous churn in and out of the marketplace is higher among younger enrolees (as one would expect considering that at younger age labour market shocks and changes to household structure are more frequent). If this was the case, the higher participation of young invincibles obtained in my counterfactuals might not only increase average elasticity in a static sense, but also lower incentives for "invest and harvest" pricing strategies (which I do not model).

3.4.2. Cost estimates. The full set of cost estimates is reported in Supplemental Appendix Table S6. Table 4 summarizes the key parameters governing heterogeneity in medical spending across buyers who differ in age and willingness-to-pay for actuarial value, and the differences in average costs across age groups for Bronze and Silver plans.

The estimate of η^{Age} derived from the Medical Expenditure Panel Survey is equal to 0.038 (Supplemental Appendix Table S5). This indicates that, on average, one year of age corresponds to approximately 3.8% higher expected medical spending. While age is observed, and partially accounted for by the regulatory age rating adjustments, willingness-to-pay for actuarial value varies unobservably conditional on age.

The parameter η^{WTP} shows that this unobservable dimension of preferences for insurance is positively correlated with medical spending. Table 4 shows that the point estimate of η^{WTP} is equal to 0.08, statistically significant at any conventional level. This implies that a $100 increase in $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$ corresponds to approximately 8% higher expected medical spending. Given the range of $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$ shown in Supplemental Appendix Figure S1a and Table 3, even conditioning on age, income, and year, willingness-to-pay for actuarial value can vary by more than $500, corresponding to 40% higher expected cost.

The estimates of $\eta(\mathbf{z},\theta)$ are the distinguishing feature of a selection market: average and marginal cost curves for a given plan *jmt* are not constant, varying as a function of base premiums. Holding base premiums fixed at the observed levels, the bottom of Table 4 summarizes the

---

[20] Using discontinuities in subsidies in the pre-ACA Massachusetts marketplace, Finkelstein *et al*. (2019) find enrolment dropping about 25% for every $40 increase in monthly premium. Applying a nested logit demand model to data from California and Washington, Saltzman (2019) estimates that a $100 increase in all premiums would induce 3.3–3.7% reduction in marketplace enrolment. In Tebaldi *et al*. (2023) we adopt a nonparametric approach and estimate that, if all 2014 monthly premiums increased by $10, the probability of enrolment in Covered California would have been 0.018–0.067 lower.

TABLE 4
Summary of cost estimates

Parameters of $\eta(\mathbf{z},\theta)=\eta^{\text{Age}}z^{\text{Age}}+\eta^{\text{WTP}}\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$			Estimator, N. Obs.	Data Source	Region FE	Year FE	Insurer FE
Age	η^{Age}	0.0379 (0.0021)	NLLSQ, N = 20,171	2014-17 MEPS	Y	Y	N
WTP for 10% AV increase ($100/year)	η^{WTP}	0.0803 (0.0104)	NLLSQ, N = 1,026	2016-19 RRF	Y	Y	Y

Insurer Expected Average Cost at Observed Premiums	Age 26–31	Age 32–37	Age 38–43	Age 44–49	Age 50–55	Age 56–61	Age 62–64
Bronze Enrolees	1, 030 (136)	1, 421 (169)	1, 861 (203)	2, 581 (247)	3, 647 (272)	5, 334 (263)	7, 503 (240)
Silver Enrolees	1, 311 (137)	1, 821 (164)	2, 361 (205)	3, 336 (220)	4, 742 (229)	7, 571 (201)	11, 208 (364)

Note: The top panel shows the estimates of the two parameters of the function $\eta(\mathbf{z},\theta)$, governing the heterogeneity in expected medical spending across individuals. The full set of non-linear least squares estimates is reported in Supplemental Appendix Table S6. The bottom panel shows the estimated average cost across Bronze and Silver enrolees, conditional on different age groups. Standard errors in parentheses, obtained as the empirical standard deviation of cost estimates obtained across 100 independent random draws of demand parameters (using the estimated variance–covariance matrix).

value of expected average claims for Bronze and Silver plans, conditioning on the seven age bins used for demand estimation. These estimates depend on $\eta(\mathbf{z},\theta)$, but also on $\boldsymbol{\phi}$, which collects year, region, and insurer-specific cost parameters (c.f. equation (7)).

For Bronze plans, expected average claims are equal to $1,030 per-person, per-year when the enrolee is aged between 26 and 31, $1,421 when between 32 and 37, almost $1,900 when between 38 and 43, and progressively increasing to more than $7,500 for the oldest group, aged between 62 and 64. Silver plans have higher average claims, reflecting both higher actuarial value ($AV_j^S=70\%$, instead of 60%) but also a different risk selection: enrolees of Silver plans have higher $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$. As a result, the average claims of Silver plans when enroling someone aged between 26 and 31 are $1,311, 27% higher than the estimate for Bronze plans, and 9% higher than the difference that would be explained by the increased actuarial value, holding risk selection fixed. This would be $1,202, computed as $1,030 × $\frac{0.7}{0.6}$.

The relative difference between Silver and Bronze expected average claims is increasing with age, reflecting the larger premium differences following the ACA rating regulations. When selecting a Silver plan, someone older than 50 must have unobservably higher $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$ relative to someone younger making the same choice. Among enrolees who are 56 or older, average claims for those selecting a Silver plan are between $7,600 and $11,200, 42–49% higher than the claims for those selecting a Bronze plan.

The relevance of heterogeneity and adverse selection in this application is highlighted in Figure 5: higher willingness-to-pay corresponds to higher expected cost. This relationship is steeper for older individuals. Among those under 35, an increase in willingness-to-pay from approximately zero to $1,000 corresponds to a cost increase from $1,000 to slightly more than $2,000. When considering individuals aged 35–64, the same difference in preferences corresponds to a cost increase from less than $2,000 to almost $6,000. Even conditioning on a specific

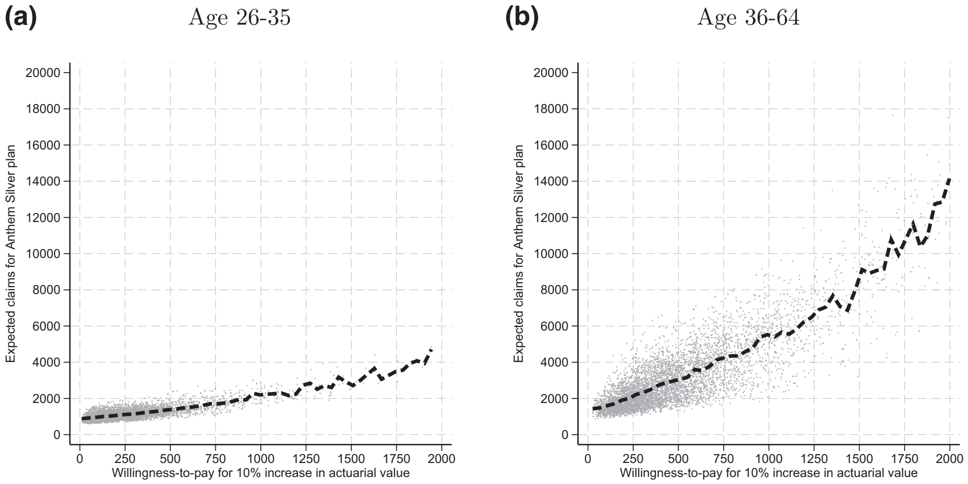

Figure 5

Empirical relationship between preferences and expected cost

Note: The figure illustrates the joint distribution of willingness-to-pay for coverage and expected cost obtained after combining demand and cost estimates. The graph is generated by randomly drawing 10,000 individuals from $G(\mathbf{z},\theta)$. For each draw, I compute willingness-to-pay for a 10% increase in actuarial value ($\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$), and expected cost if the individual enrols in a Silver plan, offered by Anthem ($\kappa_{jmt}(\mathbf{z},\theta)$, where j is Anthem's Silver plan in mt). The figure then consists of a scatter plot of these quantities, overlaying this with a local polynomial smoothing of the two quantities. The left panel is conditional on $z^{\text{Age}} \leq 35$, the right panel is conditional on $z^{\text{Age}} > 35$.

value of cost, there is significant heterogeneity in preferences, and vice versa. The joint distribution summarized in Figure 5 is the key primitive one needs to study market design in a health insurance marketplace.

4. EQUILIBRIUM AND MARKET CONDUCT

Before considering counterfactual policy design, it is necessary to model expected profits incorporating ACA regulations, and to seek empirical support for alternative assumptions about insurers' conduct. I assume that insurers understand the regulations (including subsidy adjustments to premiums and floors to subsidized premiums) and have perfect foresight of competitors' premiums.

4.1. *Rating regulations, risk adjustment, and profits*

Each insurer f offers the plans in the set $\mathcal{J}(f)$ in region m, year t. The expected profit of insurer f in mt is a function of the base premiums $\mathbf{b}_{fmt} = \{b_{jmt}\}_{j\in\mathcal{J}(f)}$. Expected total revenues for each product $j \in \mathcal{J}(f)$ are equal to

$$R_{jmt}(\mathbf{b}_{fmt}, \mathbf{b}_{-fmt}) = \int \text{Adjustment}(z^{\text{Age}}) b_{jmt} q_{jmt}(\mathbf{z},\theta) dG_{mt}(\mathbf{z},\theta);$$

where $q_{jmt}(\mathbf{z},\theta)$ depends on $(\mathbf{b}_{fmt}, \mathbf{b}_{-fmt})$, including age adjustments and subsidies, as shown in (2). Expected total costs are instead equal to

$$TC_{jmt}(\mathbf{b}_{fmt}, \mathbf{b}_{-fmt}) = \int \kappa_{jmt}(\mathbf{z},\theta) q_{jmt}(\mathbf{z},\theta) dG_{mt}(\mathbf{z},\theta).$$

To model risk adjustment I follow the ACA formula (see *e.g.* Pope *et al.*, 2014; Saltzman, 2021), as described in details in Supplemental Appendix S3. For every plan $j \in \mathcal{J}(f)$, the risk adjustment transfer is

$$RA_{jmt}(\mathbf{b}_{fmt}, \mathbf{b}_{-fmt}) = Q_{jmt} \underbrace{\frac{\sum_k R_{kmt}}{\sum_k Q_{kmt}}}_{\substack{\text{average premium} \\ \text{in region-year}}} \left(\text{Relative Risk}_{jmt} - \text{Relative Adjustment}_{jmt}\right).$$

In words, the per-enrolee risk adjustment transfer to plan j in region-year mt is the product of average premium in the region and a difference between a relative risk measure and a relative premium measure.

The risk adjustment formula is constructed to ensure that transfers sum to zero. Plans receive positive transfers if they cover costlier-than-average individuals, after controlling for actuarial value differences and premium adjustments. The other plans face negative transfers, which are larger when enrolees are, on average, less risky, after controlling for actuarial value and premium adjustments.

Expected profits for insurer f in region-year mt combine the above definitions and account for multi-plan insurers: omitting the dependence on $(\mathbf{b}_{fmt}, \mathbf{b}_{-fmt})$ to simply the notation,

$$\Pi_{fmt} = \sum_{j \in \mathcal{J}(f)} R_{jmt} - TC_{jmt} + RA_{jmt}.$$

Different subsidy designs imply different R, TC, and RA functions, by altering the relationship between $(\mathbf{b}_{fmt}, \mathbf{b}_{-fmt})$ and the composition and risk selection of individuals choosing different plans.

4.2. *Evidence on insurers' conduct*

I consider two alternative models of insurer conduct: static multi-product Nash pricing (as in Bundorf *et al.*, 2012; Starc, 2014; Decarolis *et al.*, 2020; Saltzman, 2021; Curto *et al.*, 2021), and perfect competition à la Azevedo and Gottlieb (2017), in which every plan breaks even in expectation as adopted recently by Dickstein *et al.* (2024).[21] Although I compute counterfactuals under both assumptions, I am in the position to investigate whether the data supports one over the other.

Formally, multi-product Nash pricing requires that, for every insurer f, the following FOC are satisfied for every $j \in \mathcal{J}(f)$, every m, and every t:

$$\frac{\partial \Pi_f}{\partial b_{jmt}} = \sum_{k \in \mathcal{J}(f)} \frac{\partial R_{kmt}}{\partial b_{jmt}} - \frac{\partial TC_{kmt}}{\partial b_{jmt}} + \frac{\partial RA_{kmt}}{\partial b_{jmt}} = 0. \tag{11}$$

Perfect competition requires that, for every jmt,

$$\Pi^{\mathrm{AG}}_{jmt} = R^{\mathrm{AG}}_{jmt} - TC^{\mathrm{AG}}_{jmt} + RA^{\mathrm{AG}}_{jmt} = 0. \tag{12}$$

---

[21] Future work could consider even more complex models of imperfect competition between insurers, allowing for strategies to be dynamic, or firms uncertainty about demand and cost functions (see *e.g.* Saltzman and Lucarelli, 2021).

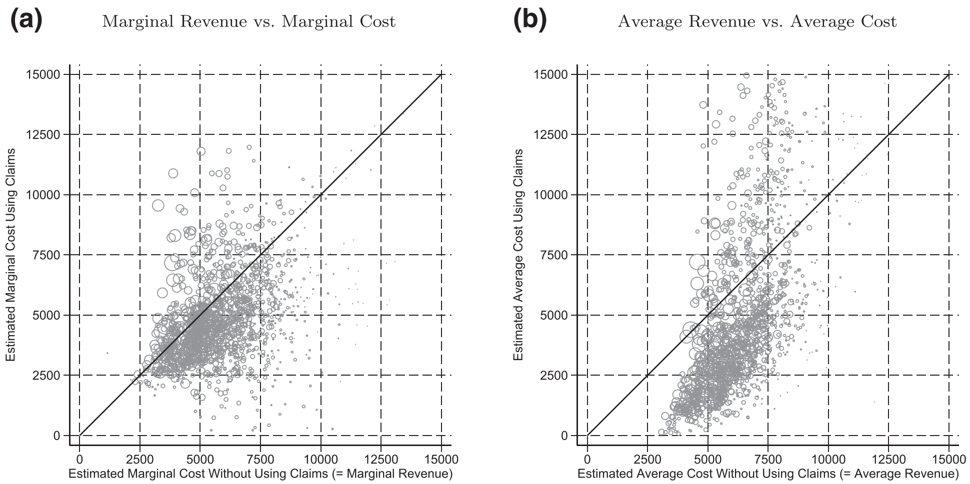

FIGURE 6

Multi-product Nash pricing versus perfect competition

*Note:*The left panel shows the comparison between per-enrolee risk-adjusted marginal costs estimated assuming multi-product Nash-in-Prices without using claims (these are equal to marginal revenue), and per-enrolee risk-adjusted marginal costs estimated using observed claims. The right panel shows the comparison between per-enrolee risk-adjusted average costs estimated assuming perfect competition without using claims (these are equal to average revenue), and per-enrolee risk-adjusted average costs estimated using observed claims. Markers are weighted by plan enrolment, each observation is a *jmt* combination.

In this expression, the superscript AG indicates that the demand function $q_{jmt}(\mathbf{z},\theta)$ is modified to let an infinitesimal fraction of "behavioural" buyers choose a given plan independently from changes in premiums or other characteristics.[22]

Figure 6 compares estimated and model-predicted marginal and average costs under alternative conduct assumptions. This is not a formal test, but it shows that observed data and estimated primitives are more consistent with multi-product Nash pricing than average-cost pricing. A formalization of this procedure, in which—rather than imposing supply assumptions during estimation—the researcher compares alternative models of conduct before running counterfactuals, represents an important venue for future work. For the case of markets without adverse selection, statistical tests to discriminate between models of conduct are known since (Bresnahan, 1987).

The comparison between the two models relies, albeit somewhat implicitly, on the possibility to identify cost curves in a selection market imposing supply-side assumptions, rather than observing costs directly. While here I discuss my findings informally, a new, formal and self-contained identification result is provided in Supplemental Appendix S1.

In Figure 6a, the horizontal axis corresponds to the per-enrolee marginal revenue for every *jmt* combination in the data. Nash pricing predicts that this would be equal to per-enrolee risk-adjusted marginal cost, following equation (11). The vertical axis corresponds to the estimate of this quantity for every *jmt*. It is important to recall that (11) has not been used as a moment or constraint for the estimation of demand and cost. The resulting scatter plot is concentrated around the 45-degree line. The enrolment-weighted average difference between per-enrolee

---

[22] I assume that a fraction of individuals equal to 0.001 chooses iid uniformly across the J plans. This ensures equilibrium existence (c.f. Azevedo and Gottlieb, 2017). Profits in this case are "almost" zero, rather than zero, as it will be the case in Tables 5 and 6.

marginal revenue and per-enrolee risk-adjusted marginal cost is \$362.99 (95%-C.I.: [272.95, 453.02]). The enrolment-weighted average ratio $\frac{\partial \Pi_f}{\partial b_{jmt}}/R_{jmt}$ is 0.049 (95%-C.I.: [0.034, 0.064]).

For comparison, Figure 6b repeats the same procedure to explore the discrepancy between average revenue and risk-adjusted average cost. Perfect competition predicts that the two quantities would be equal, and the distribution should be close to the 45-degree line. As shown in the figure, relatively to Figure 6a this seems not to be the case. For a large number of *jmt* combinations estimated risk-adjusted average cost is significantly lower than average revenue, providing evidence against perfect competition. The enrolment-weighted average difference between R_{jmt}/Q_{jmt} and $(TC_{jmt}+RA_{jmt})/Q_{jmt}$ is \$1,314.48 (95%-C.I.: [1,170.31, 1,458.64]). The enrolment-weighted average ratio Π_{jmt}/R_{jmt} is 0.26 (95%-C.I.: [0.237, 0.283]), corresponding to a departure from the model assumption 5 times as large as under Nash pricing.

One additional piece of evidence in support of modelling insurers as not perfectly competitive is provided by the estimated medical loss ratio. Despite not imposing a constraint in estimation, I calculate average medical loss ratio at the observed base premiums to be approximately equal to 0.9 (Table 6). This is above the minimum value of 0.8 mandated by the ACA, while still lower than the perfect competition value of one.

A static oligopoly model seems to perform well. However, one could note how the extent to which marginal revenues are, on average, slightly higher than marginal cost is consistent with the omission of dynamic considerations. If consumers have inertia, as documented in Drake *et al.* (2022) and Saltzman *et al.* (2021), static marginal revenues do not account for a the (positive) continuation value of keeping current enrolees. As discussed above, since younger individuals have higher churn and have been estimated to have lower inertia, the policies increasing their participation that I discuss below could also reduce dynamic pricing incentives.

5. SUBSIDY DESIGN AND EQUILIBRIUM OUTCOMES

5.1. *Price-linked subsidies versus vouchers*

I begin by comparing equilibrium under ACA subsidies and equilibrium under fixed vouchers: subsidies that do not adjust endogenously with base premiums. Jaffe and Shepard (2020) call the ACA design a “price-linked subsidy”: the market sponsor determines the maximum premium individuals should pay, and adjusts subsidies to insurers’ decisions accordingly. One alternative is to use an “equivalent” voucher: the subsidy received by every individual is fixed to the (price-linked, endogenous) amount received under the ACA. This varies then by age, income, region, and year, but it is not adjusted in equilibrium.

The transition from a price-linked subsidy to a fixed, equivalent voucher increases the own-premium semi-elasticity for the second cheapest Silver plan in the region-year. Under the ACA design, when this plan increases its base premium buyers do not face premium increases, the only effect is to lower other plans’ premiums. Under Nash pricing, switching to an equivalent voucher implies that the second cheapest Silver plan has incentives to charge lower premiums, and this effect should be larger in less-competitive, more-concentrated markets.

Jaffe and Shepard (2020) discuss this mechanism formally for the case of single-plan insurers, in which the subsidy-setting plan is the cheapest; this was the case in the pre-ACA Massachusetts marketplace. As anticipated in their appendix, the main difference in the ACA context is that insurers offer multiple plans, and that subsidies are determined to target the second cheapest Silver, rather than the cheapest Bronze.

Table 5 shows how market outcomes vary when adopting ACA price-linked subsidies or equivalent vouchers. The left panel shows results obtained assuming multi-product Nash pricing, the right panel assumes perfect competition. In the latter case, outcomes do not vary

TABLE 5

From ACA price-linked subsidies to equivalent vouchers

	Multi-Product Nash pricing				Perfect Competition			
	2–3 insurers 27 region-years		4–7 insurers 49 region-years		2–3 insurers 27 region-years		4–7 insurers 49 region-years	
	ACA subsidy	Equivalent voucher	ACA subsidy	Equivalent voucher	ACA subsidy	Equivalent voucher	ACA subsidy	Equivalent voucher
Share enroled	0.32	0.36	0.28	0.29	0.27	0.27	0.28	0.28
2nd cheapest Silver b_j	4, 127	2, 998	2, 709	2, 559	2, 387	2, 387	2, 116	2, 115
Share in Bronze plans	0.15	0.14	0.13	0.13	0.16	0.16	0.14	0.14
Medical-loss ratio	0.82	0.8	0.89	0.84	1	1	1	1
ΔCS_i relative to ACA	–	90	–	30	–	0	–	1
Average subsidy	5, 705	4, 187	3, 249	3, 258	2, 713	2, 709	2, 223	2, 211

Note: Simulated market outcomes under alternative subsidy designs and different region-year markets. The left panel corresponds to multi-plan Nash pricing, where equilibrium is simulated in every region-year by finding the vector of base premiums $\mathbf{b}_{mt}$ that minimizes the distance between the left- and right-hand side of equation (11). The right panel corresponds to perfect competition à la Azevedo and Gottlieb (2017), where equilibrium is simulated in every region-year by finding the vector of base premiums $\mathbf{b}_{mt}$ that minimizes the distance between the left- and right-hand side of equation (12). The ACA subsidy corresponds to the regulations described in (1) in Section 2. The Equivalent Voucher corresponds to setting subsidies equal to the level of the ACA subsidy, and then computing equilibrium removing price-linked adjustments of subsidies to the second cheapest Silver plan in a region-year pair. Share enroled and second cheapest Silver base premium are computed as averages across region-years, weighted by number of eligible individuals. The share in Bronze plans, medical-loss ratio, and average subsidy are computed as averages across region-years, weighted by enrolment. ΔCS_i indicates the average, per-person annual consumer surplus, which is reported in differences from the equilibrium under ACA price-linked subsidies.

across the two subsidy designs (equilibrium premiums depend only on enrolees expected costs): price-linked subsidies are non-distortionary in perfectly competitive markets. Under Nash pricing, adopting equivalent vouchers affect equilibrium outcomes, since it implies a lower second cheapest Silver base premium.

The price distortion due to linking subsidies to insurers' decisions is larger markets that are more concentrated. In small regions (2–3 insurers), second cheapest base premiums drop by 27%, from \$4,127 to \$2,998; in larger regions, with more than four participating insurers, the drop is smaller, from \$2,709 to \$2,559 (−5.5%). Cheaper Silver plans lead to a lower share of buyers choosing a (high deductible) Bronze plan.

Accounting for adjustments to all premiums, and consequent changes in plan selection and composition of enrolment pools, the Nash-pricing equilibrium under equivalent vouchers implies slightly higher marketplace enrolment, increasing from 0.32 (0.28) to 0.36 (0.29) in small (large) regions. The corresponding increase in annual per-person consumer surplus relative to the ACA design is between \$30 and \$90. In regions with less than four insurers average subsidies drop from \$5,705 to \$4,187; in larger regions they increase slightly from \$3,249 to \$3,258 due to changed in plan selection. Insurer profitability is also higher, as medical-loss ratio drops from 0.82 (0.89) to 0.8 (0.84) in small (large) regions.

Despite differences in the specific policy and market structure, the comparisons between equilibrium under ACA price-linked subsidies and vouchers are similar to the results in Jaffe and Shepard (2020). They argue that fixing vouchers to a specific level requires regulators to have prior knowledge of insurers' costs yet show that—for reasonable levels of uncertainty about costs—vouchers perform better than price-linked subsidies. My results imply that, under the ACA, adopting a system of vouchers calibrated to the early years of the marketplaces would lead to sizable gains in terms of lower premiums and government spending.

5.2. *More subsidies for the young invincibles*

The second counterfactual subsidy design amounts to providing additional enrolment incentives to the so-called "young invincibles"; in what follows this group consists of individuals aged between 26 and 35 (see *e.g.* Levine and Mulligan, 2017). Since these buyers are, at the same time, cheaper to cover and more price sensitive, lowering their (subsidized) premiums ignites a series of adjustments to a new equilibrium. Insurers lower base premiums, due to the average cost reduction and—under Nash pricing—increase in elasticity. Lower premiums lead to higher enrolment and higher consumer surplus. Importantly, since premiums across demographic groups are linked by rating regulations (which are hold fixed), the gains from higher subsidies to young individuals can be as large as to allow lowering subsidies for older individuals, while still keeping all buyers better off, increasing profits, and reducing per-buyer government spending.

There are many alternative ways to measure the benefit of higher subsidies to the young invincibles, and here I consider two. First, one can maintain a price-linked design, and lower the maximum affordable amount (c.f. Section 2, equation (1)) for young individuals. Second, using (equivalent) vouchers, one can increase vouchers for the "young", while lowering vouchers for the "old". For each alternative, the first-order, "off-equilibrium" effect of changing policy while holding base premiums fixed will be different than the equilibrium effect, which accounts for endogenous pricing behaviour.

Panel (a) of Table 6 summarizes how marketplace outcomes respond to changing the ACA price-linked design by lowering the maximum affordable amount for young invincibles by 30%. In equilibrium, the effect is to increase enrolment in all demographic groups, as well as annual per-person consumer surplus, while average cost and average subsidies are lower. Despite slight differences in magnitude, the results are qualitatively similar under alternative models of insurer conduct.

Using vouchers, the way in which alternative subsidy designs impact equilibrium outcomes is more intuitive. This is illustrated in panel (b) of Table 6, where ACA-equivalent vouchers are modified by raising annual under-35 vouchers by \$600, while lowering over-35 vouchers by \$100. Holding base premiums fixed, young invincibles would be better off, while older buyers worse off (the enrolment share for this group drops by 0.01 as they face higher premiums). In equilibrium, however, the reduction in base premiums following the larger enrolment share of under-35 individuals implies that all buyers are better off.

Considering Nash pricing, under-35 enrolment increases from 0.28 to 0.39, and over-35 enrolment from 0.32 to 0.33; despite receiving smaller vouchers, subsidized premiums of over-35 buyers are \$76 lower. The younger composition of enrolees translates in average costs that, in equilibrium, are 7.6% lower than under the ACA-equivalent voucher. Per-person consumer surplus increases by \$104 per-year, while average per-enrolee subsidies are \$68 lower. Profits are also higher since the increase in enrolment dominates the reduction in markups. The result by which the alternative vouchers represent an improvement for all buyers while not increasing average subsidies is robust to assuming perfect competition.

Figure 7 illustrates the relationship between age and changes in annual, per-person consumer surplus resulting from changing vouchers as in panel (b) of Table 6. The dash line corresponds to Nash pricing, while the solid line corresponds to the equilibrium simulations under perfect competition. In the left panel, base premiums are held fixed to the ACA-voucher equilibrium: under-35 experience a net gain, while over-35 are worse off. However, as shown in Figure 7b, at the new equilibrium the change in consumer surplus of over-35 switches sign: this group is now

TABLE 6

Counterfactual subsidy design: shifting generosity toward "young invincibles"

Panel (a): Lowering MAA for under-35 by 30%

	Multi-product Nash			Perfect Competition		
	ACA MAA	Counterfactual MAA		ACA MAA	Counterfactual MAA	
	Equilibrium	Off-equilibrium	Equilibrium	Equilibrium	Off-equilibrium	Equilibrium
Share enroled:						
26–35	0.26	0.33	0.33	0.26	0.32	0.32
36–64	0.3	0.3	0.3	0.29	0.29	0.29
Premium paid:						
26–35	1, 571	1, 265	1, 311	1, 756	1, 438	1, 440
36–64	1, 693	1, 693	1, 764	2, 009	2, 009	2, 014
Average cost ($/year)	4, 357	4, 112	4, 136	4, 192	3, 984	3, 987
Average revenue ($/year)	4, 946	4, 824	4, 842	4, 202	4, 106	3, 995
Medical-loss ratio	0.9	0.87	0.87	1	0.97	1
Per-person CS ($/year)	771	815	799	733	771	774
Average subsidy ($/year)	3, 632	3, 614	3, 542	2, 288	2, 324	2, 208
Total profits ($ million)	2, 117	2, 781	2, 694	35	454	28

(continued)

TABLE 6
Continued

Panel (b): Increasing under-35 voucher by $600/year while lowering over-35 voucher by $100/year

	Multi-product Nash			Perfect Competition		
	ACA-voucher	Counterfactual voucher		ACA-voucher	Counterfactual voucher	
	Equilibrium	Off-equilibrium	Equilibrium	Equilibrium	Off-equilibrium	Equilibrium
Share enroled:						
26–35	0.28	0.37	0.39	0.26	0.36	0.39
36–64	0.32	0.31	0.33	0.29	0.28	0.31
Premium paid:						
26–35	1, 565	1, 097	1, 012	1, 754	1, 202	1, 066
36–64	1, 660	1, 737	1, 584	2, 005	2, 100	1, 830
Average cost ($/year)	4, 207	3, 929	3, 889	4, 191	3, 873	3, 815
Average revenue ($/year)	5, 041	4, 860	4, 704	4, 200	4, 027	3, 818
Medical-loss ratio	0.84	0.81	0.83	1	0.96	1
Per-person CS ($/year)	810	851	914	734	778	864
Average subsidy ($/year)	3, 412	3, 375	3, 344	2, 278	2, 297	2, 300
Total profits ($ million)	3, 145	3, 812	3, 580	31	590	12

Note: Simulated market outcomes under alternative subsidy designs; for details on equilibrium computation, see note to Table 5. Panel (a) shows the effect of lowering the maximum affordable amount for individuals under-35 by 30%, holding fixed the other regulations as set under the ACA. Panel (b) compares the ACA-equivalent voucher to an alternative design in which vouchers for individuals under-35 are $600 higher, while vouchers for individuals over-35 are $100 lower. The Off-equilibrium columns show how outcomes vary when the subsidy design is changed, but base premiums are held fixed to the level of the ACA maximum affordable amount Equilibrium, and ACA-voucher Equilibrium, respectively. Total profits sum up profits across all insurers and year. Enrolment shares and annual per-person CS are computed as averages across region-years, weighted by number of eligible buyers. Other outcomes are enrolment-weighted averages.

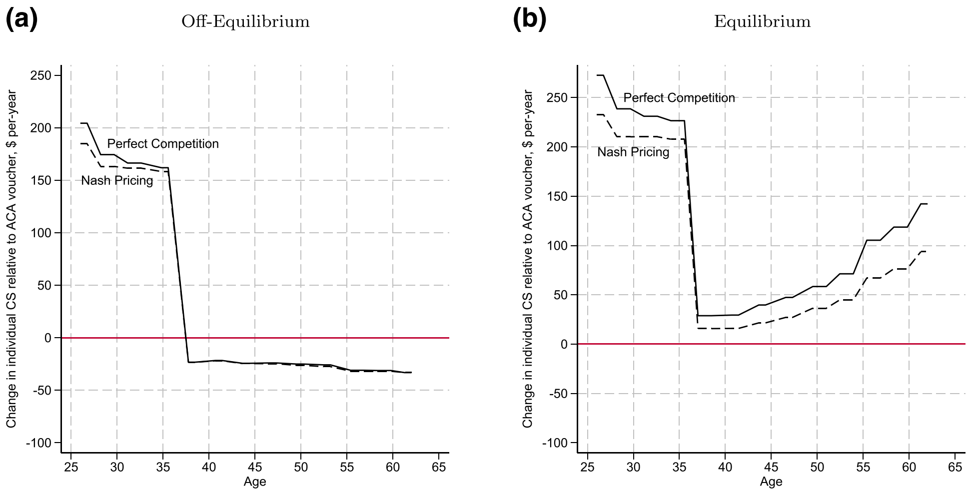

FIGURE 7

ΔConsumer surplus by age: +$600 under-35 voucher, -$100 over-35 voucher

Note: Average annual change in per-person consumer surplus when replacing ACA-equivalent vouchers with vouchers that are $600 higher for the under-35, and $100 lower for the under-35. The left panel holds base premiums fixed to the equilibrium under ACA-equivalent vouchers, the right panel corresponds to the new equilibrium. The solid lines correspond to perfect competition à la Azevedo and Gottlieb (2017), the dashed lines correspond to Nash pricing.

better relative to the ACA-voucher equilibrium, by an annual amount varying between $10 and 100.[23]

6. CONCLUSION

Expanding coverage while limiting public costs is one of the main goals of government-sponsored health insurance. If individuals who value insurance less and are more responsive to premiums are also less risky, a subsidy design in which premiums are equal for all individuals can be worse than a design in which premiums vary across types. Adjusting subsidies to observables that predict preferences and cost can lead to equilibria in which all enrolees pay lower premiums, coverage and profits are higher, and average subsidies are lower.

After discussing this point, the article measured the potential gains from introducing age adjustments to ACA subsidies using data from the California marketplace regulated under the recent healthcare reform. The data support oligopoly pricing over imperfect competition. Following the significant differences in preferences and cost across age groups, equilibrium simulations suggest that shifting subsidy generosity toward young uninsured would lower costs and premiums, increasing profits and consumer welfare. Whether this policy is desirable is a matter of political economy beyond the scope of my investigation.

---

[23] Due to the way in which rating adjustments amplify premium changes for older buyers, mid-aged individuals—while still better off—benefit the least from the alternative design. However, once established that everyone would gain, other alternatives in which vouchers are adjusted more granularly by age could smooth changes in consumer surplus across groups, while still ensuring lower premiums and lower average subsidies. Ultimately, design decisions depend on welfare weights, which here are not needed to argue that a design would improve upon the status-quo in terms of enrolment, profits, and consumer surplus.

To implement alternative subsidy schemes and to consider other market design and regulatory questions—*e.g.* the role of a public option, different risk adjustment models, or quality regulations—future work could extend the model to account for dynamic or behavioural aspects, and for the key role played by healthcare providers. Access to richer data, including measures of health risk and healthcare utilization at the individual level, would facilitate the calculation of optimal policy parameters by researchers and policymakers.

APPENDIX

A. DEMAND MODEL: PARAMETRIC ASSUMPTIONS

The premium coefficient $\alpha_t(\mathbf{z})$ is allowed to vary across year, and across seven 6-years-wide age bins, and linearly with income. The coefficient on actuarial value $\beta_t(\mathbf{z},\theta)$ is log-normally distributed with year-age-bin-specific parameters.

Letting $A^1=\{26,\dots,31\}$, $A^2=\{32,\dots,37\}$, $\dots$, $A^6=\{56,\dots,61\}$, $A^7=\{62,63,64\}$,

$$\alpha_t(\mathbf{z})=\begin{cases}\alpha_t^{0,1}+\alpha_t^{1,1}z^{\text{Inc}} & \text{if } z^{\text{Age}}\in A^1\\ \alpha_t^{0,2}+\alpha_t^{1,2}z^{\text{Inc}} & \text{if } z^{\text{Age}}\in A^2\\ \dots & \\ \alpha_t^{0,7}+\alpha_t^{1,7}z^{\text{Inc}} & \text{if } z^{\text{Age}}\in A^7\end{cases};$$

all parameters are year-specific.

The coefficient on actuarial value is log-normally distributed with year-age-bin-specific parameters:

$$\beta_t(\mathbf{z},\theta)=\begin{cases}e^{\beta_t^1+\sigma_t^1\theta}, & \text{if } z^{\text{Age}}\in A^1\\ \dots & \\ e^{\beta_t^7+\sigma_t^7\theta}, & \text{if } z^{\text{Age}}\in A^7\end{cases},\quad \text{where}\quad \theta\sim G(\theta)=\mathcal{N}(0,1);$$

$\mathcal{N}$ indicates the standard normal distribution, θ and $\mathbf{z}$ are independent:

$$G_{mt}(\mathbf{z},\theta)=G_{mt}(\mathbf{z})G(\theta).$$

The term $\boldsymbol{\mu}_t(\mathbf{z})\mathbf{x}_{jmt}$ is equal to

$$\boldsymbol{\mu}_t(\mathbf{z})\mathbf{x}_{jmt}=\begin{cases}\mu_t^{0,1}+\mu_t^{1,1}z^{\text{Inc}}+\mu_t^{2,1}z^{\text{Age}}+\mu_t^{3,1}\text{HMO}_{jmt}+\boldsymbol{\mu}_t^{4,1}\text{Insurer}_{jmt} & \text{if } z^{\text{Age}}\in A^1\\ \dots & \\ \mu_t^{0,7}+\mu_t^{1,7}z^{\text{Inc}}+\mu_t^{2,7}z^{\text{Age}}+\mu_t^{3,7}\text{HMO}_{jmt}+\boldsymbol{\mu}_t^{4,7}\text{Insurer}_{jmt} & \text{if } z^{\text{Age}}\in A^7\end{cases};$$

this allows the value of marketplace coverage to vary piecewise linearly by year, age, and income, and the value of each product to vary—with year-age-bin parameters—with the type of provider network and insurer brand. Lastly, I let γ_t to be a cubic function of ξ_{jmt}, specific to every year and every age bin:

$$\boldsymbol{\gamma}_t(\xi_{jmt};\mathbf{z})=\begin{cases}\gamma_t^{1,1}\xi_{jmt}+\gamma_t^{2,1}\xi_{jmt}^2+\gamma_t^{3,1}\xi_{jmt}^3 & \text{if } z^{\text{Age}}\in A^1\\ \dots & \\ \gamma_t^{1,7}\xi_{jmt}+\gamma_t^{2,7}\xi_{jmt}^2+\gamma_t^{3,7}\xi_{jmt}^3 & \text{if } z^{\text{Age}}\in A^7\end{cases}.$$

B. ROBUSTNESS TO MORAL HAZARD

The cost estimates in Table 4 and the simulation results in Section 5 maintained the assumption of no moral hazard (see *e.g.* Einav and Finkelstein, 2018). This assumption is dictated by the lack of data to identify correlation between willingness-to-pay and spending separately from the causal effect of coverage generosity on spending. In the model of Section 3, allowing spending to increase with actuarial value impacts the estimates of η^{WTP} and other cost parameters.

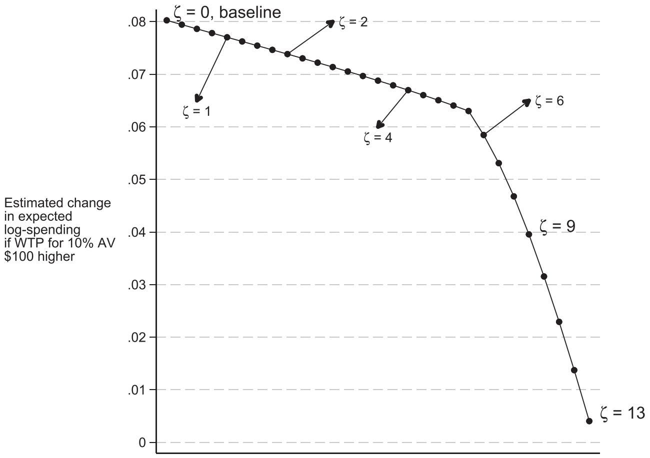

FIGURE 8

Estimated η^{WTP} varying assumptions on moral hazard

Note: The figure shows the estimated value of the adverse selection parameter η^{WTP} for different values of the moral hazard parameter ζ. The main results in the paper are obtained assuming $\zeta = 0$ (no moral hazard). The ACA risk adjustment model corresponds to $\zeta = 1$. $\zeta = 4$ (with results shown Table 7) corresponds to "400% ACA risk adjustment moral hazard".

Therefore, although the results above rely primarily on the fact that young uninsured individuals are generally healthy, the quantifications in Section 5 could be sensitive to different assumptions on moral hazard.

To address this, I re-estimate cost parameters and simulate policy counterfactuals under varying degrees of moral hazard. For reference, the ACA risk adjustment model (Pope *et al.*, 2014) assumes that medical spending increases, on average, by 3% when the individual is covered under a Silver plan (without cost-sharing reductions) relative to the spending under a Bronze plan; by 8% when covered under a Gold plan, and by 15% when covered under a Platinum plan. These moral hazard parameters are consistent with the findings of Lavetti *et al.* (2023), who estimate that when cost-sharing reductions increase actuarial value from 70% to 87% (94%) total spending is 13% (19%) higher.

Formally, I let the expected claims associated with individual *i* enroled in plan *j*, in region *m*, year *t* be equal to $\kappa^{\text{MH}}_{jmt}(\mathbf{z}_i, \theta_i) = AV^S_j L^{\text{MH}}_{jmt}(\mathbf{z}_i, \theta_i)$, with medical spending augmented for moral hazard defined as

$$L^{\text{MH}}_{jmt}(\mathbf{z}_i, \theta_i) = \left(1 + \zeta \times \chi_{ij}\right) L_{jmt}(\mathbf{z}_i, \theta_i), \tag{13}$$

where $\chi_{ij} = 0$ if $AV^D_{ij} < 70\%$, $\chi_{ij} = 0.03$ if $AV^D_{ij} \in [70\%, 75\%]$, $\chi_{ij} = 0.08$ if $AV^D_{ij} \in (75\%, 80\%]$, and $\chi_{ij} = 0.15$ if $AV^D_{ij} > 80\%$. $L_{jmt}(\mathbf{z}_i, \theta_i)$ is defined in equation (5). If $\zeta = 0$, the model is identical to the one in Sections 3 and 5. Varying ζ, one can explore the sensitivity of my findings to the presence of moral hazard. When $\zeta = 1$, the model sets moral hazard to the level assumed by the ACA risk adjustment formula.

Figure 8 shows the estimates of η^{WTP} varying ζ. From the baseline level of$\eta^{\text{WTP}} = 0.08$ obtained when $\zeta = 0$, setting $\zeta = 1$ reduces this estimate by 6% ($\eta^{\text{WTP}} = 0.075$). The estimates of η^{WTP} remain above 0.06 as long as the level of moral hazard is lower than six times the level assumed by the ACA risk adjustment formula. To obtain $\eta^{\text{WTP}} = 0$, which would indicate the absence of adverse selection, one would need to set ζ greater than 13, which seems quite unrealistic.

Table 7 explores the robustness of the results in Table 6 to alternative values of ζ. Considering the change in outcomes relative to the ACA-voucher equilibrium, the gains from increasing vouchers for young invincibles while lowering vouchers for older buyers remain present when assuming $\zeta = 1, 2$, or 4. Under perfect competition, the magnitude of the effects is almost invariant to ζ. Under Nash pricing, magnitudes are smaller when assuming larger degrees of moral hazard. However, even when setting $\zeta = 4$ the counterfactual vouchers increase enrollment for all buyers while reducing average subsidies.

Acknowledgments. Special thanks go to my Ph.D. advisor Liran Einav, and my dissertation committee Tim Bresnahan, Matt Jackson, and Mark Duggan. I also thank the editor, Thomas Chaney, and four anonymous referees for the

TABLE 7

Alternative assumptions on moral hazard and effect of age adjustments to vouchers

	Multi-product Nash						Perfect Competition					
	Change relative to ACA-voucher equilibrium +\$600 under-35 voucher, -\$100 over-35 voucher						Change relative to ACA-voucher equilibrium +\$600 under-35 voucher, -\$100 over-35 voucher					
Assumption on moral hazard:	26–35 enrolment	36–64 enrolment	26–35 premium	36–64 premium	Average CS	Average subsidy	26–35 enrolment	36–64 enrolment	26–35 premium	36–64 premium	Average CS	Average subsidy
$\zeta = 0$	0.117	0.013	−552	−76	105	−68	0.133	0.022	−689	−175	129	0
$\zeta = 1$	0.114	0.011	−542	−50	98	−39	0.131	0.022	−684	−172	129	13
$\zeta = 2$	0.103	0.003	−477	18	78	−27	0.129	0.02	−675	−155	123	4
$\zeta = 4$	0.107	0.006	−495	−27	83	−48	0.128	0.02	−666	−149	121	−17

Note: The table shows how the results of panel (b) in Table 6 vary when allowing medical spending to respond to coverage generosity (moral hazard). For each value of ζ, cost parameters are estimated replacing L_{jmt} from equation (5) with L_{jmt}^{MH} from equation (13), and equilibrium simulations are obtained with the new cost parameters. For each outcome, the results in the table correspond to the difference between the ACA-voucher equilibrium column and the counterfactual voucher equilibrium column in Table 6.

helpful suggestions. This article benefited from comments by many participants at seminars and conferences, and particularly from discussions and conversations with Francesco Decarolis, Martin Hackmann, Ben Handel, Kate Ho, and Mark Shepard.

Data availability statement

The data and code underlying this article are available in Zenodo, at https://doi.org/10.5281/zenodo.10456091.

Supplementary Data

Supplementary data are available at *Review of Economic Studies* online.

Funding

Research partly supported by the Kapnick Foundation (Stanford University) and by the Becker Friedman Institute Health Economics Initiative (University of Chicago).

REFERENCES

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

ATAL, J. P. (2019), *Lock-In in Dynamic Health Insurance Contracts: Evidence From Chile* (Mimeo: University of Pennsylvania).

ATAL, J. P., FANG, H., KARLSSON, M., *et al.* (2022), *German Long-Term Health Insurance: Theory Meets Evidence* (Mimeo: University of Pennsylvania).

AZEVEDO, E. M. and GOTTLIEB, D. (2017), “Perfect Competition in Markets with Adverse Selection”, *Econometrica*, **85**, 67–105.

BERRY, S., LEVINSOHN, J. and PAKES, A. (1995), “Automobile Prices in Market Equilibrium”, *Econometrica*, **63**, 841–890.

BERRY, S. T. and WALDFOGEL, J. (1999), “Public Radio in the United States: Does it Correct Market Failure or Cannibalize Commercial Stations?”, *Journal of Public Economics*, **71**, 189–211.

BLUMENTHAL, D., COLLINS, S. R. and FOWLER, E. J. (2020), “The Affordable Care Act at 10 Years–Its Coverage and Access Provisions”, *New England Journal of Medicine*, **382**, 963–969.

BRAMMLI-GREENBERG, S., GLAZER, J. and SHMUELI, A. (2018), “Regulated Competition and Health Plan Payment under the National Health Insurance Law in Israel–The Unfinished Story”, in *Risk Adjustment, Risk Sharing and Premium Regulation in Health Insurance Markets*, 365–395.

BRESNAHAN, T. F. (1987), “Competition and Collusion in the American Automobile Industry: The 1955 Price war”, *The Journal of Industrial Economics*, **35**, 457–482.

BROWN, J., DUGGAN, M., KUZIEMKO, I., *et al.* (2014), “How Does Risk Selection Respond to Risk Adjustment? New Evidence from the Medicare Advantage Program”, *American Economic Review*, **104**, 3335–3364.

BUNDORF, M. K., LEVIN, J. and MAHONEY, N. (2012), “Pricing and Welfare in Health Plan Choice”, *American Economic Review*, **102**, 3214–3248.

BUNDORF, M. K. and PAULY, M. V. (2006), “Is Health Insurance Affordable for the Uninsured?”, *Journal of Health Economics*, **25**, 650–673.

CHAN, D. and GRUBER, J. (2010), “How Sensitive Are Low Income Families to Health Plan Prices?”, *American Economic Review*, **100**, 292–296.

CHANDRA, A., GRUBER, J. and MCKNIGHT, R. (2011), “The Importance of the Individual Mandate–Evidence from Massachusetts”, *New England Journal of Medicine*, **364**, 293–295.

CICALA, S., LIEBER, E. M. J. and MARONE, V. (2019), “Regulating Markups in US Health Insurance”, *American Economic Journal: Applied Economics*, **11**, 71–104.

COLOMBO, F. and TAPAY, N. (2004), “Private Health Insurance in OECD Countries” (15).

COX, N., HANDEL, B., KOLSTAD, J., *et al.* (2015), “Messaging and the Mandate: The Impact of Consumer Experience on Health Insurance Enrollment Through Exchanges”, *American Economic Review*, **105**, 105–109.

CUESTA, J. I., NOTON, C. and VATTER, B. (2019), “Vertical Integration Between Hospitals and Insurers” (Available at SSRN 3309218).

CURTO, V., EINAV, L., LEVIN, J., *et al.* (2021), “Can Health Insurance Competition Work? Evidence from Medicare Advantage”, *Journal of Political Economy*, **129**, 570–606.

DECAROLIS, F. (2015a), “Medicare Part d: Are Insurers Gaming the low Income Subsidy Design?”, *American Economic Review*, **105**, 1547–1580.

DECAROLIS, F. (2015b), “The Unintended Effects of the Medicare Part D low Income Subsidy”, *Health Policy*, **119**, 597–603.

DECAROLIS, F., POLYAKOVA, M. and RYAN, S. P. (2020), “Subsidy Design in Privately Provided Social Insurance: Lessons from Medicare Part D”, *Journal of Political Economy*, **128**, 1712–1752.

DE VEN, W. P. M. M. V. and SCHUT, F. T. (2008), “Universal Mandatory Health Insurance in the Netherlands: A Model for the United States?”, *Health Affairs*, **27**, 771–781.

DIAMOND, R., DICKSTEIN, M. J., MCQUADE, T., *et al.* (2018), "Insurance Without Commitment: Evidence from the ACA Marketplaces" (Technical Report, National Bureau of Economic Research).

DICKSTEIN, M. J., DUGGAN, M., ORSINI, J., *et al.* (2015), "The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the Affordable Care Act", *American Economic Review*, **105**, 120–125.

DICKSTEIN, M. J., HO, K. and MARK, N. D. (2024), "Market Segmentation and Competition in Health Insurance", *Journal of Political Economy*, **132**, 96–1481.

DOMURAT, R., MENASHE, I. and YIN, W. (2021), "The Role of Behavioral Frictions in Health Insurance Marketplace Enrollment and Risk: Evidence from a Field Experiment", *American Economic Review*, **111**, 1549–1574.

DRAKE, C., RYAN, C. and DOWD, B. (2022), "Sources of Inertia in the Individual Health Insurance Market", *Journal of Public Economics*, **208**, 104622.

EINAV, L. and FINKELSTEIN, A. (2018), "Moral Hazard in Health Insurance: What we Know and how we Know it", *Journal of the European Economic Association*, **16**, 957–982.

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

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

ERICSON, K. M. M. and STARC, A. (2015), "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange", *Review of Economics and Statistics*, **97**, 667–682.

FANG, H. and KRUEGER, D. (2022), "The Affordable Care act after a Decade: Its Impact on the Labor Market and the Macro Economy", *Annual Review of Economics*, **14**, 453–494.

FINKELSTEIN, A., HENDREN, N. and SHEPARD, M. (2019), "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts", *American Economic Review*, **109**, 1530–1567.

FLEITAS, S. (2020), "Who Benefits When Inertia is Reduced? Competition, Quality and Returns to Skill in Health Care Markets" (CEPR Discussion Paper 14292).

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

GERUSO, M., LAYTON, T. J., MCCORMACK, G., *et al.* (2023), "The two Margin Problem in Insurance Markets", *The Review of Economics and Statistics*, **105**, 237–2572.

GERUSO, M., LAYTON, T. and PRINZ, D. (2019), "Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges", *American Economic Journal: Economic Policy*, **11**, 64–107.

GRAETZ, I., MCKILLOP, C. N., KAPLAN, C. M., *et al.* (2018), "Lessons Learned from the Affordable Care Act: The Premium Subsidy Design may Promote Adverse Selection", *Medical Care Research and Review*, **75**, 762–772.

GRAVES, J. A. and GRUBER, J. (2012), "How did Health Care Reform in Massachusetts Impact Insurance Premiums?", *American Economic Review*, **102**, 508–513.

GRUBER, J. (2010), "Massachusetts Points the way to Successful Health Care Reform", *Journal of Policy Analysis and Management*, **30**, 184–192.

HACKMANN, M. B., KOLSTAD, J. T. and KOWALSKI, A. E. (2015), "Adverse Selection and an Individual Mandate: When Theory Meets Practice", *American Economic Review*, **105**, 1030–1066.

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

HANDEL, B. and HO, K. (2021), "The Industrial Organization of Health Care Markets", in *Handbook of Industrial Organization* (Vol. 5, Elsevier) 521–614.

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

HINDE, J. M. (2017), "Incentive (less)? The Effectiveness of tax Credits and Cost-Sharing Subsidies in the Affordable Care Act", *American Journal of Health Economics*, **3**, 346–369.

HOLLY, A., GARDIOL, L., DOMENIGHETTI, G., *et al.* (1998), "An Econometric Model of Health Care Utilization and Health Insurance in Switzerland", *European Economic Review*, **42**, 513–522.

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

KREIER, R. and ZWEIFEL, P. (2010), "Health Insurance in Switzerland: A Closer Look at a System Often Offered as a Model for the United States", *Hofstra Law Review*, **39**, 89.

LAVETTI, K. J., DELEIRE, T. and ZIEBARTH, N. R. (2023), "How Do Low-Income Enrollees in the Affordable Care Act Marketplaces Respond to Cost-Sharing?", *Journal of Risk and Insurance*, **90**, 155–1831.

LEVINE, D. and MULLIGAN, J. (2017), "Mere Mortals: Overselling the Young Invincibles", *Journal of Health Politics, Policy and Law*, **42**, 387–407.

MAHONEY, N. (2015), "Bankruptcy as Implicit Health Insurance", *American Economic Review*, **105**, 710–746.

MAHONEY, N. and WEYL, E. G. (2017), "Imperfect Competition in Selection Markets", *Review of Economics and Statistics*, **99**, 637–651.

MARONE, V. R. and SABETY, A. (2022), "When Should There Be Vertical Choice in Health Insurance Markets?", *American Economic Review*, **112**, 304–342.

MCFADDEN, D. (1973), "Conditional Logit Analysis of Qualitative Choice Behavior".

MCGUIRE, T. G., SCHILLO, S. and VAN KLEEF, R. C. (2020), “Reinsurance, Repayments, and Risk Adjustment in Individual Health Insurance: Germany, the Netherlands, and the US Marketplaces”, *American Journal of Health Economics*, **6**, 139–168.

MILLER, K. S., PETRIN, A., TOWN, R., *et al.* (2022), “Optimal Managed Competition Subsidies” (Working Paper 25616, National Bureau of Economic Research).

MILLER, T. P. (2017), “Examining the Effectiveness of the Individual Mandate under the Affordable Care Act”, *American Enterprise Institute, Statement Before the House Committee on Ways and Means Subcommittee on Oversight*, **6**, 1–12.

NEVO, A. (2001), “Measuring Market Power in the Ready-To-eat Cereal Industry”, *Econometrica*, **69**, 307–342.

ORSINI, J. and TEBALDI, P. (2017), “Regulated Age-Based Pricing in Subsidized Health Insurance: Evidence from the Affordable Care Act” (Working Paper, Becker Friedman Institute for Research in Economics).

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

PAULY, M. V. (1970), “Efficiency in the Provision of Consumption Subsidies”, *Kyklos*, **23**, 33–57.

POLYAKOVA, M. and RYAN, S. P. (2019), “Subsidy Targeting with Market Power” (Technical Report, National Bureau of Economic Research).

POPE, G. C., BACHOFER, H., PEARLMAN, A., *et al.* (2014), “Risk Transfer Formula for Individual and Small Group Markets under the Affordable Care Act”, *Medicare & Medicaid Research Review*, **4**, 1–23.

ROOS, A.-F. and SCHUT, F. T. (2012), “Spillover Effects of Supplementary on Basic Health Insurance: Evidence from the Netherlands”, *The European Journal of Health Economics*, **13**, 51–62.

ROSSE, J. N. (1970), “Estimating Cost Function Parameters Without Using Cost Data: Illustrated Methodology”, *Econometrica: Journal of the Econometric Society*, **38**, 256–275.

ROTHSCHILD, M. and STIGLITZ, J. E. (1976), “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information”, *The Quarterly Journal of Economics*, **90**, 629–649.

RUGGLES, S., GENADEK, K., GOEKEN, R., *et al.* (2015), *Integrated Public Use Microdata Series: Version 6.0 [Dataset]* (Minneapolis: University of Minnesota).

RYAN, C. (2023), *Mergers in the Presence of Adverse Selection* (Mimeo: Penn State University).

SALTZMAN, E. (2019), “Demand for Health Insurance: Evidence from the California and Washington ACA Exchanges”, *Journal of Health Economics*, **63**, 197–222.

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

SALTZMAN, E. and LUCARELLI, C. (2021), “Firm Learning in a Selection Market”, in *10th Annual Conference of the American Society of Health Economists* (ASHECON).

SALTZMAN, E., SWANSON, A. and POLSKY, D. (2021), “Inertia, Market Power, and Adverse Selection in Health Insurance: Evidence from the ACA Exchanges” (Working Paper 29097, National Bureau of Economic Research).

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

SMITH, J. C. and MEDALIA, C. (2014), *Health Insurance Coverage in the United States: 2013* (USA: US Department of Commerce, Economics and Statistics Administration, Bureau of the Census).

SOMMERS, B. D., SHEPARD, M. and HEMPSTEAD, K. (2018), “Why did Employer Coverage Fall in Massachusetts after the ACA? Potential Consequences of a Changing Employer Mandate”, *Health Affairs*, **37**, 1144–1152.

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

SUMMERS, L. H. (1989), “Some Simple Economics of Mandated Benefits”, *The American Economic Review*, **79**, 177–183.

TEBALDI, P., TORGOVITSKY, A. and YANG, H. (2023), “Nonparametric Estimates of Demand in the California Health Insurance Exchange”, *Econometrica*, **91**, 107–146.

VEIGA, A. (2023), “Price Discrimination in Selection Markets”, *Review of Economics and Statistics*, 1–45.

WAGSTAFF, A. and VAN DOORSLAER, E. (2000), “Chapter 34 Equity in Health Care Finance and Delivery”, *Handbook of Health Economics*, **1**, 1803–1862.

WALDFOGEL, J. (2003), “Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets”, *RAND Journal of Economics*, **34**, 557–568.