NBER WORKING PAPER SERIES

ESTIMATING EQUILIBRIUM IN HEALTH INSURANCE EXCHANGES: PRICE COMPETITION AND SUBSIDY DESIGN UNDER THE ACA

Pietro Tebaldi

Working Paper 29869
http://www.nber.org/papers/w29869

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
March 2022

Research partly supported by the Kapnick Foundation and by the Becker Friedman Institute. Six anonymous referees and Pinelopi Goldberg provided numerous helpful comments that lead to substantial improvements in this article. Special thanks go to my thesis advisor, Liran Einav, and to Tim Bresnahan, Matt Jackson, and Mark Duggan for their support and guidance. I also thank the many participants at seminars and conferences for their feedback, and I am grateful to Sofronis Clarides, Martin Hackmann, Ben Handel, Kate Ho, Mark Shepard, and Jonathan Williams for the insightful discussions. Hanbin Yang, Honglin Li, and Matteo Saccarola provided outstanding research assistance. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2022 by Pietro Tebaldi. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA
Pietro Tebaldi
NBER Working Paper No. 29869
March 2022
JEL No. I13,I18,L98

ABSTRACT

In government-sponsored health insurance, subsidy design affects market outcomes. First, holding premiums fixed, subsidies determine insurance uptake and average cost. Insurers then respond to these changes, adjusting premiums. Combining data from the first four years of the California ACA marketplace with a model of insurance demand, cost, and insurers' competition, I quantify the impact of alternative subsidy designs on premiums, enrollment, costs, public spending, and consumer surplus. Younger individuals are more price sensitive and cheaper to cover. Increasing subsidies to this group would make all buyers better off, increase market participation, and lower average costs and average subsidies.

Pietro Tebaldi
Department of Economics
Columbia University
IAB, MC 3308
New York, NY 10027
and NBER
pt2571@columbia.edu

An online appendix is available at http://www.nber.org/data-appendix/w29869

1 Introduction

Welfare losses from adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976; Einav, Finkelstein, and Cullen, 2010a), 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 (Colombo and Tapay, 2004). However, the way in which policy-relevant market outcomes vary in response to alternative subsidy schemes is yet to be fully understood.

A recent, large-scale example of the use of means-tested subsidies in government-sponsored health insurance is found in the low-income subsidy—premium tax credit—introduced by the 2010 US health care reform (Patient Protection and Affordable Care Act; ACA). Since 2014, under this program the Federal government spent around \$40 billion per-year to provide health insurance subsidies to more than 10 million US citizens.[1] Knowledge of the relationship between subsidy design and outcomes such as coverage levels and public spending is critical to evaluate the success of the ACA, and for the design of similar programs in the future.

The goal of this article is to develop an empirical framework to analyze the dependence of equilibrium outcomes on how subsidies interact with three important features of private health insurance markets: demand from subsidized individuals, insurers' competition, and adverse selection generated by the correlation between willingness-to-pay and expected health cost. Characteristics of demand determine the extent to which subsidies alter insurance enrollment decisions. Pricing incentives of competing insurers react to these changes in demand, but also to corresponding changes in expected cost driven by differences in risk selection.

The framework is then used to measure the potential benefits of modifying the generosity of ACA subsidies, providing additional incentives for the participation of young adults. To a first approximation, the ACA determined that subsidy-eligible individuals—who have income lower than four times the federal poverty level (FPL)—must spend less than a "maximum affordable amount" for health insurance. This amount is increasing in income, but does not vary with age. Moreover, if insurers vary (pre-subsidy) premiums for one age group, they must also adjust the premiums for other age groups in fixed proportions. Figure 1a illustrates schematically the resulting ACA design, which implies that subsidies are more generous for older enrollees.

[1] Comprehensive statistics on enrollment and subsidies are released regularly by the Congressional Budget Office. See e.g. https://www.cbo.gov/taxonomy/term/45/recurring-reports, or https://www.cbo.gov/system/files/2019-04/55094-CoverageUnder65_0.pdf, last accessed on January 22, 2022.

Figure 1: *ACA subsidy design and the impact of age-adjustments to subsidy generosity*

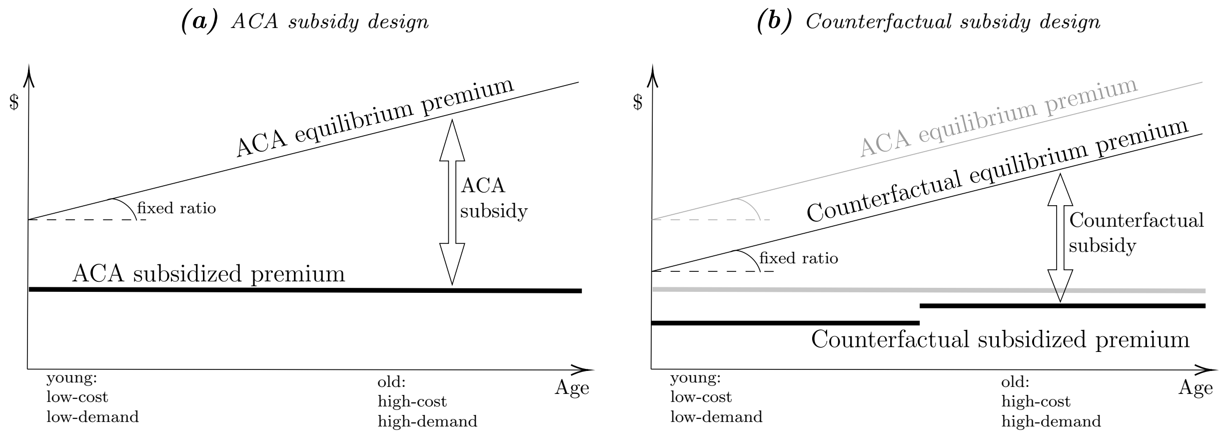

Younger individuals are, on average, less willing to pay for health insurance, more price sensitive, and cheaper to cover. Therefore, a counterfactual subsidy design that favors this group as illustrated in Figure 1b would, upon increasing the proportion of young enrollees, lower average cost and increase average elasticity of demand. This in turn would place downward pressure on equilibrium premiums, and reduce subsidy-spending per-buyer, while also ensuring that *all* enrollees, including the older ones, pay less than under the original ACA design.

To quantify this mechanism I introduce an empirical model of equilibrium pricing in ACA marketplaces. The model combines estimates of demand and cost with the details of ACA regulations, including adjusted community rating rules, risk adjustment, and the design of premium subsidies. In terms of insurers' conduct, I consider two possible alternatives: static Nash oligopoly pricing and perfect competition. Although the latter represents a relevant benchmark, I show that Nash pricing is more consistent with observed patterns in the data. In terms of subsidy design, I compare the ACA model—in which government discounts are calculated as a function of premiums (c.f. "price-linked" subsidies Jaffe and Shepard, 2020)—to the use of fixed vouchers. Despite small differences in quantifications, the alternative modelling and design choices imply similar first-order effects of shifting subsidy generosity toward young adults on enrollment, consumer surplus, and public spending.

The estimates of demand are obtained using individual-level premiums and enrollment data from one of the largest ACA marketplaces, Covered California, observed during the 2014-2017 period, which I combine with survey measures of uninsurance and subsidy eligibility by age, income, and geographic region. The raw data highlights the essential variation underlying the results: subsidized premiums are approximately constant in age, while older individuals are more likely to enroll. A mixed-logit discrete choice model of insurance demand leads to precise measures of preferences and demand heterogeneity by age. I allow preferences for plan generosity, premium, insurer, and type of provider network to vary by year of enrollment, age, and income. Moreover, I allow for unobserved heterogeneity across buyers in willingness-to-pay for coverage.

Identification of demand parameters relies on two aspects of ACA regulations. First, discrete variation in cost-sharing reductions induces discontinuities in the actuarial value of Silver plans at three income thresholds (see also Hinde, 2017; Lavetti, DeLeire, and Ziebarth, 2019). Second, community rating restrictions point directly toward a Waldfogel instrumental variable strategy (c.f. Berry and Waldfogel, 1999; Waldfogel, 2003). The ACA allows insurers to set only one (base) premium for every plan in a given geographic region, and pre-determined pricing schedules are then used to translate this 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; Polyakova and Ryan, 2019). 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, a flexible control function allows me to correct for premium endogeneity when estimating demand.

The resulting demand estimates display large heterogeneity across buyers of different age. On average, relative to older groups individuals younger than 35 value insurance generosity less than half. At the same time, they are twice as responsive to premium changes. These estimates align with the growing literature on health insurance demand among low-income adults, including Chan and Gruber (2010); Ericson and Starc (2015); Finkelstein, Hendren, and Shepard (2019), while also adding rich observed and unobserved heterogeneity, and considering a large marketplace regulated by the ACA (see also Panhans, 2019; Saltzman, 2019, 2021; Tebaldi, Torgovitsky, and Yang, 2019; Dickstein, Ho, and Mark, 2021).

To estimate expected insurance costs, the model employs year-plan-level average claims covering a large portion of plans observed in the enrollment data (as in Bundorf, Levin, and Mahoney, 2012), and individual-level healthcare spending information from the Medical Expenditure Panel Survey. Expected annual medical spending can vary across individuals and plans. For simplicity, and not having access to individual level claims data, the baseline model rules out moral hazard, while it accounts for adverse selection by letting expected medical for spending for a specific individual to vary observably with age, and unobservably with willingness-to-pay for insurance generosity. The main 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. (2019).

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 spending. An increase in willingness-to-pay (for ten percentage points in actuarial value) of $500 per year implies 35% higher medical spending. At the observed premiums, insurers' average expected cost for a Bronze (Silver) enrollee aged between 26-31 is $1148 ($1435) per year. If the enrollee is older, aged between 38-43, the cost increases to $1969 ($2504) per year. If the enrollee is aged between 56-61, the estimates imply that the cost is $5523 ($7491) per year.

The simulations of equilibrium under alternative subsidy designs and assumptions about insurer conduct combine demand and cost estimates with (adjusted) community rating, risk adjustment (see also Saltzman, 2021), and the design of premium subsidies. The results show large potential gains from increasing subsidies for young individuals. As long as each $1 decrease in subsidies for those aged 36-64 is compensated by a $4 increase in subsidies for individuals under 35, *all buyers are better off*, facing lower subsidized premiums, and experiencing higher consumer surplus. Marketplace coverage increases, while average cost and average subsidies are lower. The extent to which these "Pareto improvements" are possible represents a promising direction for the design of subsidies in ACA marketplaces, and is informative for similar regulatory contexts.

My results speak directly to the growing body of work analyzing the effect of different regulations in government-sponsored health insurance markets. Among these studies, the literature on competition and market design in Medicare Advantage and Medicare Part D is more mature, including analyses of subsidies in Decarolis (2015); Decarolis, Polyakova, and Ryan (2020); Curto, Einav, Levin, and Bhattacharya (2021); Miller, Petrin, Town, and Chernew (2022). Beyond considering a different population and regulatory environment, the model in this article includes heterogeneity in individual expected cost that is not accounted for by risk adjustment (see also Brown, Duggan, Kuziemko, and Woolston, 2014).

These studies, as well as theoretical work in Mahoney and Weyl (2017); Veiga (2020), and references therein, provide welfare considerations by aggregating across groups of individuals who are better or worse off. Here, instead, the counterfactual subsidy designs show situations in which all buyers are better off relative to the status-quo. This is essential to avoid concerns for reclassification risk (c.f. Handel, Hendel, and Whinston, 2015); the proposed alternatives are unanimous improvements for risk-averse individuals, since premiums are lower in any possible state of the world.

The US health insurance market for individuals under-65 has been analyzed primarily through the lenses of the Massachusetts healthcare reform, which served as a blueprint for the ACA. Evidence of heterogeneity in preferences and cost across age groups and the role of imperfect competition is the emphasis of Ericson and Starc (2015). Jaffe and Shepard (2020) measure the distortion of price-linked subsidies relative to fixed vouchers; my simulations extend this comparison to the ACA setting, in which insurers offer multiple plans and subsidies are linked to the second cheapest Silver, rather than cheapest plan in the market. Finkelstein et al. (2019) measure adverse selection, show that low-income adults value insurance less than expected cost, and discuss fiscal externalities (see also Mahoney, 2015) as a rationale for providing premium subsidies. In the ACA context, Saltzman (2019, 2021)—who also uses data from California—focuses on the role of individual mandates and risk adjustment, while Panhans (2019) provides evidence of adverse selection in Colorado. Polyakova and Ryan (2019) use aggregate enrollment data across many states to measure how the composition of the uninsured population affects the incidence of subsidies on consumers. My work contributes to this literature by introducing a promising alternative for the design of ACA subsidies, and by measuring its equilibrium effects using novel demand and cost estimates with rich demographic and unobserved heterogeneity.

2 ACA Marketplaces and Subsidy Design

2.1 Institutional Background and Regulations

As of 2013, 17 percent of US 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 fifty states. ACA marketplaces operate separately across states, but they all follow similar institutions and regulations as mandated by the federal reform.[2]

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, Duggan, Orsini, and Tebaldi, 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.

[2] States can choose between instituting their own marketplace, relying on the federal platform, or adopting a state-federal partnership model; the political divide between Republican and Democratic parties in supporting the reform led to different implementations.

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 6460 of title 10 of the California Code of Regulations; May 21, 2014.

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.[3]

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 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,

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

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_j^i equal to the product of b_j and the corresponding age adjustment, an increasing function of Age_i.[4]

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

Premium Subsidies: Although R_j^i is the amount collected by the insurer, enrolled 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 (MAA). 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 MAA 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 enroll 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 enrolled 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 1-1.5 times the FPL, 88% for income levels between 1.5-2 times the FPL, and 74% for income levels between

[4] 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 3b.

2-2.5 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.[5]

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 enrollment 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 enrollees 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.[6]

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%.[7] I do

[5] At the end of 2017, the Trump administration interrupted the funding of cost-sharing reductions, after a legal dispute over the appropriation of federal funds: c.f. *House v. Burwell*, *House v. Price*.

[6] 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, 2015; Decarolis et al., 2020).

[7] Federal reinsurance was mandatory between 2014-2016, 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 2014-2016. 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 was lifted by the Trump administration in 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.

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 enrollment, 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 honored (Diamond, Dickstein, McQuade, and Persson, 2018, discuss the relationship between medical spending and interruptions of premium payments).

2.2 Counterfactual Subsidy Design

The design of ACA subsidies is such that, for all subsidized individuals, premiums vary by income, while for a given income the level of premiums is age-invariant: $\overline{P}$ does not depend on age in (1) above. I argue that, holding fixed community rating regulations, a subsidy scheme that leads to equal premiums across individuals who differ in insurance willingness-to-pay *and* risk can be strictly worse, *for all consumers*, than a scheme in which subsidies are such that premiums differ across types.

To see this in a stylized model, consider a single (monopolist) insurer setting the premium p for a given (exogenous) coverage option. An individual is characterized by a scalar-valued observable type z. A type determines preferences and insurable cost: $q(p; z)$ is the probability that a type-z buyer purchases coverage when facing premium p, and $c(z)$ is the corresponding expected cost incurred by the insurer. The mass of type-z potential buyers in the population is $G(z)$.

The government provides a subsidy $s(z)$ to every type-z individual who chooses to purchase coverage. Taking $s(\cdot)$ as given, the insurer solves

$$\max_p Q^s(p)\,(p - AC^s(p))\,,$$

where quantity and average costs are

$$Q^s(p) \equiv \int q(p - s(z); z) dG(z);$$
$$AC^s(p) \equiv (Q^s(p))^{-1} \int q(p - s(z); z) c(z) dG(z).$$

The insurer sets p solving

$$p = AC^s(p) + \left[-\frac{Q^s(p)}{\frac{dQ^s(p)}{dp}} \left(1 - \frac{dAC^s(p)}{dp} \right) \right]; \tag{2}$$

the term in square brackets is the price-cost markup, which depends on the (inverse) elasticity of demand, and on a "selection correction" that takes into account the slope of the average cost curve. Adverse selection implies that $dAC^s(p)/dp > 0$. By choosing the subsidy function $s(\cdot)$, the government affects equilibrium price $p^\star(s)$, enrollment, and welfare.[8]

Let $s(z) = \overline{s}$ for all z; all individuals then face the same premium $p^\star(s) - \overline{s}$. Even "behind the veil of ignorance", i.e. considering ex-ante expected utility when z is still unknown, there is no reclassification risk (c.f. Handel et al., 2015). Individuals do not face premium uncertainty, and—denoting with $V(p; z)$ the money-metric indirect utility for type z when premium is p—the average consumer expected utility is simply $\mathbb{E}_z [V(p^\star(s) - \overline{s}; z)]$. An alternative subsidy scheme $\hat{s}(\cdot)$ such that $\hat{s}(z)$ varies by z creates reclassification risk. Relative to $s(z) = \overline{s}$, if $p^\star(s) - \overline{s} < p^\star(\hat{s}) - \hat{s}(z)$ for some z, for these types $V(p^\star(\hat{s}) - \hat{s}(z); z) < V(p^\star(s) - \overline{s}; z)$. Then, depending on G, and on the curvature of V, one cannot rule out that $\mathbb{E}_z [V(p^\star(\hat{s}) - \hat{s}(z); z)] < \mathbb{E}_z [V(p^\star(s) - \overline{s}; z)]$.

However, assuming that higher z imply higher costs, higher demand, and lower semi-elasticity of demand, it may be possible to find a non-constant $\hat{s}(z)$ for which $p^\star(s) - \overline{s} > p^\star(\hat{s}) - \hat{s}(z)$ for all z, and for which average per-enrollee subsidies are lower. If this is the case, for any G and any V, $\mathbb{E}_z [V(p^\star(\hat{s}) - \hat{s}(z); z)] > \mathbb{E}_z [V(p^\star(s) - \overline{s}; z)]$. Even if individuals face subsidized premiums that may vary with z, these are *always lower* than the amounts paid under the constant subsidy $\overline{s}$. In this scenario, the alternative subsidy scheme $\hat{s}$ is an improvement over s.

To build the alternative $\hat{s}(\cdot)$, one can increase by $\overline{\Delta}$ the subsidy for low-z types, and decrease by $\underline{\Delta}$ the subsidy for high-z types. Given a value $\hat{z}$, $\hat{s}(z) = \overline{s} + \overline{\Delta}$ for all $z \leq \hat{z}$; $\hat{s}(z) = \overline{s} - \underline{\Delta}$ for all $z > \hat{z}$. Relative to s, $\hat{s}$ implies lower average cost and higher semi-elasticity of demand, since the share of low-z types in the enrollment pool is higher. If the difference $\frac{dAC^s(p^\star(s))}{dp} - \frac{dAC^{\hat{s}}(p^\star(s))}{dp}$ is negative, or—if positive—not too large, the equilibrium pre-subsidy premium under $\hat{s}$ is lower than under s: $p^\star(\hat{s}) < p^\star(s)$. This would also hold true in a perfectly competitive market in which $p = AC^s(p)$, since $AC^{\hat{s}}(p) < AC^s(p)$ for all p.

This result is quite intuitive: by increasing participation of low-cost, high-elasticity types, the government puts downward pressure on premiums. Importantly, if $p^\star(s) -$

[8] For an extensive discussion of equilibrium existence in a market with adverse selection, I refer the reader to Azevedo and Gottlieb (2017) and references therein. When I simulate equilibrium under perfect competition, I use their result directly, ensuring existence by allowing an infinitesimal fraction of buyers to be randomly assigned across products, rather than responding to premiums. When simulating equilibrium under imperfect competition, I follow the empirical industrial organization literature of optimal pricing by multi-product firms (Bresnahan, 1987; Nevo, 2001). In the insurance context, this has been adopted widely (see e.g. Bundorf et al., 2012; Starc, 2014; Decarolis et al., 2020; Saltzman, 2021; Curto et al., 2021).

__Table 2:__ Equilibrium and Alternative Subsidy Designs: Numerical Example

Model parameters:

	z^1	z^2	z^3	z^4
a^z	6.25	6.5	7	7.25
b^z	0.15	0.12	0.1	0.075
$c(z)$	60	100	120	140

$G(z) = 1000$ for all z

$q(p;z) = \exp(a^z - b^z p) \,/\, (1 + \exp(a^z - b^z p))$

Initial subsidy design s: $\overline{s} = 70$

z	z^1	z^2	z^3	z^4
$s(z)$	70	70	70	70

Alternative subsidy design $\hat{s}$: $\hat{z} = z^2$, $\overline{\Delta} = 10$, $\underline{\Delta} = 20$

z	z^1	z^2	z^3	z^4
$s(z)$	80	80	50	50

Equilibrium outcomes under subsidy design s, setting $p = p^\star(s)$:

					$p = p^\star(s)$	$Q^s(p)$	$AC^s(p)$	$\frac{dQ^s(p)}{dp}/Q^s(p)$	$\frac{dAC^s(p)}{dp}$	RHS of FOC in (2) minus p	Subsidy per-enrollee
$p^\star(s) - \overline{s}$	85	85	85	85							
$q(p^\star(s) - \overline{s}; z)$	0.002	0.025	0.185	0.709	155	921	135	-0.04	0.28	0	70

Off-equilibrium outcomes under subsidy design $\hat{s}$, holding $p = p^\star(s)$:

					$p = p^\star(s)$	$Q^{\hat{s}}(p)$	$AC^{\hat{s}}(p)$	$\frac{dQ^{\hat{s}}(p)}{dp}/Q^{\hat{s}}(p)$	$\frac{dAC^{\hat{s}}(p)}{dp}$	RHS of FOC in (2) minus p	Subsidy per-enrollee
$p^\star(s) - \hat{s}(z)$	75	75	105	105							
$q(p^\star(s) - \hat{s}(z); z)$	0.007	0.078	0.030	0.352	155	466	131	-0.06	0.45	-15	55

Equilibrium outcomes under subsidy design $\hat{s}$, setting $p = p^\star(\hat{s})$:

					$p = p^\star(\hat{s})$	$Q^{\hat{s}}(p)$	$AC^{\hat{s}}(p)$	$\frac{dQ^{\hat{s}}(p)}{dp}/Q^{\hat{s}}(p)$	$\frac{dAC^{\hat{s}}(p)}{dp}$	RHS of FOC in (2) minus p	Subsidy per-enrollee
$p^\star(\hat{s}) - \hat{s}(z)$	42	42	72	72							
$q(p^\star(\hat{s}) - \hat{s}(z); z)$	0.499	0.817	0.458	0.867	122	2640	109	-0.03	0.57	0	65

Note: See example below.

$p^\star(\hat{s}) > \underline{\Delta}$, $V(p^\star(\hat{s}) - \hat{s}(z); z) > V(p^\star(s) - \overline{s}; z)$ for all z, $\mathbb{E}_z\left[V(p^\star(\hat{s}) - \hat{s}(z); z)\right] > \mathbb{E}_z\left[V(p^\star(s) - \overline{s}; z)\right]$ for any G and any V. If, moreover, $\overline{\Delta} < \underline{\Delta}$, average per-enrollee subsidy spending is lower under $\hat{s}$.

Example: To see this mechanism at work through a simple example, consider a market with primitives summarized in Table 2, there are four types $z = z^1, z^2, z^3, z^4$, and the model parameters are set so that higher z implies higher cost, higher demand, and lower semi-elasticity of demand. In the equilibrium under a subsidy scheme $s(z) = 70$ for all z, the premium is $p^\star(s) = 155$, and all buyers pay 85. Probability of enrollment among types z^1 and z^2 is lower than 0.03, while z^3 and z^4 enroll with probability 0.18 and 0.7, respectively. Overall enrollment is 23% of the 4000 potential buyers; average cost is 135, and the average subsidy per-enrollee is (trivially) equal to 70.

Alternatively, consider the scheme $\hat{s}$, where the subsidy of z^1 and z^2 is increased by $\overline{\Delta} = 10$, while the subsidy for z^3 and z^4 is lowered by $\underline{\Delta} = 20$. The first-order effect—holding premium fixed to $p^\star(s)$—is to make z^1 and z^2 better-off, while z^3 and z^4 are worse-off, relative to the design s. However, $p^\star(s)$ is not an equilibrium premium under the design $\hat{s}$: average cost is lower, while semi-elasticity of demand and the derivative of average cost are both higher. The insurer has incentives to set a premium lower than $p^\star(s)$, since the difference between left- and right-hand-side of (2) is -15.

The new equilibrium is $p^\star(\hat{s}) = 122$; with total enrollment 2640 (+186% relative to $\hat{s}$), average cost 109, and average per-enrollee subsidy equal to 65. Critically, *all types are better off*, since they face subsidized premiums that are lower than 85. Types z^1 and z^2 pay 42, while z^3 and z^4 pay 72. Therefore, under $\hat{s}$ buyers are unambiguously better relative to design s, and government spending per-enrollee is lower. Profits are also higher, increasing from 18420 to 34320.

To evaluate alternative designs in a specific context, one needs estimates of market primitives, and a model that must include specific regulations and competition between differentiated insurers.

3 Data Sources and Summary Statistics

3.1 Enrollment Files

Covered California provided me with individual-level enrollment 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 enrollment in Covered California during my observation period, for a total of 3.72 million individuals. Excluding dependents, who under the ACA can be as old as 25, the coverage decisions for this group are simpler, and easier to analyze. Moreover, since off-exchange plans are not eligible for subsidies, excluding the unsubsidized population mitigates concerns that enrollment files may miss many individuals purchasing coverage outside the marketplace.

The top-left panel of Table 3 summarizes the enrollment data. Average age among subsidized adults in Covered California is 45.8 (with standard deviation 11.7), while average income is 214.5 (63.9) percent of the FPL. Individuals pay, on average, \$1470 (\$1264) per-person, per-year, receiving subsidies that are, on average, more than 2.5 times as large. In terms of metal tier, 24% of enrolled individuals choose a Bronze plan, while 68% choose a Silver plan. Gold and Platinum plans are selected more rarely.

Figure 2a plots how insurer revenue, subsidized premium, and the difference between Bronze and Silver premium vary across enrollees of different age. The average

Table 3: *Summary statistics*

	Individual-level data (person-year)					
	Enrolled (Covered CA) N = 3719273		Eligible (ACS draws) N = 13265960		Surveyed (MEPS) N = 20171	
Age	45.8	(11.7)	44	(11.4)	43.8	(11)
Income (FPL %)	214.5	(63.9)	233.7	(75.4)	257.2	(81.1)
Annual Premium	1470	(1264)	-	(-)	-	(-)
Annual Subsidy	3967	(2643)	-	(-)	-	(-)
Medical Spending	-	(-)	-	(-)	4111	(12900)
Choose Bronze (0/1)	0.242	(0.428)	-	(-)	-	(-)
Choose Silver (0/1)	0.681	(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 = 1382		Base prem. quantity-weighted (Covered CA) N = 1382		Avg. claims quantity-weighted (RRF) N = 1026	
By insurer:						
Anthem (76 region-years)	0.059	(0.106)	3062	(638)	3814	(750)
Blue Shield (76 region-years)	0.06	(0.098)	3218	(625)	4140	(1846)
Health Net (33 region-years)	0.048	(0.09)	2614	(306)	3260	(1240)
Kaiser (69 region-years)	0.073	(0.094)	3245	(649)	4212	(2008)
Other 9 insurers	0.026	(0.054)	2605	(603)	2315	(1755)
By metal tier:						
Bronze	0.068	(0.071)	2468	(364)	2197	(902)
Silver	0.138	(0.132)	3125	(538)	3921	(1201)
Gold	0.009	(0.017)	3679	(689)	4847	(1543)
Platinum	0.007	(0.007)	4192	(759)	9063	(3526)

Note: The table summarizes data sources. In the Enrolled panel, each observation is an individual in the Covered California enrollment 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 100-400% FPL. The panels of Market shares and Base premiums report summary statistics from the Covered California enrollment sample. The Average claims panel summarized the 2016-2019 rate review filings matched to the Covered California sample. Standard deviations in parentheses.

amount collected by the issuers increases in age, from \$3000 per-year on average at 26 to over \$8000 for buyers older 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 \$1200 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 2b. Average

Figure 2: *Premiums by Age and Income*

(a) *Average Premium by Age* ***(b)*** *Average Premium by FPL*

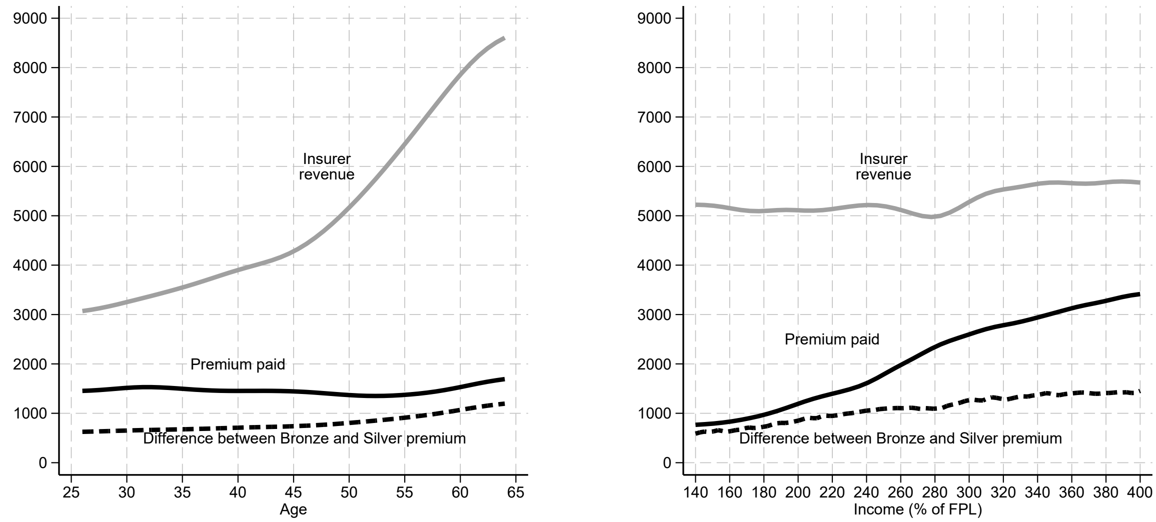

Note: The figure illustrates the relationship between average revenue collected by the insurer (gray line), average subsidized premium paid by the individual (black line), and average difference between Bronze and Silver premiums for the individual (dashed line), as a function of age (left panel) and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

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 3 summarizes market shares at the plan level (insurer-year-region-metal-network; N=1382), 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 4.8-7.3%, 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, 13.8% of enrollees in a region-year pair, about twice as large as the average share of Bronze plans. A Gold or Platinum plan covers, on average, less than 1% of the market.

3.2 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.[9] As in Bundorf et al. (2012); Saltzman (2021), while I observe enrollment 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 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 have to 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 enrollment files using HIOS-14 (a plan-insurer identifier), enrollment year, and metal tier information. The resulting sample of plans for which I observe a measure of realized average claims consists of 1,026 unique insurer-region-year-tier-network combinations, which covers 74% of the 1382 plans I observe in the enrollment data and use in my analysis.[10] In terms of enrollment, the sample of plans for which I observe RRF information covers 76% of the 3.7 million individuals included in my enrollment sample.

The bottom-right of Table 3 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 enrollment pools, regional heterogeneity, and differences in firms' cost functions. Costs vary widely across metal tiers. A Bronze plan records, on average, claims amounting to $2197 per-enrollee, per-year (with standard deviation $902). This compares to Silver plans, with average claims for $3921 ($1201) per-enrollee-year, and Gold plans, with average claims for $4847 ($1543). Platinum plans register much higher claims, with an average of $9063 per-enrollee per-year.

3.3 Survey Data

3.3.1 American Community Survey

I construct measures of potential buyers by age, income, rating region, and enrollment 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); Tebaldi et al. (2019).

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

[10] Some plans change HIOS-14 code over time, or leave the marketplace. When this is the case, I cannot match enrollment 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. Nevertheless, to the best of my knowledge, the RRF remains the best publicly available data source reporting average claims in the California marketplace, as it provides richer heterogeneity than other, state-level sources (e.g. medical-loss-ratio filings as used in Saltzman, 2021).

For every year between 2013-2016, I use the corresponding 5-year ACS sample to measure potential marketplace enrollees for the following enrollment 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.[11] Finally, I merge enrollees 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 enrollment years, which I then match to the enrollment file.[12]

As shown in Table 3, eligible buyers are, on average, two years younger and higher income (+20% FPL) relative to marketplace enrollees. Figure 3a shows more details of the relationship between age and the share of potential buyers choosing to purchase marketplace coverage, measured after combining enrollment files with the ACS. The monotone relationship between age and enrollment is evident: the average enrollment probability among under-40 individuals is between 0.22-0.25, this then increases with age until 0.38 for individuals aged between 60-64. Relating this pattern to the fact that average premium paid does not increase in age (Figure 2a) 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.

3.3.2 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 US population. I focus on individuals who are privately insured, with age and household income in the same range as the observations in the enrollment data. The resulting sample of 20171 individuals is summarized in Table 3. Average annual medical spending is equal to $4111, with standard deviation $12900. In the next section this data is used to estimate a parame-

[11] 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.

[12] 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 150-155% 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 enrollees in the same year-region-demographic combination, I measure a total marketplace share conditional on these observables equal to 0.25.

Figure 3: *Enrollment, MEPS Expenditure, and Rating Adjustments by Age*

(a) *Enrollment by Age* ***(b)*** *Expenditure and Rating Adjustments by Age*

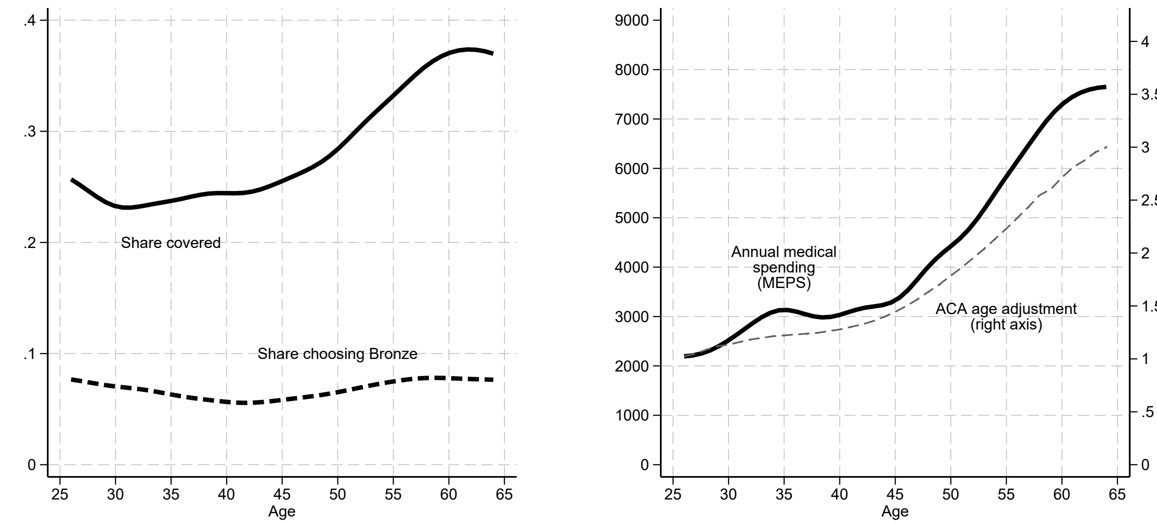

Note: The solid (dashed) black line in the left panel illustrates the relationship between age and the probability of choosing a marketplace (Bronze) plan, measured in the Eligible sample. Local polynomial with Gaussian kernel; bandwidth=2. The solid black line in the right panel illustrates the relationship between age and annual medical expenditure in the Medical Expenditure Panel Survey; Gaussian kernel with bandwidth=2. The dashed gray 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.

ter describing the relationship between age and total medical spending conditional on being insured, controlling for differences across years and MEPS geographic areas.

Figure 3b 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 MEPS the ratio of medical spending between the two age groups is higher than 3.5. Average medical spending is slightly higher than \$2000 at 26, approximately \$4000 at 47 and higher than \$7500 after 60.

4 Empirical Model

4.1 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 $\mathbf{b}_{mt} = \{b_{1mt}, ..., b_{Jmt}\}$, 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

$$q_{jmt}(\mathbf{z}, \theta) = \frac{e^{-\alpha_t(\mathbf{z}_i)P_j(\mathbf{b}_{mt}, \mathbf{z}_i) + \delta_{jmt}(\mathbf{z}, \theta)}}{1 + \sum_{k=1}^{J} e^{-\alpha_t(\mathbf{z}_i)P_k(\mathbf{b}_{mt}, \mathbf{z}_i) + \delta_{kmt}(\mathbf{z}, \theta)}}. \tag{3}$$

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

$$Q_{jmt} = \int q_{jmt}(\mathbf{z}, \theta) dG_{mt}(\mathbf{z}, \theta). \tag{4}$$

The difference between the demand model in (4) and standard discrete choice models with heterogeneous consumers (e.g. Nevo, 2001) lies in how rating regulations and subsidies determine enrollment responses to insurers' pricing decisions.

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

$$\begin{aligned}\frac{\partial Q_{jmt}}{\partial b_{kmt}} &= \int \frac{\partial q_{jmt}(\mathbf{z},\theta)}{\partial b_{kmt}} dG_{mt}(\mathbf{z},\theta) \\ &= \sum_{\ell=1}^{J} \int \frac{\partial P_\ell(\mathbf{b}_{mt},\mathbf{z})}{\partial b_{kmt}} \left(\alpha_t(\mathbf{z}) q_{jmt}(\mathbf{z},\theta) q_{\ell mt}(\mathbf{z},\theta)\right) dG_{mt}(\mathbf{z},\theta). \end{aligned} \tag{5}$$

Equation (5) highlights how changes in base premiums do not affect enrollment 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 $\ell \neq k$, $\frac{\partial P_\ell(\mathbf{b}_{mt},\mathbf{z})}{\partial b_{kmt}} = 0$.

4.2 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

$$\kappa_{jmt}(\mathbf{z}_i,\theta_i) = AV_j^S L_{jmt}(\mathbf{z}_i,\theta_i), \quad \text{where} \quad L_{jmt}(\mathbf{z}_i,\theta_i) = e^{\phi_{jmt}+\eta(\mathbf{z}_i,\theta_i)}. \tag{6}$$

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 (6) does not allow expected medical spending to vary with coverage generosity, ruling out "moral hazard" (c.f. Einav and Finkelstein, 2018). In Section 7 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

$$AC_{jmt} = \frac{1}{Q_{jmt}} \int \kappa_{jmt}(\mathbf{z},\theta) q_{jmt}(\mathbf{z},\theta) dG_{mt}(\mathbf{z},\theta), \tag{7}$$

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 $\mathbb{E}\left[\ln(\nu) | G(\mathbf{z},\theta), \mathbf{x}, \boldsymbol{\xi}, \mathbf{b}\right] = 0$.

4.3 Identification

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

$$\eta(\mathbf{z}, \theta) = \eta^{\text{Age}} z^{\text{Age}} + \eta^{\text{WTP}} \frac{\beta_t(\mathbf{z}, \theta)}{\alpha_t(\mathbf{z})}, \quad \text{and} \quad \phi_{jmt} = \phi_t^1 + \phi_m^2 + \boldsymbol{\phi}^3 \text{Insurer}_{jmt}. \tag{8}$$

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.

4.3.2 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 (4) and (7). 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,

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

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):

$$b_{jmt} = \lambda^{35} \int \mathbf{1}\left[z^{\text{Age}} \leq 35\right] dG_{mt}(\mathbf{z}) + \lambda^{\text{Tier}} + \lambda^{\text{Year}} + \lambda^{\text{Insurer}} + \xi_{jmt}. \tag{9}$$

Regression results and F-statistics are reported in Appendix Table A1, the variation in the instrument and the corresponding variation in b_{jmt} are illustrated in Appendix Figure A1. The first stage OLS estimate of the effect of age-composition of potential buyers on base premium is $\widehat{\lambda}^{35} = -5208$, 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. (2019) 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 enroll 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 Appendix Figure A1, 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.

4.3.3 Cost Identification

To identify cost parameters the structure of the data in my application is similar to the one in Bundorf et al. (2012): I observe (and estimate) demand at the individual level, while realized costs are measured at the plan level. To capture selection, my model allows costs to vary within plan across individuals who differ in age and unobservable willingness-to-pay for coverage $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$.

The MEPS data allows me to calibrate the parameter η^{Age}, which governs the age evolution of average annual medical spending when insured. 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 MEPS macro area. The parameter η^{Age} is very robust across specifications and estimated precisely; see Appendix Table A5.

Identification of cost heterogeneity across buyers with different preferences relies instead on the correlation between plan-average medical spending (total claims adjusted for actuarial value, AC_{jmt}/AV_j^S) and composition of enrollment in terms of $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$. Assuming that $\mathbb{E}\left[\ln(\nu)|G(\mathbf{z},\theta),\mathbf{x},\boldsymbol{\xi},\mathbf{b}\right]=0$, variation in participating plans, and variation in demographics of potential buyers across region-years, lead to variation in the composition of buyers that can be used to identify η^{WTP}. Equation (8) restricts the way in which insurer, year, and region affect medical spending. Given these restrictions, after controlling for insurer, year, and regional effects, 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 identifies η^{WTP}. Intuitively, if claims are higher for plans covering a larger share of individuals with high $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$, $\eta^{\text{WTP}}>0$, and vice versa.

4.4 Estimation Results

Estimation follows the steps detailed in Appendix B.

4.4.1 Demand Estimates

The full set of demand parameters is reported in Appendix Tables A2 and A3. Appendix Table A4 shows the impact of the control function on demand estimates. Omitting $\widehat{\xi}_{jmt}$ would lead to estimates of premium coefficients between one and two percent lower, and to estimates of willingness-to-pay between four and ten percent lower.

Table 4 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 \$263 among those aged between 26-31, to \$343 between 38-43, \$526 between 50-55, reaching the average value of \$892 among those aged between 62-64. This average increase is accompanied by a larger variance: the standard deviation at 26-31 (32-37) is \$210 (\$232), while at 56-61 (62-64) it is more than twice as large, equal to \$516 (\$616).

Increasing all annual premiums by \$120 (third row of Table 4) is equivalent to

Table 4: *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 for 10% AV increase	262.9	305.3	343.4	399.7	526.2	722.6	892.2
	(10.1)	(15.9)	(15.9)	(14.1)	(13.4)	(15.6)	(20.3)
St. Dev. of WTP for 10% AV increase	209.7	232.3	260.3	296.1	387.8	515.7	616.1
	(5.9)	(8.4)	(7.6)	(6.9)	(7.4)	(8.8)	(11.3)
% Change in Enrollment if +\$120/year in all Premium	−6.916	−6.527	−6.078	−5.79	−4.671	−3.69	−3.104
	(0.191)	(0.201)	(0.188)	(0.158)	(0.114)	(0.087)	(0.078)
% Change in Silver Enrollment if +1% in all Silver Premiums	−2.048	−2.047	−1.8	−1.942	−1.774	−1.546	−1.364
	(0.077)	(0.073)	(0.067)	(0.052)	(0.051)	(0.031)	(0.025)
Control Function: Year-Specific Cubic Polynomial of First-Stage Residuals	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	2588265	2265465	2003948	1944898	2013681	2039277	889550

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 Appendix Tables A2 and A3. 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.

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%, compared to 6.5% among those aged between 32-37, and 6% among those aged between 38-43. The extensive margin response to a change in all premiums is much smaller for older buyers. Conditional on age being between 56-61, if all premiums increase by \$120 enrollment drops by 3.7%. For the oldest age bin, 62-64, I estimate that average extensive margin semi-elasticity is equal to 3.1%.

Appendix Figure A2 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.[13]

The fourth row of Table 4 shows the estimates of the elasticity of Silver enrollment to Silver premiums. This is calculated as the percent change in enrollment in Silver plans

[13] Using discontinuities in subsidies in the pre-ACA Massachusetts marketplace, Finkelstein et al. (2019) find enrollment 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 enrollment. In Tebaldi et al. (2019) we adopt a nonparametric approach and estimate that, if all 2014 monthly premiums increased by \$10, the probability of enrollment in Covered California would have been 0.018-0.067 lower.

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 1.8-2%, while for older individuals this is between 1.4-1.8%.

4.4.2 Cost Estimates

The full set of cost estimates is reported in Appendix Table A6. Table 5 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 MEPS is equal to 0.038 (Appendix Table A5). 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 5 shows that the point estimate of η^{WTP} is equal to 0.07, 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 7% higher expected medical spending. Given the range of $\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$ shown in Appendix Figure A2a and Table 4, even conditioning on age, income, and year, willigness-to-pay for actuarial value can vary by more than \$700, corresponding to 50% 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 5 summarizes the 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 (8)).

For Bronze plans, expected average claims are equal to \$1148 per-person, per-year when the enrollee is aged between 26-31, \$1507 when between 32-37, almost \$2000 when between 38-43, and progressively increasing to more than \$7000 for the oldest group, aged between 62-64. Silver plans have higher average claims, reflecting both higher actuarial value ($AV_j^S = 70\%$, instead of 60%) but also a different risk selection: enrollees 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 enrolling someone aged between 26-31 are \$1435, 25% higher than the estimate for Bronze plans, and 7.2% higher than the difference that would be explained

Table 5: 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=20171	2014-17 MEPS	Y	Y	N
WTP for 10% AV increase ($100/year)	η^{WTP}	0.0699 (0.0152)	NLLSQ, N=1026	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 Enrollees	1148 (217)	1507 (266)	1969 (320)	2613 (387)	3744 (441)	5523 (436)	7294 (416)
Silver Enrollees	1435 (223)	1922 (241)	2504 (326)	3355 (371)	4919 (345)	7491 (247)	10274 (329)

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 Appendix Table A6. The bottom panel shows the estimated average cost across Bronze and Silver enrollees, 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).

by the increased actuarial value, holding risk selection fixed. This would be $1339, computed as $\$1148 \times \frac{0.7}{0.6}$.

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 enrollees who are 56 or older, average claims for those selecting a Silver plan are between $7500-$10300, 35-40% higher than the claims for those selecting a Bronze plan.

The relevance of heterogeneity and adverse selection in this application is highlighted in Figure 4: 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 $1000 corresponds to a cost increase from $1000 to slightly more than $2000. When considering individuals aged 35-64, the same difference in preferences corresponds to a cost increase from $2000 to $6000. Even conditioning on a specific value of cost, there is significant heterogeneity in preferences, and vice versa. The joint distribution summarized in Figure 4 is the key primitive one needs to study market design in a health insurance marketplace.

Figure 4: *Empirical Relationship Between Preferences and Expected Cost*

(a) *Age 26-35* ***(b)*** *Age 36-64*

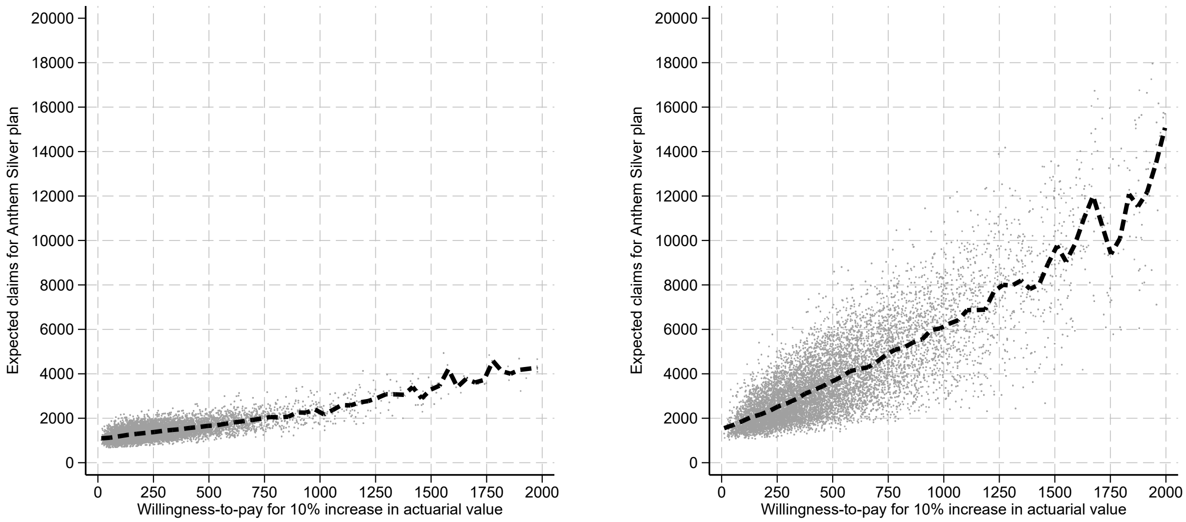

Note: The figure illustrates the joint distribution of willingness-to-pay for coverage and expected cost obtained after combining demand and cost estimates. The graph is generated by randomly drawing 10,000 individuals from $G(\mathbf{z}, \theta)$. For each draw, I compute willingness-to-pay for a 10% increase in actuarial value ($\frac{\beta_t(\mathbf{z},\theta)}{\alpha_t(\mathbf{z})}$), and expected cost if the individual enrolls 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 $z^{\text{Age}} > 35$.

5 Equilibrium and Market Conduct

To consider counterfactual policy design, it is necessary to model expected profits accounting for ACA regulations, and to make assumptions about insurers' conduct.

5.1 Rating Regulations, Risk Adjustment, and Expected 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 (3). 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 Appendix C. For every plan $j \in \mathcal{J}(f)$, the risk adjustment transfer is

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

In words, the per-enrollee 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 enrollees 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.

5.2 Insurer Conduct

To analyze the equilibrium effect of different designs I consider two alternative models of insurer conduct: 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. This has been adopted recently by Dickstein et al. (2021).

The choice between the two alternative models implies relevant trade-offs. Multi-product Nash pricing captures market power when plans differ horizontally and vertically, at the cost of increased computational complexity.[14] Perfect competition is a

[14] 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

relevant baseline case, as originally considered by Einav et al. (2010a) for the empirical analysis of equilibrium in a selection market. Computation of equilibrium is simpler and faster, and existence is guaranteed by the main theorem in Azevedo and Gottlieb (2017). In the context of my application, using both models shows the robustness of the effect of alternative subsidy designs to different assumptions on insurer behavior. The mechanism by which more generous subsidies to younger individuals lower market average cost is common to both. Differences in markups are only relevant when considering Nash pricing.

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,

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

In this expression, the superscript AG indicates that the demand function $q_{jmt}(\mathbf{z}, \theta)$ is modified to let an infinitesimal fraction of "behavioral" buyers choose a given plan independently from changes in premiums or other characteristics.[15]

Before presenting results for both conduct models, it is possible to use the estimates of demand and cost to verify whether the data provides supporting evidence in favor of one of the two alternatives. To do this, Figure 5 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.[16]

In Figure 5a, the horizontal axis corresponds to the per-enrollee marginal revenue for every jmt combination in the data. Nash pricing predicts that this would be equal to per-enrollee 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

and Lucarelli, 2021).

[15] 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).

[16] 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), yet rarely used.

Figure 5: *Multi-Product Nash Pricing vs. Perfect Competition*

(a) *Marginal Revenue vs. Marginal Cost* ***(b)*** *Average Revenue vs. Average Cost*

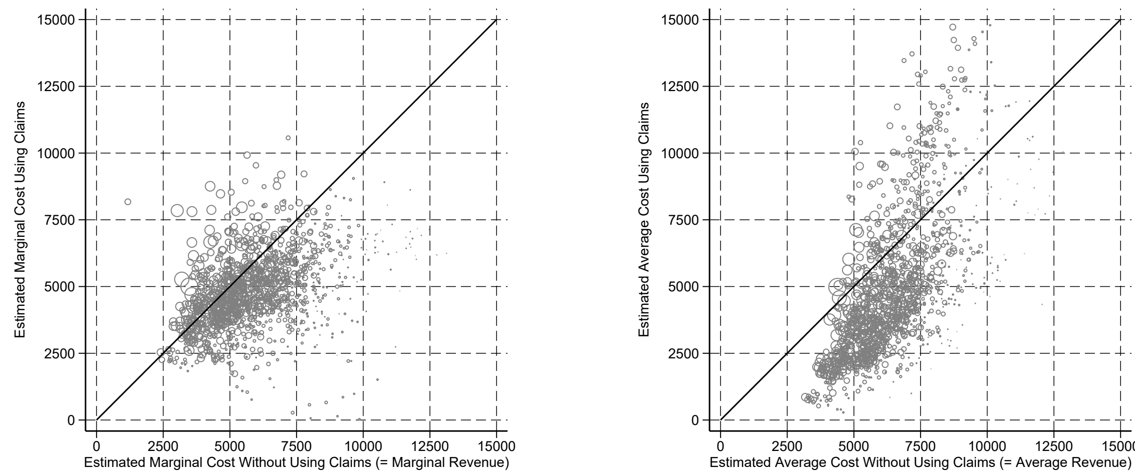

Note: The left panel shows the comparison between per-enrollee risk-adjusted marginal costs estimated assuming multi-product Nash-in-Prices without using claims (these are equal to marginal revenue), and per-enrollee risk-adjusted marginal costs estimated using observed claims. The right panel shows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

demand and cost. The resulting scatter plot is concentrated around the 45-degree line. The enrollment-weighted average difference between per-enrollee marginal revenue and per-enrollee risk-adjusted marginal cost is \$293.46 (95%-C.I.: [217.09, 369.84]). The enrollment-weighted average ratio $\frac{\partial \Pi_f}{\partial b_{jmt}}/R_{jmt}$ is 0.035 (95%-C.I.: [0.023, 0.048]).

For comparison, Figure 5b 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 5a 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 enrollment-weighted average difference between R_{jmt}/Q_{jmt} and $(TC_{jmt} + RA_{jmt})/Q_{jmt}$ is \$1331.12 (95%-C.I.: [1236.73, 1425.50]). The enrollment-weighted average ratio Π_{jmt}/R_{jmt} is 0.245 (95%-C.I.: [0.229, 0.260]), corresponding to a departure from the model assumption 14 times as large as under Nash pricing.

One additional piece of evidence in support of modeling insurers as not perfectly competitive is provided by the estimated medical-loss ratio (MLR). Despite not imposing a constraint in estimation, I calculate average MLR at the observed base premiums to be approximately equal to 0.85. This is above the minimum value of 0.8 mandated by the ACA, while still lower than the perfect competition value of one.

6 Subsidy Design and Equilibrium Outcomes

6.1 Price-Linked Subsidies vs. Vouchers

The first comparison is between 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 6 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 across the two subsidy designs (equilibrium premiums depend only on enrollees 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 13.2%, from \$3646 to \$3164; in larger regions, with more than four participating insurers, the drop is smaller, from \$2769 to \$2623 (-5.3%). 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 enrollment pools, the Nash-pricing equilibrium under equiv-

Table 6: *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 enrolled	0.315	0.324	0.266	0.273	0.274	0.274	0.268	0.268
2nd cheapest Silver b_j	3646	3164	2769	2623	2716	2715	2177	2175
Share in Bronze plans	0.171	0.166	0.154	0.142	0.181	0.181	0.153	0.153
Medical-loss ratio	0.862	0.831	0.842	0.820	0.998	0.998	0.998	0.998
ΔCS_i relative to ACA	-	28	-	24	-	0	-	1
Average subsidy	5070	4694	3388	3345	3428	3419	2339	2313

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 enrolled 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 enrollment. ΔCS_i indicates the average, per-person annual consumer surplus, which is reported in differences from the equilibrium under ACA price-linked subsidies.

alent vouchers implies slightly higher marketplace enrollment, increasing from 0.315 (0.266) to 0.324 (0.273) in small (large) regions. The corresponding increase in annual per-person consumer surplus relative to the ACA design is between \$24-\$28. In regions with less than four insurers average subsidies drop from \$5070 to \$4694; in larger regions from \$3388 to \$3345. Insurer profitability is also higher, as medical-loss ratio drops from 0.862 (0.842) to 0.831 (0.820) 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.

6.2 Equilibrium Effects of More Subsidies for the Young Invincibles

The second counterfactual subsidy design amounts to providing additional enrollment incentives to the so-called "young invincibles"; in what follows this group consists of individuals aged between 26-35 (see e.g. Levine and Mulligan, 2017). As argued in Sections 1 and 2, 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, more desirable equilibrium. Insurers lower base premiums, due to the average cost reduction and—under Nash pricing—increase in elasticity. Lower premiums lead to higher enrollment 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, 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 MAA (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 behavior.

Panel (a) of Table 7 summarizes how marketplace outcomes respond to changing the ACA price-linked design by lowering the MAA for young invincibles by 30%. In equilibrium, the effect is to increase enrollment 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. Without accounting for endogenous adjustments, premiums for older buyers are not affected by the different design. Therefore, off-equilibrium only the young invincibles are better off. In the new equilibrium, older buyers also benefit from the alternative subsidy design, as they face lower premiums and enroll more.

Using vouchers, the way in which alternative subsidy designs impact equilibrium outcomes follows more closely the mechanism discussed in Section 2. This is illustrated in panel (b) of Table 7, 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 enrollment share for this group drops by 0.01 as they face higher premiums). In equilibrium, however, the reduction in base premiums following the larger enrollment share of under-35 individuals implies that all buyers are better off.

Considering Nash pricing, under-35 enrollment increases from 0.25 to 0.348, and over-35 enrollment from 0.304 to 0.312; despite receiving smaller vouchers, subsidized

Table 7: Counterfactual Subsidy Design

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 enrolled:						
26-35	0.243	0.308	0.313	0.242	0.299	0.300
36-64	0.295	0.295	0.302	0.284	0.284	0.285
Premium paid:						
26-35	1655	1322	1296	1768	1432	1435
36-64	1704	1704	1643	1902	1902	1911
Average cost	4534	4301	4322	4423	4216	4225
Per-person CS	731	774	792	708	745	748
Average subsidy	3828	3807	3813	2594	2620	2511

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 enrolled:						
26-35	0.250	0.331	0.348	0.242	0.327	0.350
36-64	0.304	0.294	0.312	0.284	0.275	0.298
Premium paid:						
26-35	1630	1154	1086	1767	1232	1131
36-64	1675	1751	1633	1901	1989	1805
Average cost	4410	4144	4139	4422	4129	4112
Per-person CS	756	794	840	708	747	808
Average subsidy	3698	3656	3642	2573	2577	2573

Note: Simulated market outcomes under alternative subsidy designs; for details on equilibrium computation, see note to Table 6. 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 MAA Equilibrium, and ACA-voucher Equilibrium, respectively. Enrollment shares and annual per-person CS are computed as averages across region-years, weighted by number of eligible buyers. The other outcomes are enrollment-weighted averages.

premiums of over-35 buyers are \$42 lower. The younger composition of enrollees translates in average costs that, in equilibrium, are 6% lower than under the ACA-equivalent voucher. Per-person consumer surplus increases by \$84 per-year, while average per-enrollee subsidies are \$53 lower. 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 6: *ΔConsumer Surplus by Age: +$600 Under-35 Voucher, -$100 Over-35 Voucher*

(a) *Off-Equilibrium* ***(b)*** *Equilibrium*

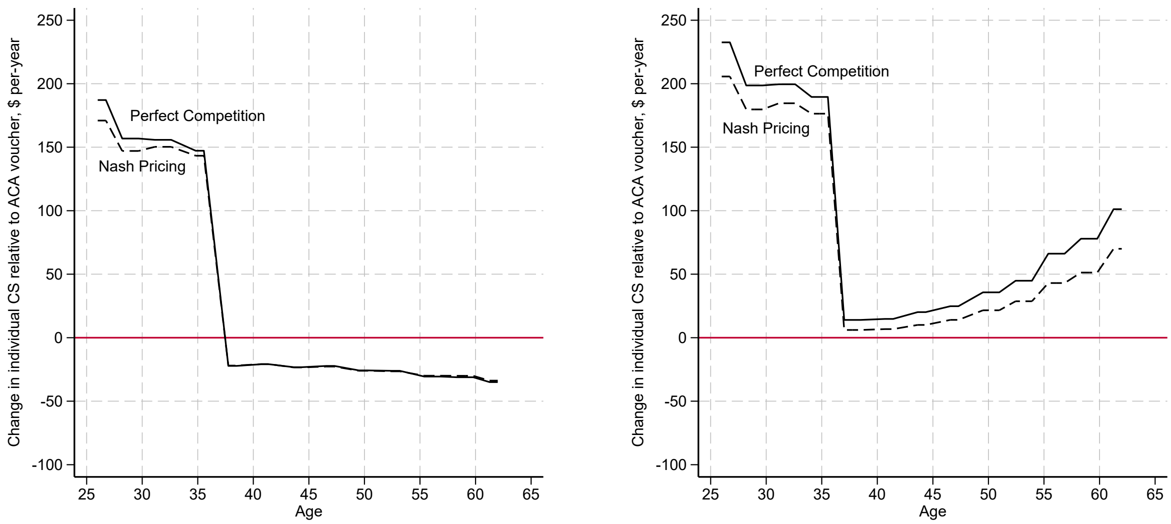

Note: Average annual change in per-person consumer surplus when replacing ACA-equivalent vouchers with vouchers that are $600 higher for the under-35, and $100 lower for the under-35. The left panel holds base premiums fixed to the equilibrium under ACA-equivalent vouchers, the right panel corresponds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

Figure 6 illustrates the relationship between age and changes in annual, per-person consumer surplus resulting from changing vouchers as in panel (b) of Table 7. 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 6b, at the new equilibrium the change in consumer surplus of over-35 switches sign: this group is now better relative to the ACA-voucher equilibrium, by an annual amount varying between $10-100.[17]

The alternatives considered so far set specific values for counterfactual subsidy designs. To explore the relationship between changes in young and old vouchers more generally, I simulate 3724 equilibria—one for every design, region, and year—over a grid of adjustments. Letting Y_{mt}^{ACA} be a market outcome under the ACA voucher, Y_{mt}^{Δ} the same outcome under alternative vouchers holding base premiums fixed, and $Y_{mt}^{\Delta\star}$

[17] 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 a Pareto improvement 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.

Table 8: *Effect of Varying Voucher Generosity by Age Group, Equations* (13) *and* (14)

	Multi-product Nash				Perfect Competition			
	Off-equilibrium		Equilibrium		Off-equilibrium		Equilibrium	
	+$100/year in 26-35 voucher	+$100/year in 36-64 voucher	+$100/year in 26-35 voucher	+$100/year in 36-64 voucher	+$100/year in 26-35 voucher	+$100/year in 36-64 voucher	+$100/year in 26-35 voucher	+$100/year in 36-64 voucher
Share enrolled:								
26-35	+0.014	0.000	+0.016	+0.001	+0.014	0.000	+0.018	+0.001
36-64	0.000	+0.009	+0.003	+0.011	0.000	+0.009	+0.004	+0.012
Premium paid:								
26-35	−77.5	0.0	−94.1	−21	−84.7	0.0	−106.3	−26.5
36-64	0.0	−81.2	−23.1	−86.1	0.0	−90.3	−32.9	−106.9
Average cost	−43.8	−4.2	−47.7	−9.8	−48.5	−5.2	−54.4	−16.9
Per-person CS	+9.3	+16.7	+17.1	+20.5	+9.7	+16.7	+20	+22.2
Average subsidy	+8.7	+59.8	+4.1	+52	+16.3	+61.2	+11.7	+53.4

Note: Coefficient estimates from equations (13) and (14), interpolating linearly each equilibrium outcome over a grid of voucher adjustments. Δ26-35 voucher varies in $\{0, 400, 500, 600, 700, 800, 900\}$, and Δ36-64 voucher varies in $\{0, -75, -100, -125, -150, -200, -250\}$. Equilibrium is computed for each subsidy design, conduct assumption, region, and year, as described in the note to Table 6.

the outcome at the new equilibrium, I estimate

$$Y_{mt}^{\Delta} - Y_{mt}^{\text{ACA}} = \hat{\rho}^{26\text{-}35}\Delta\text{26-35 voucher} + \hat{\rho}^{36\text{-}64}\Delta\text{36-64 voucher}, \tag{13}$$

$$Y_{mt}^{\Delta\star} - Y_{mt}^{\text{ACA}} = \hat{\hat{\rho}}^{26\text{-}35}\Delta\text{26-35 voucher} + \hat{\hat{\rho}}^{36\text{-}64}\Delta\text{36-64 voucher}. \tag{14}$$

Table 8 reports the results, which describe the linear approximation of the relationship between market outcomes and changes in annual vouchers across age groups.

Accounting for equilibrium responses, each $1 reduction in 36-64 vouchers can be compensated by a $4 increase in 26-35 vouchers. This guarantees that the lower subsidy for those older than 35 is more than compensated by the endogenous reduction in premiums triggered by the larger subsidy for the younger group. Average cost and consumer surplus are, respectively, lower and higher. Importantly, a $1 reduction in over-35 vouchers lead to $52 lower average subsidy, while the compensating $4 increase in 26-35 voucher only increases average subsidy by less than $17. The net effect is to lower average subsidy outlays, while ensuring that all individuals face lower premiums.

7 Robustness to Moral Hazard

The cost estimates in Table 5 and the simulation results in Section 6 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 (Einav, Finkelstein, and Levin, 2010b). In the model of Section 4, allowing spending to increase with actuarial value impacts the estimates of η^{WTP} and other cost parameters. Therefore, although the results above rely primarily on the fact that young uninsured individuals are generally healthy, the quantifications in Section 6 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. (2019), 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 enrolled 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

$$L^{\text{MH}}_{jmt}(\mathbf{z}_i, \theta_i) = (1 + \zeta \times \chi_{ij})\, L_{jmt}(\mathbf{z}_i, \theta_i), \tag{15}$$

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 (6). If $\zeta = 0$, the model is identical to the one in Sections 4 and 6. 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.

Appendix Figure A3 shows the estimates of η^{WTP} varying ζ. From the baseline level of $\eta^{\text{WTP}} = 0.07$ obtained when $\zeta = 0$, setting $\zeta = 1$ reduces this estimate by 3.7% ($\eta^{\text{WTP}} = 0.067$). The estimates of η^{WTP} remain above 0.06 as long as the level of moral hazard is lower than four 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 $\zeta = 13$, which seems quite unrealistic.

Table 9 explores the robustness of the results in Table 7 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

Table 9: *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 enrollment	36-64 enrollment	26-35 premium	36-64 premium	Average CS	Average subsidy	26-35 enrollment	36-64 enrollment	26-35 premium	36-64 premium	Average CS	Average subsidy
$\zeta = 0$	0.098	0.008	−544	−42	84	−56	0.108	0.013	−636	−96	100	0
$\zeta = 1$	0.095	0.006	−524	−30	78	−31	0.107	0.013	−628	−90	98	−8
$\zeta = 2$	0.095	0.004	−481	−3	73	−35	0.106	0.012	−621	−88	98	−14
$\zeta = 4$	0.088	0.004	−487	−7	69	−52	0.106	0.012	−617	−83	97	−32

Note: The table shows how the results of panel (b) in Table 7 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 (6) with L_{jmt}^{MH} from Equation (15), 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 7.

smaller when assuming larger degrees of moral hazard. However, even when setting $\zeta = 4$ the counterfactual vouchers make all buyers better off while reducing average subsidies.

8 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 dominated by a design in which premiums vary across types. Adjusting subsidies to observables that predict preferences and cost leads to equilibria in which all consumers are better off, coverage is 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. Following the significant differences in preferences and cost across age groups, equilibrium simulations suggests that the proposed adjustments would lead to improvements in equilibrium outcomes.

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 behavioral aspects, and for the role of healthcare providers. Access to richer data, including measures of health and healthcare utilization at the individual level, would facilitate the calculation of optimal policy parameters by researchers and policymakers.

References

AKERLOF, G. A. (1970): "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *The Quarterly Journal of Economics*, 84, 488–500. 1

AZEVEDO, E. M. AND D. GOTTLIEB (2017): "Perfect competition in markets with adverse selection," *Econometrica*, 85, 67–105. 10, 27, 28, 31, 34

BERRY, S. T. AND J. WALDFOGEL (1999): "Public radio in the United States: does it correct market failure or cannibalize commercial stations?" *Journal of Public Economics*, 71, 189–211. 3, 20

BRESNAHAN, T. F. (1987): "Competition and collusion in the American automobile industry: The 1955 price war," *The Journal of Industrial Economics*, 457–482. 10, 28

BROWN, J., M. DUGGAN, I. KUZIEMKO, AND W. WOOLSTON (2014): "How does risk selection respond to risk adjustment? New evidence from the Medicare Advantage Program," *American Economic Review*, 104, 3335–64. 4, 8

BUNDORF, M. K., J. LEVIN, AND N. MAHONEY (2012): "Pricing and welfare in health plan choice," *American Economic Review*, 102, 3214–48. 3, 10, 15, 21, 27

BUNDORF, M. K. AND M. V. PAULY (2006): "Is health insurance affordable for the uninsured?" *Journal of Health Economics*, 25, 650 – 673. 1

CHAN, D. AND J. GRUBER (2010): "How Sensitive Are Low Income Families to Health Plan Prices?" *American Economic Review*, 100, 292–96. 3

COLOMBO, F. AND N. TAPAY (2004): "Private health insurance in OECD countries," . 1

CURTO, V., L. EINAV, J. LEVIN, AND J. BHATTACHARYA (2021): "Can health insurance competition work? evidence from medicare advantage," *Journal of Political Economy*, 129, 570–606. 4, 10, 27

DECAROLIS, F. (2015): "Medicare part d: Are insurers gaming the low income subsidy design?" *American Economic Review*, 105, 1547–80. 4, 8

DECAROLIS, F., M. POLYAKOVA, AND S. P. RYAN (2020): "Subsidy design in privately provided social insurance: Lessons from Medicare Part D," *Journal of Political Economy*, 128, 1712–1752. 4, 8, 10, 27

DIAMOND, R., M. J. DICKSTEIN, T. MCQUADE, AND P. PERSSON (2018): "Insurance without Commitment: Evidence from the ACA Marketplaces," Tech. rep., National Bureau of Economic Research. 9

DICKSTEIN, M. J., M. DUGGAN, J. ORSINI, AND P. TEBALDI (2015): "The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the Affordable Care Act," *American Economic Review*, 105, 120–25. 5

DICKSTEIN, M. J., K. HO, AND N. D. MARK (2021): "Market Segmentation and Competition in Health Insurance," Tech. rep., National Bureau of Economic Research. 3, 27

EINAV, L. AND A. FINKELSTEIN (2018): "Moral hazard in health insurance: what we know and how we know it," *Journal of the European Economic Association*, 16, 957–982. 19, 35

Einav, L., A. Finkelstein, and M. R. Cullen (2010a): "Estimating welfare in insurance markets using variation in prices," *The quarterly journal of economics*, 125, 877–921. 1, 28

Einav, L., A. Finkelstein, and J. Levin (2010b): "Beyond testing: Empirical models of insurance markets," *Annu. Rev. Econ.*, 2, 311–336. 36

Ericson, K. M. M. and A. Starc (2015): "Pricing regulation and imperfect competition on the massachusetts health insurance exchange," *Review of Economics and Statistics*, 97, 667–682. 3, 5, 6

Finkelstein, A., N. Hendren, and M. Shepard (2019): "Subsidizing health insurance for low-income adults: Evidence from Massachusetts," *American Economic Review*, 109, 1530–67. 3, 5, 12, 15, 23

Geruso, M. and T. Layton (2020): "Upcoding: evidence from Medicare on squishy risk adjustment," *Journal of Political Economy*, 128, 984–1026. 8

Handel, B., I. Hendel, and M. D. Whinston (2015): "Equilibria in health exchanges: Adverse selection versus reclassification risk," *Econometrica*, 83, 1261–1313. 4, 10

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. 3

Jaffe, S. and M. Shepard (2020): "Price-linked subsidies and imperfect competition in health insurance," *American Economic Journal: Economic Policy*, 12, 279–311. 2, 5, 30, 31

Lavetti, K. J., T. DeLeire, and N. R. Ziebarth (2019): "How do low-income enrollees in the Affordable Care Act Marketplaces respond to cost-sharing?" Tech. rep., National Bureau of Economic Research. 3, 4, 21, 36

Levine, D. and J. Mulligan (2017): "Mere mortals: Overselling the young invincibles," *Journal of Health Politics, Policy and Law*, 42, 387–407. 32

Mahoney, N. (2015): "Bankruptcy as Implicit Health Insurance," *American Economic Review*, 105, 710–46. 1, 5

Mahoney, N. and E. G. Weyl (2017): "Imperfect competition in selection markets," *Review of Economics and Statistics*, 99, 637–651. 4

McFadden, D. (1973): "Conditional logit analysis of qualitative choice behavior," . 18

Miller, K. S., A. Petrin, R. Town, and M. Chernew (2022): "Optimal Managed Competition Subsidies," Working Paper 25616, National Bureau of Economic Research. 4

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. 8

Nevo, A. (2001): "Measuring market power in the ready-to-eat cereal industry," *Econometrica*, 69, 307–342. 10, 18

Orsini, J. and P. Tebaldi (2017): "Regulated age-based pricing in subsidized health insurance: Evidence from the Affordable Care Act," *Becker Friedman Institute for Research in Economics Working Paper*. 3, 6, 20

Panhans, M. (2019): "Adverse selection in ACA exchange markets: evidence from Colorado," *American Economic Journal: Applied Economics*, 11, 1–36. 3, 5

Pauly, M. V. (1970): "Efficiency in the provision of consumption subsidies," *Kyklos*, 23, 33–57. 1

Polyakova, M. and S. P. Ryan (2019): "Subsidy targeting with market power," Tech. rep., National Bureau of Economic Research. 3, 5, 20

Pope, G. C., H. Bachofer, A. Pearlman, J. Kautter, E. Hunter, D. Miller, and P. Keenan (2014): "Risk transfer formula for individual and small group markets under the Affordable Care Act," *Medicare & Medicaid Research Review*, 4. 4, 8, 27, 36, iv

Rothschild, M. and J. E. Stiglitz (1976): "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 90, 630–49. 1

Ruggles, S., K. Genadek, R. Goeken, J. Grover, and M. Sobek (2015): "Integrated Public Use Microdata Series: Version 6.0 [dataset]," *Minneapolis: University of Minnesota*. 15

Saltzman, E. (2019): "Demand for health insurance: Evidence from the California and Washington ACA exchanges," *Journal of Health Economics*, 63, 197–222. 3, 5, 8, 23

——— (2021): "Managing adverse selection: underinsurance versus underenrollment," *The RAND Journal of Economics*, 52, 359–381. 3, 4, 5, 8, 10, 15, 27, iv

Saltzman, E. and C. Lucarelli (2021): "Firm Learning in a Selection Market," in *10th Annual Conference of the American Society of Health Economists*, ASHECON. 27

Smith, J. C. and C. Medalia (2014): *Health insurance coverage in the United States: 2013*, US Department of Commerce, Economics and Statistics Administration, Bureau of the Census. 5

Starc, A. (2014): "Insurer pricing and consumer welfare: Evidence from medigap," *The RAND Journal of Economics*, 45, 198–220. 8, 10, 27

Summers, L. H. (1989): "Some simple economics of mandated benefits," *The American Economic Review*, 79, 177–183. 1

Tebaldi, P., A. Torgovitsky, and H. Yang (2019): "Nonparametric estimates of demand in the california health insurance exchange," Tech. rep., National Bureau of Economic Research. 3, 15, 23

Veiga, A. (2020): "Community Rating in Markets for Lemons," *Available at SSRN 3545479*. 4

Wagstaff, A. and E. van Doorslaer (2000): "Chapter 34 Equity in health care finance and delivery," *Handbook of Health Economics*, 1, 1803 – 1862. 1

Waldfogel, J. (2003): "Preference Externalities: An Empirical Study of Who Benefits Whom in Differentiated-Product Markets," *RAND Journal of Economics*, 34, 557–68. 3, 20