Subsidy Targeting with Market Power*

Maria Polyakova[†] Stephen P. Ryan[‡]

September 5, 2019

Abstract

Public welfare programs commonly link their benefits to observable characteristics of potential recipients. We show that when the provision of the subsidized public benefit is contracted out to imperfectly competitive firms, targeting subsidies using observable characteristics of consumers can distort both efficiency and equity of market allocations. Using a model of supply and demand, we illustrate the underlying mechanisms and quantify the distortions empirically in the context of means-tested subsidies in privately provided health insurance markets under the Affordable Care Act. In this empirical application, market power increases the welfare loss from subsidy targeting, vis-a-vis income-invariant subsidies, by 52%.

JEL classification: I11, I18, L22, D44, H57

Keywords: Tagging; Subsidies; Health Insurance; ACA; Subsidy Incidence

*We gratefully acknowledge support from the Agency for Healthcare Research and Quality (R03-HS024800). We are grateful to David Card, Pietro Tebaldi, Mark Shepard, and Keith Ericson for their discussions of the paper, and seminar participants at Daniel McFadden 80th Birthday Conference, Junior Health Economics Summit, Ohlstadt Workshop, Stanford University, London School of Economics, American Society of Health Economists, Yale University, University of Connecticut, Columbia University, Bates White Life Sciences Symposium, University of Munich, REStud Conference, Rice University, University of Utah, University of Leuven, Georgia State University, UNC Chapel Hill, NBER Health Care Meetings, and NBER Public Economics Meetings for helpful comments. We thank Alan Jaske and Lynn Hua for excellent research assistance.

[†]Stanford University School of Medicine and NBER (maria.polyakova@stanford.edu)

[‡]Washington University Olin Business School, NBER, and CESifo (stephen.p.ryan@wustl.edu).

1 Introduction

Public welfare programs have a long history of linking their benefits to observable characteristics of potential recipients, such as age, income, or employment status. Such *tagging* (Akerlof, 1978) may improve targeting of public dollars to the most needy recipients, but observable characteristics may be imperfect measures of need, or worse, individuals may try to alter their observable characteristics (the so-called *masquerading* effect) or distort their behavior in order to qualify for the benefit. The extensive theoretical and empirical literature studying the costs and benefits of tagging has almost exclusively focused on these demand-side distortions, assuming that benefits are provided by a benevolent government. However, governments have recently and increasingly turned to profit-maximizing firms to provide targeted government benefits. In this paper, we investigate the welfare consequences of strategic firms pricing in the presence of tagging.

Adding market power to the supply side of a public benefit provision in the presence of taxes or subsidies that are tagged to observables has the potential to generate substantial efficiency distortions above and beyond the well-documented masquerading effects. The intuition is simple: tagging introduces heterogeneity in subsidies across consumers and markets, and, all else equal, firms have incentives to raise prices in markets where consumers receive more generous subsidies. In the presence of market power, these incentives are not dissipated by competition. This combination of market power and tagging creates a demographic externality that can generate perverse equilibrium outcomes. For example, if consumer subsidies are computed on the basis of income, the near-poor end up paying more for identical products in markets with many poor consumers.

We explore this issue empirically on the example of the new ACA Health Insurance Marketplace market. Public health insurance has been increasingly provided by private insurers, and this new market launched in 2014 is no exception. As in all publicly funded, but privately provided health insurance markets, there is a key question of how much the government should pay insurers. This question is central for the efficiency of these markets, but is still very poorly understood.

The Marketplaces provide a fruitful empirical laboratory for understanding the effects of subsidy tagging. Public funds play a significant role in this setting - the majority of enrollees receive a subsidy in the form of a tax credit for the payment of their insurance premiums. These tax credits depend on consumers' age and income, thus following a traditional approach of conditioning a public benefit on consumer observables. Such categorical tagging with strategic insurers on the supply side that can perfectly foresee the distribution of tagged observables generates a significant potential for efficiency and allocative distortions. Moreover, in the ACA Marketplace setting, baseline subsidy levels depend on price quotes (or "bids") submitted by insurers. This feature of the market further mutes any disincentive to strategically take advantage of the tagging structure.

In this paper we set out to quantify the potential efficiency and allocative distortions that may be stemming from tagging in the presence of market power and the price-linked subsidization mechanism. We start by formulating and estimating the model of demand for ACA Marketplace plans. We utilize the unique institutional setting of the Marketplaces to implement a novel identification strategy in our demand estimates. We use sharp discontinuities in consumer-facing prices across income-age bins that are generated by premium regulation. In our empirical setting, consumers of the same age face different prices for the same product in the same market if their incomes lie above or below pre-determined income thresholds; further, consumers with the same household income face different prices for the same product if their are one year of age apart. This structure of the data allows us to identify the coefficient on the price parameter in the utility function under a semiparametric demand specification, where we estimate product-specific utility levels. The latter flexibility in demand specification is important in our setting with highly multi-dimensional products.

Using this novel strategy, we estimate reasonable levels of marginal utility of income and intuitive substitution patterns. We then proceed to derive a profit function for insurers on this market, trying to balance the institutional and especially regulatory detail with the computational tractability of the model. We arrive at first-order conditions that allow us to recover marginal costs for each product-market combination. Our estimates from the inversion of the first order condition at the product-market level are highly consistent - both in terms of levels and relative ranking of product - with accounting data at the product level.

With these estimates in hand, we analyze the welfare characteristics of the observed allocation under income tags. We find that (per capita) consumer surplus varies substantially across local geographies. We also find intuitive patterns in the distribution of consumer surplus across income and age groups - surplus decreases with income (as subsidies go down), and it increases with age. As expected, consumers that receive highest premium and cost-sharing subsidies enjoy the highest consumer surplus. We find a total consumer surplus of roughly $29 billion, which exceeds the government spending on subsidies of about $22 billion, suggesting that a dollar of subsidies generates more than a dollar of surplus or roughly breaks even if we take into account the cost of raising the public funds. To calculate the incidence of subsidies, we compare simulated allocations with and without subsidies; these allow us to track whether subsidy funds accrue primarily to consumers or insurers; we can also assess which socio-demographic groups among consumers benefit the most and the least.

In our subsequent counterfactual analyses, we consider the efficiency and allocative implications of alternative subsidization rules that either do not use categorical tagging or alter its structure. We consider several types of counterfactuals. First, to assess the distortions that arise from the combination of subsidy tagging and market power, we simulate an environment, where the insurance benefit is provided by the benevolent social planner (in practice, we impose that insurers are forced to price at marginal cost) and subsidies are administered like in the observed tagged system. Next, we consider how imperfectly competitive supply side would interact with alternative subsidization mechanisms. We consider mechanisms that keep some version of categorical tagging, but changes tags - for example, one policy option that is currently being considered by Congress is to tag subsidies to age rather than income. We further consider mechanisms that completely remove categorical tags and provide flat subsidies instead. These flat subsidies could either be regional - for example, county-specific vouchers - or national. All these mechanisms correspond to policy proposals that are being actively considered in the ACA Marketplace.

We find that under the observed subsidization mechanism about 40% of subsidies accrue to insurers and 60% accrue to consumers. We further estimate that subsidy tagging reduces total surplus relative to a subsidy policy that pays on average the same amount, but does not vary the subsidy across consumers based on their observables. Without tagging total surplus increases by $6B for $4B in government spending, 67% of which accrues to consumers. Without tagging, the insurance program generate a positive return on a dollar of public spending, while with tagging the returns are slightly negative when we consider nominal government spending and do not account for any opportunity cost of public funds in paying the same consumers in other ways. While removing tagging increases the overall size of the pie, it is not a pareto improving policy - the change in the mechanism leads to a re-allocation of surplus across geographies and demographic groups.

In a simulation of a perfectly competitive market, we find that introducing subsidization increases enrollment dramatically from 30 percent to 63 percent , but this gain in insurance coverage comes at a substantial deadweight loss. In this case, the government spends $32 billion in subsidies to generate $23 billion in consumer surplus. Reversely, when we simulate a market with an without market power (removing all subsidies), we find that market power reduces welfare by $2 billion and decreases enrollment from 30 to 11 percent. Overall, we find that the highest consumer surplus is generated in a competitive environment with means-tested subsidies, while the highest welfare is generated in a competitive environment without subsidies. This is consistent with the idea that the marginal consumers attracted into the market by subsidization have a relatively low willingness to pay for insurance coverage.

Our analysis relates to several literatures. First, the paper is closely related to the large theoretical and empirical literatures on cash-based and in-kind subsidization policies in various public programs (Currie and Gahvari 2008 provide a comprehensive overview; Allcott et al. 2015 and Lieber and Lockwood 2017 are among recent empirical applications). We add to the rich conceptual literature on optimal tagging of taxes and subsidies - Akerlof (1978) and subsequent theoretical literature - by suggesting the important role of imperfectly competitive supply side in settings where the government outsources public benefit provision to private firms. Traditionally, the literature on targeted public programs has assumed perfectly competitive markets, where consumers may *masquerade* their eligibility for any type of public benefits, but there are no strategic firms that could exploit the information on the targeting mechanisms. Increasingly, the literature has started documenting empirically - in many diverse contexts - of how firms that interact with consumers who receive public funding may respond strategically (e.g. Cellini and Goldin, 2014 on Pell grants; Rothstein, 2010 on EITC, among others).

Through our empirical application we contribute to a subset of this literature that has focused on health insurance. This strand of literature has investigated the effects of tax subsidies on employer-provided health insurance, for example in Gruber and Washington (2005); in the classic illustration of an adverse selection spiral, Cutler and Reber (1998) discuss the role of subsidy design (by the employer) in employer-sponsored plans. Enthoven (2011) and Frakt (2011) discuss some of the key conceptual points and the policy debate on the funding of publicly-funded, privately-run insurance. Conceptually and methodologically, our paper is closest to Curto et al. (2015), Tebaldi (2017), Decarolis (2015), Decarolis et al. (2016), Ho et al. (2015), Wu (2016) and Jaffe and Shepard (2017) that explore the questions about subsidies, competition, and market design, and strategic insurer behavior in the context of Medicare Advantage, Covered California, Medicare Part D, and Massachusetts Health Insurance Exchange, respectively. These papers focus on the idea that subsidies linked to prices of insurers may distort allocations.

We contribute to this literature by asking whether the tagging of subsidies to consumer observables may generate allocative distortions in the presence of market power. The possibility that subsidy tagging may be strategically exploited by firms is related to a literature outside of health insurance that has highlighted how firms may strategically respond to "targeted" public finds that may accrue to their consumers. Our paper further contributes to a rapidly growing literature that studies various aspects of the Affordable Care Act, and especially the launch and performance of the ACA Health Insurance Exchanges.

The paper proceeds as follows. Section 2 gives a brief primer on the institutional setting and describes our data sources. Section 3.2 discusses descriptive patterns in the data that are suggestive of the tagging-related distortion that we hypothesis in this setting. Sections 4.1 and 4.2 lay out the demand and supply models, respectively. Section 4.4 reports model estimation results. Section 5 proceeds to discuss the efficiency properties of observed and counterfactual subsidization mechanisms. Section 6 briefly concludes.

2 Economic Environment and Data

Our empirical application is the market for non-group health insurance contracts in the US that was created by the Affordable Care Act in 2010 and started its operation in 2014. The program allows individual consumers to purchase health insurance plans for themselves and their families. Enrollment is voluntary, although individuals that do not have any health insurance face annual penalties that have been increasing from 2014 onwards. Insurance plans on this market are complex, highly dimensional products. All plan are classified into one "metal" tier: Bronze, Silver, Gold, Platinum, and Catastrophic. These metal tiers reflect the average generosity of plans - the fraction of costs a plan would cover for a standardized population. In addition to metal labels, the plans have varying cost-sharing arrangements such as deductibles, co-insurance, and co-pays, and varying restrictions on provider networks and scope of pharmaceutical coverage.

Insurers in this market are not allowed to price-discriminate based on individual health risks, but they are allowed to set different premiums depending on individual's age, smoking status, and family composition. While several US states have created their own Marketplace programs, most states (37) use an online federal platform www.healthcare.gov to facilitate the purchase of insurance; we focus on these states in our analysis. The 37 federally-facilitated states encompass 2,566 counties with about 9 million enrollees. Within each state, counties are aggregated into "rating areas" - if a collection of counties is in the same rating area, all plans offered in these counties have to charge the same prices across different counties within the same rating area. Insurers do not have to offer plans in all counties, however. Despite the complexity of the geographic arrangements, it is helpful to think about county-level markets in this setting.

One of the key aspects of the ACA Marketplace that we focus on in this paper is the provision of subsidies for consumers with low incomes. The subsidy system is complex and consists of several pieces. We focus on subsidies that reduce annual premiums that consumers are responsible for. These subsidies are based on a classic "tagging" principle - individuals with lower incomes receive higher subsidies. In addition, subsidy levels are anchored to full prices ("bids") charged by insurers.

Premium subsidies are known as (Advanced) Premium Tax Credits - PTC - they can be paid (directly to the insurance company on consumer's behalf) at the start of the year based on projected income and be then adjusted when consumers file taxes if actual income differs from the projection. Consumers can also choose to forgo receiving advanced credit and instead claim the full amount ex post in their tax return. The PTC is calculated in two steps. First, the "MAGI" measure of income (converted to a percent of federal poverty level - FPL) determines the maximum dollar amount that the consumer should be paying for insurance premiums. Call this amount "CAP." The CAP is based on a non-linear sliding schedule. For example, if individual's income is 200% of FPL, then he or she should be spending no more than 6.34% on income on health insurance premiums. At 400% FPL, the CAP is equal to infinity as individuals with income above 400% FPL are not subsidized.

In the second step, the regulator records the bid of the Silver plan (for each market) that has the second-lowest bid in the market. Call this premium SLSP. If CAP is greater than SLSP in the county where a consumer resides, then this consumer gets no subsidy. If CAP is less than SLSP, the consumer gets a PTC that equal to the difference between the applicable SLSP and the CAP.[1]

We combine several sources of data for our analysis. We use data from 2015 - the second year of Marketplace operations - and focus our analysis on ACA Exchanges that use the federal healthcare.gov platform, as the best data is available for this year and for that part of the market. We observe detailed choice sets that consumers faced in each geographic market in premium and plan structure files that have been released by CMS and are available on the agency's web page. CMS has also released enrollment data at county-metal level, at plan level, and at county-insurer level. Kaiser Family Foundation has generously provided us with a dataset that records the potential size of the market at a fine geographic level. Finally, we use the 2015 edition of the American Community Survey (ACS). ACS data allows us to create a representative sample of uninsured individuals in each county, for whom we observe

[1]While we abstract from family-level analysis in our estimation due to lack of data, in practice income at the point of enrollment is estimated based on "tax" family composition, household income and which members of the family are getting coverage.

income, age, race, and gender.

Table 1 summarizes the key data points on the choice sets that individuals face, enrollment, and demographics. In 2015, consumers could choose among on average 39 plans offered by three large national insurers and a number of smaller firms. The annual pre-subsidy premiums for a 40-year old in these plans ranged from $\$2,500$ to $\$5,700$ with an unweighted average of about $\$3,800$. The average number of potential enrollees per market was close to $8,000$, although markets differed dramatically in their size, ranging from fewer than 100 potential enrollees to more than $500,000$. On average across markets, 62% of potential enrollees chose not to purchase a Marketplace plan; among those that did purchase, Silver plans were by far the most popular, accounting for almost 70% of choices conditional on enrollment. About 32% of potential enrollees are eligible for the most generous cost-sharing support. Potential enrollees are on average 38 years old, 83% white, with an average income of 250% FPL. On average, these consumers qualify for $\$2,120$ in premium subsidies.

3 Theory and Preliminary Evidence

3.1 Targeting and Market Power: Stylized Framework

In this section, we highlight the basic intuition that characterizes the relationship between subsidy targeting and market power. To motivate why governments may target subsidies on the basis of income, we first begin with a simple model of a benevolent government that aims to allocate in-kind subsidies in order to maximize social welfare subject to a fixed budget constraint. In this simple framework, We show that the optimal subsidy is differentiated across consumer types due to differences in the curvature of the utility function. Then, taking the idea that subsidies may be differentiated by observable consumer type, we model the incentives of a profit-maximizing firm that provisions the good to the consumers. We show that the combination of subsidy targeting and market power interacts in a complex fashion to distort market outcomes.

Suppose that the economy consists of a benevolent government, a single good w, and a unit mass of consumers. Consumers belong to one of two observable types, $t \in \{H, L\}$, and are endowed with bounded, increasing, and convex utility functions, $U_t(w)$: $U(w) < \infty$, $U_i'(w) > 0$ and $U_i''(w) < 0$. Let η and $1 - \eta$ denote the fraction of consumers within each type H and L, respectively. The government has a fixed amount of the good equal to G and wants to allocate the good among the two consumer types to maximize social welfare.

Formally, the problem is:

$$\max_{w_L, w_H} (1-\eta)U_L(w_L) + \eta U_H(w_H) \text{ s.t. } w_L + w_H = G. \quad (1)$$

Substituting in the budget constraint and taking derivatives, the optimal (interior) solution is characterized by the following first-order condition:[2]

$$(1-\eta)U_L'(w_L) = \eta U_H'(G - w_L). \quad (2)$$

Intuitively, the optimal allocation equates share-weighted marginal utilities across consumer types. From this simple expression, one can also see that the optimal split will, in general, be unequal at both the group- and individual-level. In this sense, the social planner finds it desirable to *target* allocations (i.e. subsidies) as a function of consumer type, which is a ubiquitous feature of public assistance programs.

Next, taking the idea that subsidy structure constant from the first part of the model, we introduce a profit-maximizing firm to act as an intermediary to provide the good.[3] Following the design of many publicly-subsidized and privately-provided markets around the world, we allow that firm to set a single price, b, after observing both the subsidy structure and the distribution of consumer types. However, the firm is not allowed to price discriminate.

A consumer of type t purchases one unit of the good from the firm if this consumer's willingness to pay, v_t, is greater than the effective (post-subsidy) price $p_t = \max\{0, b - z_t\}$, where each consumer type has a type-specific subsidy, z_t. Denote the share of consumers of each type purchasing the good as $s_t(p_t)$. We assume that for any given price p, (weakly) more H type consumers are willing to buy the good as $s_H(p) > s_L(p)$, for any p. We also assume that H consumers have a flatter demand curve, $\frac{\partial s_H}{\partial p} < \frac{\partial s_L}{\partial p}$ for the same ∂p.

We assume that the good can be provided at the same marginal cost c to both types of consumers.[4] The firm chooses bid b to maximize the following profit function:

$$\pi(b; z) = (b - c)(s_H(p_H(b))\eta + s_L(p_L(b))(1-\eta)). \quad (3)$$

[2]When an interior solution does not exist, it implies that one consumer type gets all of the subsidy, which is the most extreme form of inequality.

[3]One motivation for doing so is that the firm (or, in the case of regulated competition, many firms) can provide the good at lower cost or higher quality than the government.

[4]This assumption can be relaxed to allow for differences in costs across types. Allowing for costs to vary doesn't change the intuition of the mechanism, but adds another degree of freedom into the first-order condition, making algebraic expressions less transparent. We hence use the constant marginal cost case for the stylized discussion in this section.

Assuming an interior solution, the first order condition for the firm's choice of b takes the following form:[5]

$$s_H(p_H(b))\eta + s_L(p_L(b))(1-\eta) \\ + (b-c)\left(\eta \frac{\partial s_H(p_H(b))}{\partial p_H(b)} \frac{\partial p_H(b)}{\partial b} + (1-\eta)\frac{\partial s_L(p_L(b))}{\partial p_L(b)} \frac{\partial p_L(b)}{\partial b}\right) = 0. \quad (4)$$

This is the standard first-order condition for a monopolist selling to both types at a uniform price with one key difference: the term $\frac{\partial p_t(b)}{\partial b}$, which does not exist in unsubsidized markets, captures the role of the subsidy design. The key economic force of our model is that, under targeted subsidies, the response of the effective consumer price to changes in b varies across consumer types:

1. For the unsubsidized consumers of type H, $\frac{\partial p_H}{\partial b} = 1$, as an increase in the bid translates one-to-one into the price they pay.

2. Consumers of type L, however, receive a subsidy $z_L > 0$. If the firm's current price is above that subsidy, $b > z_L$, then any additional increase is passed on to the consumers, $\frac{\partial p_L}{\partial b} = 1$. However, if $b < z_L$, then consumers pay zero additional premiums for a small increase in b and hence $\frac{\partial p_L}{\partial b} = 0$.

As a result, targeted subsidies may induce a distortion in the aggregate demand curve perceived by the firm relative to the true underlying demand curve. Critically, the distortion is asymmetric, as the targeted subsidies only change the demand curve for type L. However, since the firm is forced by regulation to charge a single price, the presence of the subsidized type L generically influences the price paid by type H. In the context of targeted subsidies, We call this effect the *demographic externality.* The sign and magnitude of this effect depends on both the level of the subsidy and the relative demand curves for the two types. Increasing the proportion of type L in the market has a nuanced effect. While there are more subsidies in the market that the firm will try to capture by increasing price, it also faces different demand curves across the two types. So decreasing η has two effects: a subsidy effect and a demand curve effect.

We summarize the comparative statics of the model in three propositions that we will be able to test empirically in Section 5:

[5] In the empirical application we will observe that firms may be pricing at a corner solution due to restrictions on their profit margins via a regulation known as a minimum loss ratio regulation (MLR). As a result, the MLR can in practice mute the conceptual externality effects that we discuss in this section. To highlight the economic mechanism, however, we assume an interior solution in this exposition.

Proposition 1 *Introducing subsidy targeting by giving subsidy z to consumers L has an ambiguous effect on consumers H. Consumers H experience a negative externality if z is large enough, leading the insurer to increase bids to the level of z, taking advantage of the inelastic demand by consumers L. Consumers H experience a positive externality if z is small enough that it shifts the market towards consumers L with a higher marginal utility of income, preserving the price responsiveness of these consumers on the margin.*

Proposition 2 *Removing subsidy targeting by giving subsidy z to previously unsubsidized consumers H generates a negative externality on incumbent subsidized consumers L if consumers H have a less elastic demand function that is shifted out by subsidies.*

Proposition 3 *Under subsidy targeting, as the share of unsubsidized consumers η decreases, the list price b decreases if subsidized consumers have a more elastic demand function and the subsidy z is small enough. For a large enough z, as the share of unsubsidized consumers η decreases, the list price b increases and approaches the subsidy level.*

In the next section, we will examine first empirical evidence that is consistent with these comparative statics. Then, in the counterfactual analyses of Section 5 we will be able to simulate how targeted subsidies affect market equilibria, and whether the incidence of subsidies falls onto consumers or producers. First, in the absence of market power—when the product is offered at marginal cost—we will quantify the extent of the efficiency-equity tradeoff between targeted subsidies that are available only to low income beneficiaries and non-means-tested subsidies. Second, we will estimate how this trade-off changes when we incorporate market power and require the firm to set one list price for all consumer types.

3.2 Descriptive Evidence

Before proceeding to the empirical model of supply and demand in section 4, we start by investigating descriptive patterns in the data. We first test whether the observed relationships between prices and demographics of local markets are consistent with the comparative static in Proposition 1 above.

We have two related empirical predictions. First, list prices charged by insurers for second-lower silver plans should be lower in markets with more consumers that are not eligible for subsidies. Since subsidies in ACA markets are linked to the list prices of the second-lowest silver plans, for these plans the subsidy is such that $p(b) = min(p, b)$ and $s = b - p$ for a pre-determined p. Hence, we expect prices for these plans to be lower in places with a higher share of consumers with income above 400% FPL. Second, if list prices are lower in places with a higher share of consumers not eligible for subsidies, the effective prices paid by *partially* subsidized consumers, should also be lower in the same markets. Consumer prices are determined as the difference between the income-related cap on premiums and the list of the second lowest silver plans. For consumers with higher incomes, the cap is higher and can in fact sometimes exceed the list price of the second lowest silver (2LSP) plan. In that case, consumers face the list price of the 2LSP rather than the premium cap. As 2LSP goes down, the price for these consumers also goes down. Hence, the price faced by partially subsidized consumers may vary across markets if the list price of the 2LSP is low enough.

We test whether both of these empirical predictions are consistent with the data in Figure 1. Panel (a) plots the relationship between individual list prices (we take the price for a 40 year old adult - since prices are scaled with a mechanical step function, it does not matter which age we pick) and the estimated share of potential consumers with household income above 400% FPL. We estimate this share from ACS data. Given the significant variation in the regulatory and cost environment between states, we first residualize both the horizontal and the vertical axis onto state fixed effects (and then add back the mean). We then group 2,505 counties into 20 equally sized bins along the x-axis. We see large differences in the list prices for 2LSP across the ventiles of the share of potential consumers that are ineligible for subsidies. While list prices are around $5,000$ per year in markets with 10% or fewer of such potential consumers, they are 400 or 8% lower in markets where more than a third of potential enrollees are not eligible for subsidies. The relationship is statistically very pronounced, with a linear slope of 155 dollar decrease for each 10 percentage point increase in the share of unsubsidized consumers.

Next, we test whether markets with a higher share of potential consumers that are not eligible for premium subsidies exhibit lower effective premiums for those potential consumers that are eligible for partial subsidies. This relationship would exist in the data if insurer list prices in markets with more elastic demand were falling faster than subsidies, so that net premiums were lower.

Panel (b) of Figure 1 suggests that this relationship holds empirically. Here we plot the same x-axis of the figure we plot the share of potential Marketplace enrollees per county that have income above 400% FPL and hence are not eligible for premium subsidies. We group 2,566 counties into 20 bins equally-spaced by the x-axis value. For each county, we calculate the average effective premiums that potential consumers with incomes above 250% FPL and below 400% FPL would have faced. These consumers are eligible to receive premium subsidies, but typically their subsidies only cover a portion of the premium. We plot the average of these county-level effective premiums within each bin on the y-axis.[6] A clear pattern emerges - individuals that are poor, but not too poor to receive full subsidies face higher premiums in markets that have fewer “elastic” consumers.

Panel (b) of Figure 1 illustrates how the underlying demographic composition of the market affects enrollment. The x-axis of this figure is the same as in Panel (a). On the y-axis, we plot the county-level share of potential consumers in the income bracket between 250% and 400% FPL that purchase any plan on the ACA Marketplaces. Consistent with these individuals facing higher premiums in markets with fewer “elastic” consumers, we observe lower enrollment in these markets.

We formalize these relationships in Table **??**. We start with a regression that captures the same relationship as Panel (a) of Figure 1: how effective premiums for second-lowest cost silver plans among potential consumers in the income bracket between 250% and 400% FPL vary with the share of potential consumers with income above 400% FPL in their county c and state s. To focus on comparable consumers across counties, we control for consumers' age (a) and income (w), as well as state fixed effects. The coefficient of interest is β that measures the correlation between the share of elastic consumers (σ) in a county and premiums that consumers with partial subsidies face.

$$p_{i(cs)} = \beta \sigma_{cs} + \sum_a \alpha_a \mathbf{1}\left[age_{i(cs)} = a\right] + \kappa w_{i(cs)} + \gamma_s + \epsilon_{i(cs)} \tag{5}$$

We estimate two versions of this specification: at the individual level on the ACS sample of 206,064 potential consumers with income between 250% and 400% FPL, clustering standard errors at the county level. The individual level allows us to precisely control for age fixed effects and income. Similar to Panel (a) of Figure 1, we estimate a negative relationships: individuals in markets with 10 percentage point more of elastic consumers face on average $45 lower premiums. In Column (4) we report an aggregated version of this relationship at the county level, we control for average age in the county rather than age fixed effects. Here we find a slightly attenuated coefficient, suggesting that on average an increase in the share of unsubsidized consumers by 10 percentage points leads to $19 lower premiums.

Columns (1) to (3) of Table **??** report the results of enrollment regressions. For each county, we calculate the share of potential enrollees with income between 250% and 400%

[6] Both the x and y axis are residualized to account for fixed differences across states and for the exact income level of potential consumers.

FPL that purchased a plan on the ACA Marketplace.[7] In column (1) we regress this share of inside option enrollment on average effective second-lowest cost silver plan premium for this group of potential consumers. We expect this linear demand estimation to be biased, as observed premiums are the equilibrium outcome of market interactions. Indeed, we find that the coefficient on premiums is positive and noisy. Next, we instrument premiums with the share of "elastic" consumers in each market. Column (2) reports the reduced form of this specification, illustrating that enrollment is higher in markets that have a higher share of "elastic" demand (the same relationship that we observed in Panel (b) of Figure 1). The 2SLS specification of demand in Column (3) produces a more meaningful estimate of the demand slope, suggesting that the inside share decreases by 0.3 percentage points (off the mean of 0.58) for each $100 increase in the effective annual premiums for the second lowest cost silver plans.

Taken together, these relationships provide strong suggestive evidence for insurers' strategic response to income-tagging of subsidies. Insurers have an incentive to raise price in places where more potential consumers are heavily subsidized, and an incentive to lower prices in place where more consumers are paying full premiums. As subsidies are tagged to observable income, insurers have nearly perfect information on which market is going to be more elastic and are likely to respond strategically to this information.

4 Empirical Model

4.1 Demand

We formulate and estimate a random utility model of demand for health insurance plans on ACA Marketplaces. Utility takes a semi-nonparametric form. Each consumer i gets utility of ϕ_{ij} from buying plan j and pays a premium p_{ij} that lowers i's utility by α_i utils per dollar. As consumers typically buy insurance for everyone in their family who needs coverage ("coverage" family), demand operates at the family level. Following the institutional design of the ACA Marketplaces, we impose that each member of the "coverage" family buys the same plan. We accomplish this by assuming that each consumer maximizes the average utility of their coverage family. Consumers pick plans that give their families the highest utility at the lowest price, or choose not to participate in the market.

[7] We observe the numerator of the share in the data released by CMS. We compute the denominator by applying the county-level share of individuals with income between 250% and 400% FPL in the ACS data to the total market size in each county.

Formally, we posit that individual i in family f in market t chooses plan j from a set of choices J. The set of choices that each family faces depends on the family's geographic location, family income, and age distribution of family members that require insurance coverage. It is helpful to think about geographic markets as operating at the county level, since, as Fang and Ko (2018) demonstrate, insurers effectively make their strategic decisions at the county level.[8] The indirect utility function as observed by each consumer takes the following form:

$$u_{ijt} = -\alpha_i p_{ijt} + \phi_{ijt} \tag{6}$$

Where p_{ijt} is individual-specific price that consumer i, who lives in market t, faces for plan j. This price is individual-specific because consumers of different ages with different family incomes face different prices for the same plan j in market t.[9] ϕ_{ijt} is the amount of utility a consumer gets from plan jt. This utility can vary across consumers for two reasons. First, consumer valuation of the same plan may differ due to heterogeneity in preferences. Second, the same plan j may have different characteristics for different consumers, as consumers with lower incomes get more generous coverage when they buy certain plans.

We make a few assumptions about α_i and ϕ_{ijt} to arrive at an empirically-tractable version of this utility function that still preserves the semi-nonparametric approach. First, we replace individual-specific α_i with a coarser set of parameters that vary across nine demographic groups d. The demographic groups are all permutations of three age categories: age under 25, age between 25 and 40, age above 40, and three income categories: income under 200% FPL, income between 200% and 400% FPL, income above 400% FPL. Second, we decompose the utility that a consumer gets from plan j into several additively separable components. The first component ψ_a captures the average level of utility that consumers get from purchasing any insurance plan. We allow this intercept parameter to vary across three age groups a (same as the age groups above), so as to capture the idea that older consumers may value insurance more, all else equal.[10] The second component captures the deviations in the

[8] In practice, counties are aggregated into *service areas* that are collections of counties (one or more) where plan j is offered; and *rating areas* that are collections of counties where plan j has to offer the same price in all counties if it chooses to operate in these counties. Service areas and rating areas need not overlap. We account for the exact detail of rating and service areas in estimation.

[9] In theory premiums also vary depending on the consumer's smoking status. In what follows, we abstract from the differences in premiums for smoking and non-smoking enrollees. Smoking is self-reported and insurers have no obvious mechanism to verify whether an enrollee smokes.

[10] We further allow for a separate intercept for the group of consumers with income under 100% FPL. While this group of consumers should not be participating in ACA Marketplaces, as they are commonly eligible for Medicaid and are not eligible for ACA subsidies, we observe some very few enrollees from this group in the data - a separate intercept for this group allows the model to rationalize very low, but non-zero,

generosity of a plan j that a consumer may face with low enough income. This component is captured through the actuarial value (AV) characteristic of a plan - the AV of silver plans is adjusted up for lower income consumers. Finally, we include a plan-specific constant δ_j for each plan j that captures the average utility that a consumer gets from purchasing plan j. The remaining difference between ϕ_{ijt} and these three components is captured in the part of the utility function that is observed by the consumer, but is not observable to the econometrician - ϵ_{ijt}. This random individual-plan specific shock to the utility function is assumed to be distributed Extreme Value Type 1, which leads to a logit discrete choice model. The extensive literature on the statistical properties of discrete choice models has demonstrated that this set-up is extremely flexible and can approximate any random utility function (McFadden and Train, 2000). To summarize, the empirical version of the utility function becomes:

$$u_{ijt} = -\alpha_{d(i)} p_{ijt} + \psi_{a(i)} + \gamma AV_{ij} + \delta_j + \epsilon_{ijt} \tag{7}$$

To close the model we assume that individuals choose plan j that maximizes their family's average utility across all possible choices, or they choose not to enroll, which gives a normalized utility of zero. Formally, i chooses j if $\frac{1}{|f|}\sum\limits_{i\in f} u_{ijt} > \frac{1}{|f|}\sum\limits_{i\in f} u_{ikt}$ for all k in J such that k is not equal to j.

The variation in premiums that we observe in the data does not stem from experimental assignment. Policies that govern subsidy levels in this market, however, deliver a set of natural experiments that allow us to pin down the marginal utility of income parameters α_d. Consider two consumers of the same age in a given market. Without subsidies, these consumers would face the same prices for a given plan j. In the presence of subsidies, however, two consumers face different effective prices for the same plan j if their incomes are different. The variation in subsidies across income levels depends on a pre-specified formula that generates kinks in the relationship between income and effective premiums, as we discuss in Section 2. In the data, we observe - for each of more than 2,000 markets - how the share of consumers purchasing any insurance plan varies within a market, within an age group across adjacent income levels. Similarly, for consumers of the same age, we observe different prices - and different enrollment shares across consumers of adjacent ages. This policy-related variation in prices for the same plan j across income and age groups allow us to identify α_d parameters. These institutional features of the market allow us to improve substantially on the more common identification approach of using instrumental variables for

inside share for this group.

prices that rely on the firm's pricing or product characteristics in other geographic markets (Hausman, 1996; Berry et al., 1995; Berry and Haile, 2016).

4.2 Supply

4.2.1 Profit function

Insurers on ACA Exchanges decide which geographic markets to enter, how to design their plans, and how to price them. In this paper we are interested in the conceptual question of how subsidy targeting to lower-income consumers may affect equilibrium prices, conditional on entry and contract design decisions; hence, we keep insurers' entry and product design fixed. Modeling price-setting in this market poses a significant challenge, as pricing is constrained by an array of regulatory provisions. We first start with a brief accounting of payment flows in the market. We then discuss the assumptions we make to get an empirically tractable supply-side model that allow us to focus on the role of subsidy targeting.

For each consumer i of age a, plan j collects revenue that consists of several pieces. First, plan j collects premium p_{ij} from the consumer. For consumers that do not receive premium subsidies, the premium is equal to insurer's full list price for consumers of age a, b_j^a. For consumers that are eligible for subsidies, the insurer collects $p_{ij} < b_j^a$ from the consumer as well as a subsidy from the federal government. Together, the premium and the subsidy add up to b_j^a.[11] Second, the insurer may collect revenue from three risk-equalization programs that we describe below.

On the cost side, both realized and expected costs that plan j incurs for consumer i differ across consumers and plans. Let the total (i.e. out of pocket and insurer payment) expected healthcare spending of consumer i in plan j be h_{ij}. This spending depends on consumer's underlying health risk, which we denote with r_i, as well as the features of plan j. The features of plan j may affect h_{ij} either by changing consumer demand for healthcare, e.g. through moral hazard, or plan j may simply have different negotiated prices for the same services. We denote j's contract features, including any negotiated price provisions, with ϕ_j. Then, h_{ij} is a function of r_i and ϕ_j. Plan j's expected cost for consumer i is not equal to h_{ij}. Instead, the plan expects to pay only a portion of h_{ij}, net of consumer cost-sharing. Consumer cost-sharing, in turn, is either paid directly by the enrollee or can be paid by the government in the form of cost-sharing subsidies. The source of payment doesn't affect insurer's cost per

[11] If the subsidy is higher than the bid, the consumer pays zero and does not receive the cash value of the "unused" subsidy. In practice, the subsidy operates as a tax credit; the estimated level of the credit is reconciled during tax filing.

se; however, insurers' costs may go up if cost-sharing subsidies induce additional demand for healthcare services. As eligibility for cost-sharing subsidies depends on individual income, we can write that the plan's expected cost for enrollee i is $c_{ij}(r_i, \phi_j, D_i)$, where D_i denotes consumer i' s income.

Prior to any risk-equalization programs, plan j's expected profit for consumer i of age a is then:

$$\pi_{ij}(b_j^a) = b_j^a - c_{ij}(r_i, \phi_j, D_i), \tag{8}$$

Suppose that for any plan j, there is a baseline plan-specific cost c_j^a of covering an average enrollee of age a. Then, we can re-write c_{ij} as the sum of the plan-specific cost and an idiosyncratic (but predictable) individual cost component: $c_{ij}(r_i, \phi_j, D_i) = c_j^a + \tilde{c}_{ij}(r_i, \phi_j, D_i)$. Using this notation, we can re-write the expected profit of plan j from enrolling individual i of age a to be:

$$\pi_{ij}(b_j^a) = b_j^a - c_j^a - \tilde{c}_{ij}(r_i, \phi_j, D_i), \tag{9}$$

The individual-specific expected cost term $\tilde{c}$ allows for the presence of advantageous or adverse selection that is a function of plan characteristics ϕ_j. Three programs exist on ACA Marketplaces (within the time horizon we study) that are aimed at equalizing expected insurers' costs across all enrollees. The aim of these programs is to reduce the incentives for active cream-skimming by insurers and ameliorate the consequences of adverse selection of sicker consumers into more generous plans. It is easier to think about these programs as affecting insurers' costs; however, in practice, the programs constitute revenue streams. The first program - risk adjustment - generates lump-sum payments or lump-sum collections from a plan, depending on whether the plan has enrollees whose risk is above or below the average in the market. Second, the reinsurance program transfers additional revenue to insurers to cover expenditures on particularly high-cost consumers. Finally, insurers may receive funds from or be required to pay into a so-called risk corridor program. This last program attempts to reduce the ex post volatility in realized profits relative to the ex ante risk pool.

Intuitively, the idea of the risk-equalization programs is to create transfers that exactly offset the idiosyncratic cost component $\tilde{c}_{ij}(r_i, \phi_j, D_i)$, so that every enrollee has the same expected cost in the insurer's profit function. The reinsurance program effectively gives insurers additional individual-specific revenue for individuals with particularly high $\tilde{c}_{ij}(r_i, \phi_j, D_i)$, so as to reduce the impact of this term on insurer's profit function. We can then consider the difference between this additional revenue and the idiosyncratic cost-component as the net idiosyncratic cost that is relevant for insurer's decision-making. Denote this difference with

η_{ij}. Now let the (positive or negative) lump-sum risk-adjustment payment to the insurer be R_j. This term is a function of risk types r_i of all individuals that enroll in a plan and is not individual-specific, but it aims to offset the sum of η_{ij} across all i's in cases where the *expected* (or in other words, predictable) individual-specific deviations in risk across consumers add up to a positive or a negative quantity. Let $H_j = \sum_{i \in j} \eta_{ij} - R_j$ denote any residual selection. If the risk-equalization programs fully offset the ex ante net idiosyncratic shocks, this term would be zero.

In the ACA environment, H_j may not be zero. The direction of this term is likely to primarily be a function of ϕ_j, or contract generosity. This follows from the fact that individuals with the lowest incomes, who are also likely to have the highest expected costs (as has been shown by a voluminous literature on health-income gradient), receive subsidies that significantly compress the variation in both prices of plans that these consumers face. Hence, the key differences across plans that are likely to drive selection lie in non-pecuniary plan features of ϕ_j, such as physician networks, formulary breadth, and chronic condition management. As without data on individual enrollment and expected costs in Marketplace plans, we can only estimate the profit function up to the $H_j(\phi_j)$ constant, we simplify the notation by assuming that $H_j(\phi_j) = 0$ in subsequent discussion. We come back to this assumption in our welfare computation and test the sensitivity of the subsidy incidence analysis in Section 5 to the assumption about the levels of $H_j(\phi_j)$.[12] Importantly, the assumption of $H_j(\phi_j) = 0$ has no bearing on the first order condition in prices, since as we argued above price-based risk screening, conditional on ϕ_j, is unlikely to be quantitatively important (in other words, the derivative of $H_j(\phi_j)$ with respect to prices is zero).[13] We do not explicitly incorporate the ex post risk-corridor transfers into the model - these payments can be interpreted as a reduction in insurers' fixed cost of purchasing private re-insurance policies and should not affect insurers' pricing incentives on the margin.

Rewriting the profit function using the share notation, where s_j^a denotes the share of plan j among consumers of age a and M^a denotes the number of consumers of age a on the market, we get:

$$\pi_j(b_j) = \sum_a s_j^a M^a b_j^a - \sum_a s_j^a M^a c_j^a \tag{10}$$

[12] In general, to the best of our knowledge, there exists no empirical evidence that would allow assessing the precision of risk-adjustment in the ACA market. From other markets that employ risk-equalization policies, we know that while risk-equalization leaves scope for residual selection, it goes a long way to reducing the differences in costs in expectation.

[13] We have illustrated this point empirically in an earlier paper that studied subsidization mechanisms in the context of Medicare Part D (Decarolis, Polyakova and Ryan, forthcoming).

According to ACA statutes, insurers have to follow a statutory age schedule for their bids that constraints age-specific underwriting. This restriction allows us to simplify the problem further. Let there be a fixed set of age-specific multipliers that apply to bids. We assume that the same multipliers apply to expected baseline costs, capturing how healthcare costs increase with age.[14] Let the multiplier vector be τ^a. The profit equation for plan j then becomes:

$$\pi_j(b_j) = (b_j - c_j) \sum_a s_j^a M^a \tau_j^a \tag{11}$$

At the insurer level, we aggregate across all plans j offered by insurer f:

$$\pi_f(\mathbf{b}) = \sum_{j \in f} \left[(b_j - c_j) \sum_a s_j^a M^a \tau_j^a \right] \tag{12}$$

Finally, to close the model we introduce the last feature of the institutional environment - medical loss ratio (MLR) regulation that has been documented to be binding for the majority of insurance contracts in this market (Cicala et al., forthcoming). The MLR regulation stipulates that insurers in the ACA market spend at least 80% of their revenue on healthcare claims and quality improvement, constraining the markups to be at most 25%, and requiring insurers to rebate extra revenue consumers. We impose this restriction when inverting the first order condition to recover marginal costs. Under this restriction, the insurer maximizes profits by choosing a bid b_j for each plan j in its portfolio.

4.2.2 First order conditions

Insurers choose bids that maximize their profits taking into account the actions of other firms. The first order condition for a one plan-firm implied by the profit function in 12 is:

$$\frac{\partial \pi_f}{\partial b_j} = (b_j - c_j) \sum_a \frac{\partial s_j^a}{\partial b_j} M^a \tau_j^a + \sum_a s_j^a M^a \tau_j^a = 0. \tag{13}$$

For an insurer that offers more than one plan in a market, the set of j first-order conditions accounts for own and cross-price elasticities of demand, taking the vector form given by $S - \Omega(B - C) = \mathbf{0}$, where row j of vector S is given by: $S_j = \sum_a s_j^a M^a \tau^a$ and row j of

[14]The assumption of the same multipliers on costs and bids simplifies the problem computationally, but can conceptually be relaxed (Tebaldi, 2017), allowing costs to follow a different slope with respect to age than the statutory age-specific multipliers. Examining age-cost gradients in commercially insured populations, we found that the discrepancies are likely to be the largest in the oldest population that comprises the smallest share of Marketplace enrollment, and are thus unlikely to qualitatively affect our results.

vector $(B - C)$ is given by $(B - C)_j = (b_j - c_j)$, while row k, column j of matrix Ω is:

$$\Omega_{kj} = -\sum_{a} \frac{\partial s_j^a}{\partial b_k} M^a \tau^a \tag{14}$$

for plans k and j offered by firm f. This gives us j equations in j unknowns for each insurer, as want to recover costs c_j that are unknown up to the scaling factor τ. We invert Equation 13 and compute the baseline marginal cost c_j for each plan as a function of observed equilibrium prices and the elasticity of demand that is given by the demand parameters from Section 4.1.

The key term of the first order condition is the derivative of the (age-specific) share with respect to the (age-specific) bid: $\frac{\partial s_j^a}{\partial b_j^a}$. We drop the age superscripts to simplify notation in what follows, as age scaling is given by regulation and age markets are additive in our set up.[15] The share derivative reflects how much the demand for plan j changes when this plan increases its bid by a small amount. Unlike in a standard product-market setting, this term captures the complex relationship between premiums and bids within the ACA Marketplaces. Bids and premiums are linked via the premium subsidy mechanism, so formally:

$$\frac{ds_j(p_j, p_{-j})}{db_j} = \frac{\partial s_j}{\partial p_j} * \frac{\partial p_j}{\partial b_j} \tag{15}$$

Recall that the subsidy is a function of the bid set by the second-lowest cost silver plan and consumer's family income. The second-lowest cost silver plans face a different set of incentives relative to other plans, as under the observed subsidy regime a change in their bids affects not only their own prices, but also the subsidies, and hence consumer premiums, of all plans.[16] To account for the idea that the first order condition of a Bertrand game do not capture these additional incentives, we do not use inverted marginal costs for these plans.[17] Instead, following (Decarolis et al., forthcoming), we use other plans to impute the marginal costs of SLSPs. we project the estimates of marginal costs for the non-2LSPs on a vector of their characteristics and use this hedonic projection to predict the marginal costs of the SLSPs.

Since many consumers in the ACA market do not pay the full bid b_j for any plan, it follows that the change in the share of plan j in response to a small increase in bid b_j

[15] In practice, age markets interact through plan- or insurer-level risk-equalization policies and the MLR constraint.

[16] Formally, these plans have another terms in 15 that is non-zero: $\frac{\partial s_j}{\partial p_{-j}} * \frac{\partial p_{-j}}{\partial b_j}$.

[17] Note that since marginal costs of each plan are separable in the inverted first-order condition, the inversion of marginal costs for non-SLSPs is not affected by the SLSP.

depends crucially on the design of the subsidy mechanism and the composition of consumer incomes in the market. As can be seen from Equation 12, the total market share of plan j is a weighted sum of the shares in each income-age bin (weighted by the share of consumers in each bin). In follows that the derivative of the plan's share with respect to its own bid, and hence its pricing decisions, depend on the share of each consumer type in each market. This feature of the problem generates the key mechanism that is the focus of our paper: the subsidy-income linkage creates a demographic externality, as insurers have an incentive to change prices in markets where the subsidy structure either shifts out the intercept or reduces the slope of the demand function for a significant share of potential consumers.

4.3 Efficiency Metric

We define a welfare function (W) that consists of three pieces: consumer surplus (CS), insurer profits (Π), and government subsidies (G):

$$W = CS + \Pi - \lambda G, \tag{16}$$

where λ is the social cost of raising public revenues, which we assume to be 30 cents on a dollar of public spending. Following Williams (1977) and Small and Rosen (1981), surplus for consumer i with a vector of marginal utilities θ_i takes the following form:

$$CS(\theta_i) = \frac{1}{\alpha_i}\left[\gamma + \ln\left[1 + \sum_{j=1}^{J} \exp(v_{ij}(\theta_i))\right]\right], \tag{17}$$

where γ is Euler's constant, and v_{ij} is the deterministic component of utility for person i (recall that this is, in return, is the average utility within a family) for plan j and is equal to utility net of the idiosyncratic ϵ term.[18] We integrate out over the empirical distribution (as observed in the ACS) of ages, income, and family composition to obtain average annual per capita consumer surplus:

$$CS = \int CS(\theta) dF(\theta). \tag{18}$$

Producer surplus Π is computed following equation 12. We assume that any risk-equalization payments, including risk corridors, contribute to cost equalization and are already captured in marginal cost estimates, implying that they do not separately enter the

[18] Euler's constant is the mean value of the Type I Extreme Value idiosyncratic shock under the standard normalizations in the logit model, and is approximately equal to 0.577.

profit function.

Government spending G captures three parts. Nominal spending includes subsidies for insurance premiums as well as subsidies for cost-sharing reduction. The former are computed either from the data or are adjusted in simulations of Section 5. Cost-sharing reduction (CSR) spending is held at observed levels. Specifically, using CMS data we compute the average per capita spending on CSR subsidies by consumer type, based on income brackets. In all counterfactual simulations, we then assign this average spending level to each consumer who falls into the respective income bracket and who enrolls in a plan where cost-sharing reduction is available. In the final step, we account for the fact that when a consumer enrolls into an ACA plan, the government likely saves some money on this consumer; for example, if a consumer enrolls in a formal insurance plan, this consumer is then unlikely to benefit from any public payments for uncompensated care. Following the Kaiser Family Foundation 2013 report on public spending on uncompensated care for the uninsured, we assume that the government saves \$1,827 per capita in public funds for each consumer enrolled in an ACA insurance plan.

4.4 Estimation results

4.4.1 Demand Parameters

We use a simulated method of moments procedure to estimate utility function parameters. As we do not observe plan-market-level enrollment that would give rise to a more standard model akin to (Berry et al., 1995), we adapt the estimation routine to incorporate enrollment moments at varying level of aggregation. We first use market (i.e. county) level enrollment moments that report enrollment at three different levels: (i) the share of consumers enrolled in plans of different metal levels; (ii) the share of consumers of different age intervals purchasing any ACA plan; (iii) the share of consumers of different income levels purchasing any ACA plan. These moments give us several thousand cross-sectional restrictions on the underlying demand function, including demographic interactions. We use these market-level moments to estimate α_d, ψ_a, and γ. Next, we use plan-level enrollment moments to recover δ_j for every plan.

Panel A of Table 5 reports the results. We find intuitive patterns for the variation in the marginal utility of income parameter across demographic groups. A one dollar increase in price has on average a larger impact on the utility of poorer consumers and younger consumers, for whom we would expect a dollar of insurance to constitute a higher share of their annual budget. The relationship between the overall value of insurance and age is non-linear. While consumers above the age of 40 value any insurance more than consumers aged 25 to 40, the demand by youngest consumers - below age 25 - exhibits an even higher valuation. In general, while the patterns are mostly intuitive, we are cautious about the interpretation, as the consumers in the model are assumed to be maximizing average family utility, hence the marginal utility of income parameters capture family level preferences. Family-level demand could, for example, exhibit a higher valuation of insurance by younger consumers, stemming from the valuation of their parents rather than individuals themselves. This would lead to a high estimated value of insurance at young age, as the younger group includes children, whose parents may place a high value on having insurance for their child. As would also be expected, we find that consumers get significant utility from purchasing plans with a higher level of coverage as measured by the actuarial value, conditional on other characteristics of plans held fixed as captured by the plan's δ that measure the relative attractiveness of plans to consumers on average. Together, the coefficients suggest that, for example, consumers over age 40 with income over 400% FPL value insurance coverage by a plan with the highest plan fixed effect that on average pays 70% of expenditures (maximum δ_j for a silver plan equals 0.34) at about \$2,891, which lies within the support of the distribution of list premiums for a 40 year old consumer for silver plans that ranges from \$2,391 to \$9,057 in the data.

We investigated several measures of the model's in-sample fit. We find that the model is able to match the county-by-metal-tier moments well. The between observed and simulated moments are centered at zero and are not concentrated in any particular geographic area of the country. Appendix Figure A.2 illustrates one set of moments used for the estimation and the geographic distribution of the model fit. In Panel (a) we report the county-level market share of Silver plans. In panel (b) we report the average in-sample difference between the data and the model's prediction of county-by-metal-level enrollment shares across all 2,566 counties that are used in the estimation. We closely match these aggregated enrollment shares with the model being able to capture a substantial amount of variation in the data.

4.4.2 Cost

Panel B of Table 5 reports the results of a hedonic regression that projects marginal cost estimates from the first order condition inversion onto plan characteristics for non-SLS plans. The estimates are intuitive, in that we find that more generous benefit design is associated with higher marginal costs. The average (baseline, for a 20 year-old) non-SLSP marginal cost from the inversion procedure is circa $1,940, with a standard deviation of $540. Moving a plan's actuarial value by 1 basis point, while keeping the "metal level" the same increases the marginal cost by $19. Moving from a gold to a silver plan decreases costs by ca. $500 ($300 metal label affect and plus the actuarial value adjustment of $192). The regression also includes measures of out of network coverage, whether a plan is HSA-eligible, whether a plan covers some common benefits, whether it offers management of common chronic conditions, and insurer fixed effects. We use this regression equation to impute marginal costs for SLSPs. As in the marginal cost inversion, we impose the MLR constraint on marginal cost predictions.

Figure A.3a in the Appendix illustrates the resulting distribution of estimated marginal costs for the baseline age group of 20-year old consumers. We plot the distribution separately for Bronze and Gold plans. We observe two pronounced patters. First, there is substantial heterogeneity in costs within a metal level. This is not surprising, since plans on the ACA Marketplaces are extremely diverse, with some plans being offered by large national insurers and some by local co-operatives. Second, there is substantial differences in costs between more and less generous plans, as we already saw in Panel B of Table 5. This is intuitive, as mechanically gold plans cover 80 percent of consumers' healthcare expenditures on average, while the Bronze plans cover only 60 percent. We would expect that the ratio of costs between these plans is on average 1.3, which is consistent with the shift in the distribution that we observe.

Appendix Figure A.3b compares our estimates of marginal costs from the first-order condition inversion (and projection for the SLSPs) to plan-level accounting costs as reported by plans to CMS. The accounting costs are measured with error, as insurers are allowed to report their costs equally split across their plans rather than providing a true plan-level attribution of costs. Moreover, the reported accounting costs do not include some ex post cost reconciliation, such as, for example, MLR repayments. Nevertheless, the accounting cost data provide a valuable informational signal, as they are likely to on average provide an accurate ordinal ranking of plans from least to most expensive, and also give a general sense of cost levels in the market. As can be seen in Appendix Figure A.3b, as we would expect given the ex post cost reconciliation that characterizes our institutional environment, our estimates of marginal cost are on average lower than reported accounting cost, although they have the same general order of magnitude. Further, we observe a very strong correlation between accounting costs and marginal costs, which supports the idea that we are accurately able to differentiate more and less expensive plans. [19]

4.4.3 Demographic Externality

The combination of estimated demand and cost parameters allows us to illustrate the demographic externality mechanism, as described conceptually in Section 3.1, in the context of our empirical model. We perform two counterfactual simulations that map into the comparative static predictions of Section 3.1.

The first simulation increases the number of subsidized consumers without changing the distribution of marginal utility of income in the population. For each market, we set subsidies for consumers with income above 400% FPL as if these consumers had income of 151% FPL. This in practice means that we "endow" the 400% + FPL consumers with income of 151% FPL, but do not change the parameters of their utility function. This means that in each market the share of consumers with subsidies increases, while the share of unsubsidized consumers goes to zero. The results of this first simulation are reported in Figure 2 and are marked as "case A." This figure reports how average premiums and consumer surplus changes for *directly unaffected* consumers, i.e. those with incomes between 150% and 400% FPL when their neighbors with income above 400% start getting subsidies. As we would expect in the presence of market power, insurers take advantage of the fact that in the 400%+ FPL market segment, consumers now face lower prices for any given list price. Insurers increase list prices in response. The increase in list prices, in turn, affects consumers of all income levels. In particular, plans now become more expensive for "unaffected" consumers with incomes below 400% FPL. As the light dashed line marked with "A" in the figure illustrates, the average annual effective premium for consumers with income under 400% FPL that remain on the market increases by 10$. Consumer surplus, marked with grey circles, in turn declines by $0 to $20, depending on how much consumers of different income levels are exposed to the increase in price and how much that affects their utility function. While average changes in prices and consumer surplus are quantitatively relatively small, they are highly unequally distributed across geographic areas. Panel A of Table 3 illustrates that some counties are completely unaffected by the simulated change, while some others experience more than

[19] Related work in this areas has pursue a different approach of directly using accounting costs as inputs into the counterfactual exercises, avoiding the inversion of the first-order conditions (see for example, Tebaldi, 2017). We do not pursue this strategy in our context, given that accounting costs are not observed at the product-market level and may capture several levels of ex-post accounting of cash flows through risk-equalization mechanisms, making it hard to know what exactly is being measured. In practice, the decision on which approach to pursue appears to not be consequential for the subsequent analyses, given the strong correlation between these measures.

times the average loss in consumer surplus. This counterfactual simulation cleanly illustrates a simple mechanism - subsidizing one group of consumers in a market with market power, all else equal, increases prices and decreases welfare for other consumers. Thus, subsidized consumers exert a negative demographic externality on other consumers.[20]

In the second exercise we simulate a scenario that is more likely to account for cross-sectional variation in prices across different markets at a point in time. In this scenario, we not only assign consumers with income of 400%+ FPL the subsidies of lower income consumers, but we also endow the higher income consumers with the marginal utility of income parameter of 151% FPL consumers. In other words, we make 400%+ FPL consumers look identical to 151% FPL consumers. This would be equivalent to moving from a county that had some fraction of unsubsidized consumers with 400%+ FPL to a county that had no 400% FPL consumers. Relative to the previous scenario, the effects are more nuanced. While the firms now face more subsidized consumers, which pushes prices up ("subsidy effect"), the firms also face much more elastic consumers, as we estimate a substantially higher marginal utility of income parameter for the lower consumers, which pushes prices down ("elasticity effect"). "Case B" in Figure 2 illustrates that the second effect dominates in our empirical setting (although this is not a general result; in general, as we show in Section 3.1, the direction of price change is ambiguous). In our context, moving to an environment with more subsidized, but highly elastic consumers, decreases list prices. As the dashed line marked with "B" in 2 illustrates, the annual average consumer-facing prices for consumers that are not directly affected by our simulation go down by $20-$30. This decline in premiums leads to an increase in consumer surplus among consumers with incomes between 150% and 400% FPL, whose subsidies or utility functions are not directly manipulated in the simulation. As Panel B in Table 3 again illustrates, the effect is highly heterogeneous across areas. In fact, while some areas experience a high increase in consumer surplus - in other words, here the "elasticity" effect dominates - other areas experience price increases and losses in consumer surplus, so the "subsidy" effect dominates. On average, however, in this simulation the lowest-income consumers exert a positive demographic externality on other consumers in the market.

We next examine how the economic forces of demographic externality can affect aggregate efficiency and equity properties of market allocations. In our counterfactual analyses in Section 5, we will simulate what happens when we replace means-tested with income-invariant

[20] Importantly, this effect is not special to insurance markets and does not depend on the fact that insurance contracts are pooling risks, which is different from the age-based externality from risk pooling as documented in Tebaldi (2017).

subsidies, which effectively changes the share of subsidized consumers across different markets.

5 Subsidy Design and Welfare

Efficiency of observed subsidy payments We begin our analysis by examining the economic costs and benefits of subsidy payments as observed in the data. To facilitate this computation we first re-simulate the allocation under observed subsidies in our model. The simulation serves two purposes. First, it allows us to establish a baseline that differences out any model simulation error. Second, it allows us to compute an allocation that shuts down the SLSP mechanism, while preserving observed subsidy levels. A baseline without the SLSP mechanism is useful, because in subsequent counterfactual analyses we will want to isolate the effects of subsidy levels and subsidy targeting without the SLSP mechanism, so a comparable baseline is required to interpret these counterfactual analyses.[21] As expected, the resulting simulated allocation is close to the allocation observed in the data.

Column (1), row (1) of Table 4 reports the surplus calculation results in this baseline scenario. We estimate that the baseline consumer surplus in the market amounts to $38 billion. This corresponds to $1,570 in annual surplus per capita, on average, among all potential consumers. Consumer surplus levels vary substantially across different areas of the country and across different socio-demographic groups. Figure A.4 in the Appendix illustrates the geographic variation in surplus that ranges from $621 in Brazos county of Texas to $6,260 in the McHenry County of Illinois. The socio-demographic variation is also substantial. Older consumers in poorest households enjoy the highest level of consumer surplus. Surplus falls monotonically with age, but is non-monotonic with respect to income. As the highest income consumers in this market are estimated to have the lowest marginal utility of income, consumer surplus declines with income for incomes up to 300% FPL, but then starts rising with income. Youngest consumers with household income in the 250-300% FPL range get the least amount of consumer surplus on this market. Consistent with lower elasticity of demand among higher income consumer, rows 14 to 26 of Column 1 suggest that consumers with income above the subsidization threshold of 400% FPL constitute a non-trivial share of enrollees that choose to buy insurance (26%). Enrollment is distributed

[21] We implement this simulation by assigning imputed marginal costs to SLSPs and letting all plans set prices according to the first order condition in 13, taking observed subsidy levels as given. In practice, this implies that SLSPs end up being the only plans for which prices change relative to observed prices, as for other plans we recover back observed prices from the first order conditions.

more evenly across age groups, with older beneficiaries accounting for a higher share of enrollment, which we would expect if *ceteris paribus* older individuals value insurance more. Surplus monotonically falling with age is consistent with the observation in Tebaldi (2017), arguing that the structure of ACA subsidies does not encourage enrollment of the youngest consumers; however, to the extent that the healthiest consumers are both younger and higher income, the subsidy structure does not fully discourage enrollment of the healthiest group.

Producer surplus amounts to $3.9 billion (Column 1, row 2). Under the observed allocation as simulated in our model, the government is spending $30 billion in premium and cost-sharing subsidies, which is broadly consistent with the subsidy spending reported by the Congressional Budget Office[22] Thus in total, we conclude that, nominally, the program that attracts about 41% of potential enrollees (row 17) generates $1.39 of surplus for $1 of nominal government spending, as the nominal government spending is $12 billion lower than the sum of consumer and producer surplus. In row (5) we add one additional piece of information - an estimate of how much the government would have likely spent on the same consumers had they not enrolled in ACA Marketplaces and instead, for example, used unsubsidized care. We estimate that the savings of public funds on the same set of consumers amounts to $16 billion, which implies that the net additional government outlays for premium and cost-sharing subsidies on the Marketplaces is ca. $14 billion. Taking into account this foregone spending and the cost of raising public funds, we arrive at a total return of $2.33 on a dollar of public funds under observed subsidy levels (row 9).

Incidence of observed subsidy payments Column 2 of Table 4 allows us to assess the incidence of observed subsidies. In this column, we characterize the allocation that would result in the absence of premium subsidies, holding other features of the market (such as market power and cost-sharing reduction subsidies) fixed. Without premium subsidies consumer surplus drops from $38 to $27 billion.

Table 4 reports the full set of counterfactual simulations. We start in Panel (a) with a counterfactual that removes subsidies, but preserves the market power in the market. The comparison of this counterfactual and the baseline surplus levels allows us to compute the incidence of premium subsidies between consumers and insurers. We estimate that adding (income-tagged) subsidies at the level observed in the data generates an addition surplus of $11 billion (from baseline of $10 billion) for $23B in (nominal, without opportunity cost)

[22]CBO reports $39 billion in net premium and cost-sharing subsidy spending for 2016, which includes spending in non-federally facilitated states (https://www.cbo.gov/sites/default/files/114th-congress-2015-2016/reports/51385-healthinsurancebaselineonecol.pdf)

government spending. This again highlights the negative return on the nominal government spending in the program. Importantly, however, subsidization significantly increases insurance coverage - the share of potential consumers who purchase any coverage increases from 11% to 37%. Out of $11 billion in additional total surplus, $4 billion accrues to insurers, suggesting a consumer incidence of subsidies of approximately 65%.

Efficiency cost of market power in the absence of subsidies The counterfactual simulations in columns (2) and (5) of Table 4 highlight the role of market power. The two columns compare allocations with and without market power when there are no subsidies. Removing market power increases total surplus by $6 billion from the baseline of $36 billion. In addition, $2 billion of surplus is re-allocated from insurers to consumers, thus increasing consumer surplus by a total of $5 billion and consumer enrollment from 14 to 20 percent. The effects of market power are heterogeneous across locations and result in a redistribution of surplus across the country.

Efficiency cost of subsidies in the absence of market power To further understand the contribution of subsidies, we pursue two more comparisons. In Panel (c) of Table 4 we report the results of two counterfactuals that simulate a classic public finance analysis. Assuming perfect competition, we ask how subsidies contribute to welfare. This comparison allows us to estimate the pure deadweight loss from subsidization. To implement the idea of perfect competition, we force insurers to price at the estimated marginal costs. Hence, in both counterfactuals, producer surplus is zero. The simulation suggests that at marginal cost pricing, subsidies generate an additional consumer surplus of only $11B for $32B extra in (nominal) government spending. Hence, two thirds of nominal government spending constitute a deadweight loss. However, the program achieves high enrollment rates. Without subsidies and marginal cost pricing our simulated enrollment is 30% of potential consumers. With (income-tagged) subsidies enrollment increases to 63 percent. The deadweight loss is not surprising in the presence of such large enrollment increase, as the marginal consumers attracted by increasingly generous subsidies have an increasingly declining willingness to pay for insurance.

Subsidy targeting in the presence of market power We next consider the central question of the paper, which is the effect of tagging subsidies to income in the presence of market power. Panel (b) of Table 4 compares the surplus generated in the program under tagged subsidies (at observed levels) vis-a-vis flat uniform subsidies that are equal to the average observed subsidy with tagging (this effectively holds government spending constant for the same number of enrollees). Our simulation suggests that removing tagging substantially increases the surplus and enrollment in the program. Moving to flat subsidies generates an additional \$ 6 billion in surplus for \$4 billion in additional government spending. 67% of the increase in surplus accrues to consumers. Insurance coverage increases substantially from 37% to 48%. Without accounting for the cost of public funds, with flat subsidies the program generates positive returns to \$1 in nominal government spending.

The allocation with flat subsidies, however, is not Pareto improving over the tagging mechanism. Flat subsidies generate a re-allocation of surplus from previously highly subsidized consumers to previously less subsidized consumers. As both panels in Figure 5 illustrate, flat subsidies lead to a re-allocation of surplus from several states in the Midwest, Northeast and Southwest to parts of Texas, Nebraska, Georgia, and North Carolina.

6 Conclusion

Traditionally, targeted benefits have been provided directly by the government. As a result, the vast majority of the literature has modelled the "supply" side in these settings as a benevolent social planner. Increasingly, however, governments continue funding social insurance and welfare programs, but relegate the actual provision of the benefits to private markets. In this paper we have argued that adding market power to the supply side of a public benefit in the presence of taxes or subsidies that are targeted on observables has the potential to generate change the efficiency and equity properties of allocations in unintended ways. The intuition is simple. If a firm knows the income composition of each market, it will adjust prices so as to take advantage of means-tested subsidies received by consumers. As long as the firm is required to set one price per market for its good, means-tested subsidies lead to a demographic externality that in an imperfectly competitive environment is not dissipated by competition.

References

Akerlof, G., "The Economics of "Tagging" as Applied to the Optimal Income Tax, Welfare Programs, and Manpower PlanPlan," *American Economic Review*, 1978, *68(1)*, 8–19.

Allcott, H., C. Knittle, and D. Taubinsky, "Tagging and Tageting of Energy Efficient Subsidies," *American Economic Review Papers and Proceedings*, 2015, *105(5)*, 187–191.

Berry, S. and P. Haile, "Identification in Differentiated Products Markets," *Annual Review of Economics*, 2016, *8*, 27–52.

Berry, Steve, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63(4)*, 841–889.

Cellini, S. and C. Goldin, "Does Federal Studen Aid Raise Tuition? New Evidence on For-Profit Colleges," *American Economic Journal: Economic Policy*, 2014, *6(4)*, 174–206.

Cicala, Steve, Ethan Lieber, and Victoria Marone, "Regulating Markups in U.S. Health Insurance," *American Economic Journal: Applied Economics*, forthcoming.

Currie, Janet and Firouz Gahvari, "Transfers in Cash and In-Kind: Theory Meets the Data," *Journal of Economic Literature*, 2008, *46(2)*, 333–383.

Curto, V., L. Einav, J. Levin, and J Bhattacharya, "Can Health Insurance Competition Work? Evidence from Medicare Advantage," *NBER Working Paper No. 20818*, 2015.

Cutler, David M. and Sarah J. Reber, "Paying For Health Insurance: The Trade-Off Between Competition And Adverse Selection," *Quarterly Journal of Economics*, 1998, *113(2)*, 433–466.

Decarolis, F., M. Polyakova, and S. Ryan, "Welfare Effects of Supply-Side Regulation in Medicare Part D," *NBER Working Paper No. 21298*, 2016.

Decarolis, Francesco, "Medicare Part D: Are Insurers Gaming the Low Income Subsidy Design?," *American Economic Review*, 2015, *105(4)*, 1547–1580.

_ , Maria Polyakova, and Stephen Patrick Ryan, "Subsidy Design in Privately-Provided Social Insurance: Lessons from Medicare Part D," *Journal of Political Economy*, forthcoming.

Enthoven, Alain C., "Reforming Medicare by Reforming Incentives," *New England Journal of Medicine*, 2011, *364(21).*

Fang, Hanming and Ami Ko, "Partial Rating Area Offering in the ACA Marketplaces: Facts, Theory and Evidence," Working Paper 25154, National Bureau of Economic Research October 2018.

Frakt, Austin, "The key difference between exchanges and the Ryan-Rivlin plan," http://theincidentaleconomist.com/wordpress/the-key-difference-between-exchanges-and-the-ryan-rivlin-plan/ March 2011.

Gruber, J. and E. Washington, "Subsidies to Employee Health Insurance Premiums and the Health Insurance Market," *Journal of Health Economics*, 2005, *24*, 253–276.

Hausman, J., *The Economics of New Goods*, University of Chicago Press,

Ho, Kate, Joseph Hogan, and F. Scott Morton, "The Impact of Consumer Inattention on Insurer Pricing in the Medicare Part D Program," *NBER Working Paper No. 21028*, 2015.

Jaffe, S. and M. Shepard, "Price-Linked Subsidies and Health Insurance Markups," *NBER Working Paper No. 23104*, 2017.

Lieber, E. and L. Lockwood, "Targeting with In-kind Transfers: Evidence from Medicaid Home Care," *University of Michigan Retirement Research Center Working Paper, WP 2017-359.*, 2017.

McFadden, D. and K. Train, "Mixed MNL Models for Discrete Response," *Journal of Applied Econometrics*, 2000, *15*, 447–470.

Rothstein, J., "Is the EITC as Goog as an NIT? Conditional Cash Transfers and Tax Incidence," *American Economic Journal: Economic Policy*, 2010, *2(1)*, 177–208.

Small, Kenneth A. and Harvey S. Rosen, "Applied Welfare Economics with Discrete Choice Models," *Econometrica*, 1981, *49(1)*, 105–130.

Tebaldi, P., "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," *University of Chicago Working Paper*, 2017.

Williams, H.C.W.L, "On the formation of travel demand models and economic evaluation measure of user benefit," *Environment and Planning*, 1977, *9(3)*, 285–344.

Wu, Yufei, "Supply Response to Consumer Inertia: Strategic Pricing in Medicare Part D," *MIT Unpublished Mimeo*, 2016.

Figure 1: Descriptive evidence: η and market power

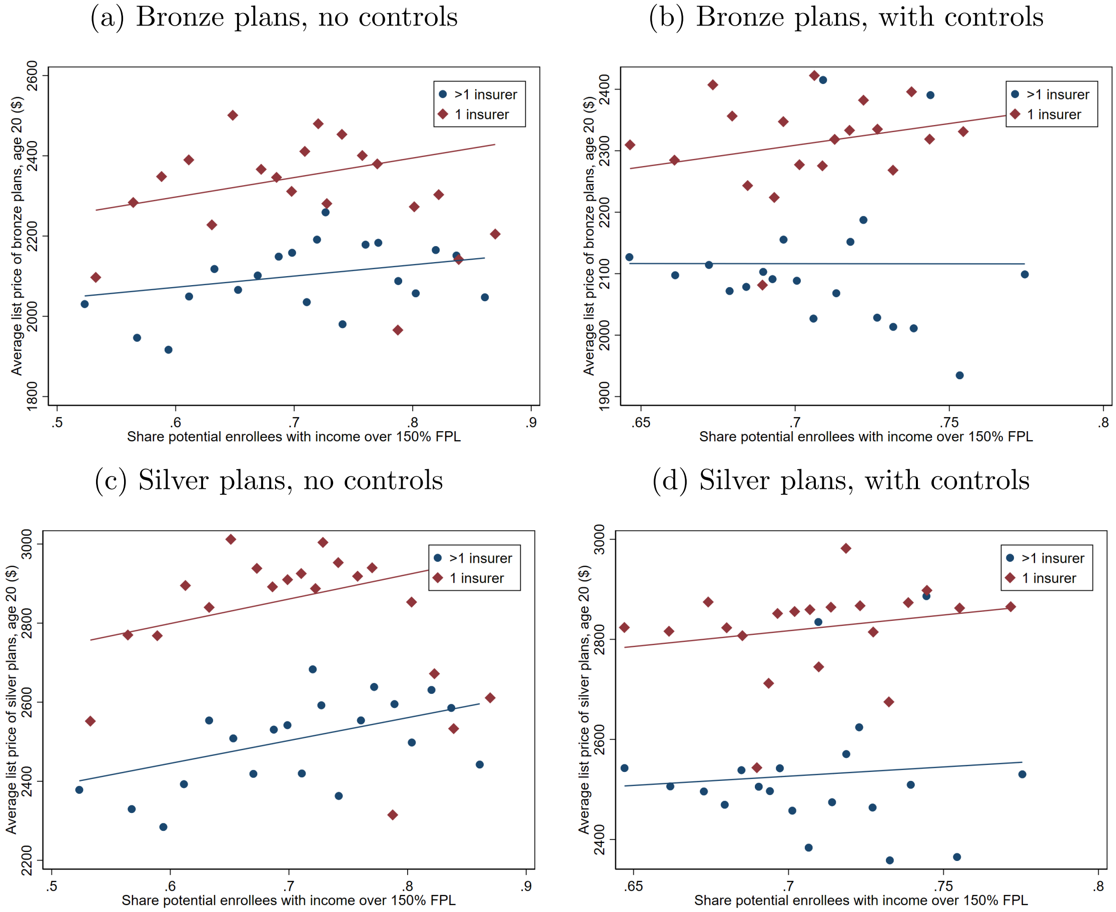

(a) Bronze plans, no controls

(b) Bronze plans, with controls

(c) Silver plans, no controls

(d) Silver plans, with controls

Notes: The Figures illustrates the relationship between the demographic composition of markets (counties) and list prices for bronze (Panels a, b) and silver plans (Panels c, d). The demographic composition is measured as a share of *potential* consumers with income above 150% FPL in a market (x-axis). These consumers are most likely to face a marginal change in the consumer-facing price in response to a marginal change in the a plan's list price and hence correspond to parameter η in the model of Section 3.1. The y-axis in each panel plots the average list price for a 20 year old consumer across all plans of the corresponding metal level in a county. The counties are aggregated in 20 equally-sized bins, so that each point represents the average list price, averaged across all counties in a bin. In each panel, prices are reported separately for counties that are characterized by a monopoly insurer (diamonds) versus more competitive counties with more than one insurer serving the market. The lines mark linear fits. Panels (b) and (d) include the following control variables: average age in the market, average income in the market, share of women, and share of white.

Figure 2: Demographic externality with targeted subsidies and market power

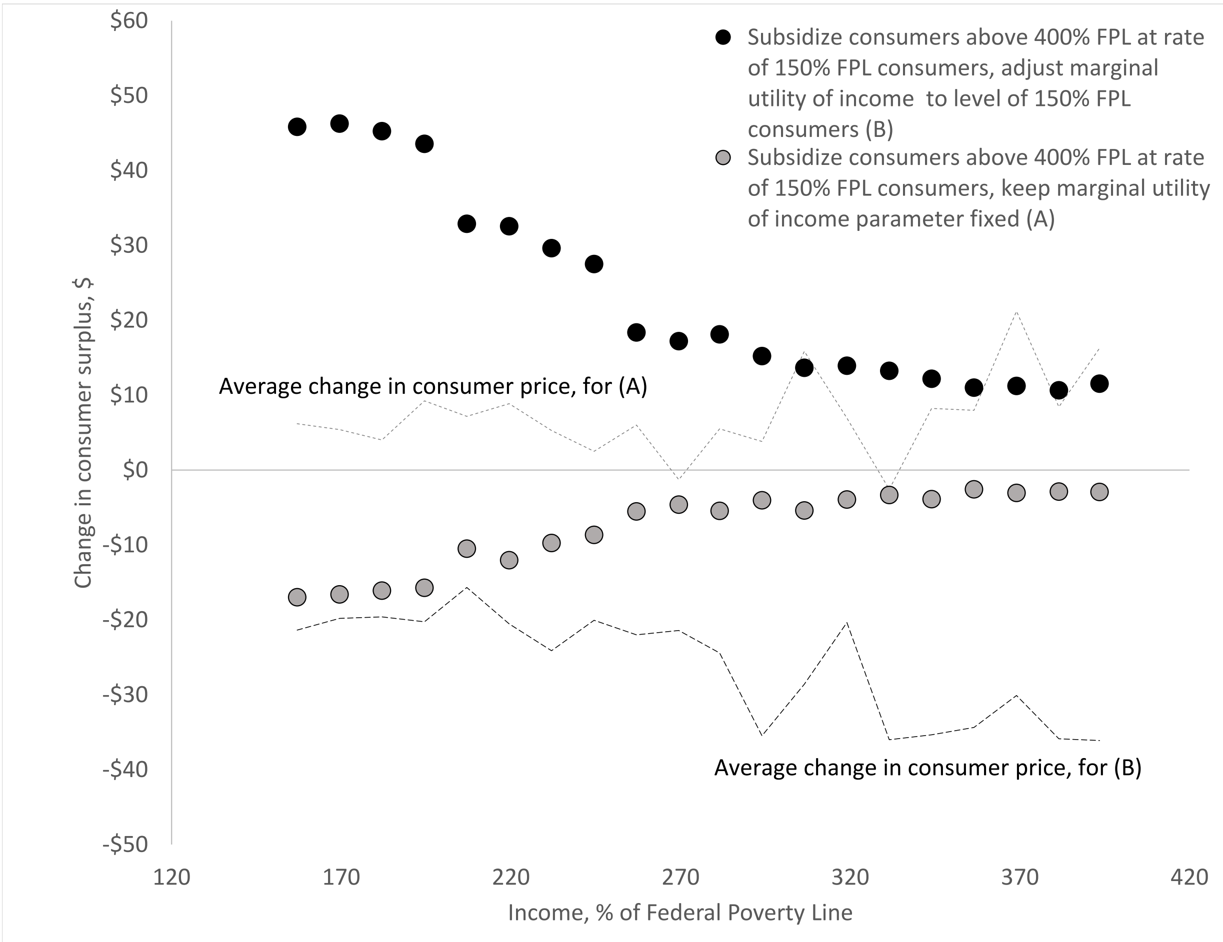

Notes: Figure reports estimated change in consumer surplus and consumer-facing prices (y-axis, in dollars) by income level (x-axis) in two counterfactual cases that capture the "demographic externality". The first counterfactual simulation (case A) changes income of consumers with true income of above 400% FPL to be 151% FPL. This change results in these consumers now receiving subsidies at the same rate as 151% FPL consumers. The counterfactual simulation holds everything else constant, including subsidies of other consumers and all utility function parameters, and allows firms to reprice their plans. Consumers with (true) income between 150% and 400% FPL are affected by price changes. As can be seen in the lighter dashed line, effective prices paid by consumers (that stay in the market) go up, while consumer surplus (grey circles) goes down. In another simulation - Case B - we additionally change the marginal utility of income parameter for consumers with true income above 400% FPL, assigning them the utility parameter of consumers with 151% FPL. Reverse price and consumer surplus patterns that are observed in this case are recorded in the darker dashed line (prices) and black circles (consumer surplus).

Figure 3: Geographic incidence of means-tested subsidies

(a) Under perfect competition

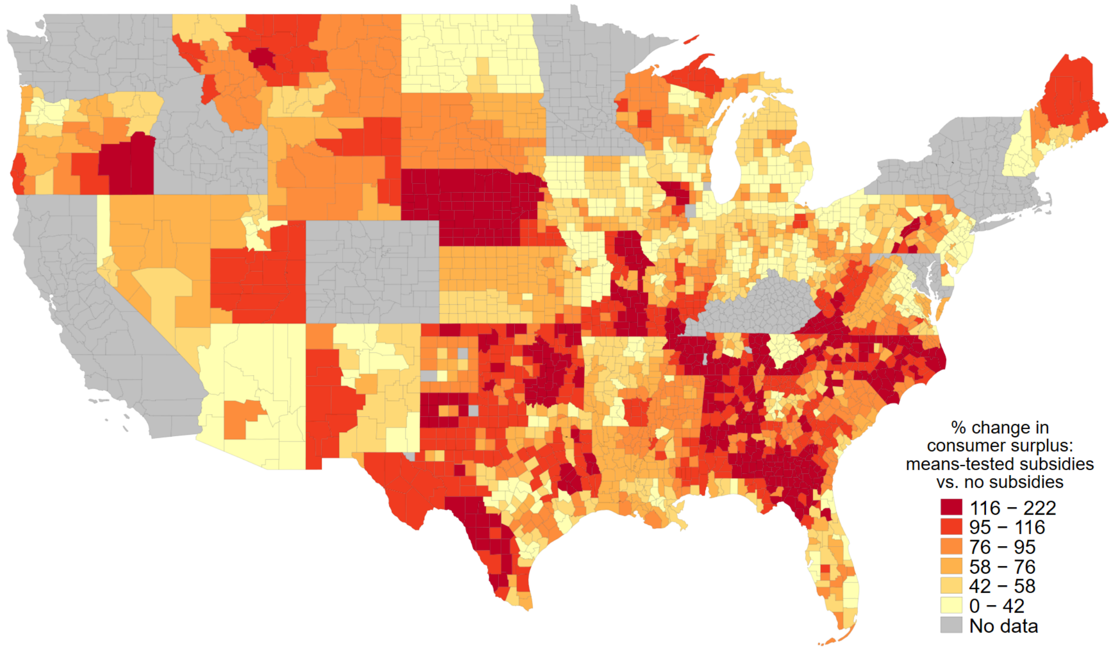

(b) Under imperfect competition

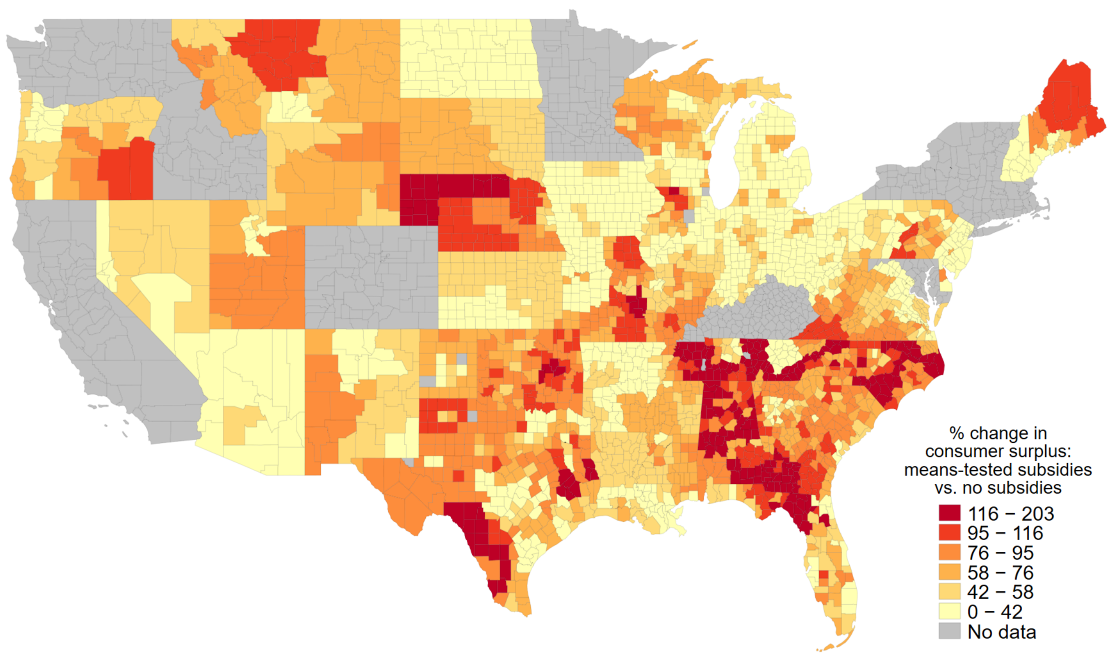

Notes: Figure reports the percent change in average consumer surplus for moving consumers from an environment with no subsidies to means-tested subsidies, as observed in the data. (Panel a) forces firms to price at the marginal cost. In (Panel b), we allow insurers to adjusted prices, re-simulating the observed market environment. The latter simulation shuts down the second-lowest silver plan part of the observed algorithm, in order to make the simulation comparable to counterfactual subsidy mechanisms.

Figure 4: Flat subsidies vs. means-tested subsidies: efficiency effects

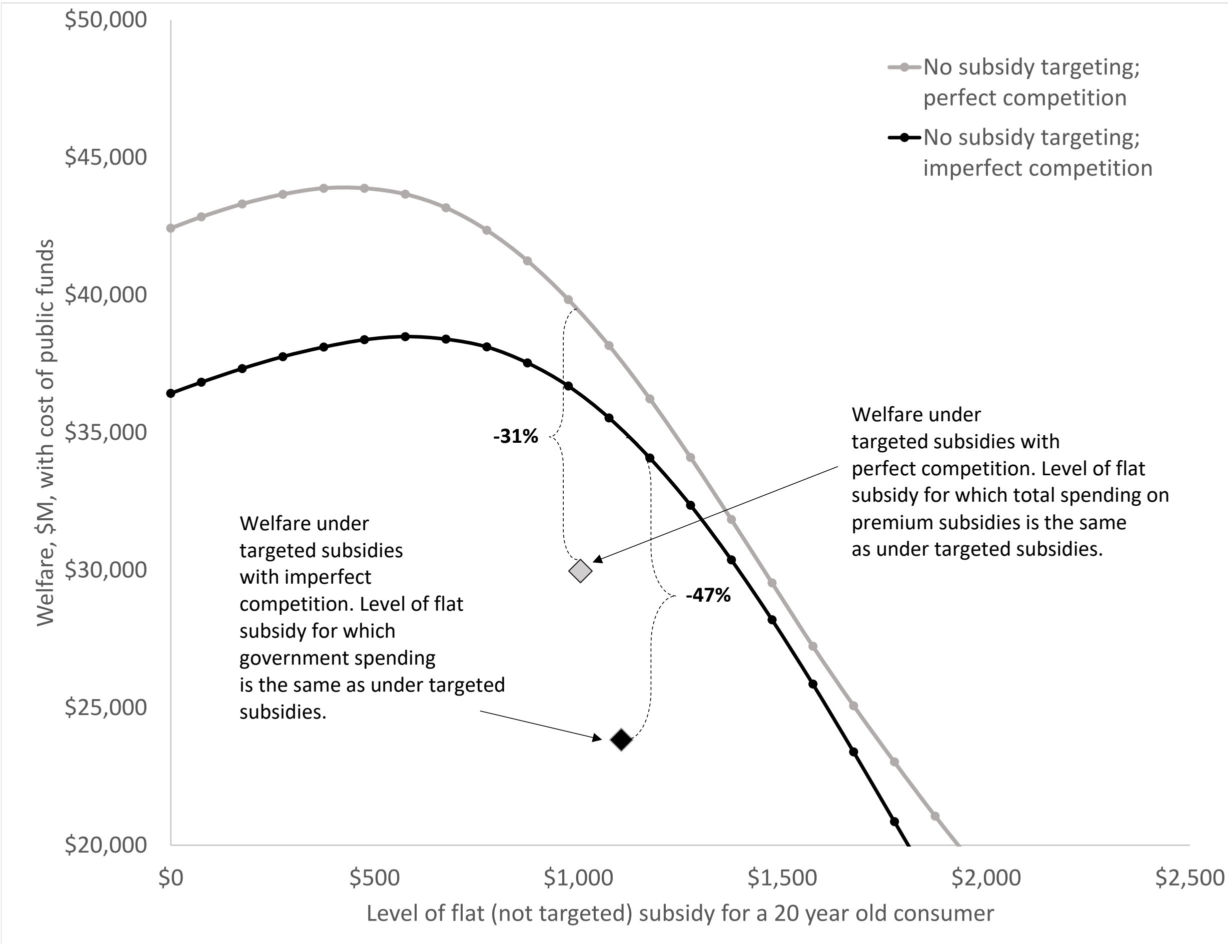

Notes: Figure reports estimated total welfare (y-axis), including the opportunity cost of government spending and the cost of raising public funds, under subsidies that do not vary with consumer income *within* an age group. The x-axis marks the value of the flat voucher that is offered to a 20-year old consumer. Consumers of all other ages receive vouchers that are scaled by the premium age curve as observed in the data. The grey line marks welfare estimates in counterfactuals that shut down market power in the market, setting baseline (for a 20 year olds) premiums to be equal to estimated marginal costs. The black line marks cases that preserve market power. Two diamonds mark the two levels of baseline (for a 20 year old) flat subsidies that lead to the same level of nominal government spending (premium and cost-sharing reduction subsidies combined) as under means-tested subsidies, in the perfectly competitive (grey diamond) and the imperfectly competitive (black diamond) environment. The y-axis value for both points record the level of total welfare achieved under means-tested subsidies. The values correspond to row (7), Columns (1) and (4) in Table 4. Curly brackets show the difference in welfare that is achieved - for the same level of government spending - between means-tested and flat (age-adjusted) subsidies. We observe that the welfare loss from subsidy means-testing is higher in the presence of market power.

Figure 5: Flat subsidies vs. means-tested subsidies: distributional effects

(a) Perfect Competition

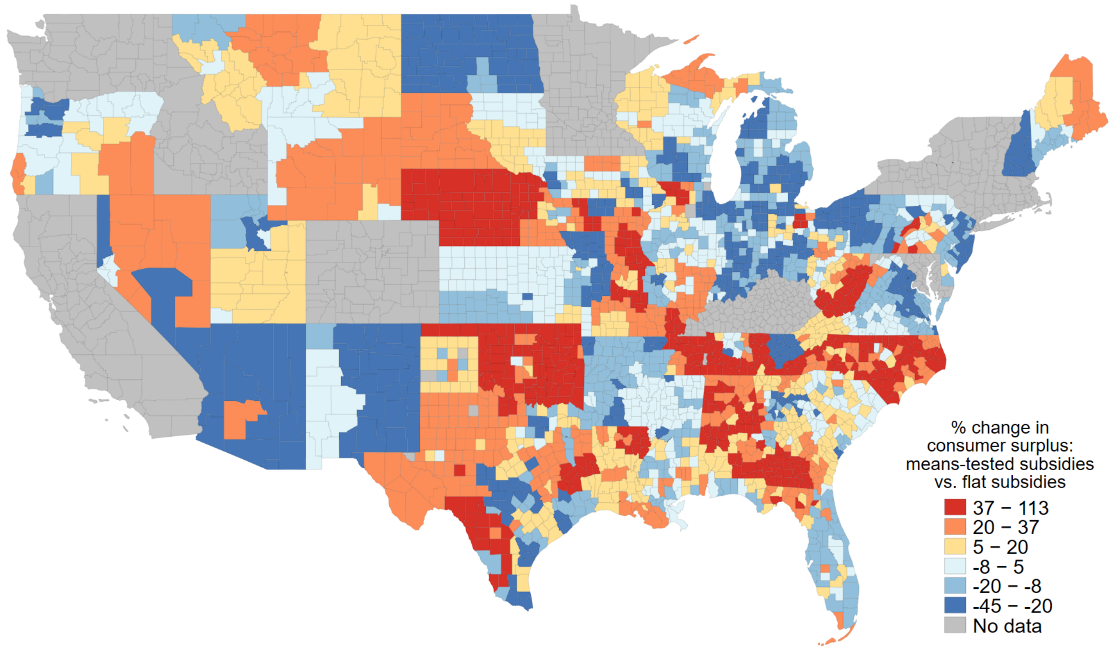

(b) Imperfect Competition

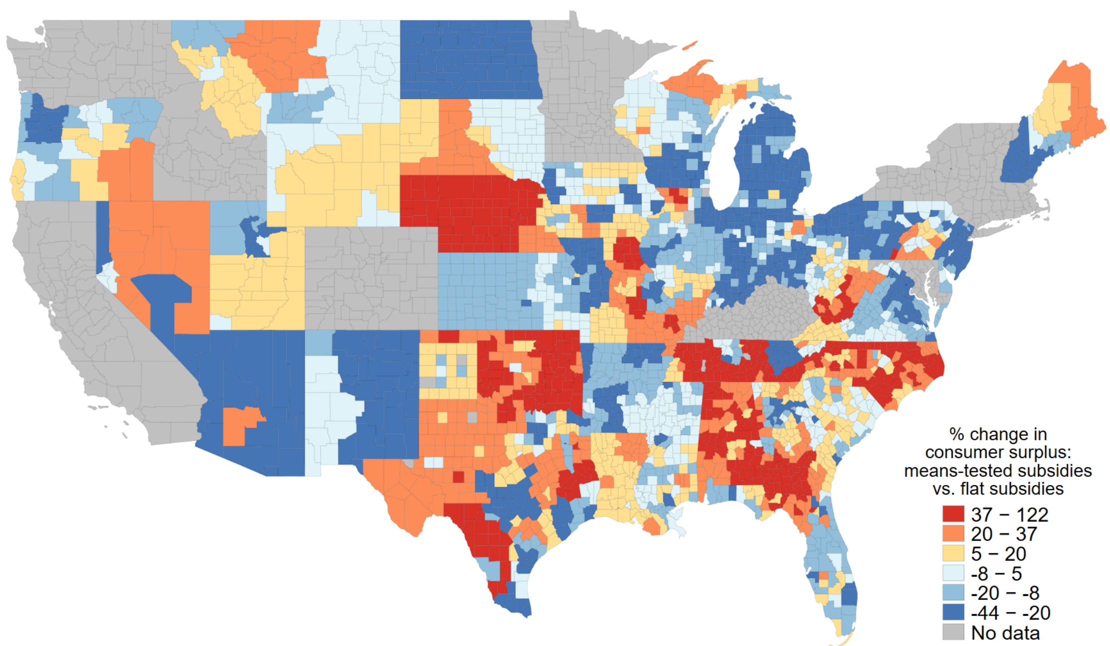

Notes: Figure reports the percent change in average consumer surplus from moving consumers from means-tested to flat (age-adjusted) subsidies, preserving total nominal government spending on subsidies. The baseline flat subsidy for a 20-year old that results in the same government spending as under means-tested subsidies is $1,005 under perfect competition (Panel a) and $1,105 under imperfect competition (Panel b).

Table 1: Summary statistics

	Mean‡ (1)	Std.Dev. (2)	10th pctile (3)	90th pctile (4)
A. Choice set				
(1) Number of plans	21	13	8	37
(2) Number of insurers	2.2	1.1	1	4
(3) Average annual premium (age 40), $	5,106	902	3,978	6,351
B. Enrollment				
(1) Market size‡‡	7,867	25,756	479	15,671
(2) Share outside option	0.6	0.2	0.4	0.8
(3) Share bronze plans	0.09	0.05	0.04	0.2
(4) Share silver plans	0.3	0.1	0.2	0.4
(5) Share gold plans	0.01	0.02	0	0.03
(6) Market-level enrollment	3,536	13,798	168	6,411
(7) Plan-level enrollment^	3,165	12,040	39	6,353
C. ACS Sample of Potential Consumers				
(1) Age	40	2.5	37	43
(2) Share women	0.5	0.04	0.4	0.5
(3) Share white	0.9	0.1	0.7	1.0
(4) Income in % FPL	262	36	212	309
(5) Annual premium subsidy, $^^	3,301	1,293	1,791	4,988

‡ Across counties

‡‡ Based on Kaiser Family Foundation estimates

^ Mean, Std. Dev., 10th and 90th percentiles for plan enrollment are reported across plans, not across counties

^^ Reports average individual-level subsidy, which is computed as the average subsidy within a coverage family

Notes: Panels A and B report the distribution of choices and enrollment in federally-facilitated ACA Marketplaces in year 2017. Choice set statistics (**Panel A**) are based on data from Health Insurance Marketplace Public Use Files, released by the Center for Medicare and Medicaid Services as well as the Center for Consumer Information and Insurance Oversight. Enrollment statistics (**Panel B**) are based on county and plan-level enrollment data released by the Center for Medicare and Medicaid Services. Demographic data in **Panel C** are based on the public use sample of the American Community Survey for year 2017. Potential enrollees in the ACS sample were defined as individuals who did not have active employer-sponsored insurance, were not enrolled in any type of public health insurance coverage, and were not eligible for insurance under Medicaid expansion in those states that expanded Medicaid. Annual premium subsidies were imputed using the ACS records of income and tax family composition following instructions for 2017 IRS Form 8962 (Premium Tax Credit).

Table 2: Model Estimates

	Mean	Consumer type - age dimension		
		Age<25	Age 25-40	Age >40
	(1)	(2)	(3)	(4)
Panel A: Parameters of utility function				
Coefficient on premium, $000 (α)				
Income FPL <200	-	-4.75	-1.47	-2.33
	-	(0.30)	(0.10)	(0.20)
Income FPL > 200 and FPL < 400	-	-4.71	-0.98	-2.94
	-	(0.33)	(0.06)	(0.22)
Income FPL > 400	-	-1.68	-0.33	-0.41
	-	(0.11)	(0.03)	(0.19)
Age-specific intercepts	-	1.39	-2.40	0.00
	-	(0.10)	(0.16)	0.00
Actuarial Value	16.45			
	(1.07)			
Average plan-level utility (plan fixed effects; 2,851 plans)	-11.25			
Std. Dev plan-level utility (plan fixed effects; 2,851 plans)	3.00			
Panel B: Marginal cost projection				
Actuarial value	1.90			
	(0.10)			
PPO	0.18			
	(0.01)			
Catastrophic	-0.0045			
	(0.09)			
Bronze	-0.80			
	(0.03)			
Silver	-0.58			
	(0.02)			
Gold	-0.28			
	(0.01)			
Platinum	reference			
New plan	-0.023			
	(0.00)			
Mean dependent variable (inverted MC, $000)	1.94			
Standard deviation dependent variable	0.54			
R-squared	0.83			
N	49222			

Notes: **Panel A** reports non-linear least squares parameter estimates for the demand model described in Section 4.1. The NLLS objective function minimizes the squared distance between estimated and real age- and income-specific enrollments in each market. Bootstrapped standard errors are reported in parantheses. The model includes, but does not report an intercept for consumers with income below 100% FPL. **Panel B** reports the results of a hedonic regression that projects marginal cost estimates - obtained via the inversion of the first order condition - on plan characteristics, for plans other than the second lowest silver plan. The model includes, but does not report: indicators for plan's HSA eligiblity, out of network and out of country coverage, presence of a national network; measures of quantity limits and coverage exclusions; indicators for the requirement of pregnancy notices, referrals to specialits, presence of a wellness program, offers of chronic condition and pregnancy management; indicators for coverage of most common services; insurer fixed effects.

Table 3: Heterogeneous impacts of demographic externality

	State	County	Change in consumer surplus, $	State	County	Change in consumer surplus, $
	(1)	(2)	(3)	(4)	(5)	(6)
	Panel A: Subsidize consumers above 400% FPL at rate of 150% FPL consumers, keep marginal utility of income parameter fixed					
	Negative impact, top 10[^]			Negative impact, bottom 10		
(1)	GA	Walton County	-0.016	PA	Adams County	-145.9
(2)	MO	Linn County	-0.011	NE	Keya Paha County	-131.6
(3)	MO	Sullivan County	-0.011	NE	Rock County	-131.6
(4)	MO	Putnam County	-0.011	NE	Holt County	-131.6
(5)	OH	Richland County	-0.006	NE	Boyd County	-131.6
(6)	FL	Liberty County	-0.005	PA	York County	-130.8
(7)	OH	Crawford County	-0.005	NE	Antelope County	-130.6
(8)	WV	Wyoming County	0.000	NE	Madison County	-130.6
(9)	WV	Mingo County	0.000	NE	Colfax County	-130.6
(10)	WV	McDowell County	0.000	NE	Platte County	-130.6
	Panel B: Subsidize consumers above 400% FPL at rate of 150% FPL consumers, adjust marginal utility of income to 150% FPL consumers					
	Positive impact, top 10			Negative impact, bottom 10		
(1)	WI	Langlade County	236.1	WV	Calhoun County	-4.3
(2)	WI	Waushara County	228.4	WV	Roane County	-4.3
(3)	WI	Iron County	209.7	WV	Clay County	-3.0
(4)	TX	San Jacinto County	185.7	WV	Fayette County	-2.4
(5)	WI	Ozaukee County	185.6	OH	Hocking County	-1.9
(6)	WI	Washington County	185.6	WV	Raleigh County	-1.4
(7)	FL	Palm Beach County	185.4	WY	Natrona County	-1.0
(8)	WI	Waukesha County	182.1	AL	Etowah County	-0.5
(9)	FL	Pinellas County	172.1	AZ	Pima County	0.0
(1)	ME	Oxford County	169.9	SC	Oconee County	0.0

[^] Excludes 3,864 counties that had no change in average consumer surplus.

Notes: Table reports the change in average consumer surplus (in dollars) among consumers with incomes between 150% and 400% FPL, for most and least affected counties, for two counterfactual simulations that capture the "demographic externality". The first counterfactual simulation (**Panel A**) changes income of consumers with true income of above 400% FPL to be 151% FPL. This change results in these consumers now receiving subsidies at the same rate as 151% FPL consumers. The counterfactual simulation holds everything else constant, including subsidies of other consumers and all utility function parameters, and allows firms to reprice their plans. Consumers with (true) income between 150% and 400% FPL are affected by price changes. Simualtion reported in **Panel B** additionally changes the marginal utility of income parameter for consumers with true income above 400% FPL, assigning them the utility parameter of consumers with 151% FPL. The average effects on consumer surplus and prices for all counties in the US in these two counterfactual scenarios are reported in Figure 2.

Table 4: Surplus under counterfactual subsidy mechanisms

		Preserve market power			No market power			Demographic externality	
		Simulated baseline with targeted subsidies^	No premium subsidies	Flat subsidies with G ≈ G in (1) ($1,105 subsidy for a 20 year-old)~	Targeted premium subsidies	No premium subsidies	Flat subsidies with G ≈ G in (4) ($1,005 subsidy for a 20 year-old consumer)~	Subsidize consumers above 400% FPL at rate of 150% FPL consumers	(7) and change α_{400} to α_{150}
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1)	Consumer surplus, $M	37,892	26,382	45,634	47,133	31,300	53,822	49,524	30,863
(2)	Insurer profit, $M	3,861	1,703	4,291	-	-	-	-	4,418
(3)	**Consumer and producer surplus, $M**	41,753	28,085	49,925	47,133	31,300	53,822	49,524	35,281
(4)	Public spending on premium and cost-sharing subsidies **(G)**, $M	30,040	727	30,156	33,406	1,374	33,372	44,334	53,301
(5)	Pubilc savings on uncompensated care for un(under-)insured, $M	16,254	7,143	18,772	20,200	9,933	22,241	18,865	20,365
(6)	**Net government spending, $M**	13,786	(6,416)	11,384	13,206	(8,559)	11,131	25,469	32,936
(7)	Total surplus, including the cost of public funds, $M	23,831	36,426	35,126	29,965	42,427	39,351	16,414	(7,536)
(8)	Return on a dollar of nominal public spending, $	1.39	38.63	1.66	1.41	22.78	1.61	1.12	0.66
(9)	Return on a dollar of adjusted public spending $^^	2.33	-	3.37	2.75	-	3.72	1.50	0.82
(10)	**Characteristics of the allocation**								
(11)	Inside option enrollment, '000	8,234	2,750	8,163	10,469	4,008	10,433	9,894	10,747
(12)	Inside option enrollment, percent of total market	40.9	13.7	40.5	52.0	19.9	51.8	49.1	53.3
(13)	Composition of inside share enrollment by consumer type								
(14)	Income 0-100% FPL	1.1%	0.0%	0.1%	1.6%	0.0%	0.1%	0.9%	0.9%
(15)	Income 100-150% FPL	24.9%	7.7%	13.7%	20.1%	10.7%	14.7%	20.7%	19.1%
(16)	Income 150-200% FPL	21.3%	4.5%	11.3%	19.2%	7.1%	12.9%	17.7%	16.5%
(17)	Income 200-250% FPL	13.3%	3.9%	9.7%	14.1%	5.3%	10.9%	11.0%	10.4%
(18)	Income 250-300% FPL	6.5%	2.1%	7.6%	8.7%	3.1%	8.6%	5.4%	5.1%
(19)	Income 300-400% FPL	6.4%	3.3%	12.1%	9.9%	4.8%	13.7%	5.2%	5.1%
(20)	Income 400% FPL and above	26.4%	78.5%	45.4%	26.4%	69.0%	39.1%	39.1%	42.8%
(21)	Age 0-18	6.6%	12.2%	13.7%	8.1%	12.0%	13.0%	8.3%	8.8%
(22)	Age 19-26	10.9%	6.4%	11.6%	11.4%	7.3%	12.1%	10.5%	10.5%
(23)	Age 27-35	14.1%	14.4%	13.5%	13.9%	14.7%	13.7%	12.8%	13.5%
(24)	Age 36-45	16.7%	19.2%	17.1%	16.2%	19.3%	16.6%	15.4%	15.5%
(25)	Age 46-55	23.4%	25.5%	22.6%	22.5%	25.0%	21.9%	22.8%	22.2%
(26)	Age 55-64	28.4%	22.3%	21.6%	27.9%	21.7%	22.8%	30.2%	29.4%
(27)	Average unweighted list premium for a 20 year old consumer, $	2,410	2,425	2,422	1,940	1,940	1,940	2,417	2,332
(28)	Average premium paid, $	977	2,291	1,100	732	1,872	830	678	431
(29)	Average subsidy paid, $	3,732	-	2,653	3,271	-	2,442	4,671	5,124
(30)	Average consumer surplus, $	1,570	999	1,643	1,991	1,173	2,000	2,148	1,410

^ Simulation shuts down the second-lowest Silver plan part of the observed mechanism to allow comparability with counterfactual subsidization policies

~ Subsidies are flat within an age group. Baseline flat subsidy for a 20-year old consumer are adjusted for age using the age curve from the observed allocation

^^ Accounts for opportunity cost of public funds in the form of uncompensated care and the cost of raising public funds (λ, assumed to be 30 cents on a dollar)

Notes: Table reports the levels of consumer surplus, producer surplus, government spending, and total welfare under the observed allocation (column 1) and under counterfactual allocations (columns 2 to 8). We compute these objects using estimates of demand and marginal costs, and simulated market equilibria that allocate consumers to Marketplace insurance plans or the outside option. All simulations are performed within the ACS sample of consumers and are then scaled to the total market size. Consumer surplus is computed as discussed in Section 5. Firm profits reported in row (2) account for risk-equalization programs indirectly, as marginal cost estimates are "pricing-relevant" marginal costs, i.e. net of firms' expectations about positive or negative risk-equalization transfers. Cost-sharing reduction (CSR) subsidies in row (4) are computed by multiplying consumer-type specific average CSR values as reported by CMS for 2016 by enrollment share of each consumer type ($1,440 per year for conusmers with income under 150% FPL, $1,068 for those with income between 150% and 20)% FPL, and $144 for consumers with income between 200% and 250% FPL. Uncompensated care spending is computed at the rate of $1,827, following the Kaiser Family Foundation 2013 report on public spending on uncompensated care for the uninsured. Rows 11-27 describe consumer sorting for each allocation. Negative quantities are reported in parentheses.

Table 5: Demographic incidence of subsidies

Change in consumer surplus, %	Age<25 (1)	Age 25-40 (2)	Age >40 (3)	Age<25 (4)	Age 25-40 (5)	Age >40 (6)
	Panel A: Moving from no subsidies to means-tested subsidies					
	Under perfect competition			**Under imperfect competition**		
Income FPL <200	422.4	200.0	380.2	317.3	179.2	351.1
Income FPL > 200 and FPL < 400	148.6	200.0	203.6	58.6	179.2	126.3
Income FPL > 400	0.0	0.0	0.0	0.5	0.3	0.6
	Panel B: Moving from means-tested subsidies to flat subsidies^					
	Under perfect competition			**Under imperfect competition**		
Income FPL <200	-46.3	-37.6	-52.7	-48.4	-38.9	-57.7
Income FPL > 200 and FPL < 400	34.2	-37.6	-26.4	59.2	-38.9	-16.9
Income FPL > 400	108.1	33.9	52.1	127.0	35.9	58.4

^ Flat subsidies, adjusted for age, such that government spending is the same as under means-tested subsidies.

Notes: Table reports the percent change in average consumer surplus, by consumer type, for a set of allocations under counterfactual subsidization policies. **Panel A** reports how consumer surplus changes when we compare observed, means-tested subsidy regime with a regime where no consumers receive premium subsidies (the cost-sharing reduction subsidies are kept fixed). In the panels marked "under perfect competition," insurers are assumed to price at their marginal costs. In the panels market "under imperfect competition," insurers choose prices taking into account the subsidy structure and consumer demand. **Panel B** reports the change in consumer surplus that we simulate when moving the market from means-tested to flat (but age-adjusted) subsidies. The baseline flat subsidy for a 20-year old that results in the same nominal government spending (on premium and cost-sharing reduction subsidies), as under means-tested subsidies is $1,005 under perfect competition and $1,105 under imperfect competition. Consumer surplus is computed as discussed in Section 5. Surplus does not vary with consumer choices, but only depends on the set of available products that is held fixed and consumer-facing product prices. Thus, any changes in consumer surplus reported in this table result from price changes, either purely due to subsidy, under perfect competition, or due to subsidy and firms' adjustment of list prices, under imperfect competition.

A ONLINE APPENDIX

Figure A.1: Consumer interface on healthcare.gov

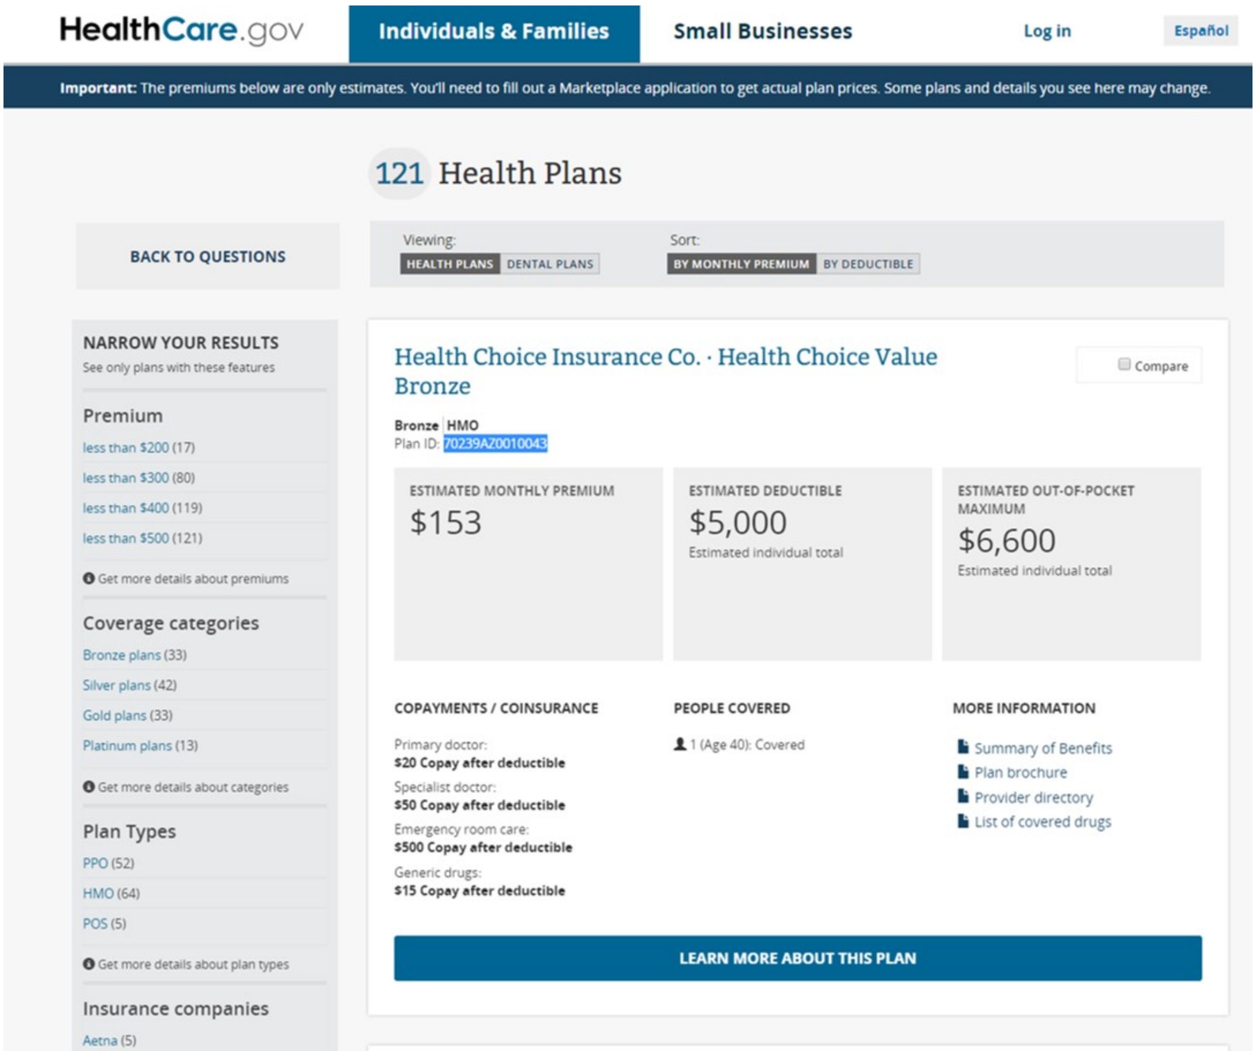

Notes: Snapshot of healthcare.gov for one of 121 plans that were offered to 40-year old individuals in Cook County, IL in 2015. The premium that individuals see on the web page incorporates their individual premium subsidy if they enter their income information during the selection process.

Figure A.2: Demand model: empirical moments and model fit

(a) Silver plan market share

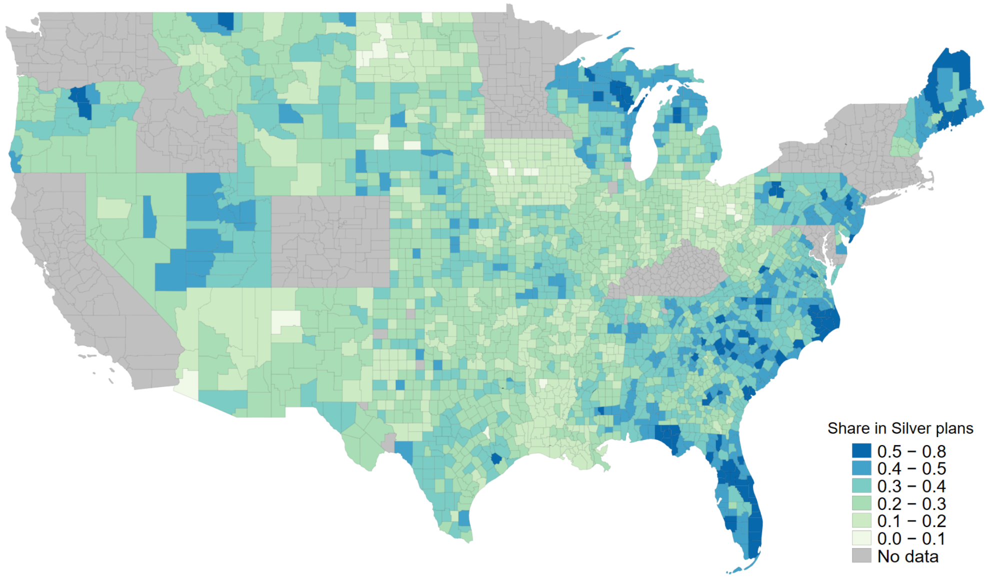

(b) Model in-sample prediction error of Silver plan market share

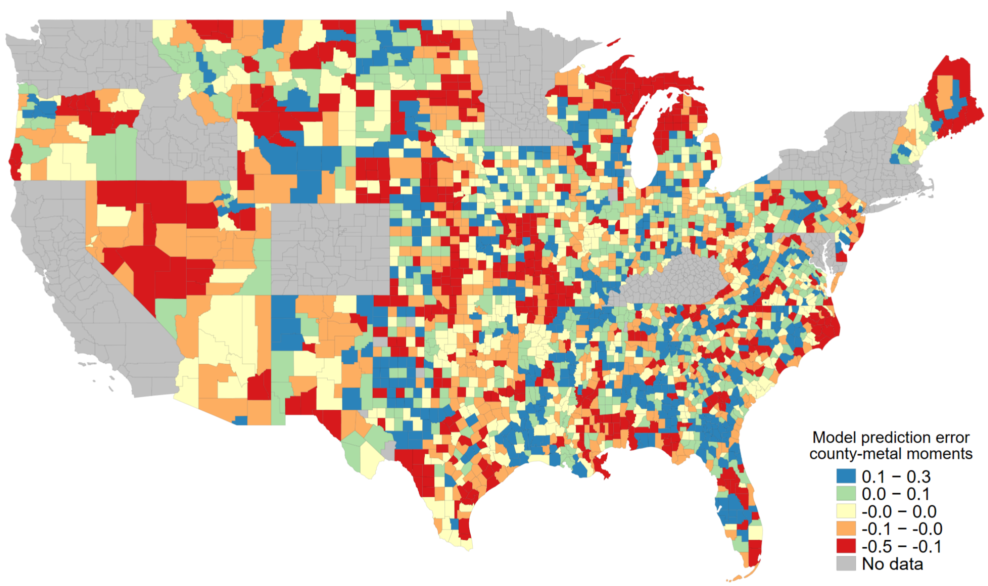

Notes: Map in panel (a) plots the share of potential consumers in each county that enrolled in a Silver plan on ACA Marketplaces. States that are marked with grey are not federally facilitated and do not enter our analysis. The counts of the pool of potential consumers (denominator) was provided by the Kaiser Family Foundation and is based on estimates from national surveys of how many people were uninsured or underinsured in each geographic region. The number of people that purchased a Silver plan (numerator) are administrative enrollee counts reported by CMS that do not account for disenrollments. Data is for year 2015. Map in panel (b) plots the difference between the observed share of enrollees in Silver plans - as pictured in Panel (a) - and the share of enrollment in Silver plans as predicted by demand model of Section 4.1.

Figure A.3: Marginal cost estimates

(a) Distribution of marginal cost estimates

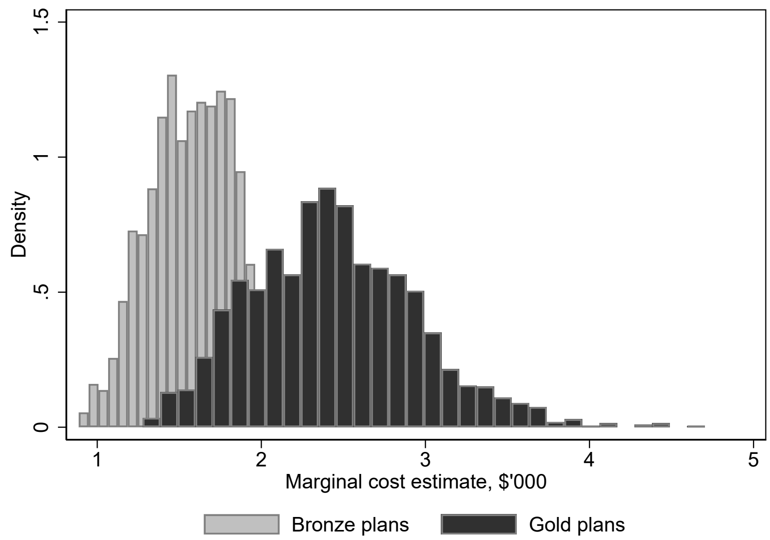

(b) Marginal cost estimates and observed accounting cost

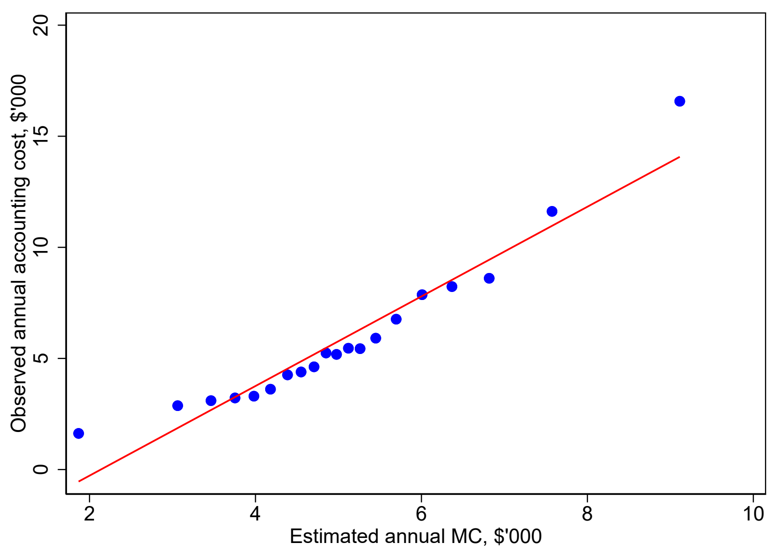

Notes: Panel (a) plots the distribution of plan-market level marginal costs as estimated in Section 4.2. The costs are plotted for a baseline, age 20, consumer. The costs are plotted separately for Bronze plans that provide the lowest amount of coverage and Gold plans that provide the highest amount of coverage for most consumers (excluding rare Catastrophic and Platinum plans). Panel (b) plots the correspondence between average estimated marginal cost (plan-market level costs were aggregated to plan-level) in each plan (x-axis) and plan-level accounting costs reported by CMS (y-axis).

Figure A.4: Geographic variation in consumer surplus under observed subsidies

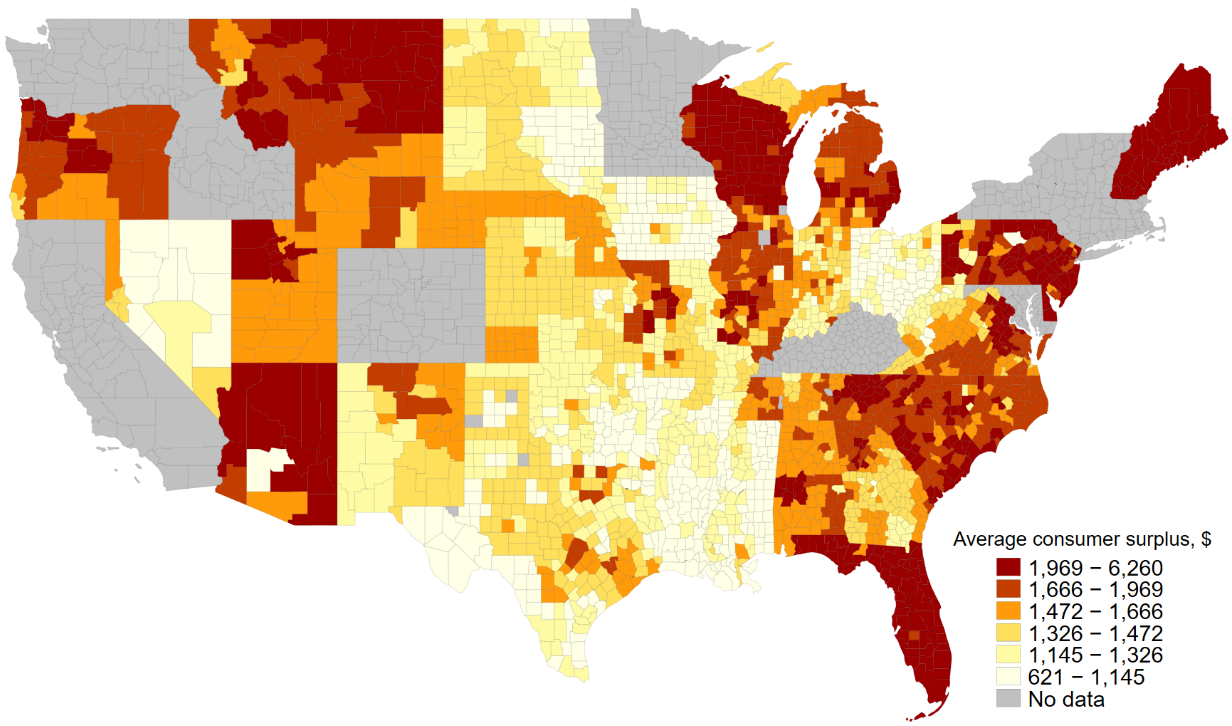

Notes: Map plots average consumer surplus per county, estimated under the observed subsidy regime.

Table 6: Surplus under counterfactual subsidy mechanisms, no MLR constraint

		Preserve market power			No market power			Demographic externality	
		Simulated baseline with targeted subsidies^	No premium subsidies	Flat subsidies with G ≈ G in (1) ($1,105 subsidy for a 20 year-old)~	Targeted premium subsidies	No premium subsidies	Flat subsidies with G ≈ G in (4) ($1,005 subsidy for a 20 year-old consumer)~	Subsidize consumers above 400% FPL at rate of 150% FPL consumers	(7) and change α_{400} to α_{150}
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1)	Consumer surplus, $M	37,892	26,382	45,634	47,133	31,300	53,822	49,524	30,863
(2)	Insurer profit, $M	3,861	1,703	4,291	-	-	-	-	4,418
(3)	**Consumer and producer surplus, $M**	41,753	28,085	49,925	47,133	31,300	53,822	49,524	35,281
(4)	Public spending on premium and cost-sharing subsidies **(G)**, $M	30,040	727	30,156	33,406	1,374	33,372	44,334	53,301
(5)	Pubilc savings on uncompensated care for un(under-)insured, $M	16,254	7,143	18,772	20,200	9,933	22,241	18,865	20,365
(6)	**Net government spending, $M**	13,786	(6,416)	11,384	13,206	(8,559)	11,131	25,469	32,936
(7)	Total surplus, including the cost of public funds, $M	23,831	36,426	35,126	29,965	42,427	39,351	16,414	(7,536)
(8)	Return on a dollar of nominal public spending, $	1.39	38.63	1.66	1.41	22.78	1.61	1.12	0.66
(9)	Return on a dollar of adjusted public spending $^^	2.33	-	3.37	2.75	-	3.72	1.50	0.82
(10)	**Characteristics of the allocation**								
(11)	Inside option enrollment, '000	8,234	2,750	8,163	10,469	4,008	10,433	9,894	10,747
(12)	Inside option enrollment, percent of total market	40.9	13.7	40.5	52.0	19.9	51.8	49.1	53.3
(13)	Composition of inside share enrollment by consumer type								
(14)	Income 0-100% FPL	1.1%	0.0%	0.1%	1.6%	0.0%	0.1%	0.9%	0.9%
(15)	Income 100-150% FPL	24.9%	7.7%	13.7%	20.1%	10.7%	14.7%	20.7%	19.1%
(16)	Income 150-200% FPL	21.3%	4.5%	11.3%	19.2%	7.1%	12.9%	17.7%	16.5%
(17)	Income 200-250% FPL	13.3%	3.9%	9.7%	14.1%	5.3%	10.9%	11.0%	10.4%
(18)	Income 250-300% FPL	6.5%	2.1%	7.6%	8.7%	3.1%	8.6%	5.4%	5.1%
(19)	Income 300-400% FPL	6.4%	3.3%	12.1%	9.9%	4.8%	13.7%	5.2%	5.1%
(20)	Income 400% FPL and above	26.4%	78.5%	45.4%	26.4%	69.0%	39.1%	39.1%	42.8%
(21)	Age 0-18	6.6%	12.2%	13.7%	8.1%	12.0%	13.0%	8.3%	8.8%
(22)	Age 19-26	10.9%	6.4%	11.6%	11.4%	7.3%	12.1%	10.5%	10.5%
(23)	Age 27-35	14.1%	14.4%	13.5%	13.9%	14.7%	13.7%	12.8%	13.5%
(24)	Age 36-45	16.7%	19.2%	17.1%	16.2%	19.3%	16.6%	15.4%	15.5%
(25)	Age 46-55	23.4%	25.5%	22.6%	22.5%	25.0%	21.9%	22.8%	22.2%
(26)	Age 55-64	28.4%	22.3%	21.6%	27.9%	21.7%	22.8%	30.2%	29.4%
(27)	Average unweighted list premium for a 20 year old consumer, $	2,410	2,425	2,422	1,940	1,940	1,940	2,417	2,332
(28)	Average premium paid, $	977	2,291	1,100	732	1,872	830	678	431
(29)	Average subsidy paid, $	3,732	-	2,653	3,271	-	2,442	4,671	5,124
(30)	Average consumer surplus, $	1,570	999	1,643	1,991	1,173	2,000	2,148	1,410

^ Simulation shuts down the second-lowest Silver plan part of the observed mechanism to allow comparability with counterfactual subsidization policies

~ Subsidies are flat within an age group. Baseline flat subsidy for a 20-year old consumer are adjusted for age using the age curve from the observed allocation

^^ Accounts for opportunity cost of public funds in the form of uncompensated care and the cost of raising public funds (λ, assumed to be 30 cents on a dollar)

Notes: Table reports the levels of consumer surplus, producer surplus, government spending, and total welfare under the observed allocation (column 1) and under counterfactual allocations (columns 2 to 8). We compute these objects using estimates of demand and marginal costs, and simulated market equilibria that allocate consumers to Marketplace insurance plans or the outside option. All simulations are performed within the ACS sample of consumers and are then scaled to the total market size. Consumer surplus is computed as discussed in Section 5. Firm profits reported in row (2) account for risk-equalization programs indirectly, as marginal cost estimates are "pricing-relevant" marginal costs, i.e. net of firms' expectations about positive or negative risk-equalization transfers. Cost-sharing reduction (CSR) subsidies in row (4) are computed by multiplying consumer-type specific average CSR values as reported by CMS for 2016 by enrollment share of each consumer type ($1,440 per year for conusmers with income under 150% FPL, $1,068 for those with income between 150% and 20)% FPL, and $144 for consumers with income between 200% and 250% FPL. Uncompensated care spending is computed at the rate of $1,827, following the Kaiser Family Foundation 2013 report on public spending on uncompensated care for the uninsured. Rows 11-27 describe consumer sorting for each allocation. Negative quantities are reported in parentheses.