In-Kind Transfers, Tagging, and Market Power: Evidence from the Affordable Care Act

Maria Polyakova and Stephen P. Ryan*

October 1, 2018
PRELIMINARY AND INCOMPLETE

Abstract

Public welfare programs have a long history of linking their benefits to observable characteristics of potential recipients, such as age, income, or employment status. We argue that this common mechanism, *tagging*, whose goal it is to improve efficiency with better targeting of in-kind transfers, may lead to substantial market distortions in an environment where the final good is provided to recipients by strategic private firms. The efficiency losses may be exacerbated when the level of transfers that subsidize the purchase of the good is anchored to the price information supplied by these firms. We explore this possibility empirically, using data on Health Insurance Marketplaces that were created in 2014 under the Affordable Care Act. We build a model of supply and demand in this new market. The estimated model primitives allow us to analyze the efficiency of market equilibria and the incidence of subsidies under the observed subsidy regime with tagging, as well as under counterfactual subsidization mechanisms.

JEL: I11, I18, L22, D44, H57

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

*Polyakova, Stanford University, CESifo, and NBER (mpolyak@stanford.edu); Ryan, Washington University in St. Louis, CESifo, and NBER (stephen.p.ryan@wustl.edu). We gratefully acknowledge support from the Agency for Healthcare Research and Quality (R03-HS024800). We are grateful to David Card and Pietro Tebaldi 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, University of Munich, REStud Conference, Rice University, and University of Utah for helpful comments. We thank Alan Jaske and Lynn Hua for excellent research assistance. Previously titled as "Subsidy Tagging in Privately-Provided Health Insurance Markets."

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, as in Medicare Part D or under the Affordable Care Act (ACA). In this paper, we investigate the welfare consequences of strategic firms pricing in 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 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 that was 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 and by which mechanism should the government pay insurers. This question is central for the efficiency of these markets, but is still very poorly understood.

The Marketplaces provide a fruitful ground 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 the tagging and price-linking subsidization mechanisms. We start by formulating and estimating the model of demand for ACA Marketplace plans. We find intuitive preference patterns, with individuals disliking higher premiums and liking more generous coverage levels. 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 a profit function that gives us first-order conditions, which in turn allow us to recover marginal cost estimates for insurance plans.

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 tagged 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 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 4.3 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 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

[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.

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 Descriptive Evidence

We start our analysis by exploring the observed relationships between premiums, enrollment and the demographic composition of the local geographic markets for ACA insurance plans. The key prediction is that effective premiums that consumer i faces depend on the income composition of similarly-aged population in this consumer's local market. For any given age group, insurers anticipate the distribution of demand elasticities in the population that depends on the level of income-tagged subsidies. For instance, in a market that only has consumers with income above 400% FPL, insurers face a fully elastic demand function, as these consumers are not eligible for subsidies. Reversely, in a market in which all potential consumers have income below 150% FPL and receive nearly full premium subsidies, insurers face zero residual demand elasticity and have an incentive to raise prices.

We test whether these predictions are consistent with the data. Using the ACS demographic sample, we calculate premium subsidies that each potential consumer would have faced if he or she chose to enroll in one of the Marketplace plans. The subsidy level is a function of individual's age and income relative to the Federal Poverty Line. We then combine the subsidy calculation with information about premiums for second-lowest cost silver plans in each geographic market. This allows us to calculate the effective premium for the second-lowest silver plan that each consumer would have faced if they chose to enroll. We focus on premiums for the second-lowest silver plans for ease of interpretation; the same exercise could be repeated for any other one plan, or any statistic of the premium distribution.

Next, we test whether in markets with more demand elasticity - which we define as those that have 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 subsidies. This relationship would exist in the data if insurer "bids" in markets with more elastic demand were falling faster than subsidies, so that net premiums were lower.

Panel (a) of Figure 3 suggests that this relationship holds empirically. On the 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.[2] 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 3 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 2. We start with a regression that captures the same relationship as Panel (a) of Figure 3: 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.

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

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 3, 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 2 report the results of enrollment regressions. For each county, we calculate the share of potential enrollees with income between 250% and 400%

[2]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.[3] 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 3). 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 Model

4.1 Demand

We formulate and estimate a random utility model of demand for health insurance plans on ACA Marketplaces. The empirical model maps characteristics of plans, such as premiums, cost-sharing rules, and non-pecuniary features, into a scalar measure of utility; consumers then pick plans that give them the highest utility.

Formally, we posit that individual i, characterized by a vector of demographic characteristics such as age and income, which we denote with D_i, chooses plan j from a set of choices J available to this individual, so as to maximizes utility. The set of choices that each consumer faces depends on their geographic location. It is helpful to think about a geographic market as a county, although there are multiple nuances on what serves as a geographic "market"

[3]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.

within the ACA Exchanges and we take these nuances into account when estimating the model. We let the indirect utility function take the following form:

$$u_{ij} = -\alpha_i p_{ij} + \beta_i X_{ij} + \epsilon_{ij} \tag{2}$$

Where p_{ij} denotes consumer i's post-subsidy premium for insurance plan j, vector X_{ij} captures observable characteristics of the plan, such as coverage generosity for various providers and services. The observable characteristics of the plans are allowed to be individual-specific. X_{ij} also includes fixed effects for insurer brand, as consumers may have strong preferences for specific insurers. Parameter α measures the marginal utility of income, or in other words the value of a dollar in "utils." Parameter vector β measures average consumer preferences for plan features. We allow for unobserved heterogeneity in these preferences. We further implicitly allow preferences to vary with observable consumer characteristics that include age and income, since prices and plan characteristics vary with consumer age and income as we describe in more detail below.

The last part of utility - ϵ_{ij} - is a random individual-plan specific shock to the utility function that accounts for the fact that there are some aspects for why a given individual may have higher or lower utility for a given plan that are known to the individual but unobserved by the researcher. We assume that this random shock is 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 close the model we assume that individuals choose plan j that maximizes their utility across all possible choices, or they choose not to enroll, which gives a normalized utility of zero. Formally, i chooses j if $u_{ij} > u_{ik}$ for all k in J such that k is not equal to j.

In our empirical setting, consumers that have the same set of plans to choose from, may face different plan characteristics; these characteristics depend on consumer demographics. First, consumers with lower incomes have reduced cost-sharing in plans with Silver metal labels, requiring us to adjust X_{ij} for these consumers. Second, insurers are allowed to age-rate premiums, so as to partially account for higher average healthcare spending at older ages, so older consumers will have higher p_{ij}. In addition, and this aspect is central for our analysis, while individuals of the same age in the same market face the same nominal premiums, effective premiums vary substantially across individuals according to their income. Effective consumer premiums depend on the level of premium subsidies that consumers receive. In the model, p_{ij} is the premium that consumer i has to pay for plan j net of consumer-specific

premium subsidies.

The variation in premiums that pins down the marginal utility of income parameter stems from insurer decisions and government policies rather than from experimental assignment. Hence, we may be concerned about a bias in our parameter estimates. For example, it is possible that there exists a characteristic of a plan j that we do not observe and do not include into the utility function, but this characteristic is observable to the individual, affects his or her choice and is also correlated with the premium of the plan. This would lead to an omitted variables bias in our estimates. To address this concern we use two instrumental variable strategies. The first strategy follows the intuition in Section 3 and is similar in spirit to the ideas in Waldfogel (2003), as also discussed in Berry and Haile (2016). We instrument prices in a any given county with the share of potential consumers in that county who do not qualify for income-tagged subsidies (i.e. consumers with annual income of more than 400% FPL). As illustrated in Section 3, there is a strong negative relationship between nominal premiums and the share of subsidy-ineligible consumers among potential consumers. The second strategy is to add spatial instrumental variables in the spirit of Hausman (1996) instruments. We instrument premiums of plans in insurer-metal level combination A in market (state and rating area combinations) X with a variable that computes the average premium that insurer-metal level combination A charges in all state and rating area combinations other than X. The logic behind this instrument is that it should capture the fundamentals of plan A's cost structure, such as for instance, its ability to negotiate with providers, and not be correlated with local demand shocks in market X. To accommodate the instrumental variables strategy into the model we use a control function approach as derived in Kim and Petrin (2010). We run first-stage specifications separately for each age level - the first instrument does not vary with age, while the second does. The first stage specifications include two instruments and all plan characteristics that also enter the utility function. We compute residuals for each first stage regression and then take a simple average of these residuals, using that average as the control function in demand estimation.

4.2 Supply

4.2.1 Profit function

Insurers in the ACA Insurance Exchange market decide which geographic markets to enter, how to design their plans, and how to price them. In this analysis, we focus on how insurers

set prices conditional on having made the entry and contract design decisions. Modeling price-setting poses a significant challenge, as pricing is constrained by an array of regulatory provisions. We start with a detailed accounting of payment flows in the market. We then discuss the assumptions we need, to make the supply-side model empirically tractable.

Consider a 40-year old individual i that lives in market t, and purchases an insurance plan j. For this individual, plan j collects revenue that consists of several pieces. First, the insurer collects a premium p_{ij} from the consumer. The premium is consumer-specific and depends on consumer's age and income.[4] For consumers that are not eligible for premium subsidies, the premium is equal to insurer's full price or the "bid." For consumers that are eligible for a subsidy, the insurer collects the premium from the consumer as well as a subsidy from the federal government that together add up to the "bid." Consumer i receives (a maximum) subsidy z_i that is a function of i's income and the bid of the second-lowest cost silver plan in market t. The consumer pays the difference between the bid and the subsidy in premiums. If the subsidy is higher than the bid, the consumer pays zero and does not receive the cash value of the "unused" subsidy.[5] Second, the insurer collects revenue from three risk-equalization programs that we describe below.

On the cost side, both the realized costs that a plan j incurs for consumer i and the ex ante costs that a plan expects to incur 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, i.e. through moral hazard, or plan j may simply have different negotiated prices for the same services. We denote j's cost-sharing characteristics or provider bargaining power 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 se; however, insurers' costs

[4] In what follows, we abstract from the differences in premiums for smoking and non-smoking enrollees, as well as from the individual versus family coverage. Smoking is self-reported upon the purchase of insurance; the insurers have no mechanism to verify whether an enrollee smokes, which likely leads to severe underreporting. We do not obesrve household-level enrollment information and simplify the model by treating each enrollee as an independent agent. This simplification is ameliorated by the fact that premiums for household vs. individual-level enrollment tend to just be scaled from the baseline premiums for single adults, suggesting that the insurer faces the same incentives across single and household consumers on the margin.

[5] In practice, the subsidy operates as a tax credit; the estimated level of the credit is reconciled during tax filing.

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.

Without any risk-equalization programs, plan j's expected profit for consumer i as a function of plan j's bid b_j and bids of all other plans b_{-j} in a given market for this consumer's age group a is equal to:

$$\pi_{ij}(b_j^a; b_{-j}^a) = p_{ij}(D_i, b_j^a, b_{-j}^a) + z_i(D_i, b_j^a, b_{-j}^a) - c_{ij}(r_i, \phi_j, D_i), \tag{3}$$

Suppose that for any plan j, there is a baseline plan-specific cost c_j^a of covering an average enrollee of a given age. Then, we can re-write the individual cost c_{ij} as the sum of the plan-specific cost and an idiosyncratic cost component: $c_{ij}(r_i, \phi_j, D_i) = c_j^a + \tilde{c}_{ij}(r_i, \phi_j, D_i)$. We further note that while the split of insurer revenue between consumer premium p_{ij} and the subsidy z_{ij} is important in determining how many people choose to enroll in the plan, once a consumer enrolles, the premium and the subsidy add up to insurer's bid. In other words, $p_{ij}(D_i, b_j^a, b_{-j}^a) + z_i(D_i, b_j^a, b_{-j}^a) = b_j^a$ for a consumer of age a with any income level. Using this notation, we can then re-write the profit of plan j from enrolling individual i of age a as follows:

$$\pi_{ij}(b_j^a; b_{-j}^a) = b_j^a - c_j^a - \tilde{c}_{ij}(r_i, \phi_j, D_i), \tag{4}$$

The individual-specific cost term allows for the presence of advantageous or adverse selection that is a function of plan characteristics ϕ_j. For example, selection into a plan would be adverse if the sum of $\tilde{c}_{ij}(r_i, \phi_j, D_i)$ across all consumers purchasing plan j, is a positive number. In other words, selection is adverse if a plan j systematically enrolls consumers with idiosyncratic costs above the plan-specific expected cost component c_j. Seelction would be advantageous if this sum were negative, and there would be no selection if the sum was close to zero.

Three programs exist on ACA Marketplaces (within the time horizon we study) aimed at equalizing expected insurers' costs of 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.

To the first-order approximation, the risk-equalization programs attempt to neutralize the idiosyncratic component $\tilde{c}_{ij}(r_i, \phi_j, D_i)$ of enrollees' ex ante costs. We incorporate these programs into our notation as follows. 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. Let this additional revenue be $\tilde{z}_{ij}(r_i, \phi_j, D_i)$. 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 $\eta_{ij} = \tilde{c}_{ij} - \tilde{z}_{ij}$. Now let the 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. We do not explicitely 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 (although they likely affect entry decisions).

Let I_j^a denote the number of consumers of age a that choose to enroll in plan j. Then, for all individuals i across all age groups a that enroll in plan j, we can write the profit of plan j as:

$$\pi_j(b_j; b_{-j}) = \sum_a I_j^a b_j^a - \sum_a I_j^a c_j^a - \sum_{i \in j} \eta_{ij} + R_j \tag{5}$$

Let $H_j(\phi_j)$ denote any residual selection that is left after risk-adjustment. In other words, $H_j(\phi_j) = \sum_{i \in j} \eta_{ij} - R_j$. If the risk-equalization programs resulted in perfect risk-adjustment that completely removed the ex ante net idiosyncratic shocks, this term would be zero. If risk adjustment is imperfect, $H_j(\phi_j)$ can be either positive or negative depending on the nature of residual selection. We assume that $H_j(\phi_j)$ is not zero, but that it depends only on ϕ_j rather than the bids *conditional* on ϕ_j.[6]

[6]This is a plausible assumption, as the costliest individuals are likely to have the lowest incomes due to a negative health-income gradient, and thus are likely to receive generous premium and cost-sharing subsidies, making plans similar in terms of their financial characteristics. It follows that for these individuals, the key differences across plans lie in non-pecuniary plan features in ϕ_j, such as physician networks, formulary breadth, and chronic condition management. 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.

Rewriting the profit function using the share notation, we get:

$$\pi_j(b_j; b_{-j}) = \sum_a s_j^a M^a b_j^a - \sum_a s_j^a M^a c_j^a + H_j(\phi_j) \tag{6}$$

Empirically, and according to ACA statutes, the bid menu of each plan j across different ages follows a fixed scaling schedule. This observation 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 multiplies apply to expected baseline costs, capturing how healthcare costs increase with age.[7] Let a multiplier vector for plan j be τ_j (we allow for plan-specific multipliers; in practice, the observed market outcomes results from fixed τ's across plans). The profit equation for plan j then becomes:

$$\pi_j(b_j; b_{-j}) = \sum_a s_j^a M^a \tau_j^a b_j - \sum_a s_j^a M^a \tau_j^a c_j + H_j(\phi_j) \tag{7}$$

or, re-arranging:

$$\pi_j(b_j; b_{-j}) = (b_j - c_j) \sum_a s_j^a M^a \tau_j^a + H_j(\phi_j) \tag{8}$$

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

$$\pi_f(\mathbf{b}) = \sum_{j \in f} \left[(b_j - c_j) \sum_a s_j^a M^a \tau_j^a + H_j(\phi_j) \right] \tag{9}$$

The insurer maximizes profits by choosing a bid b_j for each plan j in its portfolio.

4.2.2 First order conditions for pivotal and non-pivotal plans

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 9 is:

$$\frac{\partial \pi_f}{\partial b_j} = (b_j - c_j) \sum_a \frac{\partial s_j^a}{\partial b_j} M^a \tau_j^a + \sum_a s_j^a M^a \tau_j^a = 0. \tag{10}$$

For an insurer that offers more than one plan in a market, the vector notation for the set of first-order conditions becomes:

[7] The discrepancy between the statutory age-specific multipliers and the true age slope in healthcare cost is the source that can generate adverse selection. Since insurers are allowed partial age-rating of their premiums, the selection that we may be concerned about in the ACA setting is *within* age selection - this source of selection is the target of the risk-equalization programs as discussed above.

$$S - \Omega(B - C) = \mathbf{0}. \tag{11}$$

where row j of vector S is given by:

$$S_j = \sum_a s_j^a M^a \tau^a \tag{12}$$

and row j of vector $(B - C)$ is given by:

$$(B - C)_j = (b_j - c_j) \tag{13}$$

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. We invert Equation 11 and compute the marginal cost (for a 20-year old as a baseline) of 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. In practice, they interact through any plan- or insurer-level risk-equalization policy. Further, to the extent that age-specific scaling in bids does not perfectly equalize differences in age-specific costs, there is cross-subsidization across ages within a plan. 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. Recall that the subsidy is a function of the bid set by the second-lowest cost silver plan and consumer's income. These SLS plans - or what we call "pivotal" plans - face a different set of incentives, as a change in their bid affects not only their own prices, but also the subsidies, and hence consumer premiums, of all plans. In what follows, we consider the object $\frac{\partial s_j}{\partial b_j}$, separately for pivotal and non-pivotal plans. We assume a complete information game, in which pivotal plans know that they are pivotal, while non-pivotal plans know that they are not pivotal, and price accordingly.

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 depends crucially on the composition of consumer incomes in the market. As we highlight below, the overall share of plan j is a weighted sum of the shares in each income-age consumer bin (weighted by the share of consumers in each bin). In other words, the derivative of the plan's share with respect to its own bid, and hence its pricing decisions, depend crucially on the *share of elastic consumers* 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 the insures have an incentive to raise prices in markets with poorer - and consequantly more subsidized and less elastic - consumers.

Consumer prices are functions of all silver plan bids, as they depend on the subsidy, which itself is a function of the the second-lowest bid silver plan. In other words, we have: $s_j(p_j, p_{-j})$, and $p_j(b_j, b_{-j})$ as well as $p_{-j}(b_j, b_{-j})$. The overall derivative of the share function can then be written as:

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

Non-pivotal plans Consider the share derivative in Equation 15 for non-pivotal plans. For these plans, $\frac{\partial p_{-j}}{\partial b_j} = 0$ since the bid of plan j does not affect prices for other plans, conditional on the bids of other plans. The term, $\frac{\partial p_j}{\partial b_j}$ depends on the subsidy structure. Under the means-tested subsidy structure, we have the following components (within one age group) in $\frac{\partial p_j}{\partial b_j}$:

1. For consumers with no subsidy $\frac{\partial p_j}{\partial b_j} = 1$, as for these consumers a small increase in the bid translates into the same increase in consumer premium. Consumers with a zero subsidy are those with income over 400% FPL or those whose income-specific premium cap $\bar{p}_i$ is strictly above the bid of the second-lowest cost Silver plan b_j^{2LSP}

2. All other consumers receive subsidy z_i that is (at most) equal to the difference between their individual income-specific premium cap $\bar{p}_i$ and the bid of the second-lowest silver plan in their market. For these consumers $\frac{\partial p_j}{\partial b_j} = 1$ if b_j is greater than the subsidy. If b_j is less than the subsidy, then consumers pay zero out of pocket premiums and hence $\frac{\partial p_j}{\partial b_j} = 0$.

Summarizing, we get:

$$\frac{\partial p_j}{\partial b_j} = \begin{cases} 0 & \text{if } b_j < z_i \\ 1 & \text{otherwise} \end{cases} \tag{16}$$

Denote the fraction of elastic consumers within an age-geography market with η. We also denote the share of elastic consumers that buy pan j with s_j^e and the share of inelastic consumers that buy plan j with s_j^{ine}. We then get that $j's$ total market share in age group a in market t:

$$s_j = \eta s_j^e + (1-\eta) s_j^{ine} \tag{17}$$

$$\frac{ds_j(p_j, p_{-j})}{db_j} = \eta \frac{ds_j^e(p_j, p_{-j})}{db_j} + (1-\eta) \frac{ds_j^{ine}(p_j, p_{-j})}{db_j} \tag{18}$$

$$\frac{ds_j(p_j, p_{-j})}{db_j} = \eta \frac{\partial s_j^e}{\partial p_j} + (1-\eta) * 0 \tag{19}$$

$$\frac{ds_j(p_j, p_{-j})}{db_j} = \eta \frac{\partial s_j^e}{\partial p_j} \tag{20}$$

For elastic consumers that are subsidized, each consumer faces a different, income-specific premium, hence the derivative $\frac{\partial s_j^e}{\partial p_j}$ is itself a sum over income cells. Equation **??** highlights the importance of elastic consumers on the market. As η approaches 1, i.e. as the share of elastic consumers grows, the elasticity of demand increases. Reversely, when there are no elastic consumers on the market and η approaches zero, the elasticity of demand apporaches zero and the firm faces no marginal incentives to decrease prices.

For example, in the case of a logit probability at the individual level, we get:

$$s_j^e = \sum_{i=1}^{N} s_{ij}^e = \sum_{i=1}^{N} \frac{exp(\delta_j - \alpha p_{ij})}{1 + \sum_K exp(\delta_k - \alpha p_{ik})}$$

$$\frac{\partial s_j^e}{\partial p_j} = \sum_{i=1}^{N \in e} \frac{\partial s_{ij}}{\partial p_{ij}} = \sum_{i=1}^{N \in e} \alpha s_{ij}(s_{ij} - 1)$$

allowing for heterogeneity in preferneces, we get that the (age-specific) share derivative for non-pivotal plans is equal to:

$$\frac{ds_j(p_j, p_{-j})}{db_j} = (1-\eta) \sum_{i=1}^{N \in e} \alpha_i \frac{exp(\delta_{ij} - \alpha_i p_{ij})}{1 + \sum_K exp(\delta_{ik} - \alpha_i p_{ik})} \left(\frac{exp(\delta_{ij} - \alpha_i p_{ij})}{1 + \sum_K exp(\delta_{ik} - \alpha_i p_{ik})} - 1 \right)$$

Pivotal plans We proceed analogously for pivotal plans. Consider the first term of the share derivative in Equation 15 for pivotal plans. Under income-linked subsidies, we have

the following components (within one age group) in $\frac{\partial p_j}{\partial b_j}$ for pivotal (i.e. 2LSP) plans:

1. For consumers with no subsidy $\frac{\partial p_j}{\partial b_j} = 1$, as for these consumers a small increase in the bid translates into the same increase in consumer premium. Consumers with a zero subsidy are those with income over 400% FPL (in which case the income-specific premium cap is infinity) or those whose income-specific premium cap $\bar{p}_i$ is strictly above b_j (since j in this case is the second-lowest cost silver plan).

2. Subsidized consumers always pay exactly their income-specific premium cap $\bar{p}_i$ for the second-lowest silver plan, as the subsidy z_i is equal to the difference between their individual income-specific premium cap and b_j. Hence, for subsidized consumers, $\frac{\partial p_j}{\partial b_j} = 0$. Consumers are subsidized if their income-specific premium cap $\bar{p}_i$ is below b_j. Summarizing, we get:

$$\frac{\partial p_j}{\partial b_j} = \begin{cases} 0 & \text{if } \bar{p}_i < b_j \\ 1 & \text{if } \bar{p}_i > b_j \end{cases} \tag{21}$$

As with the non-pivotal plans, the share of elastic consumers matters for the pricing of second-lowest silver plans, as those consumers who receive any subsidies always pay a fixed amount for the second-lowest silver plans and hence have a completely inelastic demand.

We next consider the term $\frac{\partial p_{-j}}{\partial b_j}$. Unlike in the case of non-pivotal plans, this term is not equal to zero for pivotal plans. When pivotal plans increase their bid b_j by epsilon, the maximum subsidy that eligible consumers can get - z_i - also increases by epsilon, which in turn affects consumer prices of all other plans. Specifically,

1. For consumers with income being over 400% FPL and no subsidy, i.e. $z_i = 0$, $\frac{\partial p_{-j}}{\partial b_j}$=0, since the bid of the pivotal plan doesn't affect the prices these consumers pay for other plans

2. For consumers whose income-specific premium cap $\bar{p}_i$ is below b_j: these consumers receive subsidy z_i that is (at most) equal to the difference between p $\bar{p}_i$ and b_j , since b_j is the pivotal plan. When b_j increases by epsilon, z_i increases by epsilon. For these consumers:

   (a) If b_{-j} is below z_i, an epsilon increase in b_j increases the maximum subsidy by epsilon, but has no effect on consumer prices for b_{-j} , as consumers are paying zero premiums for these plans. Then, $\frac{\partial p_{-j}}{\partial b_j} = 0$.

(b) If b_{-j} is above z_i, then increasing the maximum subsidy z_i by epsilon results in a decrease of consumer price for plan $-j$ by epsilon, hence $\frac{\partial p_{-j}}{\partial b_j} = -1$, since consumers have to pay less of the incremental cost of more expensive plans. This condition applies by construction to all plans with bids above the second lowest silver plan bid as well as to some plans, whose bids are above z_i, but below the second-lowest silver plan.

Summarizing, we get:

$$\frac{\partial p_{-j}}{\partial b_j} = \begin{cases} 0 & \text{if } b_{-j} < z_i \text{ or } z_i = 0 \\ -1 & \text{if } b_{-j} > z_i \end{cases} \tag{22}$$

The term $\frac{\partial s_j}{\partial p_{-j}} * \frac{\partial p_{-j}}{\partial b_j}$ is a sum of these terms across each consumer type. Within each consumer type, consumer face different prices, so that the derivative $\frac{\partial s_{-j}}{\partial p_j}$ is itself a weighted sum over income cells for each consumer type. Equation 22 highlights that for a set of consumers and plans, the change in the bid of the second-lowest silver plan may decrease consumer prices for other plans and hence increase the demand for these plans conditional on the plans' own bids. We take this relationship into account when computing derivative shares numerically.

For example, in the case of a logit probability at the individual level, we get:

$$s_j = \sum_{i=1}^{N} s_{ij} = \sum_{i=1}^{N} \frac{exp(\delta_{ij} - \alpha_i p_{ij})}{1 + \sum_K exp(\delta_{ik} - \alpha_i p_{ik})}$$

$$\frac{\partial s_j}{\partial p_{-j}} = \sum_{i=1}^{N} \frac{\partial s_{ij}}{\partial p_{-ij}} = \sum_{i=1}^{N} \alpha_i s_{ij} s_{-ij}$$

and

$$\frac{\partial s_j}{\partial p_j} = \sum_{i=1}^{N} \frac{\partial s_{ij}}{\partial p_{ij}} = \sum_{i=1}^{N} \alpha s_{ij}(s_{ij} - 1)$$

Multiplying this with probability of consumer types, we get:

$$\frac{\partial s_j}{\partial p_{-j}} * \frac{\partial p_{-j}}{\partial b_j} = (Pr(b_{-j} > z_i)) * (-1) * \sum_{i=1}^{N} \alpha_i s_{ij} s_{-ij}$$

and

$$\frac{\partial s_j}{\partial p_j} * \frac{\partial p_j}{\partial b_j} = (Pr(\bar{p}_i > b_j)) * \sum_{i=1}^{N} \alpha s_{ij}(s_{ij} - 1)$$

We get the total derivative for the pivotal plan to be:

$$\frac{ds_j(p_j, p_{-j})}{db_j} = \frac{\partial s_j}{\partial p_j} * \frac{\partial p_j}{\partial b_j} + \frac{\partial s_j}{\partial p_{-j}} * \frac{\partial p_{-j}}{\partial b_j} \tag{23}$$

$$\frac{ds_j(p_j, p_{-j})}{db_j} = (Pr(\bar{p}_i > b_j)) * \sum_{i=1}^{N} \alpha s_{ij}(s_{ij} - 1) + (Pr(b_{-j} > z_i)) * (-1) * \sum_{i=1}^{N} \alpha_i s_{ij} s_{-ij} \tag{24}$$

Conditional on knowing the share derivative for each plan, marginal cost is the only unknown in the first order condition. Hence, we can estimate marginal costs by inverting equation 10 as follows:

Given the complexity of the share derivative, there is no algebraic closed-form solution for it, but we can estimate it numerically and use the calculation to derive a vector of marginal costs for each plan.

4.3 Estimation results

Demand Parameters We adapt the estimation routine for discrete choice models, to incorporate the varying levels of aggregation at which we observe enrollment. We do not observe plan-level market shares as would be common in this setting. Instead, but rather county-level market shares of plans aggregates into their metal levels: platinum, gold, silver, bronze, and catastrophic, as well as plan (but not county) level enrollment, and finally insurer-county enrollment. These moments give us several thousand cross-sectional restrictions on the underlying demand function, including demographic interactions. The intuition of the estimation is similar to a standard discrete choice model; computationally, however, the approach is different, as we rely on bottom up simulation of the model to match moments at different levels of aggregation.

Table 3 reports the simulated method of moments estimates for several demand specifications. We start with a specification in Column (I) that includes only the main parameters of the utility function: individual-specific (income-adjusted) premium, individual-specific measure of average plan generosity, the control function, and a constant capturing the value of the outside option. While very parsimonious, this specification robustly captures the key

patterns in the data. Individuals dislike higher premiums and like more coverage - consumers are willing to pay about $350 a year for a ten percentage point increase in the average generosity (actuarial value) of plans. Column (2) adds random coefficients to the specification. Column (3) adds the full vector of plan characteristics, but not the random coefficients. While the exact point estimates for the willingness to pay vary across specifications, the qualitative and quantitative patterns are very similar.

Figure 5 illustrates the in-sample fit of the demand model (I). We simulate consumer choices in the model and compute the resulting enrollment shares at the plan metal level in each market. The figure plots the histogram, across 2,566 counties, of metal-level shares in the data and in out simulation. We closely match these aggregated enrollment shares and the model is able to capture a substantial amount of variation in the data.

Cost Parameters Figure 6 illustrates the distribution of estimated marginal costs (based on elasticity estimates from demand model (I)). Given the restrictions of the inversion procedure, we estimate the marginal cost for one age group per plan and assume that marginal costs for other ages groups in the same plan move with τ multipliers. We estimate the costs for 20-year old consumers, who have the reference *tau* of 1. The figure plots the distribution of the cost estimates across plan-markets, separately for Bronze and Platinum plans. We observe two pronounced patters. First, there is substantial heterogeneity in costs within a plan type. 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, which we would expect, as mechanically platinum plans cover 90% of consumers' healthcare expenditures, while the Bronze plans cover only 60%. The ratio of 1.5 in costs between these plans that we would expect on average is consistent with the magnitude of the shift in the distribution of costs from light gray (bronze plans) to dark gray (platinum plans) that we estimate.

Figure 6 presents another piece of evidence that suggests that our marginal cost estimates are reasonable. In this figure we report the correlation between our marginal cost estimates from the inversion of the first-order condition and accounting costs. We observe plan-level (although not plan-county level) accounting costs reported by plans to CMS in their annual rate review files. These files contain a measurement error in accounting costs, as insurers are allowed to report their costs equally split across their plans rather than providing a true plan-cost attribution. Nevertheless, these data provide valuable informational signal, as they

on average sort plans by their costs. We observe a strong correlation between accounting costs and the marginal costs that we estimate from the inversion procedure. The levels of two cost measures are not comparable. While re report marginal costs for a 20-year old consumer, the accounting costs are reported at insurer or plan level and hence are averages across consumers of all ages.

5 Counterfactuals and Welfare

Welfare under Observed Allocation Before turning to counterfactual analyses of subsidy mechanisms, we first focus on understanding the incidence of the existing subsidization process. We are interested in two aspects. First, how does total consumer surplus and insurer profit compare to the total (premium) subsidy spending by the government. Second, how does the surplus from subsidies differ across socio-economic groups and across geographic locations. We start the analysis with a simple exercise of documenting the differences in per capita consumer surplus across socio-economic groups and geographies. We then proceed to analyze the differences in the distribution of surplus between consumers and producers across these groups and locations.

Following Williams (1977) and Small and Rosen (1981), surplus for consumer i with marginal utilities θ_i from plan characteristics, including the premium, takes the following form:

$$CS(\theta_i) = \frac{1}{\alpha_i}\left[\gamma + \ln\left[1 + \sum_{j=1}^{J}\exp(v_{ij}(\theta_i))\right]\right], \tag{25}$$

where γ is Euler's constant, and v_{ij} is the deterministic component of utility for person i from plan j and is equal to utility net of the idiosyncratic ϵ term.[8] We integrate out over consumer heterogeneity to obtain consumer surplus:

$$CS = \int CS(\theta)dF(\theta). \tag{26}$$

We estimate that under the observed allocation and prices consumer surplus for the federally facilitated ACA Marketplaces amounts to $15 billion. Producer surplus (where we assume that any risk-equalization payments, including the risk corridors contribute to cost equalization and do not explicitly add these to profits) amounts to $6 billion. Under

[8]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.

the observed allocation as simulated in our model, the government is spending $23 billion in subsidies, which is consistent with the subsidy spending reported by the Congressional Budget Office ($32 billion, which include spending in non-federally facilitated states). Thus in total, we conclude that nominally the program that attracts about 40% of potential enrollees generates negative return on government spending, as the subsidy payments are higher than the total surplus. It is important to emphasize the word nominal here, as we are not accounting for other spending that the government would have potentially had on the same consumers in the absence of ACA Marketplaces. The latter could range from enrolling these consumers in other public health insurance programs, such as Medicaid, or being a residual payer for uncompensated care. If those other potential expenditures on the same consumers would have been more than $2B, which is very likely, then the Marketplaces would generate a large positive net welfare gain.

Figure 8 illustrates the geographic variation in average welfare. We observe a significant degree of variation, ranging from positive average per capita welfare of circa $800 generated in Tennessee, many parts of Texas, and Iowa. At the same time we observe negative welfare generated in several southern states, Florida, and Maine. The main driver of this variation is the different amount of subsidies across areas. As we saw in Figure 2, the the areas with the lowest welfare estimates are the ones that also tend to have the lowest share of unsubsidized consumers.

Subsidy incidence and counterfactual mechanisms 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) 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%.

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 we observe in Panel (b) of Figure 10, flat subsidies lead to a re-allocation of surplus from the south-eastern states to Texas, Tennesse and several northern states.

To further understand the contribution of the market power, 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.

We report the last counterfactual simulation in Panel (d) on Table 4 . The two columns of this panel compare allocations with and without market power when there are no subsidies. Removing market power increases total surplus by $2 billion from the baseline of $10 billion. In addition, $2 billion of surplus is re-allocated from insurers to consumers, thus increasing consumer surplus by a total of $4 billion and consumer enrollment from 11 to 30 percent. Panel (d) of Figure 10 again suggests that the effects of market power are heterogeneous

across locations and result in a redistribution of surplus across the country.

6 Conclusion

Traditionally, "tagged" 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, for example, insurance benefit to private markets. While the mode of providing social insurance and welfare benefits is changing, many policies still rely on traditional "tagging", frequently implemented as a means-test. In this paper we argue that adding market power to the supply side of a public benefit 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. If a firm knows that in a particular market with low income consumers, all consumers will receive very generous subsidies, effectively rendering demand for insurance inelastic, it will have the incentives to raise prices. In the presence of market power, this demographic externality is not dissipated by competition.

References

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

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

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

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

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

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

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

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

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

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

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

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

Hausman, J. (1996). *The Economics of New Goods*, Chapter Valuation of New Goods under Perfect and Imperfect Competition, pp. 207–248. University of Chicago Press.

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

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

Kim, K. and A. Petrin (2010). Control Function Corrections for Unobserved Factors in Differentiated Product Models. *University of Minnesota*.

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

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

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

Small, K. A. and H. S. Rosen (1981). Applied welfare economics with discrete choice models. *Econometrica 49(1)*, 105–130.

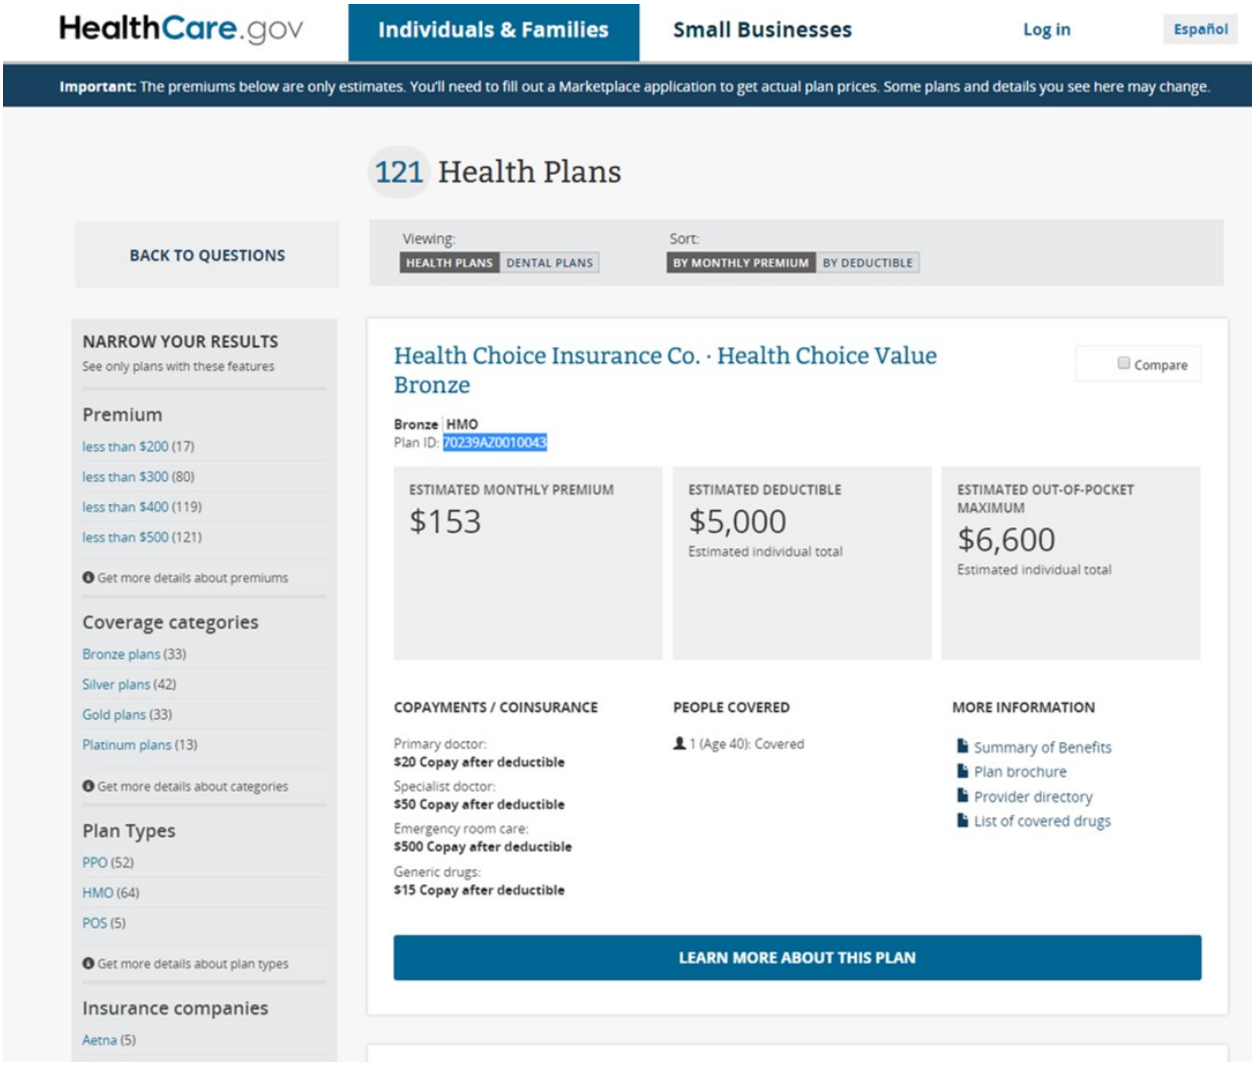
Tebaldi, P. (2017). Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA. *University of Chicago Working Paper*.

Waldfogel, J. (2003). Preference externalities: An empirical study of who benefits whom in differentiated-product markets. *RAND Journal of Economics 34*, 557–568.

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

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

Figure 1: Consumer interface on healthcare.gov



Notes: Snapshot of 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 2: Variation in the share of unsubsidized consumers across geographic markets

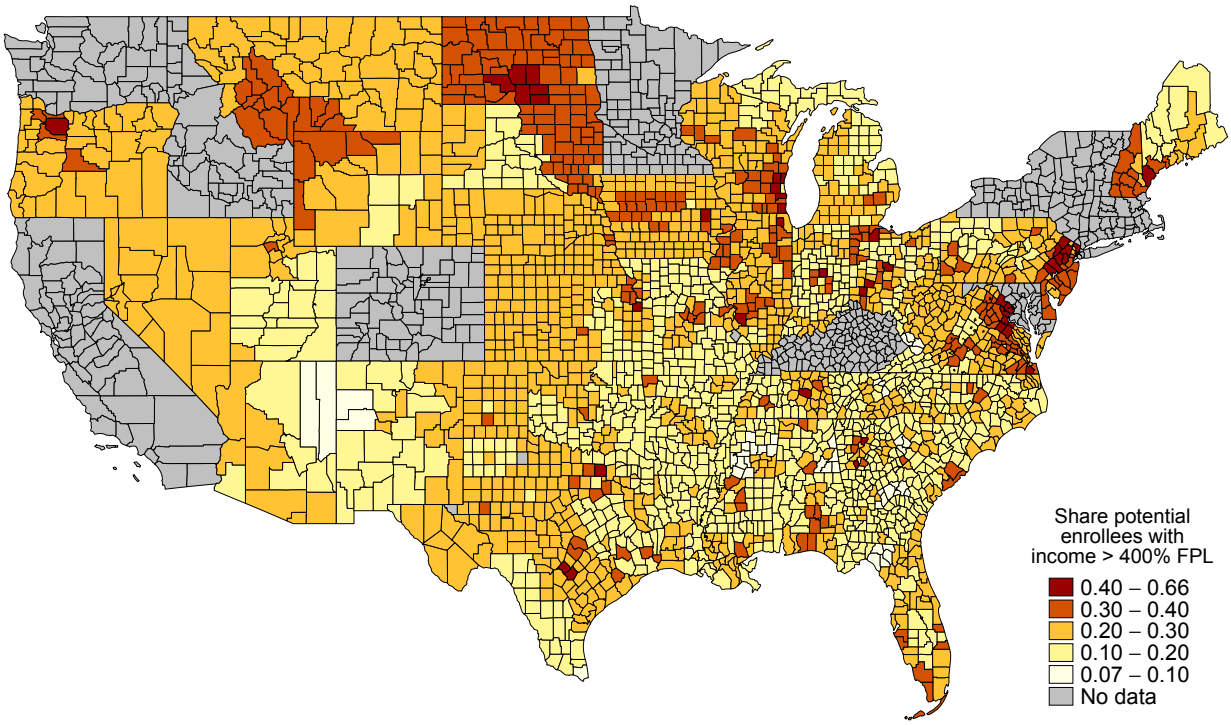
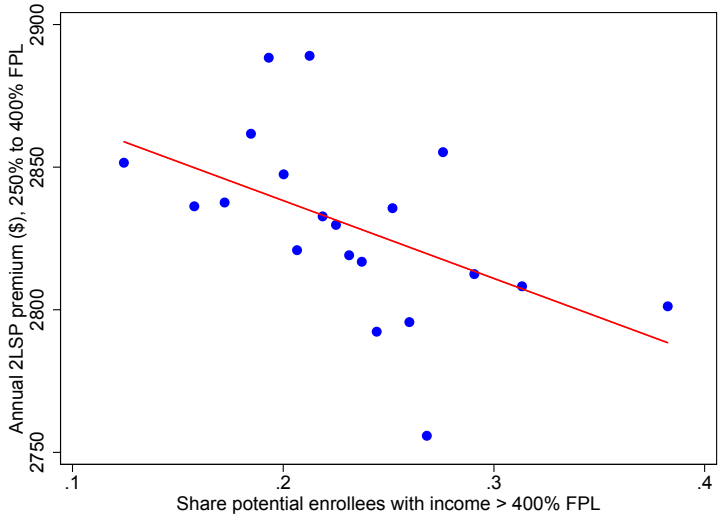
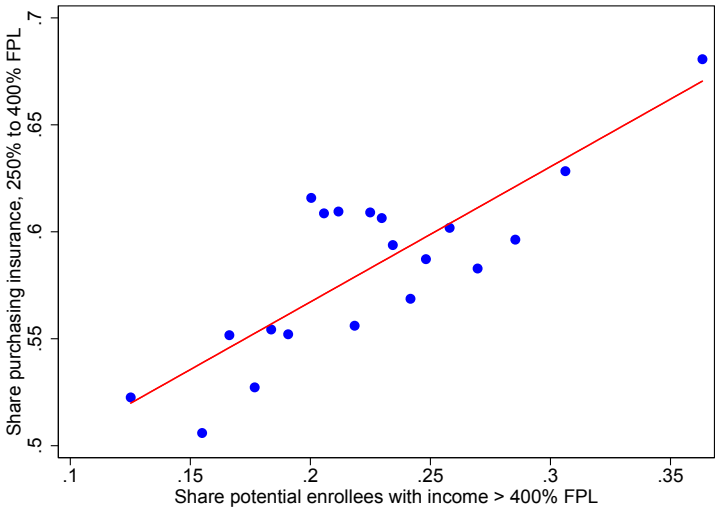


Figure 3: Variation in premiums and enrollment by the share of unsubsidized consumers

(a) Effective premiums

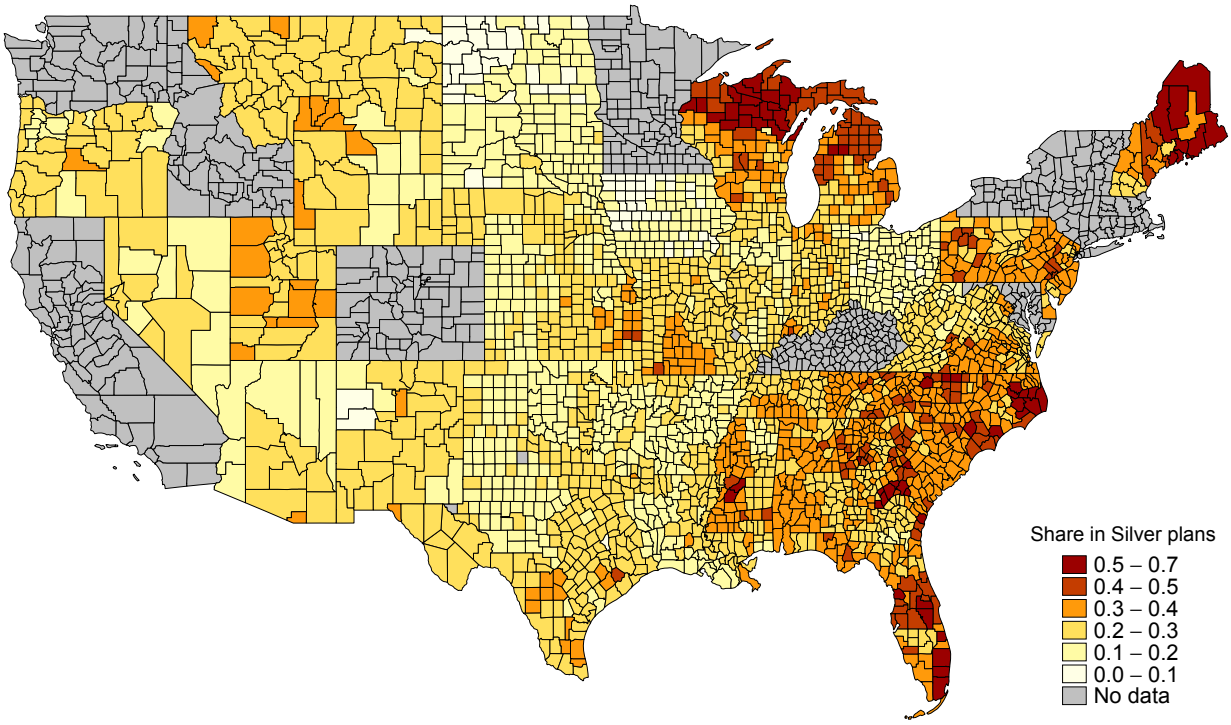


(b) Enrollment



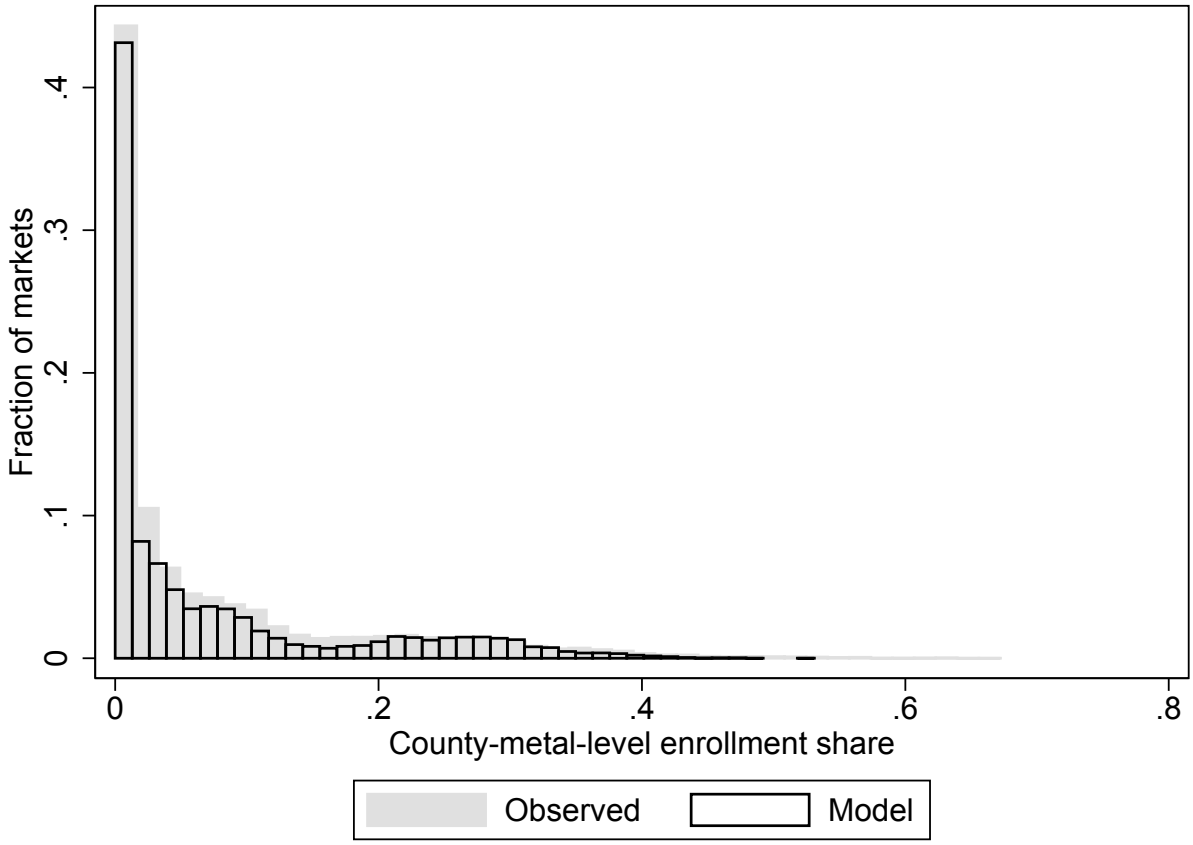
Notes: The panels illustrate the descriptive relationship between the demographic composition of counties and the market experience of consumers with income above 250% FPL and below 400% FPL. The top panel computes for each county the share of individuals that are not eligible for subsidies (income above 400% FPL) among individuals that do not have private or public insurance and are thus in the potential consumer pool for the Marketplaces. The x-axis plots the residualized (to state fixed effects) version of this measure. The y-axis plots the (residualized) effective premiums that individuals with income above 250% FPL and below 400% FPL would face if they enrolled. The bottom panel has the same x-axis. On the y-axis, it records the share of consumers with income above 250% FPL and below 400% FPL that enrolled in any Marketplace plan. County-level observations are aggregated into 20 equal-sized bins, each point in the scatterplot reports the average of the y-variable in the bin. The line marks a linear fit.

Figure 4: Empirical demand moments: share of Silver enrollment by county



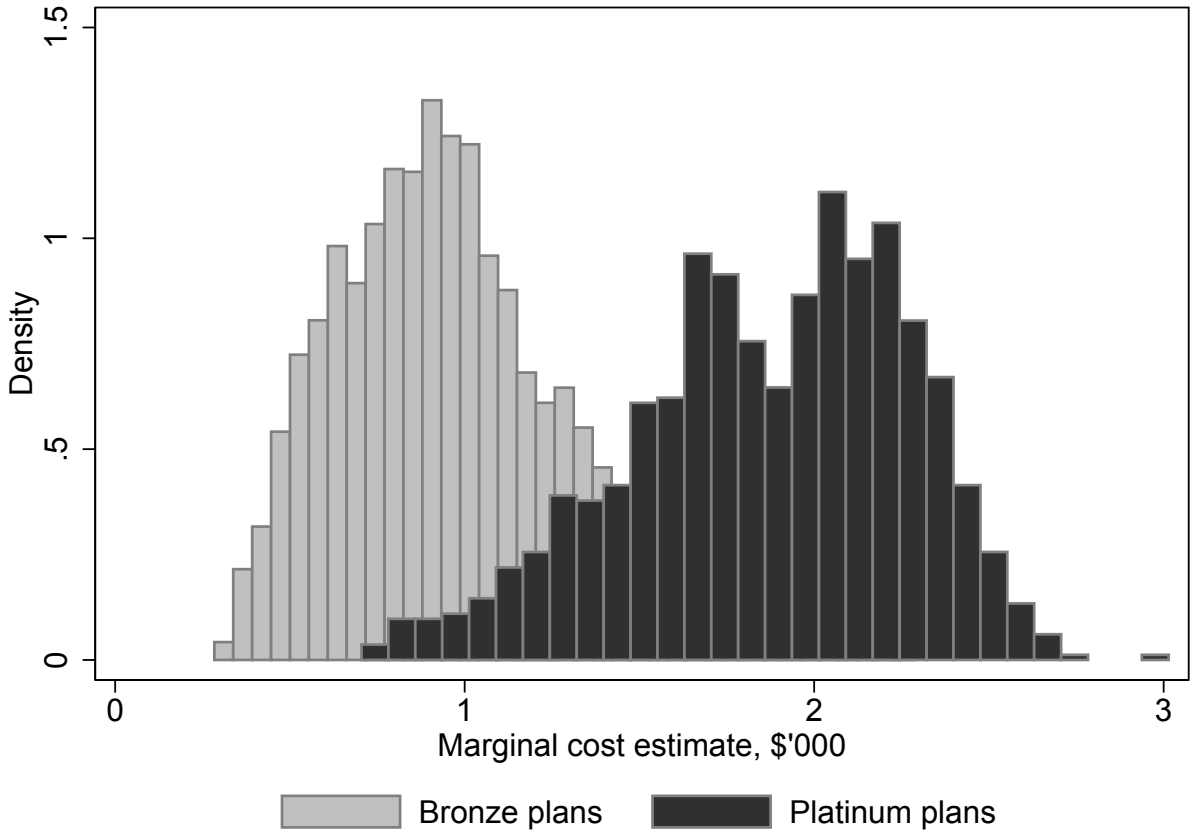
Notes: The map plots the share of potential consumers in each county that enrolled in a Silver plan on ACA Marketplaces. States that are marked with grey are not federally facilitated and do not enter our analysis. The counts of the pool of potential consumers (denominator) was provided by the Kaiser Family Foundation and is based survey metrics 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. The data is for year 2015.

Figure 5: Demand model fit: county-metal level moments



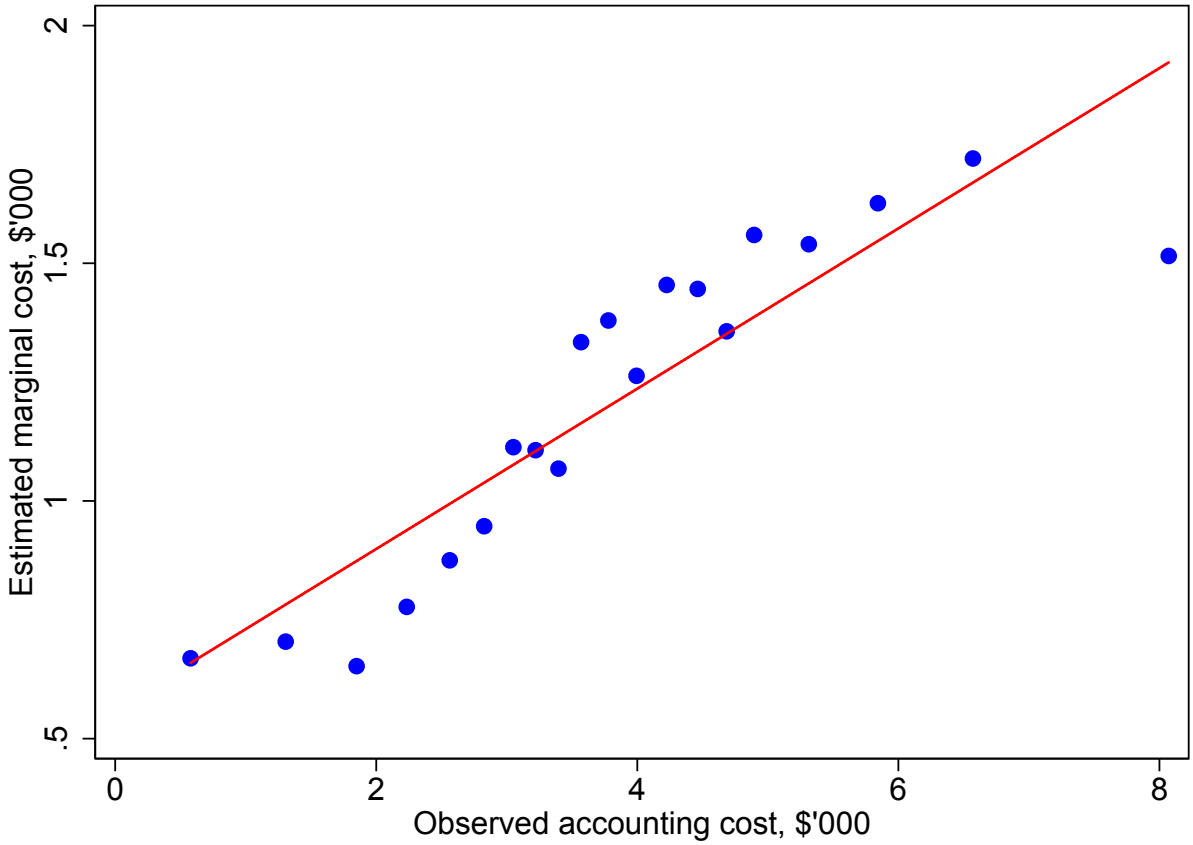
Notes: The graph compares the distribution of county-metal label enrollment shares between data and the demand model. X-axis measure enrollment share: for example, if 20% of potential enrollees in Cook County, IL purchased a Silver plan, this would be recorded as 0.2 on the x-axis. Y-axis measures the share of county-metal label combinations that fall in respective x-axis bins. There are a total of 2,566 counties and 5 metal labels, creating a total of 12,830 observations that underlie each histogram. The light-counter histogram records the distribution of county-metal label enrollment in the data. The dark-contoured histogram records the (in-sample) predictions of the demand model.

Figure 6: Marginal cost estimates from first-order-condition inversion



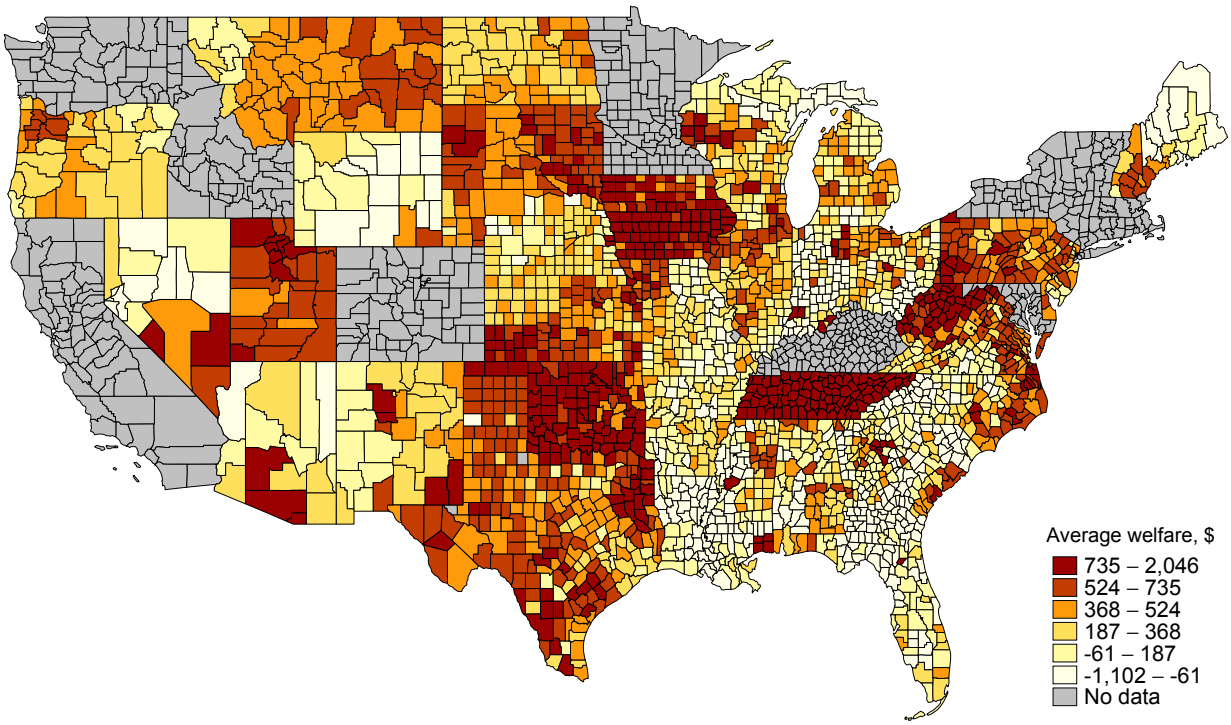
Notes: Marginal cost estimates are for a 20 year old consumer. We assume that marginal costs scale with respect to age with the same multipliers as observed prices. The x-axis is in thousands of dollars per year.

Figure 7: Inverted marginal costs correlate with accounting costs



Notes: Marginal cost estimates are reported for a 20 year old consumer. We assume that marginal costs scale with respect to age with the same multipliers as observed prices. Accounting costs are reported as plan-averages and hence average across all age groups. Accouting costs are extracted with rate review files released annually from the Centers for Medicare and Medicaid Services.

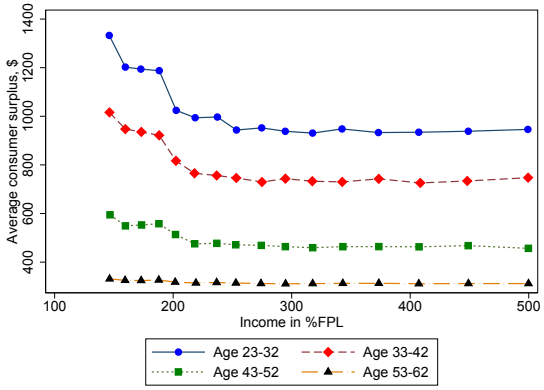
Figure 8: Distribution of welfare under the observed allocation



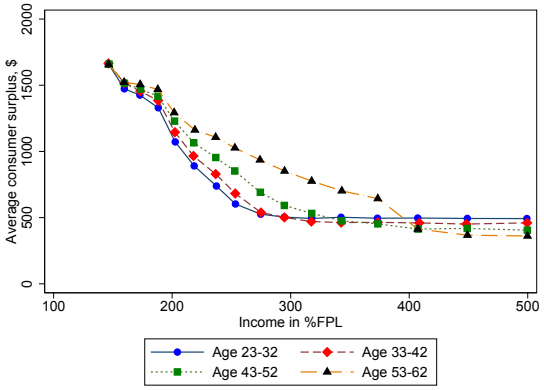
Notes: The map illustrates the average total surplus from ACA Marketplace plans in each county of the federally facilitated states. Total surplus is computed as a sum of consumer and producer surplus net of public subsidy spending. We compute consumer surplus for each potential enrollee in our representative ACS sample; the surplus depends on the choice set available to each consumer and not on whether any given consumer actually chose a plan. Aggregating up from the sample gives us the average consumer surplus in each market. We proceed similarly for producer surplus and government spending, scaling the ACS sample to the population level.

Figure 9: Demographic incidence of consumer surplus under counterfactual allocations

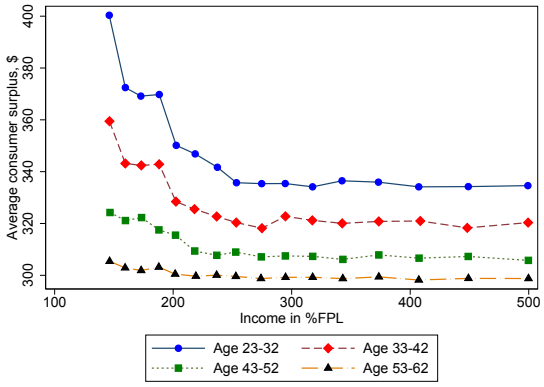
(a) Flat subsidies with market power



(b) Tagged subsidies without market power



(c) No subsidies with market power



(d) No subsidies without market power

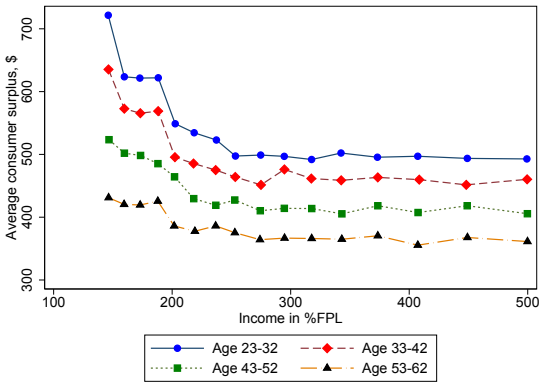
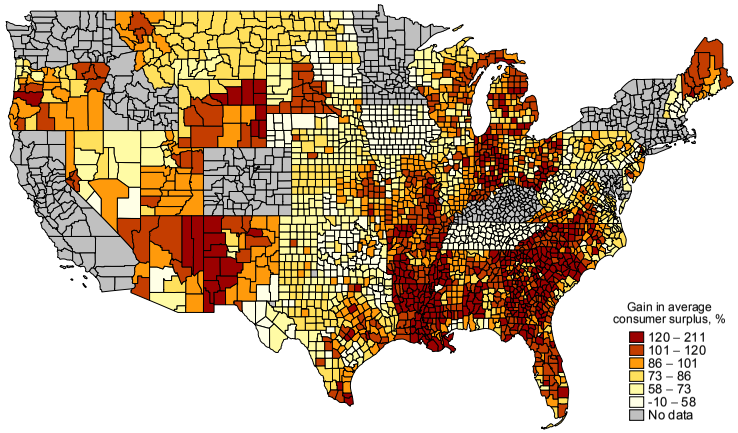
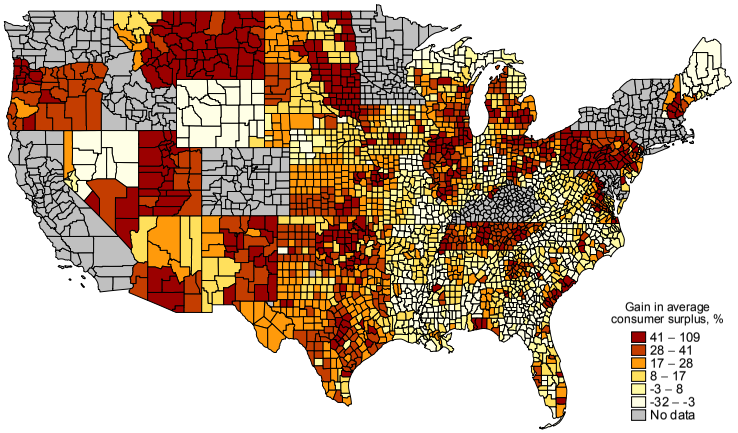


Figure 10: Change in market-level average consumer surplus under counterfactual allocations

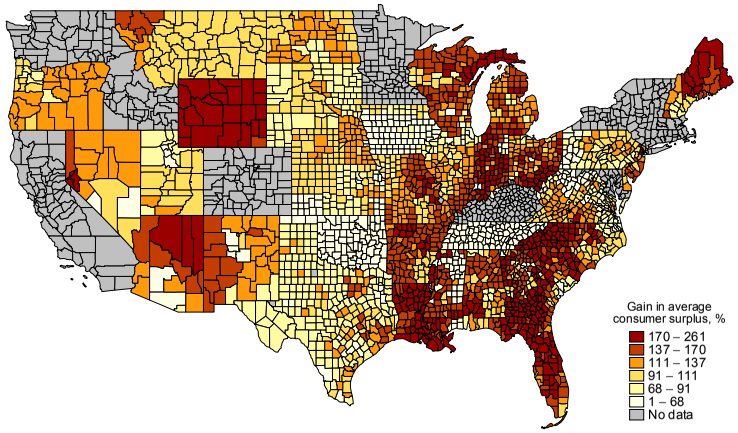
(a) No subsidies vs tagged subsidies (with market power)



(b) Tagged vs flat subsidies



(c) No subsidies vs tagged subsidies (no market power)



(d) No market power vs market power (no subsidies)

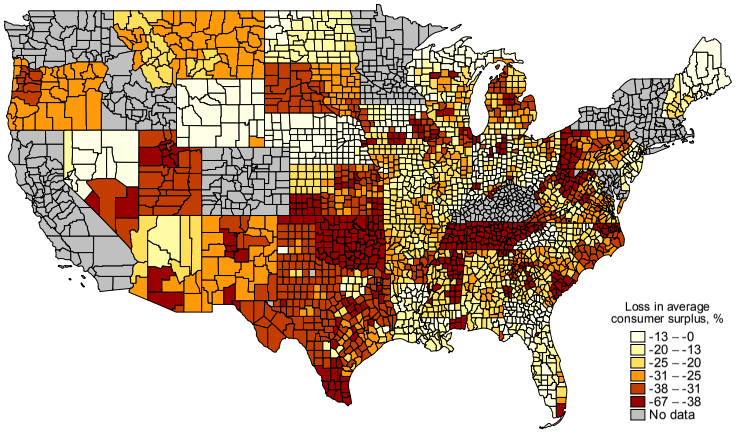


Table 1: Summary statistics

	μ	σ	min	max
Choice set				
Number of plans	39	21	6	165
Number of large insurers	3	1	1	8
Average annual premium (age 40), $	3,768	491	2,558	5,672
Average standard actuarial value	0.66	0.03	0.56	0.72
Enrollment				
Market size	7,878	25,751	32	536,531
Total enrollment	3,436	13,560	12	392,442
Share Outside Option	0.62	0.12	0.15	0.858
Share Bronze	0.09	0.04	0	0.27
Share Silver	0.26	0.1	0.01	0.67
Share Gold	0.03	0.02	0	0.12
Demographics				
Average age	38.4	2.7	27.9	45.8
Average share white	0.83	0.15	0.05	1
Average income, % FPL	249	34	167	391
Share with 94% cost-sharing reduction	0.32	0.09	0.11	0.58
Share with 87% cost-sharing reduction	0.13	0.03	0.01	0.27
Share with 73% cost-sharing reduction	0.11	0.03	0.03	0.25
Average annual premium subsidy, $	2,117	627	597	4,842

Notes: Summary statics on ACA Marketplace plans and consumers in year 2015. Choice set statistics are based on data from Health Insurance Marketplace Public Use Files that have been released by the Center for Medicare and Medicaid Services as well as the Center for Consumer Information and Insurance Oversight. Enrollment statistics are based on county and plan-level enrollment data that have been released by the Center for Medicare and Medicaid Services. Demographic data are based on American Community Survey.

Table 2: Descriptive evidence: demand of partially subsidized consumers

	(1) OLS	(2) Reduced Form	(3) 2SLS	(4) First Stage
2LSP Premium, $000	0.16		-3.27	
	(0.035)		(0.82)	
Share unsubsidized consumers		0.63		-0.19
		(0.078)		(0.038)
Mean Y	0.58	0.58	0.58	2.87
Std. dev. Y	0.24	0.24	0.24	0.22
R-squared	0.52	0.53	-	0.77
N	2,566	2,566	2,566	2,566

Notes: The table reports point estimates of regressions that formalize the relationships recorded in Panels (a) and (b) of Figure 3. Columns (1)-(3) are county-level regressions. The outcome variable in Columns (1)-(3) is the share of individuals eligible for partial subsidies - those with income between 250% and 400% FPL - that purchased any plan on ACA Marketplaces. The outcome variable in Column (4) is the effective premium for the second-lowest cost silver plans that individuals eligible for partial subsidies face. The regression in Column (4) is estimated at the individual level on the ACS sub-sample that satisfies the income restriction of 250% to 400% FPL. Column (1) measures whether the inside enrollment share of partially subsidized consumers is lower when their average 2LSP premium is higher. Column (2) measures whether the inside enrollment share of partially subsidized consumers is higher in places that have more unsubsidized consumers. Column (3) measures whether the inside share is higher when the average 2LSP premium is higher, where we instrument for premiums with the share of unsubsidized consumers. The relationship between premiums of partially subsidized consumers and the share of unsubsidized consumers is recorded in Column (4). All regressions control for state fixed effects. Column (4) includes, but doesn't report individual demographics as described in Section 3.

Table 3: Demand Estimates

	(I)	(II)	(II)
Income-Adjusted Premium Mean ($1000)	−1.94	0.57	−3.47
	(0.12)	(0.04)	(0.19)
Income-Adjusted Premium Std.dev. ($1000)		0.09	
		(0.01)	
Adjusted AV Mean	6.78	5.38	8.67
	(0.43)	(0.34)	(0.08)
Adjusted AV Std.dev		5.64	
		(0.35)	
Brand Fixed Effects	no	no	yes
Plan Characteristics	no	no	yes
Control Function	yes	yes	yes
Moment level	county-metal	county-metal	county-metal

Notes: Standard errors are reported in parentheses. The models are estimated using a simulated method of moments with an objective function that attempts to maximize the match of county-metal level enrollment shares. Individual-specific premiums are constructed by subtracting the individual-specific subsidy that we compute for each individual in the ACS sample using information about their income from the "bid" or list premium for each plan for the corresponding age of the individual. The adjusted actuarial value measure reflect the AV that each individual in the ACS sample would face in each plan depending on their income. For example, individuals with the lowest incomes receive cost-sharing subsidies such that the actuarial value of Silver plans for them become 94% rather than the standard 70%. In Specification number (III), we include the following characteristics: whether a plan is a PPO/HMO/POS/EPO, whether the plan is new to the market, whether it is eligible for HSA accounts, provides out of network and out of country coverage, has a national network of providers, applies quantity limits and exclusions on any services or drugs, requires pregnancy notice, specialist referral, offers wellness programs, disease management, asthma management, diabetes management, depression management, heart disease management, high blood pressure management, back pain management, pain management, pregnancy management, as well as 13 indicators for coverage exclusion for a set of common services.

Table 4: Counterfactual surplus simulations

(a) Subsidy incidence

	Tagged subsidies (preserve market power)	No subsidies
Consumer surplus (in $B)	15	8
Producer surplus (in $B)	6	2
Government spending (in $B)	23	0
Share inside option enrollment	37%	11%

(b) Tagging with market power

	Tagged subsidies (preserve market power)	Flat subsidies
Consumer surplus (in $B)	15	19
Producer surplus (in $B)	6	8
Government spending (in $B)	23	27
Share inside option enrollment	37%	48%

(c) Subsidy incidence without market power

	Tagged subsidies (shut down market power: p=mc)	No subsidies
Consumer surplus (in $B)	23	12
Producer surplus (in $B)	0	0
Government spending (in $B)	32	0
Share inside option enrollment	63%	30%

(d) Market power without subsidies

	Market power (no subsidies)	No market power
Consumer surplus (in $B)	8	12
Producer surplus (in $B)	2	0
Government spending (in $B)	0	0
Share inside option enrollment	11%	30%