

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

Maria Polyakova and Stephen P. Ryan*

PRELIMINARY AND INCOMPLETE:
PLEASE DO NOT CITE OR REDISTRIBUTE

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 subsidy funds, may lead to substantial market distortions in an environment where public insurance is provided by strategic private firms and the level of subsidies 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 and NBER (mpolyak@stanford.edu); Ryan, Washington University in St. Louis, CESifo, and NBER (stephen.p.ryan@wustl.edu). We thank Alan Jaske and Lynn Hua for excellent research assistance. We gratefully acknowledge support from the Agency for Healthcare Research and Quality (R03-HS024800). 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 consumer surplus varies substantially across local geographies. As expected, consumers that receive highest premium and cost-sharing subsidies enjoy the highest consumer surplus. 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. 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.

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. 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 to 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. \(2016\)](#), and [Jaffe and Shepard \(2017\)](#) that explore the questions about subsidies, competition, and market design in the context of Medicare Advantage, Covered California, Medicare Part D, and Massachusetts Health Insurance Exchange, respectively. We contribute to this specific branch of the literature by exploring the role of subsidy tagging to consumer observables.

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 hypothesize in this setting. Sections 4.1 and 4.2 lay out the demand and supply models, respectively. Section 5 reports model estimation results. Section 6 proceeds to discuss the efficiency properties of observed and counterfactual subsidization mechanisms. Section 7 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 plans 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

¹While 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. These individuals qualify for on average \$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 2 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 2 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. Column (1) reports the results of a regression that captures the same relationship as Panel (a) of Figure 2: 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.

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

We estimate this specification on the ACS sample of 206,064 potential consumers with income between 250% and 400% FPL, clustering standard errors at the county level. Similar to Panel (a) of Figure 2, we estimate a negative relationships: individuals in markets with 10 percentage point more of elastic consumers face on average \$45 lower premiums.

Columns (2) to (4) 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% FPL that purchased a plan on the ACA Marketplace.³ In column (2) 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,

²Both the x and y axis are residualized to account for fixed differences across states and for the exact income level of potential consumers.

³We 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.

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 (3) 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 2). The 2SLS specification of demand in Column (3) produces a more meaningful estimate of the demand slope, suggesting that the inside share decreases by 0.1 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 subsidized to what extent.

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 maximize 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” 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:

$$u_{ij} = -\alpha p_{ij} + \beta D_i X_{ij} + \epsilon_{ij} \tag{2}$$

Where p_{ij} denotes consumer i ’s premium for insurance plan j , vector X_{ij} captures observ-

able characteristics of the plan, such as coverage generosity for various providers and services. 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 preferences to vary with observable consumer characteristics D_i that include age, sex, race, and income.

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 plan names 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 Waldfoegel (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 something about the fundamental cost structure of plan A , 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. We choose to take average residuals across all age groups as control function, as the premium that enters the utility function is individuals specific and hence varies with age.

4.2 Supply

4.2.1 Profit function

Insurers on the ACA Marketplaces face a complex problem of deciding which geographic markets to enter, what kinds of plans to offer and how to price them. In this analysis we focus on the part of the problem when insurers set prices for their products conditional on entry and contract design decisions. Modeling price-setting in ACA markets poses a significant challenge, as we have to take into account an array of regulatory provisions. We start with a detailed accounting of payment flows in the market. We then discuss how we attempt to make the supply-side model empirically tractable.

Suppose a 40-year old, non-smoking individual i that lives in market t purchases plan j .

For this individual, plan j collects revenue that consists of several pieces. First, the insurer gets its “bid” (or list price) for a non-smoking 40-year-old in this market.⁴ Second, the insurer receives revenue from risk-equalization programs that we describe below. The bid is only partially paid by the individual. Consumer i receives subsidy z_i that is a function of i ’s income and the bid of the second-lowest cost silver plan on market t . The consumer then pays the difference between the bid and the subsidy as what we call an “effective premium.” If the subsidy is higher than j ’s bid, the consumer doesn’t pay any additional premiums, but also doesn’t receive any rebates.⁵

On the cost side, the ex-post costs that plans experience and ex ante costs that plans expect differ for each enrollee. Let the overall expected healthcare spending of consumer i be h_i . This spending depends on consumer’s underlying health risk, which we denote with r_i , as well as the cost-sharing features of plan j , which could generate moral hazard. We denote j ’s cost-sharing characteristics with ϕ_j . Then, h_i is a function of r_i and ϕ_j . Plan j ’s cost for consumer i is not equal to h_i . Instead, the plan pays only a portion of h_i , 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 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 cost for enrollee i is $c_i(r_i, \phi_j, D_i)$, where D_i denotes demographics and includes income.

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

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

Suppose that for any plan j , there is a baseline plan-specific cost c_j 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: $c_{ij}(r_i, \phi_j, D_i) = c_j + \tilde{c}_{ij}(r_i, \phi_j, D_i)$. We further note that while the split of insurer revenue between consumer premium p_{ij} and subsidy z_{ij} is important in determining how many people choose to enroll in the plan, once

⁴In what follows, we abstract from bid differences for smoking and non-smoking enrollees, as well as from individual versus family coverage, as we do not have information on enrollment across these groups. This simplification is ameliorated by the fact that frequently, premiums for smoking adults and family enrollments are scaled versions of the baseline premiums for single, non-smoking adults.

⁵In practice, the subsidy operates as a tax credit, so it is estimated at the time of purchase given the information about income, and can be adjusted during tax filing.

a consumer enrolls, the premium and the subsidy add up to insurer's bid. In other words, $p_{ij}(D_i, b_j, b_{-j}) + z_i(D_i, b_j, b_{-j}) = b_j$. Using this notation, we can then re-write the profit of plan j from enrolling individual i of age a as follows:

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

The individual-specific cost term allows for the presence of advantageous or adverse selection (depending on its sign) that is a function of plan characteristics ϕ_j . This set-up does not allow for additional selection purely on premiums conditional on all components of ϕ_j .

Three programs exist on ACA Marketplaces aimed at equalizing the costs of all enrollees from insurers' perspective, so as to reduce the incentives for cream-skimming and ameliorate the consequences of adverse selection into more generous plans. It is easier to think about these programs as affecting insurers' costs; however, in practice, the programs often 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 or pay into a so-called risk corridor program. This program attempts to reduce the ex post volatility in profits that may result from enrolling a particularly good or bad risk pool.

While these programs may be imperfect, to the first order approximation, they attempt to equalize the idiosyncratic portions of enrollees' costs ex ante. 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 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 enrolled in a plan. Finally, let Γ_j denote a risk-corridor transfer to the plan.

Let I_j^a denote the number of people 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:

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

Let $H_j(\phi_j)$ denote any residual selection that is left after risk-adjustment and risk-corridors. In other words, $H_j(\phi_j) = \sum_{i \in j} \eta_{ij} - R_j - \Gamma_j$. Under the assumption of perfect risk-equalization programs that completely neutralize differences in ex ante net idiosyncratic shocks, this term would be zero. Without this assumption, $H_j(\phi_j)$ can be either positive or negative depending on the nature of residual selection. If we keep the assumption that $H_j(\phi_j)$ depends only on ϕ_j , it does not enter the first order conditions. In that case, it is not crucial to assume anything about the sign of H . For our subsequent analysis we let $H_j(\phi_j)$ be zero, or in other words, assume that $\sum_{i \in j} \eta_{ij} = R_j + \Gamma_j$. The profit then becomes:

$$\pi_j(b_j; b_{-j}) = \sum_a I_j^a b_j^a - \sum_a I_j^a c_j^a \quad (6)$$

Rewriting this using share notation, we get:

$$\pi_j(b_j; b_{-j}) = \sum_a s_j^a M^a b_j^a - \sum_a s_j^a M^a c_j^a \quad (7)$$

Empirically and according to ACA statutes, insurer bids across different ages follow a fixed schedule that doesn't appear to change substantially over time. This observation allows us to simplify the problem further. We assume that there is a fixed set of age-specific multipliers that apply to bids and costs. To reduce the computational burden, we take these multipliers from the data. Let a multiplier vector for plan j be τ_j . The profit equation for plan j then becomes:

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

or, re-arranging:

$$\pi_j(b_j; b_{-j}) = (b_j - c_j) \sum_a s_j^a M^a \tau_j^a \quad (9)$$

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

$$\pi_f(\mathbf{b}) = \sum_{j \in f} \left[(b_j - c_j) \sum_a s_j^a M^a \tau_j^a \right] \quad (10)$$

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

4.2.2 First order conditions for pivotal and non-pivotal plans

Within a standard Bertrand-Nash game, insurers choose bids that maximize their profits taking into account the actions of other firms. The plan-level first order condition implied by the profit function in 10 is:

$$\frac{\partial \pi_j}{\partial b_j} = (b_j - c_j) \sum_a \frac{\partial s_j^a}{\partial b_j} M^a \tau_j^a + \sum_a s_j^a M^a \tau_j^a = 0. \quad (11)$$

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

$$S - \Omega(B - C) = \mathbf{0}. \quad (12)$$

where

$$S_j = \sum_a s_j^a M^a \tau^a \quad (13)$$

$$(B - C)_j = (b_j - c_j) \quad (14)$$

and

$$\Omega_{kj} = \begin{cases} -\sum_a \frac{\partial s_j^a}{\partial b_k} M^a \tau^a & \text{if } \{j, k\} \in f, \\ 0 & \text{else,} \end{cases} \quad (15)$$

The key term of the first order condition is the derivative of the (age-specific) share with respect to the bid: $\frac{\partial s_j^a}{\partial b_j}$. This derivative reflects how much the demand of age group a for plan j changes when this plan changes 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. In our setting, bids and premiums are related via premium subsidy. This subsidy is a function of a bid set by the second-lowest cost Silver plan. These 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 subsidies that consumers can apply to all plans. We assume a complete information game, in which pivotal plans know that they are pivotal and

set bids accordingly.

The share derivative consists of two parts: the response of consumer premiums to a change in the bid and the response of enrollment to the change in consumer premium. Consider a change in consumer premiums in response to a slight increase of plan j 's bid. For non-pivotal plans, when the bid goes up by epsilon, all consumer prices (i.e. for all income levels) weakly go up by epsilon, except in those cases where consumer prices remain zero because the new bid is still below the subsidy level. For pivotal plans, when the bid goes up by epsilon, its own price doesn't change (because the subsidy is adjusted exactly by epsilon), but prices of all other plans weakly go down by epsilon. The response of enrollment to the change in consumer premiums follows from the demand model.

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 11 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.

5 Estimation results

5.1 Demand Parameters

We significantly adapt the standard estimation routine for discrete choice of models to incorporate different levels of aggregation of enrollment data that we observe. In particular, we do not observe the market shares of each plan, 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 results of increasingly rich demand specifications. We start with a very basic specification in Column (1) that includes only the individual-specific premium, the individual-specific measure of average generosity, and the control function. This specification captures the key patterns. Individuals dislike higher premiums and like more coverage

- consumers are willing to pay almost \$400 a year for a ten percentage point increase in average generosity of plans. Column (2) adds brand and state fixed effects to the specification. Column (3) further adds the full vector of plan characteristics. Conditional on these characteristics, consumer willingness to pay for 10 percentage points of average generosity increases to \$500.

Table 4 reports the fit of all models aggregated to metal level enrollment across the whole US - all models closely match these aggregated enrollment shares. Figure 4 illustrates the county-metal level fit of the richest specification; the model is able to capture a substantial amount of variation in the data.

6 Counterfactuals and Welfare

6.1 Welfare under Observed Subsidization Mechanism

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:

$$CS(\theta_i) = \frac{1}{\alpha_i} \left[\gamma + \ln \left[1 + \sum_{j=1}^J \exp(v_{ij}(\theta_i)) \right] \right], \quad (16)$$

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.⁶ We integrate out over consumer heterogeneity to obtain consumer surplus:

$$CS = \int CS(\theta) dF(\theta). \quad (17)$$

⁶Euler'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.

Figure 7 illustrates the average level of consumer surplus by county. Surplus levels are very different across the country, as would be expected with subsidy tagging. Areas where more enrollees are eligible for premium and cost-sharing subsidies capture higher consumer surplus. Figure 8 illustrates that within geographic areas, the distribution of surplus varies substantially between different income-age groups. Surplus is decreasing with income - individuals that have income that makes them ineligible for subsidies get consumer surplus that is about 20% of the surplus for lowest-income individuals. Moreover, for groups with partial subsidies, surplus is much higher for older consumers. These surplus distribution patterns are again consistent with targeted distribution of public funds for the tax credits in this program.

- To get the right marginal costs weighted by enrollment, we need to simulate outcomes and then assign average marginal costs for those plans, then compare that graph to the bin-scatter against accounting costs

7 Discussion

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 providing social insurance and welfare benefits is changing, many policies still rely on traditionally “tagged” policies. 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. Consider a social insurance example. Suppose 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. This knowledge generates a strong incentive to set higher prices for insurance in this poorer market; in the presence of market power, these incentives are 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.
- 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.
- 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.
- 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.

Figure 1: Consumer interface on healthcare.gov

The screenshot displays the HealthCare.gov website interface. At the top, there are navigation tabs for 'Individuals & Families' (selected), 'Small Businesses', 'Log in', and 'Español'. A dark blue banner contains an important notice: 'Important: The premiums below are only estimates. You'll need to fill out a Marketplace application to get actual plan prices. Some plans and details you see here may change.'

The main content area is titled '121 Health Plans'. Below this, there are filters for 'Viewing:' (Health Plans, Dental Plans) and 'Sort:' (By Monthly Premium, By Deductible). A 'BACK TO QUESTIONS' button is located on the left.

The left sidebar, titled 'NARROW YOUR RESULTS', includes sections for 'Premium' (with filters for less than \$200, \$300, \$400, and \$500), 'Coverage categories' (Bronze, Silver, Gold, Platinum plans), 'Plan Types' (PPO, HMO, POS), and 'Insurance companies' (Aetna).

The main plan detail card for 'Health Choice Insurance Co. · Health Choice Value Bronze' (Plan ID: 70239AZ0010043) shows the following key metrics:

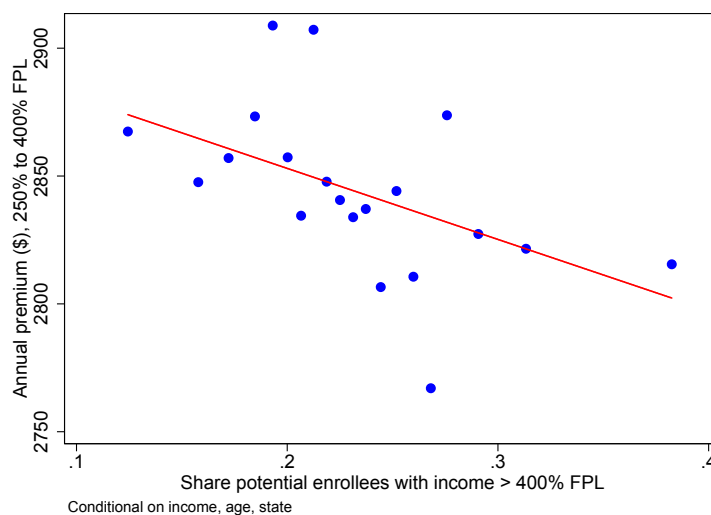
- ESTIMATED MONTHLY PREMIUM:** \$153
- ESTIMATED DEDUCTIBLE:** \$5,000 (Estimated individual total)
- ESTIMATED OUT-OF-POCKET MAXIMUM:** \$6,600 (Estimated individual total)

Additional details include 'COPAYMENTS / COINSURANCE' (Primary doctor: \$20 Copay after deductible; Specialist doctor: \$50 Copay after deductible; Emergency room care: \$500 Copay after deductible; Generic drugs: \$15 Copay after deductible) and 'PEOPLE COVERED' (1 (Age 40): Covered). A 'MORE INFORMATION' section lists links for Summary of Benefits, Plan brochure, Provider directory, and List of covered drugs. A prominent blue button at the bottom of the card reads 'LEARN MORE ABOUT THIS PLAN'.

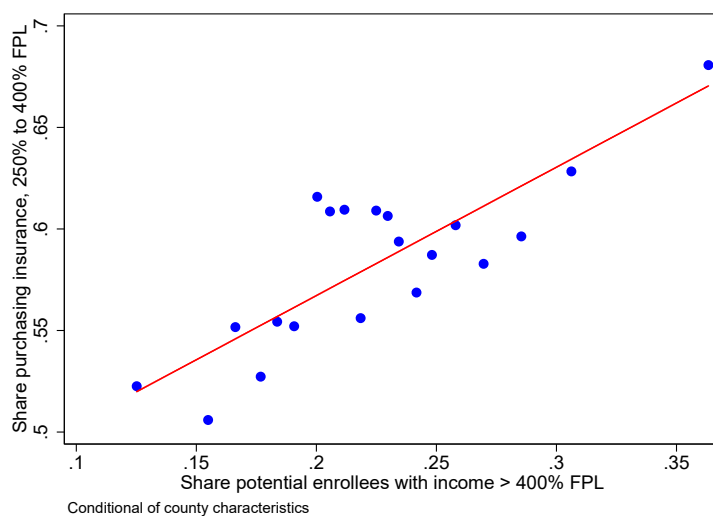
Notes: Snapshot of one of 121 plans that were offered to 40-year old individuals in Cook County, IL in 2015. The premium that individuals see on the web page incorporates their individual premium subsidy if they enter their income information during the selection process.

Figure 2: Descriptive evidence: tagged subsidies and premiums

(a) Effective premiums

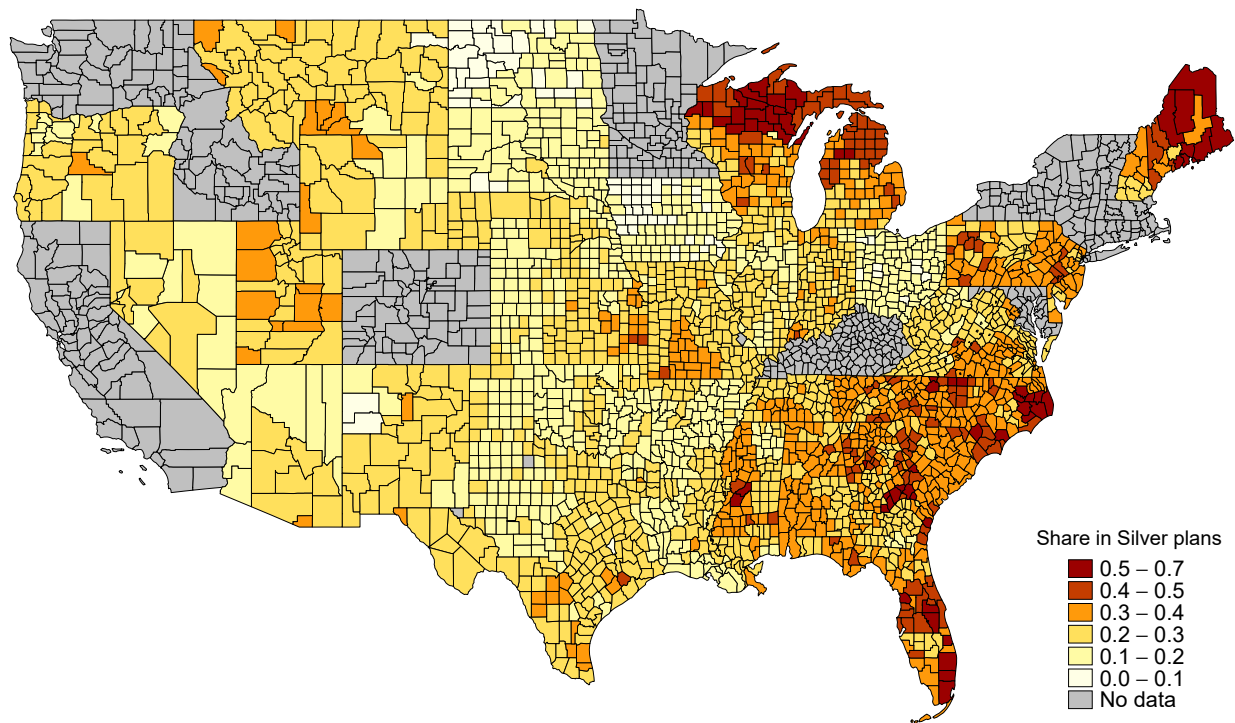


(b) Enrollment



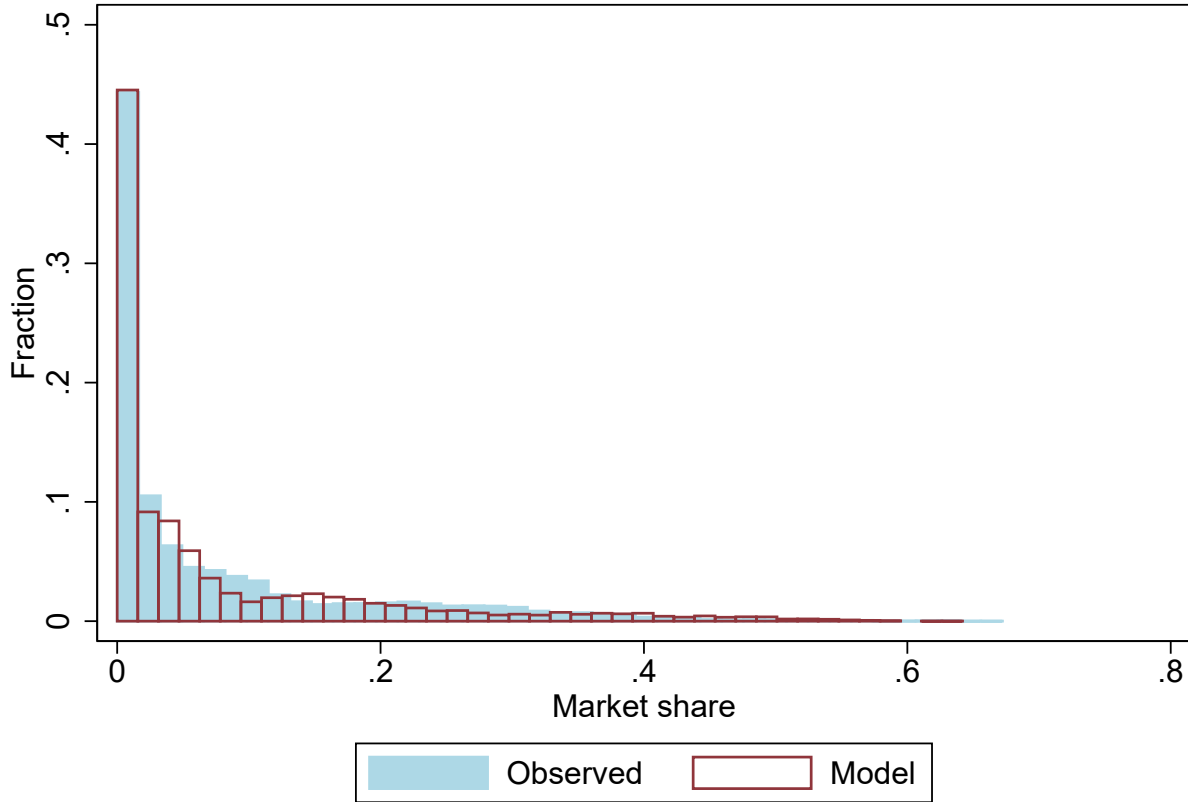
Notes: The panels illustrate the descriptive relationship between the demographic composition of counties and the market experience of consumers with income above 250% FPL and below 400% FPL. The top panel computes for each county the share of individuals that are not eligible for subsidies (income above 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 3: Empirical moments: share of Silver plans 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 generously 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 4: Demand model fit



One observation is county-metal label

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 blue-colored histogram records the distribution of county-metal label enrollment in the data. The red-contoured histogram records the (in-sample) predictions of the demand model.


Figure 5: Marginal cost estimates and accounting costs



All values restricted to >0; account. costs restricted to <10. Drops 17% of plans.

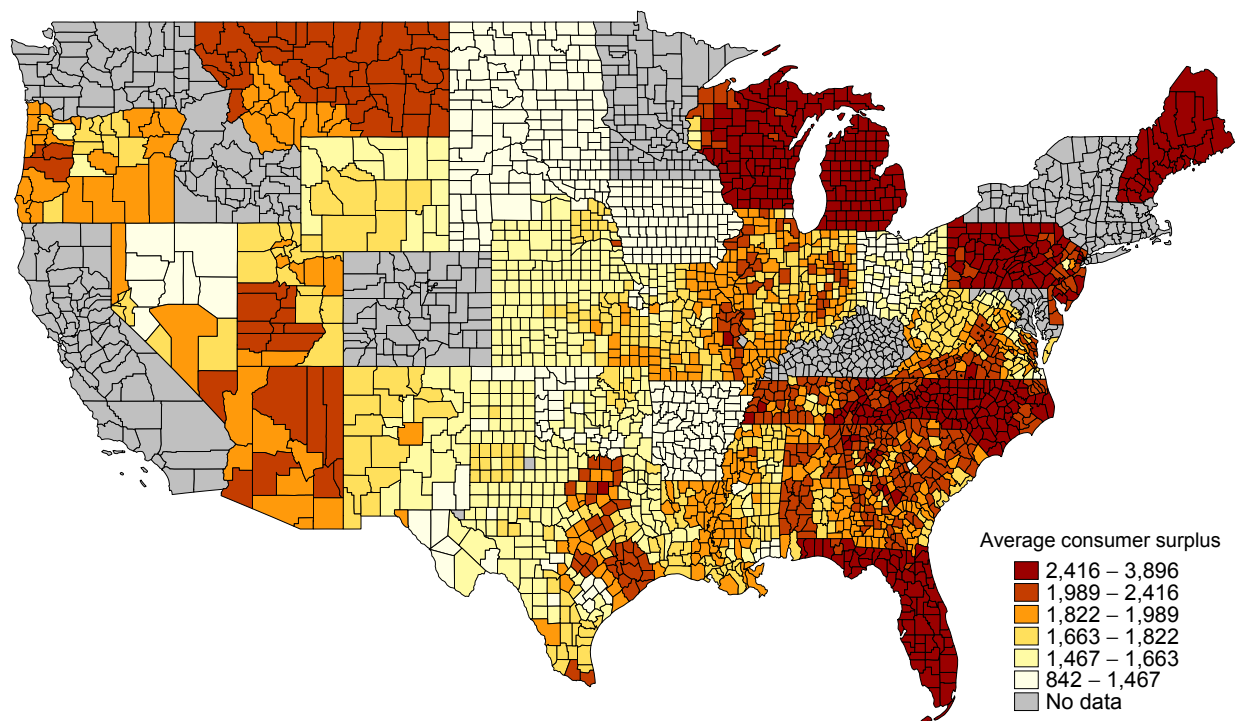
Notes: Marginal cost estimates are done for a 21 year old consumer. We assume that marginal costs scale with respect to age with the same multipliers as observed prices. Accounting costs are computed as an average claim across all enrollees (i.e. of all ages) in the plan. The costs on both axes are in thousands of dollars per year.

Figure 6: Estimated mark-ups by insurer




Notes: Marginal cost estimates are done for a 21 year old consumer. We assume that marginal costs scale with respect to age with the same multipliers as observed prices.

Figure 7: Geographic distribution of consumer surplus



Notes: The map illustrates the average annual consumer surplus from ACA Marketplace plans in each county of the federally facilitated states. Consumer surplus calculation is described in section 5. 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.

Figure 8: Distribution of consumer surplus by income and age



Notes: The figure illustrates the average annual consumer surplus from ACA Marketplace plans in by age and income bracket. Consumer surplus calculation is described in section 5. 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 by age and income bins.

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)	(2)	(3)	(4)
	Inside share	Inside share	Inside share	2LSP Premium
Share unsubsidized consumers		0.632 (0.078)		-459.2 (104.0)
2LSP premium	0.000 (0.000)		-0.001 (0.000)	
State fixed effects	Yes	Yes	Yes	Yes
Mean Y	0.583	0.583	0.583	\$ 3,396
Standard deviation Y	0.244	0.244	0.244	\$ 1,395
Number of observations	2,566	2,566	2,566	206,064

Notes: The table reports point estimates of regressions that formalize the relationships recorded in Panels (a) and (b) of Figure 2. 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)
Constant	-5.440 (0.082)	
Income-Adjusted Premium Mean (\$1000)	0.649 (0.084)	0.723 (0.063)
Income-Adjusted Premium Std.dev. (\$1000)	0.495 (0.0.254)	0.289 (0.121)
Adjusted AV Mean	5.667 (0.128)	5.343 (0.252)
Adjusted AV Std.dev	3.144 (0.505)	3.665 (0.206)
Control function	-0.101 (0.024)	-0.082 (0.021)
Brand Fixed Effects	No	Yes
Plan Characteristics	No	Yes
Objective Function	164,260	104,022

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: Aggregate Demand Fits by Metal Level

Metal Level	Empirical	(I)	(II)	(III)
Catastrophic	0.00	0.00	0.00	0.00
Bronze	0.09	0.09	0.10	0.09
Silver	0.30	0.25	0.24	0.28
Gold	0.03	0.04	0.05	0.04
Platinum	0.01	0.00	0.01	0.01

Notes: Fit simulation for each model specification in Table 3. The model is simulated for each individuals in the representative ACS sample and then aggregated to the level of metal level enrollment shares for all 37 federally facilitated states. The first column in the table reports observed enrollment shares in 2015 data.