Subsidy Targeting with Market Power*

Maria Polyakova† Stephen P. Ryan‡

September 30, 2019

Abstract

Public welfare programs commonly link their benefits to observable characteristics of potential recipients. We show that when the provision of the public benefit is contracted out to imperfectly-competitive firms, targeting subsidies using observable characteristics of consumers can distort both efficiency and equity of market allocations. Using a structural model of supply and demand, we illustrate the underlying mechanisms and quantify the distortions empirically in the context of means-tested subsidies in privately-provided health insurance markets under the Affordable Care Act. In this setting, market power increases the welfare loss from subsidy targeting, vis-a-vis income-invariant subsidies, by 33 percent.

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

Keywords: targeting; in-kind transfers; means-testing; health insurance; ACA

*We gratefully acknowledge support from the Agency for Healthcare Research and Quality (R03-HS024800). We are grateful to David Card, Pietro Tebaldi, Mark Shepard, and Keith Ericson for their discussions of the paper, and seminar participants at Daniel McFadden 80th Birthday Conference, Junior Health Economics Summit, Ohlstadt Workshop, Stanford University, London School of Economics, American Society of Health Economists, Yale University, University of Connecticut, Columbia University, Bates White Life Sciences Symposium, University of Munich, 2018 REStud Reunion Conference, Rice University, University of Utah, University of Leuven, Georgia State University, UNC Chapel Hill, SUNY Stony Brook, NBER Health Care Meetings, and NBER Public Economics Meetings for helpful comments. We thank Alan Jaske and Lynn Hua for excellent research assistance.

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

‡Washington University in St. Louis Olin Business School, NBER, and CESifo (stephen.p.ryan@wustl.edu).

1 Introduction

Public welfare programs have a long history of linking their benefits to observable characteristics of potential recipients, such as age, income, or employment status. Conceptually, targeting of transfers helps improve the allocation of public dollars to the most needy recipients. Targeting on observable characteristics, however, may imperfectly identify the most needy, or may induce recipients to distort their behavior in order to qualify for the benefit. The extensive theoretical and empirical literature studying the costs and benefits of targeted transfers has almost exclusively focused on these demand-side distortions, assuming that benefits are provided by a benevolent government.

In practice, however, governments are increasingly turning to profit-maximizing firms to provide public benefits.[1] How the government should pay private firms to administer public benefits is a central policy question. The amount of public funds distributed through subsidies is staggering: the US government expects to pay $560 billion in subsidies for privately-provided health insurance in 2019 alone, with that number growing to $1.2 trillion by 2029 (CBO, 2019a,b,c). Subsidy mechanisms directly influence the economic efficiency of public programs and their ability to achieve underlying redistributional goals; yet, relatively little is understood about the optimal design of these instruments when goods and services are provided by profit-maximizing firms. In this paper, we combine insights from the literature in public economics and industrial organization to investigate the distributional and efficiency consequences of targeting subsidies on observable characteristics of the recipients in the presence of firms with market power.

The interaction of market power on the supply side and targeted taxes or subsidies can generate distortions in equilibrium. There are two key, potentially countervailing, forces at work. First, subsidy targeting introduces heterogeneity in consumer-facing prices across markets and across consumers within markets. All else equal, firms have incentives to raise list prices in markets where consumers receive more generous subsidies. In the presence of market power, these incentives are not dissipated by competition and may lead to perverse equilibrium outcomes. For example, if consumer subsidies are means-tested and price discrimination is forbidden, as is the case in many public programs, the near-poor may end up paying more for identical products if they live in markets with many poor consumers. Sec-

[1] Notably, paying private agents to provide public benefits is particularly common in health care and health insurance. Poterba (1996) outlines the stark differences in the choice of public policy instruments between education and health care, highlighting the much more common payment to private agents—rather than direct provision—in the health domain.

ond, targeted subsidies change the relative importance of different consumers in the firm's profit-maximization problem. If subsidized consumers are also more price-sensitive, bringing these consumers into the market through subsidies that retain their elasticity of demand on the margin can generate negative pressure on prices, benefiting unsubsidized consumers. We call this link between the distribution of consumer types and equilibrium prices under subsidy targeting with market power a *demographic externality*.[2]

We explore equilibrium effects of subsidy targeting in the presence of market power empirically in the setting of publicly-subsidized, privately-provided health insurance Marketplaces created under the Affordable Care Act (ACA). The Marketplaces provide a fruitful empirical laboratory for understanding the effects of subsidy targeting in the presence of market power for several reasons. First, 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. Second, these tax credits are targeted based on consumers' household income, thus following a traditional approach of means-tested public benefits. Third, the exercise of market power is likely, as firms are both allowed to set uniform prices to maximize profits (subject to regulatory limitations) and most markets are highly concentrated with two or fewer insurers. For these reasons, we anticipate the interaction between subsidy targeting and market power to be quantitatively important.

To quantify this interaction, we formulate and estimate a structural model of supply and demand for ACA Marketplace plans. We utilize the unique institutional setting of the Marketplaces to implement a novel identification strategy in our demand estimates. As in Tebaldi et al. (2019), we leverage price variation across consumers of different ages and incomes that is generated by the subsidy mechanism. Within each market, consumers of the same age face different prices for the same product depending on their income; analogously, consumers with the same household income, but different ages, face different prices for the same product. This structure of the data allows us to identify the coefficient on the price parameter in the utility function under a semiparametric demand specification, where we estimate product-specific utility levels. This flexibility in demand specification is important in our setting with highly multi-dimensional products.

The regulatory variation in prices across age-income bins also allow us to capture the

[2] This idea is closely related to the "preference externality" as proposed in George and Waldfogel (2003) and Waldfogel (2003), in which the features of differentiated products available to a consumer depend on the preferences of the majority of consumers' neighbors. The mechanism that generates the externality in our setting, however, is different—the externality is driven by a policy instrument that exogenously changes (slope and/or levels, depending on subsidy design) of demand functions for some, but not other consumers.

underlying heterogeneity in demand elasticity across different types of consumers, by letting the marginal utility of income parameter vary across age and income brackets. Using this strategy, we estimate reasonable levels of marginal utility of income and intuitive substitution patterns for all groups of consumers. We then proceed to derive a profit function for insurers on this market, trying to balance the institutional and especially regulatory detail with the computational tractability of the model. We arrive at first-order conditions that allow us to recover marginal costs for each product-market combination. Our estimates from the inversion of the first-order condition at the product-market level are consistent—both in terms of levels and relative ranking of products—with accounting data that we observe at the product level.

With these estimates of demand and supply in hand, we compute a series of counterfactuals that help us investigate the interaction between market power and targeting. We begin with a set of stylized model simulations that illustrate the demographic externality in action in our setting. We then turn to a series of counterfactuals that modify the subsidy regime. To establish a quantitative baseline, we start by measuring the efficiency and incidence of subsidies under the observed allocation. We then compute a set of counterfactuals where we impose either perfect competition, non-targeted subsidies, or both. This allows us to assess the efficiency and distributional effects of market power and subsidy targeting, both in isolation and jointly.

To illustrate the demographic externality, we start with two conceptual exercises that show the complex interplay between targeted subsidies and the demographic distribution of consumers within a market. In our first stylized experiment we introduce subsidies for the highest-income consumers, who were previously not eligible for subsidies, at the same level as received by the lowest-income consumers. This leads to a shift out in the aggregate demand curve and induces insurers to increase prices. As a result, consumers not directly affected by the subsidy change experience a negative demographic externality; their consumer surplus falls because their neighbors now get subsidies. As discussed above, however, the demographic externality depends both on the amount of the subsidy and also on who is eligible to receive it. To show this, our second exercise endows those newly-subsidized highest-income consumers with the demand curve of the lowest-income consumers. In our setting, the lowest-income consumers have a more elastic demand function, so even though subsidies still increase in this second experiment, prices actually move down in response to a more elastic demand, implying that the newly-subsidized consumers now exert a positive demographic externality on other consumers. How these two opposing forces play out in any given market in equilibrium is an empirical question, which we turn to next.

We compute a set of counterfactual simulations that allow us to measure the efficiency of the observed equilibrium and to document the incidence of means-tested subsidies across consumers and producers under this allocation. We find that under the observed subsidy regime, the ACA Marketplace markets generate \$42 billion in consumer and producer surplus for \$30 billion in nominal spending. The net government outlays are, however, lower—we estimate that the government would have spent \$16 billion on the same consumers through other channels (such as, for example, uncompensated care) had they not been in the ACA Marketplaces. We thus estimate the net government outlays of the program to be \$14 billion. Accounting for the cost of public funds of 30 cents on a dollar, the one dollar of these \$14 billion in extra government spending invested in this market generates \$2.33 of surplus in return. Contrasting this allocation to a scenario that shuts down any public premium subsidies, we conclude that the consumers capture the majority of subsidy payments. We estimate the pass-through of the subsidy-induced surplus to be 84 percent to consumers and 16 percent to firms. The distribution of surplus is highly heterogeneous across different groups of consumers and different geographic areas. While no consumers end up being worse off on average when means-tested subsidies are introduced (relative to a situation with no subsidies), lower income consumers that are targeted by means-tested subsidies experience up to a four-fold higher increase in consumer surplus.

With this set of baseline estimates for the existing allocation, we next turn to computing a series of counterfactuals to help understand the role that market power, targeting, and their interaction play in determining equilibrium outcomes. To understand the role of market power in the observed mechanism, we first compute an equilibrium in the environment that keeps means-tested subsidies fixed, but imposes perfect competition on the supply side. When prices are set at marginal cost, average premiums paid by consumers drop by nearly \$250 (25 percent), enrollment increases by 27 percent, consumer surplus goes up by nearly \$10 billion, from \$38 to \$47 billion, while government expenditures on premium subsidies actually decrease by \$3 billion in nominal spending (net government expenditures decrease by about \$0.5 billion). The return on a dollar of public spending increases from \$2.33 to \$2.75. Overall, market power clearly manifests itself in increased prices and reduced consumer surplus in this market and it is quantitatively important.[3]

To isolate the influence of the subsidy design, we next compare the equilibrium above,

[3]A separate exercise that drops any public subsidies allows us to see the pure effect of market power even more clearly. In the absence of subsidies, list prices for a 20-year old consumer are almost \$500 higher in the presence of market power and enrollment is 31 percent lower.

which keeps means-tested subsidies fixed but imposes perfect competition, to an equilibrium where both subsidy targeting and market power are removed. We replace the targeted subsidy structure with a flat, income-invariant, voucher. We focus on the voucher level that keeps nominal government spending on subsidies the same as under perfect competition with means-tested subsidies. Switching from means-tested to flat subsidies when firms are forced to price at marginal cost has almost no effect on aggregate enrollment, but has a large effect on consumer surplus and what type of consumers choose to enroll. Enrollment patterns shift, with many more higher-income consumers enrolling under flat vouchers. Consumer surplus increases from $47 billion to $54 billion, while the net government spending decreases to $11.3 billion. The combination of these outcomes leads to a $3.72 return on each dollar of public funds, which is nearly a whole $1 (or 36 percent) higher return. In aggregate surplus terms, we estimate that moving from a flat subsidy to means-testing, keeping nominal government expenditures fixed, leads to a 24 percent decline in welfare. These results are intuitive and clearly highlight the equity-efficiency trade-off that is inherent in subsidy targeting. While spending the same amount of public funds through non-targeted subsidies can substantially increase the size of the pie, it is not a Pareto-improving policy.

Finally, we compare the loss in welfare generated by targeted subsidies in the environment with perfect competition to the loss in welfare that subsidy targeting generates when we retain market power. We again focus on the voucher level that holds government spending on premium and cost-sharing subsidies constant relative to the case with means-tested subsidies (which here is just our observed baseline). We observe similar patterns. While overall consumer enrollment stays practically the same, we estimate significant distributional differences across income and age. Far fewer lower-income consumers purchase the good, as their effective subsidy has decreased, while far more higher-income consumers do the reverse. There is also a shift in the age distribution towards more uniform enrollment across ages than the older-skewed majority in the observed baseline. With flat vouchers, both consumer surplus and insurer profit increase by nearly $9 billion to a combined $50 billion, while the net government expenditure decreases by nearly $2.5 billion. Return on a dollar of public funds increases from $2.33 to $3.37. In terms of aggregate surplus, we estimate that moving from flat to means-tested subsidies in the presence of market power leads to a 32 percent loss in welfare.

Returning to our initial motivation to understand how market power and subsidy targeting interact, we can interpret our findings in several ways. One salient point of view is to consider the policymaker's perspective when considering the equity-efficiency tradeoff involved with using targeted versus non-targeted subsidies. Without considering market power on the supply side, the policymaker would conclude that the deadweight loss generated by targeted subsidies is $9.4 billion, or 24 percent of total welfare. In contrast, accounting for the targeted subsidies interaction with market power changes that estimate dramatically; total welfare is $6.1 billion lower and the deadweight loss is $11.3 billion, or 32 percent of total welfare. This is the primary message of the paper: accounting for the presence of market power can have a large effect on the estimates of the efficiency of standard subsidy mechanisms that have long been used by policy-makers in many settings. Ignoring that interaction understates the deadweight loss of targeted subsidies by 33 percent.

Our analysis relates to several different strands of the literature. First, the paper is related to the extensive theoretical and empirical literature 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, 2018 are among recent empirical applications). Traditionally, the literature on the optimal targeting of taxes and subsidies has assumed that benefits are provided directly by the government, so only demand-side distortions can occur. We add to this literature by introducing the idea that targeted subsidies may generate additional distortions in the presence of market power on the supply side, and by quantifying the extend of this distortion in an empirical example. This idea is closely related to several recent papers that have documented how private firms strategically interact with consumers who receive public benefits in-kind or in-cash. For example, Cellini and Goldin (2014) and Fillmore (2019) have examined the "Bennett hypothesis" on the relationship between federal grants and college tuition; Rothstein (2010) empirically examined how firms set wages in the presence of the Earned-Income Tax Credit, while Stantcheva (2014) considers the theory of optimal taxation in an environment where government policies and employer decisions interact; Meckel (2019) and Goldin et al. (2018) consider the effects of food assistance programs on prices in grocery stores.

Through our empirical application we contribute to a growing literature that examines the design of subsidies in health insurance. This literature has investigated the effects of tax subsidies on employer-provided health insurance, for example in Gruber and Washington (2005); Gruber (2005); Gruber and Poterba (1996); further, 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 closely related to Curto et al. (2015), Tebaldi (2017), Tebaldi et al. (2019), Decarolis

(2015), Decarolis et al. (forthcoming), Jaffe and Shepard (2018), Miller et al. (2019) that all study the design of subsidies in health insurance markets such as Medicare Advantage, Covered California, Medicare Part D, and Massachusetts Health Insurance Exchange. These papers focus on the idea that subsidies linked to prices of insurers may distort allocations. Einav et al. (2019) consider the differences in market outcomes when public payments channeled through the demand (through subsidies) or supply (through risk-adjustment) side of the market.

A closely related work is Tebaldi (2017), who considers an externality that arises through insurers' cost in the ACA, as insurers pool risks across consumers of different ages, but cannot perfectly price discriminate across these consumers. In that setting, risk pooling generates a externality via costs, as consumers face higher prices if they are pooled together with high-cost consumers. In contrast, the demographic externality that we study is not specific to markets with selection and can arise in any market with subsidy targeting and where firms posses market power.

Finally, our paper contributes to a rapidly growing literature that studies various aspects of the Affordable Care Act, and especially the economic properties of the ACA Health Insurance Exchanges.[4]

The paper proceeds as follows. Section 2 discusses the theoretical framework. Section 3 gives a brief primer on the institutional setting and describes our data sources for our empirical application. Sections 4.1 and 4.2 lay out the empirical models of demand and supply. Section 4.4 reports estimation results. Section 5 proceeds to discuss the efficiency properties of observed and counterfactual subsidization mechanisms. Section 6 briefly concludes.

2 Stylized Framework

In this section, we highlight the basic intuition that characterizes the relationship between subsidy targeting and market power. To motivate why governments may target subsidies on the basis of observable information about consumers, such as income, we first begin with a simple model of a benevolent government that aims to allocate in-kind subsidies in order to maximize social welfare subject to a fixed budget constraint. In this simple framework, We show that the optimal subsidy is differentiated across consumer types due to differences in the

[4] Among recent studies are Frean et al. 2017; Geruso and Layton 2017; Geruso et al. 2019; Layton et al. 2015, 2016; Layton and McGuire 2017; Layton et al. 2017; Kowalski 2014; Abraham et al. 2017; Courtemanche et al. 2017; Panhans 2019; Dafny et al. 2015; Orsini and Tebaldi 2017; Sacks et al. 2018; Dickstein et al. 2015; Diamond et al. 2018.

curvature of the utility function. Then, taking the idea that subsidies may be differentiated by observable consumer type, we model the incentives of a profit-maximizing firm that provisions the good to the consumers. We show that the combination of subsidy targeting and market power interacts in a complex fashion and can distort market outcomes.

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

Formally, the problem is:

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

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

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

Intuitively, the optimal allocation equates share-weighted marginal utilities across consumer types. From this simple expression, one can immediately see that the optimal split will, in general, be unequal at both the group- and individual-level. In this sense, the social planner finds it desirable to *target* allocations (i.e. subsidies) as a function of consumer type, which is a ubiquitous feature of public assistance programs.[5]

Next, we introduce a profit-maximizing firm to act as an intermediary to provide the good.[6] Following the design of many publicly-subsidized and privately-provided markets around the world, we allow that firm to set a single price, b, after observing both the subsidy structure and the distribution of consumer types. Our model has a similar flavor to those used in the price discrimination literature, particularly with regard to the effects of enforcing a single price across multiple consumer types. Our theory extends the standard price discrimination model to allow for targeted subsidies that move the demand curves of each type differentially. We are particularly interested in the interaction of the single-price

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

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

restriction, targeted subsidies, and the exercise of market power by the firm.

A consumer of type t purchases one unit of the good from the firm if this consumer's willingness to pay, v_t, is greater than the effective post-subsidy price $p_t(b) = \max\{0, b - z_t\}$, where each consumer type has a type-specific subsidy, z_t. Without loss of generality, we will assume that the subsidy is positive for L consumers, and is zero for H consumers. Denote the share of consumers of each type purchasing the good as $s_t(p_t)$. We assume that the good can be provided at the same marginal cost c to both types of consumers.[7] The firm chooses a list price b to maximize the following profit function:

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

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

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

This is the standard first-order condition for a monopolist selling to both types at a uniform price with one key difference: the term $\frac{\partial p_t(b)}{\partial b}$ captures the role of the subsidy design by introducing the possibility of a gap between the firm's price, b, and the price that the consumer faces, p. With targeted subsidies, the response of the effective consumer price to changes in b varies across consumer types:

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

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

As a result, targeted subsidies may induce a distortion in the aggregate demand curve perceived by the firm relative to the true underlying demand curve. Critically, the distortion

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

is asymmetric, as the targeted subsidies only change the demand curve for type L. However, since the firm is forced by regulation to charge a single price, the presence of the subsidized type L generically influences the price paid by type H.

The heart of our theoretical and empirical analysis is the *demographic externality*, which we define as $\partial p_H(b)/\partial \eta$. The sign of this effect is theoretically ambiguous, as it depends on both the level of the subsidy and the relative demand curves for the two types. On one hand, holding all else equal, when η increases there is a price-raising incentive due to the presence of additional subsidies in the market. However, if the two groups have different demand curves, it could be the case that when η increases the firm now faces a more elastic aggregate demand curve overall, which incentivizes the firm to lower prices. In any given context, it is an empirical question as to which of these two forces is stronger.

The demographic externality creates an additional economic force within the market that the policymakers need to be cognizant of when designing the subsidy mechanism. The combination of subsidy targeting and market power along with a prohibition on price discrimination can lead to unintended equilibrium outcomes, such as prices being higher in markets with more poor, but heavily-subsidized, consumers. In such a scenario, removing subsidy targeting and using a universal subsidy could lead to higher social welfare at the same level of government spending. At the same time, if targeted subsidies are used, their optimal level in the presence of market power needs to take into account the demographic externality and may be different the the optimal subsidy level under public provision. Intuitively, the level of the subsidy changes the consumer surplus of consumers in group L directly, and of consumers in group H indirectly through changes in equilibrium prices.

Finally, we note that the demographic externality we describe here is conceptually distinct from other settings where prices are linked to the distribution of consumer types in a market. In an insurance context, for example, the cost of providing the service depends on which consumer types buy coverage. Under community rating and perfect competition, the equilibrium price is the average cost across all types. Hence, on the margin, low-cost consumers end up paying more for the same good if they live in markets with more high-cost types. However, this effect is driven by changes to the firm's cost function. Here, the effect we describe is driven by links across consumer types enforced on the firm's revenue function and a public policy that differentially changes consumer demand function by providing targeted subsidies.

3 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 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, administered through the tax system, that have been increasing from 2014 onward. 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 as measured by the fraction of costs a plan would cover for a standardized population. In addition to metal labels, plans vary in their cost-sharing rules, restrictions on provider networks, and the scope of procedural and pharmaceutical coverage.

While several US states have created their own ACA 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. These 37 so-called "federally-facilitated" states encompass 2,566 counties with about 9 million enrollees. Within each state, sets of counties are aggregated into broader "rating areas," which is the regulated level of aggregation at which prices are set. Specifically, while insurers do not have to offer any given plans in all counties within a rating area, if an insurer does offer a plan then they must set a single price common across all those counties. Despite the complexity of the geographic arrangements, it is helpful to think about county-level markets in this setting (see also Fang and Ko, 2018).

Insurers in this market are not allowed to price-discriminate based on individual health risks, but they are allowed to charge different premiums depending on individual's age and smoking status. Age adjustment of prices is tightly regulated, as insurers have to follow a regulatory age curve that determines how much prices can differ across consumers of different ages.[8]

One of the key aspects of the ACA Marketplace is the availability of subsidies for consumers with low incomes. The subsidy system consists of several pieces. The first type of subsidies—which are the focus of this paper—reduce annual premiums that consumer have to pay. The second type of subsidies, held constant in our analysis, are cost-sharing reduction subsidies that reduce the out-of-pocket liability that consumers may face when utilizing

[8]Insurers can choose to also underwrite on consumers' smoking status; however, as whether someone smokes is hardly verifiable and very few consumers in the data are flagged as smokers, we do not consider prices for smokers in our analysis.

healthcare in form of deductibles, co-pays, etc. Both types of subsidies are means-tested, where families with lower household incomes receive higher subsidies.

Premium subsidies are known as Advanced Premium Tax Credits (APTC) they can be paid (directly to the insurance company on consumers' behalf) at the start of the year based on projected household 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 APTC is calculated in several steps. First, the "Modified Adjusted Gross Income (MAGI)" for a tax family as reported in federal tax forms 1040 is converted to the percent of Federal Poverty Level (FPL). The level of FPL varies depending on the family composition and allows comparing incomes of families of different size using the same scale. The MAGI relative to FPL measure then determines the maximum dollar amount that the household "should be" paying for insurance premiums. Call this amount "CAP." The CAP is based on a non-linear sliding schedule. For example, if a household's income is 200 percent of FPL, then this household should be spending no more than 6.34 percent of their income on health insurance premiums in the ACA Marketplace. At 400 percent FPL, the CAP is equal to infinity as households with income above 400 percent FPL are not subsidized.

Next, to determine the level of subsidies, the regulator considers the distribution of list prices in the county where the household resides. The regulator identifies the list price of the second-cheapest Silver plan (SLSP) and sets this as the reference premium. If the family's CAP is greater than the sum of SLSP premiums that the family would need to pay for all family members that require coverage, then this household gets no subsidy. If CAP is less than SLSP, the household gets a premium tax credit that is equal to the difference between the applicable SLSP and the CAP.

We combine several sources of data for our analysis. We use data from 2017, which was the fourth year of Marketplace operations. We focus our analysis on ACA Marketplaces that use the federal healthcare.gov platform, as the most accurate 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. These are recorded in premium and plan files that have been released by the Centers for Medicare and Medicaid Services that is running the Marketplaces. 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 2017 edition of the American Community Survey (ACS) 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, enrollment, and demographics. In 2017, consumers could choose among on average 21 plans offered by 1 to 4 large national insurers and a number of smaller firms. The annual pre-subsidy premiums for a 40-year old in these plans ranged from $3,978 (10th percentile) to $6,351 (90th percentile) with an unweighted average of about $5,160. The average number of potential enrollees per market was close to 8,000 individuals, although markets differed dramatically in their size, ranging from fewer than 479 potential enrollees at the 10th percentile of counties to more than 15,000 at the 90th percentile. On average across markets, 60 percent 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 75 percent of choices conditional on enrollment. In an average market, the average plan had 3,156 enrollees, again with plan sizes varying substantially between plans with fewer than 50 and more than 6,000 enrollees. Potential enrollees were on average 40 years old, 90 percent white, with an average income of 262 percent FPL. On average, these consumers qualified for $3,301 in premium subsidies.

4 Empirical Model

4.1 Demand

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

Formally, we posit that individual i in family f in market t chooses plan j from a set of choices J. The set of choices that each family faces depends on the family's geographic location, family income, and age distribution of family members that require insurance coverage. It is helpful to think about geographic markets as operating at the county level, since, as Fang and Ko (2018) argue, insurers effectively make their strategic decisions at the county level.[9] The indirect utility function as observed by each consumer takes the following form:

$$u_{ijt} = -\alpha_i p_{ijt} + \phi_{ijt} \tag{5}$$

Where p_{ijt} is individual-specific price that consumer i, who lives in market t, faces for plan j. This price is individual-specific because consumers of different ages with different family incomes face different prices for the same plan j in market t.[10] ϕ_{ijt} is the amount of utility a consumer gets from plan jt. This utility can vary across consumers for two reasons. First, consumer valuation of the same plan may differ due to heterogeneity in preferences. Second, the same plan j may have different characteristics for different consumers, as consumers with lower incomes get more generous coverage when they buy certain plans.

We make a few assumptions about α_i and ϕ_{ijt} to arrive at an empirically-tractable version of this utility function that still preserves the semi-nonparametric approach. First, we replace individual-specific α_i with a coarser set of parameters that vary across nine demographic groups, d. The demographic groups are all combinations of three age categories: age under 25, age between 25 and 40, age above 40, and three income categories: income under 200 percent FPL, income between 200 percent and 400 percent FPL, income above 400 percent FPL. Second, we decompose the utility that a consumer gets from plan j into several additively separable components. The first component ψ_a captures the average level of utility that consumers get from purchasing any insurance plan. We allow this intercept parameter to vary across three age groups a (same as the age groups above), so as to capture the idea that older consumers may value insurance more, all else equal.[11] The second component captures the deviations in the generosity of a plan j that a consumer may face with low enough income. For low-income consumers, there are additional non-premium subsidies for silver plans that we capture through the actuarial value (AV) of a plan.[12] Finally, we

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

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

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

[12] There are adjustments to cost-sharing provisions of the health plan, e.g. the deductible, that the government partially subsidizes for low-income consumers. One can compute an equivalent actuarial value for

include a plan-specific constant δ_j for each plan j that captures the average utility that a consumer gets from purchasing plan j. The remaining difference between ϕ_{ijt} and these three components is captured by an idiosyncratic taste shocks, ϵ_{ijt}, that is observed by the consumer but not by the econometrician. This random individual-plan specific shock to the utility function is assumed to be distributed Extreme Value Type 1, which leads to a logit discrete choice model.[13] To summarize, the empirical version of the utility function becomes:

$$u_{ijt} = -\alpha_{d(i)}p_{ijt} + \psi_{a(i)} + \gamma AV_{ij} + \delta_j + \epsilon_{ijt} \tag{6}$$

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

The variation in premiums that we observe in the data does not stem from experimental assignment. However, the regulatory design of this market generates variation in prices independent of demand shocks to help us identify the marginal utility of income parameters, α_d. To see this, we first note the well-known result from Berry et al. (1995) that there exists a unique vector δ that perfectly rationalizes plan enrollments. However, that δ cannot, by itself, explain the patterns of enrollments within a market across consumers of various ages and incomes. That is where the α plays a critical role. Consider two consumers of the same age in a given market. Under ACA regulations, two consumers will face different effective prices for the same plan j if their incomes or ages are different. The variation in subsidies across income levels depends on a pre-specified formula that generates a non-decreasing relationship between income and effective premiums, as discussed in Section 3. The statutory age-adjustment curve does the same for consumers of different ages. In conjunction with data on enrollment in the any plan in a market by age and income buckets, this policy-related variation in prices for the same plan j across income and age groups allow us to identify α_d parameters; the δ explains the average enrollments in total, while the α parameters explain the rate at which consumers substitute to the outside option within and across age- and income-levels.

This identification approach is similar in spirit to that pursued in Tebaldi et al. 2019. In

these adjustments and apply them to silver plans, which are the only category of plans that are covered by this cost-sharing reduction, for consumers that qualify.

[13] McFadden and Train (2000) show that the logit model can approximate any random utility function if one is sufficiently flexible with the specification of how the utility parameters vary across individuals.

sum, despite the fact that we do not observe product-market shares to pursue a more standard demand estimation strategy (Berry et al., 1995), the institutional features of the market in fact open a novel identification strategy that contrasts to the more common identification approach of using instrumental variables constructed from product characteristics or prices in other geographic markets (Hausman, 1996; Berry and Haile, 2016).

4.2 Supply

4.2.1 Profit function

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

For each consumer i of age a, plan j collects revenue that consists of several pieces. First, plan j collects premium p_{ij} from the consumer. For consumers that do not receive premium subsidies, the premium is equal to insurer's full list price for consumers of age a, b_j^a. For consumers that are eligible for subsidies, the insurer collects $p_{ij} < b_j^a$ from the consumer as well as a subsidy from the federal government. Together, the premium and the subsidy add up to b_j^a.[14] Second, the insurer may collect revenue from three risk-equalization programs that we describe below.

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

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

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

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

$$\pi_{ij}(b_j^a) = b_j^a - c_{ij}(r_i, \phi_j, D_i), \tag{7}$$

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

$$\pi_{ij}(b_j^a) = b_j^a - c_j^a - \tilde{c}_{ij}(r_i, \phi_j, D_i), \tag{8}$$

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

Intuitively, the idea of the risk-equalization programs is to create transfers that exactly offset the idiosyncratic cost component $\tilde{c}_{ij}(r_i, \phi_j, D_i)$, so that every enrollee has the same expected cost in the insurer's profit function. The reinsurance program effectively gives insurers additional individual-specific revenue for individuals with particularly high $\tilde{c}_{ij}(r_i, \phi_j, D_i)$, so as to reduce the impact of this term on insurer's profit function. We can then consider the difference between this additional revenue and the idiosyncratic cost-component as the net idiosyncratic cost that is relevant for insurer's decision-making. Denote this difference with ν_{ij}. Now let the (positive or negative) lump-sum risk-adjustment payment to the insurer be R_j. This term is a function of risk types r_i of all individuals that enroll in a plan and is not individual-specific, but it aims to offset the sum of ν_{ij} across all i's in cases where the *expected* (or in other words, predictable) individual-specific deviations in risk across consumers add up to a positive or a negative quantity. Let $H_j = \sum_{i \in j} \nu_{ij} - R_j$ denote any residual selection. If the risk-equalization programs fully offset the ex ante net idiosyncratic shocks, this term will be zero.

In the ACA environment, H_j may not be zero. The direction of this term is likely to primarily be a function of ϕ_j, or contract generosity. This follows from the fact that individuals with the lowest incomes, who are also likely to have the highest expected costs (as has been shown by a voluminous literature on the health-income gradient), receive subsidies that significantly compress the variation in both prices of plans that these consumers face. Hence, the key differences across plans that are likely to drive selection lie in non-pecuniary plan features of ϕ_j, such as physician networks, formulary breadth, and chronic condition management. Without data on individual enrollment and expected costs in Marketplace plans, we can only estimate the profit function up to the constant $H_j(\phi_j)$; we simplify the notation by assuming that $H_j(\phi_j) = 0$ in subsequent discussion.[15] Importantly, the assumption of $H_j(\phi_j) = 0$ has no bearing on the first-order condition in prices, since as we argued above price-based risk screening, conditional on ϕ_j, is unlikely to be quantitatively important (in other words, the derivative of $H_j(\phi_j)$ with respect to prices is zero).[16] We do not explicitly incorporate the ex post risk-corridor transfers into the model; these payments can be interpreted as a reduction in insurers' fixed cost of purchasing private re-insurance policies and should not affect insurers' pricing incentives on the margin.

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

[15] 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.

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

market, we get:

$$\pi_j(b_j) = \sum_a s_j^a M^a b_j^a - \sum_a s_j^a M^a c_j^a \tag{9}$$

According to ACA statutes, insurers have to follow a statutory age schedule for their bids that constraints age-specific underwriting. This restriction allows us to simplify the problem further. Let there be a fixed set of age-specific multipliers that apply to bids. We assume that the same multipliers apply to expected baseline costs, capturing how healthcare costs increase with age.[17] Let the multiplier vector be τ^a. The profit equation for plan j then becomes:

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

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

$$\pi_f(\mathbf{b}) = \sum_{j \in f} \left[(b_j - c_j) \sum_a s_j^a M^a \tau_j^a \right] \tag{11}$$

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

4.2.2 First-order conditions

Insurers choose bids that maximize their profits taking into account the actions of other firms. The first-order condition for a single-plan firm is:

$$\frac{\partial \pi_f}{\partial b_j} = (b_j - c_j) \sum_a \frac{\partial s_j^a}{\partial b_j} M^a \tau_j^a + \sum_a s_j^a M^a \tau_j^a = 0. \tag{12}$$

For an insurer that offers more than one plan in a market, the set of j first-order conditions

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

accounts for own and cross-price elasticities of demand, taking the vector form given by $S - \Omega(B - C) = \mathbf{0}$, where row j of vector S is given by: $S_j = \sum_a s_j^a M^a \tau^a$ and row j of vector $(B - C)$ is given by $(B - C)_j = (b_j - c_j)$, while row k, column j of matrix Ω is:

$$\Omega_{kj} = -\sum_a \frac{\partial s_j^a}{\partial b_k} M^a \tau^a \tag{13}$$

for plans k and j offered by firm f. This gives us j equations in j unknowns for each insurer, as want to recover costs c_j that are unknown up to the scaling factor τ. We invert Equation 12 and compute the baseline marginal cost c_j for each plan as a function of observed equilibrium prices and the elasticity of demand that is given by the demand parameters from Section 4.1.

The key term of the first-order condition is the derivative of the (age-specific) share with respect to the (age-specific) bid: $\frac{\partial s_j^a}{\partial b_j^a}$. We drop the age superscripts to simplify notation in what follows, as age scaling is given by regulation and age markets are additive in our set up.[18] 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:

$$\frac{ds_j(p_j, p_{-j})}{db_j} = \frac{\partial s_j}{\partial p_j} * \frac{\partial p_j}{\partial b_j} \tag{14}$$

Recall that the subsidy is a function of the bid set by the second-lowest cost silver plan and consumer's family income. The second-lowest expensive silver plans face a different set of incentives relative to other plans, as under the observed subsidy regime a change in their bids affects not only their own prices, but also the subsidies, and hence consumer premiums, of all plans.[19] To account for the idea that the first-order conditions do not perfectly capture these additional incentives, we do not use inverted marginal costs for these plans.[20] Instead, following (Decarolis et al., forthcoming), we use other plans to impute the marginal costs of SLSPs. We project the estimates of marginal costs for the non-2LSPs on a vector of their characteristics and use this hedonic projection to predict the marginal costs of the SLSPs.

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

[19] Formally, these plans have another term in 14 that is non-zero: $\frac{\partial s_j}{\partial p_{-j}} * \frac{\partial p_{-j}}{\partial b_j}$.

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

4.3 Efficiency Metric

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

$$W = CS + \Pi - \lambda G, \tag{15}$$

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

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

where γ is Euler's constant, and v_{ij} is the deterministic component of utility for person i (recall that this is, in return, is the average utility within a family) for plan j and is equal to utility net of the idiosyncratic ϵ term.[21] We integrate out over the empirical distribution (as observed in the ACS) of ages, income, and family composition to obtain average annual per capita consumer surplus:

$$CS = \int CS(\theta)dF(\theta). \tag{17}$$

Producer surplus, Π, is computed following equation 11. We assume that any risk-equalization payments, including risk corridors, contribute to cost equalization and are already captured in marginal cost estimates, implying that they do not separately enter the profit function.

Government spending G captures three parts. Nominal spending includes subsidies for insurance premiums as well as subsidies for cost-sharing reduction. The former are computed either from the data or are adjusted in simulations of Section 5. Cost-sharing reduction (CSR) spending is held at observed levels. Specifically, using CMS data reports we compute the average per capita spending on CSR subsidies by consumer type, based on income brackets.[22] In all counterfactual simulations, we then assign this average spending level to each consumer who falls into the respective income bracket and who enrolls in a plan where cost-sharing reduction is available. In the final step, we account for the fact that

[21]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.

[22]The data was accessed in June 2019 at Health Insurance Marketplace Cost-Sharing Reduction Subsidies by Zip Code and County 2016.

when a consumer enrolls into an ACA plan, the government likely saves some money on this consumer; for example, if a consumer enrolls in a formal insurance plan, this consumer is then unlikely to benefit from any public payments for uncompensated care. Following the Kaiser Family Foundation and Urban Institute 2013 report on public spending on uncompensated care for the uninsured (Coughlin et al., 2014), we assume that the government saves $1,827 per capita in public funds for each consumer enrolled in an ACA insurance plan.

4.4 Estimation results

4.4.1 Demand Parameters

We use non-linear least squares to estimate utility function parameters using the approach described above. Panel A of Table 2 reports the results. We find intuitive patterns for the variation in the marginal utility of income parameter across demographic groups. A one dollar increase in price has on average a larger impact on the utility of poorer consumers and younger consumers, for whom we would expect a dollar of insurance to constitute a higher share of their annual budget. Figure 4a plots compensating variation (CV) by income and age. The relationship between the overall value of insurance and age is non-linear. While consumers above the age of 40 value any insurance more than consumers aged 25 to 40, the demand by consumers below age 25 exhibits an even higher valuation. We also find a U-shaped relationship between CV and income. This follows from the fact that subsidies are declining in income, lowering CV, while our estimates suggest generally less elastic demand as income rises, increasing CV.

As would also be expected, we find that consumers get significant utility from purchasing plans with a higher level of coverage as measured by the actuarial value, conditional on other characteristics of plans held fixed as captured by the plan's δ that measure the relative attractiveness of plans to consumers on average. Together, the coefficients suggest that, for example, consumers over age 40 with income over 400 percent FPL value insurance coverage by a plan with the highest plan fixed effect that on average pays 70 percent of expenditures (maximum δ_j for a silver plan equals 0.34) at about $2,891, which lies within the support of the distribution of list premiums for a 40 year old consumer for silver plans that ranges from $2,391 to $9,057 in the data.

In general, while the patterns are mostly intuitive, we are cautious about the interpretation, as the consumers in the model are assumed to be maximizing average family utility, hence the marginal utility of income parameters capture family level preferences.

Family-level demand could, for example, exhibit a higher valuation of insurance by younger consumers, stemming from the valuation of their parents rather than individuals themselves. This would lead to a high estimated value of insurance at young age, as the younger group includes children, whose parents may place a high value on having insurance for their child.

To assess fit, Figure 1 illustrates one set of moments used for the estimation and the geographic distribution of the model fit. In Panel (a) we report the county-level market share of Silver plans. In panel (b) we report the average in-sample difference between the data and the model's prediction of county-by-metal-level enrollment shares across all 2,566 counties that are used in the estimation. We closely match these aggregated enrollment shares, with the model being able to capture a substantial amount of variation in the data.

4.4.2 Cost

Panel B of Table 2 reports the results of a hedonic regression that projects marginal cost estimates from the first-order condition inversion onto plan characteristics for non-SLS plans. The estimates are reassuring, in that we find that more generous benefit design is associated with higher marginal costs. The average (baseline, for a 20 year-old) non-SLSP marginal cost from the inversion procedure is circa \$1,940, with a standard deviation of \$540. Moving a plan's actuarial value by 1 basis point, while keeping the metal level the same, increases the marginal cost by \$19. Moving from a gold to a silver plan decreases costs by approximately \$500 (\$300 metal-label affect and plus the actuarial value adjustment of \$192). The regression also includes measures of out of network coverage, whether a plan is HSA-eligible, whether a plan covers some common benefits, whether it offers management of common chronic conditions, and insurer fixed effects. We use this regression equation to impute marginal costs for SLSPs. As in the marginal cost inversion, we impose the MLR constraint on marginal cost predictions.

Figure 2a illustrates the resulting distribution of estimated marginal costs for the baseline age group of 20-year old consumers. We plot the distribution separately for Bronze and Gold plans. We observe two pronounced patters. First, there is substantial heterogeneity in costs within a metal level. This is not surprising, since plans on the ACA Marketplaces are extremely diverse, with some plans being offered by large national insurers and some by local co-operatives. Second, there is substantial differences in costs between more and less generous plans, as we already saw in Panel B of Table 2. This is intuitive, as mechanically gold plans cover 80 percent of consumers' healthcare expenditures on average, while the Bronze plans cover only 60 percent. We would expect that the ratio of costs between these plans is on average 1.3, which is consistent with the shift in the distribution that we observe.

Figure 2b compares our estimates of marginal costs from the first-order condition inversion (and projection for the SLSPs) to plan-level accounting costs as reported by plans to CMS. The accounting costs are measured with error, as insurers are allowed to report their costs equally split across their plans rather than providing a true plan-level attribution of costs. Moreover, the reported accounting costs do not include some ex post cost reconciliation, such as, for example, MLR repayments. Nevertheless, the accounting cost data provide a valuable informational signal, as they are likely to on average provide an accurate ordinal ranking of plans from least to most expensive, and also give a general sense of cost levels in the market. As can be seen in Figure 2b, as we would expect given the ex post cost reconciliation that characterizes our institutional environment, our estimates of marginal cost are on average lower than reported accounting cost, although they have the same general order of magnitude. Further, we observe a very strong correlation between accounting costs and marginal costs, which supports the idea that we are accurately able to differentiate more and less expensive plans. [23]

4.4.3 Demographic Externality

The combination of estimated demand and cost parameters allows us to illustrate the demographic externality mechanism, as described conceptually in Section 2, in the context of our empirical model.

The first simulation increases the number of subsidized consumers without changing the distribution of marginal utility of income in the population. For each market, we set subsidies for consumers with income above 400 percent FPL as if these consumers had income of 151 percent FPL. This means that in each market the share of consumers with subsidies increases, while the share of unsubsidized consumers goes to zero. The results of this first simulation are reported in Figure 3 and are marked as “Case A.” This figure reports how average premiums and consumer surplus changes for *directly unaffected* consumers, i.e. those with incomes between 150 percent and 400 percent FPL when their neighbors with income above 400 percent start getting subsidies. As we would expect in the presence of market power,

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

insurers take advantage of the fact that in the 400 percent and higher FPL market segment, consumers now face lower prices for any given list price. Insurers increase list prices in response, making plans now become more expensive for "unaffected" consumers with incomes below 400 percent FPL. As the light dashed line marked with "A" in the figure illustrates, the average annual effective premiums for consumers with income under 400 percent FPL increases by $10. Consumer surplus, marked with grey circles, in turn declines by up to $20 for the poorest consumers. While average changes in prices and consumer surplus are quantitatively relatively small, this masks highly unequally distributed across geographic areas. Some counties are completely unaffected by the simulated change, while some others experience more than ten times the average loss in consumer surplus. This counterfactual simulation cleanly illustrates a simple mechanism: subsidizing one group of consumers in a market with market power, all else equal, increases prices and decreases welfare for other consumers. In this case, subsidized consumers exert a negative demographic externality on other consumers. Importantly, this effect is not special to insurance markets and does not depend on the fact that insurance contracts are pooling risks, which is different from the age-based externality from risk pooling as documented in Tebaldi (2017).

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

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

5 Subsidy Design and Welfare

In this section we report a set of counterfactual model simulations that allow us to assess the quantitative importance of the posited interaction effect between subsidy targeting and market power. We start off by characterizing the observed market allocation under means-tested subsidies and market power. The efficiency properties of this allocation, as well as the incidence of economic surplus under this allocation, form the baseline for our analysis. We then examine the economic consequences of subsidy means-testing and market power in isolation, before putting them together in a set of counterfactual comparisons that allow us to isolate the additional welfare loss that is generated when a targeted subsidy is used as a policy instrument in an environment with market power.

Observed allocation We begin our analysis by examining the economic costs and benefits of subsidy payments as observed in the data. To facilitate this computation we first re-simulate the allocation under observed subsidies in our model. The simulation serves two purposes. First, it allows us to establish a baseline that differences out any model simulation error. Second, it allows us to compute an allocation that shuts down the SLSP mechanism, while preserving observed subsidy levels. A baseline without the SLSP mechanism is useful, because in subsequent counterfactual analyses we will want to isolate the effects of subsidy levels and subsidy targeting without the SLSP mechanism, so a comparable baseline is required to interpret these counterfactual analyses.[24]

[24]We implement this simulation by assigning imputed marginal costs to SLSPs and letting all plans set prices according to the first-order condition in 12, taking observed subsidy levels as given. In practice, this implies that SLSPs end up being the only plans for which prices change significantly relative to observed

Column (1) of Table 3 reports the results in this baseline scenario. We estimate that consumer surplus in the market amounts to $38 billion. This corresponds to $1,570 in annual surplus per capita, on average, among all potential consumers. Consumer surplus levels vary substantially across different areas of the country and across different socio-demographic groups. Panel (b) of Figure 4b illustrates the geographic variation in surplus that ranges from $621 in Brazos County, Texas to $6,260 in McHenry County, Illinois. The socio-demographic variation is also substantial. Panel (a) of Figure 4a plots consumer surplus by income and age. We estimate that older workers have higher surplus than younger workers for all income levels. There is a U-shaped relationship between income and surplus, with workers in the 250-300 percent FPL range having the lowest valuations for insurance products in the ACA. These surplus calculation dovetail with enrollment levels across income; rows 14 to 26 of Column 1 show that more than half of the enrollees have incomes below 250 percent FPL, while high-incomes consumers above 400 percent FPL constitute about a quarter of the market. Enrollment is distributed modestly more evenly across age groups, but enrollment rates increase monotonically with age.[25]

Producer surplus amounts to $3.9 billion (Column 1, row 2). Under the observed allocation as simulated in our model, the government is spending $30 billion in premium and cost-sharing subsidies, which is broadly consistent with the subsidy spending reported by the Congressional Budget Office.[26] The program attracts about 41 percent of potential enrollees (row 17) and generates $1.39 of surplus for $1 of within-ACA government spending, as the total outlays on subsidies are $12 billion lower than the sum of consumer and producer surplus. However, as is the case in many health settings, government expenditures need to account for the opportunity cost of spending on these same consumers if they had not enrolled in the ACA; this amount is reported in row (5). We estimate that the savings of public funds on the same set of consumers amounts to $16 billion, which implies that the net additional government outlays for premium and cost-sharing subsidies on the Marketplaces is approximately $14 billion. Taking into account this foregone spending and the cost of raising public funds, we estimate a total return of $2.33 on a dollar of public funds under observed subsidy levels (row 9).

prices.

[25]Surplus monotonically increasing with age is consistent with the observation in Tebaldi (2017), who argues that the structure of ACA subsidies does not encourage enrollment of the youngest consumers.

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

Incidence of the subsidy Column 2 of Table 3 allows us to assess the incidence of observed subsidies between consumers and producers. In this column, we compute the allocation that would result in the absence of premium subsidies, holding other features of the market (such as market power and cost-sharing reduction subsidies) fixed. Without premium subsidies consumer surplus drops from $38 billion to $26 billion, while producer surplus drops from $3.9 billion to $1.7 billion. Overall, introducing targeted subsidies as observed in the data, generates an additional surplus of $13.7 billion for $29 billion in within-ACA spending on subsidies; this is partially mitigated by $9 billion of additional public savings, leading to a net increase in government spending of $20 billion. From this perspective, the program generates negative economic value, as the return on a dollar of subsidies was only 68.5 cents. However, these additional subsidies did substantially benefit consumers, whose enrollment jumps from 14 percent without subsidies to 41 percent with subsidies. This is particularly true of the consumers below 400 percent FPL, who basically exit the market without subsidy support.

Out of $13.7 additional surplus, 84 percent accrues to consumers and 16 percent is captured by firms, which is a comparatively high pass-through for an imperfectly-competitive market. Two margins are at work here: one, the MLR constrains firm profit margins, and two, targeted subsidies disproportionately attract lower-income consumers into the market that have substantially more elastic demand. Consumers with more elastic demand both have lower consumer surplus and put pressure on insurer margins. Rows (29) and (30) of Columns (1) and (2) summarize this effect. While the average premium subsidy payment between the two regimes goes up from $0 to $3,732, the average consumer surplus increases from $999 to $1,570, a substantially lower increase in surplus as compared to government expenditures.[27]

The effects of subsidies on consumer surplus are highly heterogeneous across locations. Panel (a) of Figure 5 illustrates the geographic distribution of the estimated gain in consumer surplus under the observed mechanism relative to an economy with no subsidies. Highest changes in surplus accrue to southern parts of Texas, parts of Nebraska, and several states in the Southeastern United States, including parts of Tennessee, North Carolina, Alabama, Georgia, and Florida, where in many counties average consumer surplus increases by more than 100 percent when means-tested subsidies are introduced.

[27] The potentially puzzling phenomenon of low valuation (as measured from revealed preferences) of formal insurance by low-income consumers has been documented in prior literature. See an overview in, for example, Poterba (1996) and more recent empirical evidence in Finkelstein et al. (2019), who speculate about the role of uncompensated care and behavioral biases in accounting for the low revealed willingness to pay.

Isolating the efficiency cost of market power Before turning to the discussion of how market power interacts with subsidy design, we first assess the baseline efficiency cost of market power in this market. Counterfactual simulations in columns (1) and (4) of Table 3 compare the properties of allocations with and without market power, keeping everything else about the environment fixed. In column (4), the allocation without market power is simulated by setting prices equal to marginal costs. Removing market power increases total surplus by $6 billion from the baseline of $24 billion. In addition, $3.9 billion of surplus is reallocated from insurers to consumers, with consumer surplus increasing by nearly $10 billion (and by $420 on average) and consumer enrollment increasing from 41 to 52 percent. List premiums in the absence of market power are $450 lower on average, which attracts slightly more lower-income consumers into the market. On average, consumers lose 21 percent of per capita surplus due to the presence of market power in this market.

Isolating the efficiency cost of targeting To further understand how the subsidy structure per se affects surplus in this market, we pursue two more comparisons. In columns (4), (5) and (6) of Table 3, we assume that the market is operating under perfect competition and ask how much do subsidies contribute to welfare. This comparison allows us to estimate the extent of deadweight loss from means-tested subsidies (comparing column 5 to 4) and from means-invariant subsidies (comparing column 5 to 6). The first comparison—between having means-tested subsidies at the level observed in the ACA market and no subsidies—suggests that $32 billion in nominal premium subsidy spending generates only about $16 billion in consumer surplus. Hence, nearly half of nominal government spending constitutes a deadweight loss. The amount of deadweight loss is reduced if we take into account that the government saves on those consumers who choose to enroll in the ACA market in the presence of subsidies. In this case, the extra public spending constitutes $22 rather than $32 billion, decreasing the implied deadweight loss to 27 percent of extra public spending, or $6 billion. The comparison of the two simulations reveal that subsidies play the central role in ensuring high enrollment rates into the program. Without subsidies and with marginal cost pricing, we simulate that only 20 percent of potential consumers would enroll in this market (69 percent of these consumers coming from relatively higher-income households). With means-tested subsidies, overall enrollment increases to 52 percent and more than half of the consumers are now coming from households with income under 250 percent FPL. In sum, means-tested subsidies work as intended—they lower prices for the lowest-income consumers and thus attract them to the market. That the subsidies generate substantial deadweight loss is not surprising in the presence of such a large enrollment increase, as the marginal consumers attracted by increasingly generous subsidies have an increasingly-declining willingness-to-pay for insurance.

In column (6) we ask whether the amount of deadweight loss from subsidies under perfect competition depends on the design of the subsidy system. We estimate how much consumer surplus could be achieved, on average and in total, if we kept the level of public spending the same, but made subsidies income-invariant (although still varying by age). We find that under income-invariant subsidies, the same amount of total public spending can deliver \$6.6 billion more in total consumer surplus. This substantial gain in efficiency comes at a re-distributional cost. While the overall enrollment in the market remains the same, enrollment shifts from poorest to less poor consumers, which clearly illustrates the equity-efficiency trade-off of the subsidy design. In the next set of counterfactual simulations, we examine how these trade-offs can be amplified when the good is provisioned by strategic, imperfectly-competitive firms that can change prices in response to changes in the subsidy design.

Targeting subsidies in the presence of market power We next consider the central question of the paper, which whether the presence of market power can introduce additional distortions when consumers receive targeted subsidies. We isolate the interaction effect between targeted subsidies and market power by comparing how much efficiency loss subsidy targeting generates in the absence of market power (Columns (4) versus (6) of Table 3) to the amount of efficiency loss in the presence of market power (Columns (1) versus (3) of Table 3). Within each of this pairwise comparison, we find flat voucher levels that keep the amount of public spending fixed in equilibrium.

Figure 6 illustrates our findings. Our first result holds market power and government spending constant: we estimate that having means-tested rather than income-invariant subsidies leads to a 32 percent loss in total welfare. The black curve traces out the level of welfare that is generated in the market in the presence of market power, under different levels of flat subsidies. We note that total welfare is concave with respect to the level of flat subsidies. Highest welfare is achieved at baseline subsidy levels of around \$600–\$700, and starts falling at higher subsidy levels. At the flat subsidy of \$1,105, as we see in Column (3) of Table 3, the allocation with flat subsidies results in the same amount of public spending as with means-tested subsidies at levels as observed in the data. The level of welfare under means-tested subsidies is marked with a black diamond in Figure 6. Moving from flat subsidies at the level of $1,105 to means-tested subsidies with the same amount of government spending generates a surplus loss of 32 percent. Not all of this loss in welfare, however, is attributable to the interaction between subsidy design and market power. The light grey line that repeats the same analysis for the perfectly-competitive case illustrates that the benchmark loss in efficiency purely attributable to subsidy design is 24 percent. We conclude that market power significantly exacerbates the efficiency losses from means-testing, leading to a 33 percent higher loss in welfare.

Crucially, however, the allocation with income-invariant (flat) subsidies is not Pareto improving over the targeted subsidy mechanism. Flat subsidies generate a re-allocation of surplus from previously highly subsidized consumers to previously less subsidized consumers. This is true both under perfect and imperfect competition, with the effects amplified under imperfect competition. Panel (b) of Figure 5 illustrates the geography of losers and winners in the presence of market power. Counties in the Midwest are hurt by means-tested subsidies, while counties in the Southeast and some areas of Texas benefit. Table 4 summarizes the magnitude of changes in consumer surplus by consumer type in Panel B. Under perfect competition, we observe that moving from means-tested to flat subsidies hurts the lowest income consumers of all ages, as well as older consumers with incomes between 200 percent and 400 percent FPL who are partially subsidized under the means-tested subsidy regime. For these, consumer surplus declines by 26 to 53 percent. Young consumers with income between 200 percent and 400 percent FPL benefit from flat subsidies, as do consumers of all ages with income above 400 percent FPL, as they now also receive subsidies, their surplus increasing by 34 to 108 percent. The effects are amplified in the presence of market power. As insurers re-price their plans, low income consumers experience slightly higher losses of surplus, while higher-income consumers experience slightly higher gains. As price adjustments are constrained by the minimum loss ratio requirements, these estimates capture only the lower bound of the negative effects on the lowest income consumers, as insurers are restricted in the maximum increase in prices that they can undertake.

6 Conclusion

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

We examine this theory in the empirical context of the ACA Marketplaces. We show that market power has a substantial interaction effect with income-targeted subsidies, leading to a 33 percent higher loss in welfare. We conclude with the observation that policymakers should account for this interaction when designing subsidy mechanisms, as ignoring the effect can generate substantial biases in surplus and social returns in publicly-subsidized, privately-provided programs.

References

Abraham, Jean, Coleman Drake, Daniel W. Sacks, and Kosali Simon, "Demand for health insurance marketplace plans was highly elastic in 2014 to 2015," *Economics Letters*, 2017, *159*, 69 – 73.

Allcott, Hunt, Christopher Knittle, and Dmitry Taubinsky, "Tagging and Targeting of Energy Efficient Subsidies," *American Economic Review Papers and Proceedings*, 2015, *105(5)*, 187–191.

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

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

CBO, "Federal Subsidies for Health Insurance Coverage for People Under Age 65: 2019 to 2029," Technical Report, Congressional Budget Office Report May 2019.

_ , "Medicaid—CBO's May 2019 Baseline," Technical Report, Congressional Budget Office Report May 2019.

— , "Medicare—CBO's May 2019 Baseline," Technical Report, Congressional Budget Office Report May 2019.

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

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

Coughlin, Teresa, John Holahan, Kyle Caswell, and Megan McGrath, "Uncompensated Care for Uninsured in 2013: A Detailed Examination," Technical Report, Urban Institute May 2014.

Courtemanche, Charles, James Marton, Benjamin Ukert, Aaron Yelowitz, and Daniela Zapata, "Early Impacts of the Affordable Care Act on Health Insurance Coverage in Medicaid Expansion and Non-Expansion States," *Journal of Policy Analysis and Management*, 2017, *36* (1), 178–210.

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

Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya, "Can Health Insurance Competition Work? Evidence from Medicare Advantage," *NBER Working Paper No. 20818*, 2015.

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

Dafny, Leemore, Jonathan Gruber, and Christopher Ody, "More Insurers Lower Premiums: Evidence from Initial Pricing in the Health Insurance Marketplaces," *American Journal of Health Economics*, 2015, *1* (1), 53–81.

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

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

Diamond, Rebecca, Michael J Dickstein, Timothy McQuade, and Petra Persson, "Take-Up, Drop-Out, and Spending in ACA Marketplaces," May 2018, (24668).

Dickstein, Michael J., Mark Duggan, Joe Orsini, and Pietro Tebaldi, "The Impact of Market Size and Composition on Health Insurance Premiums: Evidence from the First Year of the Affordable Care Act," *American Economic Review*, May 2015, *105* (5), 120–25.

Einav, Liran, Amy Finkelstein, and Pietro Tebaldi, "Market Design in Regulated Health Insurance Markets:Risk Adjustment vs. Subsidies," *Unpublished mimeo*, 2019.

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

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

Fillmore, Ian, "Price Discrimination and Public Policy in the U.S. College Market," *Unpublished mimeo*, 2019.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard, "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts," *American Economic Review*, 2019, *109* (4).

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

Frean, Molly, Jonathan Gruber, and Benjamin D. Sommers, "Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act," *Journal of Health Economics*, 2017, *53*, 72 – 86.

George, Lisa and Joel Waldfogel, "Who Affects Whom in Daily Newspaper Markets?," *Journal of Political Economy*, 2003, *111* (4), 765–784.

Geruso, Michael and Timothy J. Layton, "Selection in Health Insurance Markets and Its Policy Remedies," *Journal of Economic Perspectives*, November 2017, *31* (4), 23–50.

— , **Timothy Layton, and Daniel Prinz**, "Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges," *American Economic Journal: Economic Policy*, May 2019, *11* (2), 64–107.

Goldin, Jacob, Tatiana Homonoff, and Katherine Meckel, "SNAP Benefit Issuance and Grocery Sales, Pricing, and Promotional Behavior," *Unpublished mimeo*, 2018.

Gruber, Jonathan, *Tax Policy for Health Insurance*, in Tax Policy and the Economy, Volume 19 by MIT Press, September 2005

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

— **and James M. Poterba**, *Tax Subsidies to Employer-Provided Health Insurance*, in Empirical Foundations of Household Taxation by University of Chicago Press, January 1996

Hausman, Jerry, *The Economics of New Goods*, University of Chicago Press,

Jaffe, Sonia and Mark Shepard, "Price-Linked Subsidies and Imperfect Competition in Health Insurance," *NBER Working Paper No. 23104*, 2018.

Kowalski, Amanda E, "The Early Impact of the Affordable Care Act State-By-State," October 2014, (20597).

Layton, Timothy, Ellen J Montz, and Mark Shepard, "Health Plan Payment in U.S. Marketplaces: Regulated Competition with a Weak Mandate," Working Paper 23444, National Bureau of Economic Research May 2017.

Layton, Timothy J. and Thomas G. McGuire, "Marketplace Plan Payment Options for Dealing with High-Cost Enrollees," *American Journal of Health Economics*, 2017, *3* (2), 165–191.

Layton, Timothy J, Randall P Ellis, and Thomas G McGuire, "Assessing Incentives for Adverse Selection in Health Plan Payment Systems," September 2015, (21531).

Layton, Timothy J., Thomas G. McGuire, and Anna D. Sinaiko, "Risk Corridors and Reinsurance in Health Insurance Marketplaces: Insurance for Insurers," *American Journal of Health Economics*, 2016, *2* (1), 66–95.

Lieber, Ethan M.J. and Lee M Lockwood, "Targeting with In-kind Transfers: Evidence from Medicaid Home Care," January 2018, (24267).

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

Meckel, Katherine, "Is the Cure Worse than the Disease? Unintended Effects of Payment Reform in a Quantity-based Transfer Program," *Unpublished mimeo*, 2019.

Miller, Keaton S, Amil Petrin, Robert Town, and Michael Chernew, "Optimal Managed Competition Subsidies," Working Paper 25616, National Bureau of Economic Research February 2019.

Orsini, Joe and Pietro Tebaldi, "Regulated Age-Based Pricing in Subsidized Health Insurance: Evidence from the Affordable Care Act," Working Paper, Becker Friedman Institute for Research in Economics November 2017.

Panhans, Matthew, "Adverse Selection in ACA Exchange Markets: Evidence from Colorado," *American Economic Journal: Applied Economics*, April 2019, *11* (2), 1–36.

Poterba, James M., *Government Intervention in the Markets for Education and Health Care: How and Why?*, University of Chicago Press, January

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

Sacks, Daniel W., Khoa Vu, Tsan-Yao Huang, and Pinar Karaca-Mandic, "The Effect of the Risk Corridors Program on Marketplace Premiums and Participation," Research Paper No. 18-57, Kelley School of Business April 2018.

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

Stantcheva, Stefanie, "Optimal Income Taxation with Adverse Selection in the Labour Market," *The Review of Economic Studies*, 2014, *81* (3 (288)), 1296–1329.

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

_ , **Alexander Torgovitsky, and Hanbin Yang**, "Nonparametric Estimates of Demand in the California Health Insurance Exchange," Working Paper 25827, National Bureau of Economic Research May 2019.

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

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

Figure 1: Demand model: empirical moments and model fit

(a) Silver plan market share

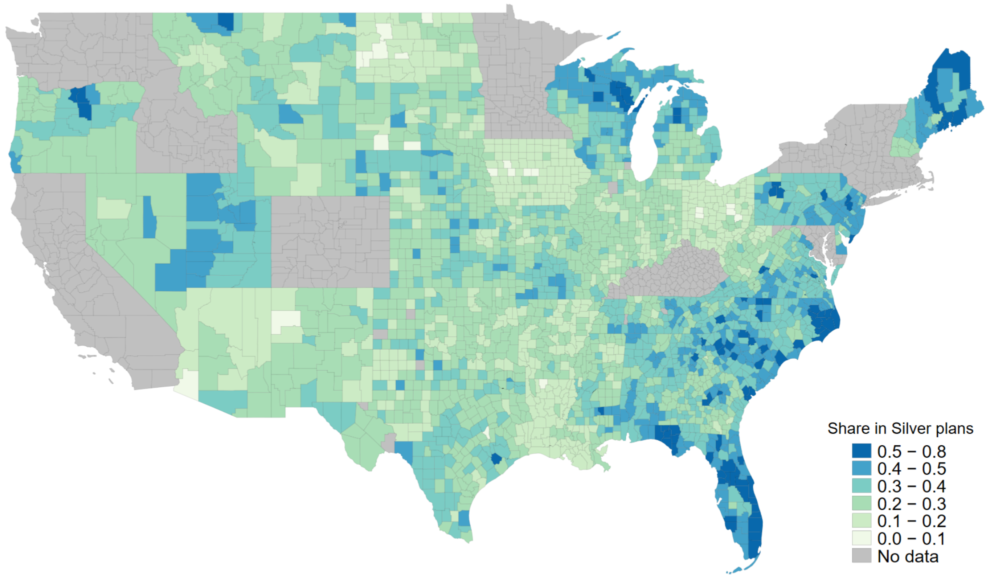

(b) Model in-sample prediction error of Silver plan market share

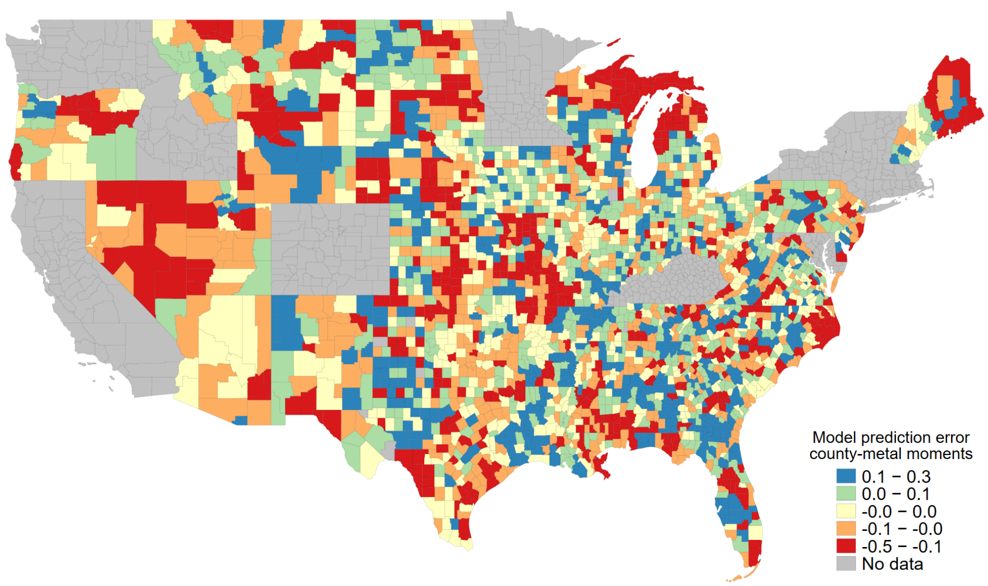

Notes: Map in panel (a) plots the share of potential consumers in each county that enrolled in a Silver plan on ACA Marketplaces. States that are marked with grey are not federally facilitated and do not enter our analysis. The counts of the pool of potential consumers (denominator) was provided by the Kaiser Family Foundation and is based on estimates from national surveys of how many people were uninsured or underinsured in each geographic region. The number of people that purchased a Silver plan (numerator) are administrative enrollee counts reported by CMS that do not account for disenrollments. Data is for year 2017. Map in panel (b) plots the difference between the observed share of enrollees in Silver plans - as pictured in Panel (a) - and the share of enrollment in Silver plans as predicted by demand model of Section 4.1.

Figure 2: Marginal cost estimates

(a) Distribution of marginal cost estimates

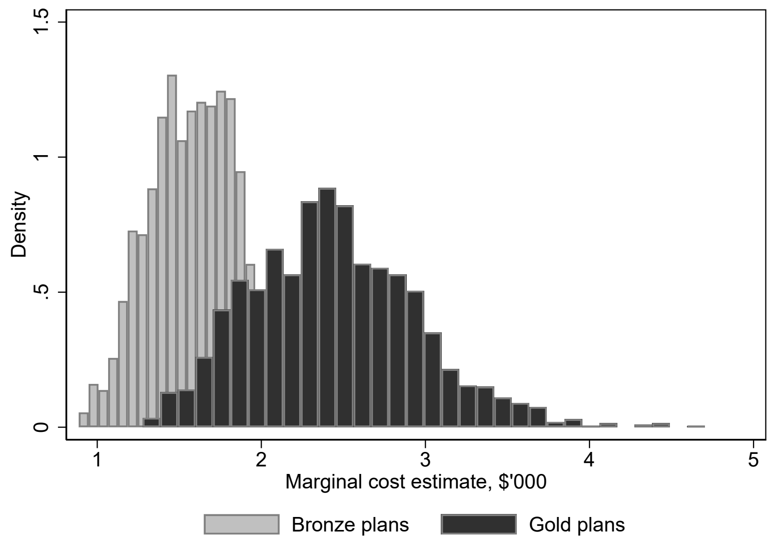

(b) Marginal cost estimates and observed accounting cost

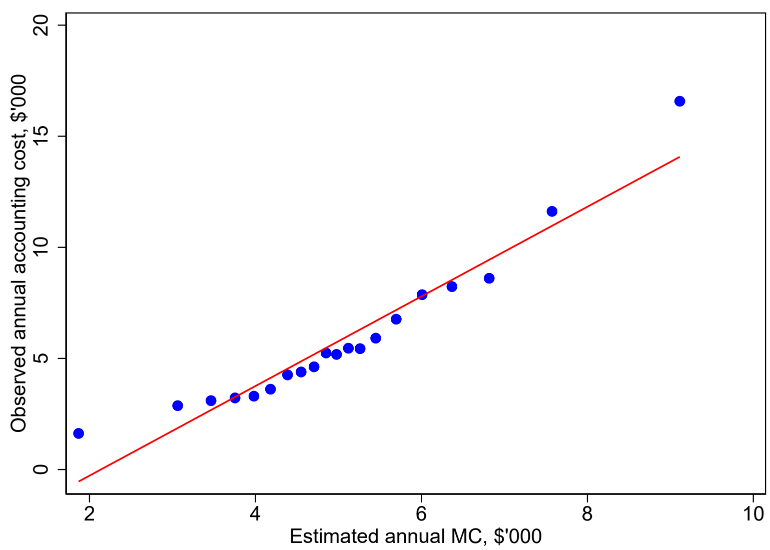

Notes: Panel (a) plots the distribution of plan-market level marginal costs as estimated in Section 4.2. The costs are plotted for a baseline, age 20, consumer. The costs are plotted separately for Bronze plans that provide the lowest amount of coverage and Gold plans that provide the highest amount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

Figure 3: Demographic externality with targeted subsidies and market power

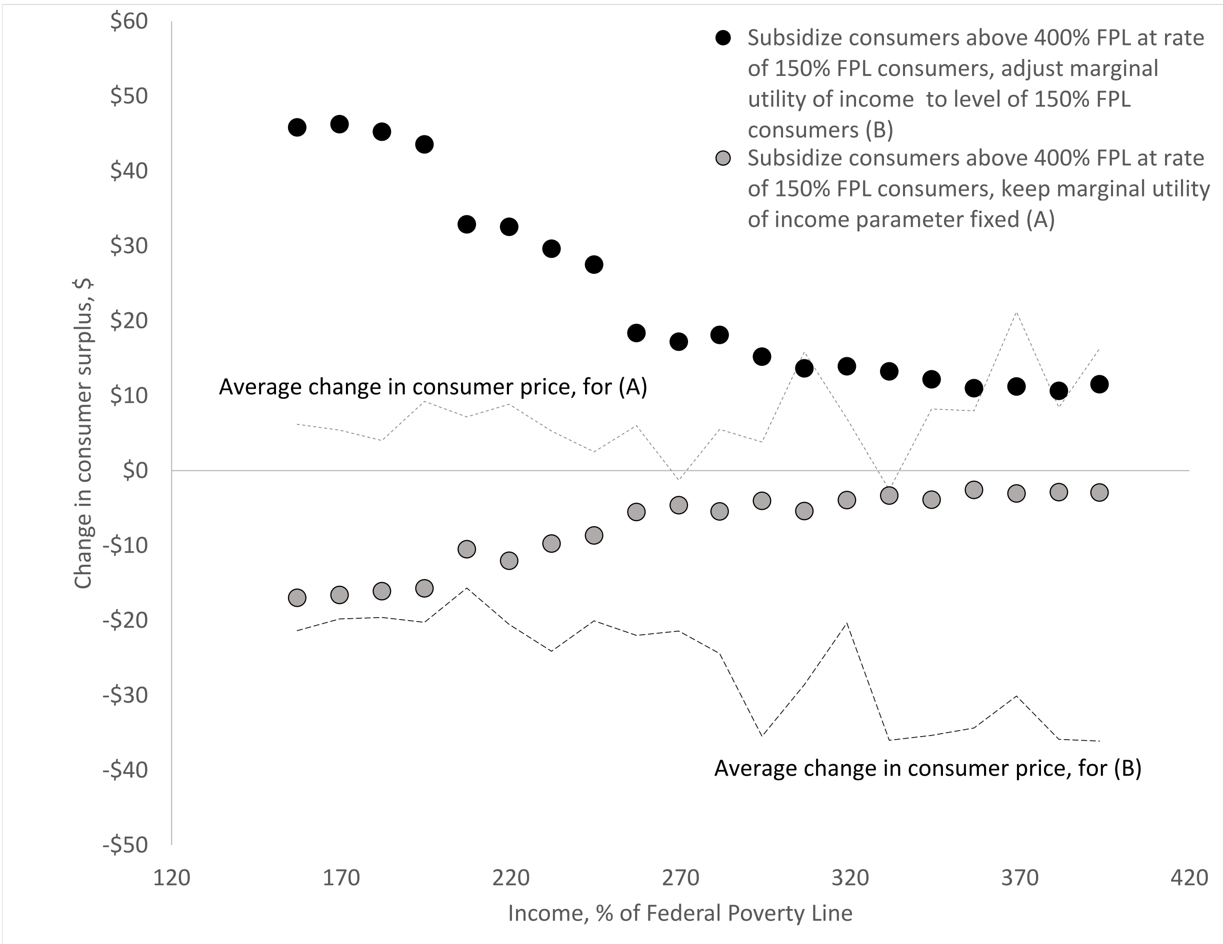

Notes: Figure reports estimated change in consumer surplus and consumer-facing prices (y-axis, in dollars) by income level (x-axis) in two counterfactual cases that capture the "demographic externality". The first counterfactual simulation (case A) changes income of consumers with true income of above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

Figure 4: Variation in consumer surplus under observed subsidies

(a) Demographic variation in consumer surplus

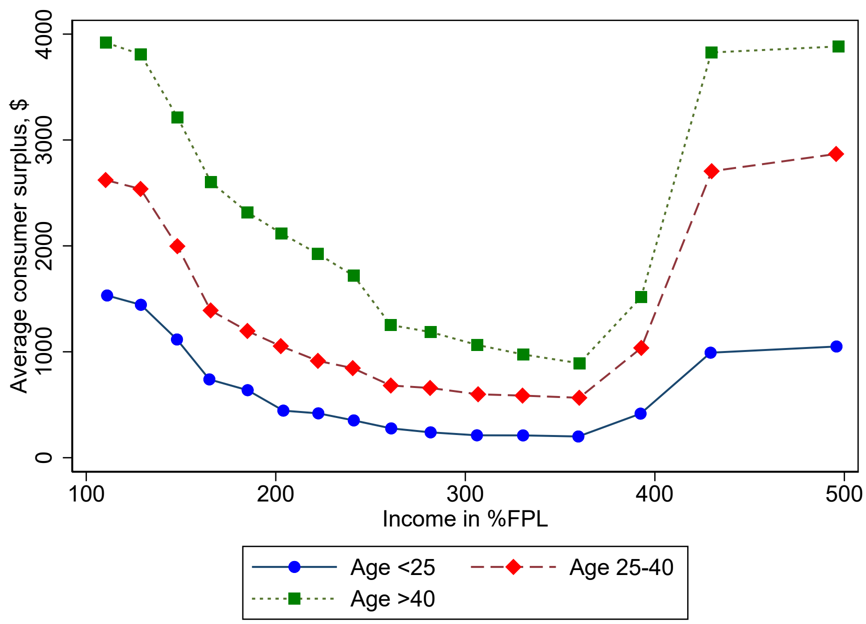

(b) Geographic variation in consumer surplus

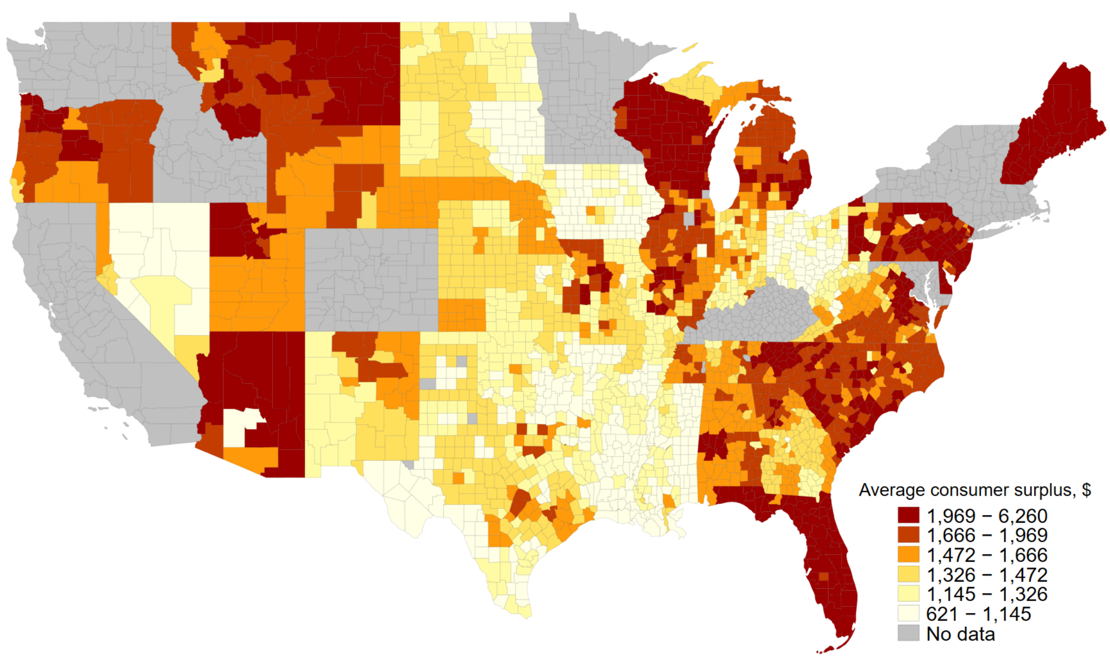

Notes: Panel (a) plots the average consumer surplus by income and age groups, as estimated under the observed subsidy regime. Map in Panel (b) plots average consumer surplus per county, also as estimated under the observed subsidy regime.

Figure 5: Geographic incidence of subsidies under imperfect competition

(a) Means-tested subsidies relative to no subsidies

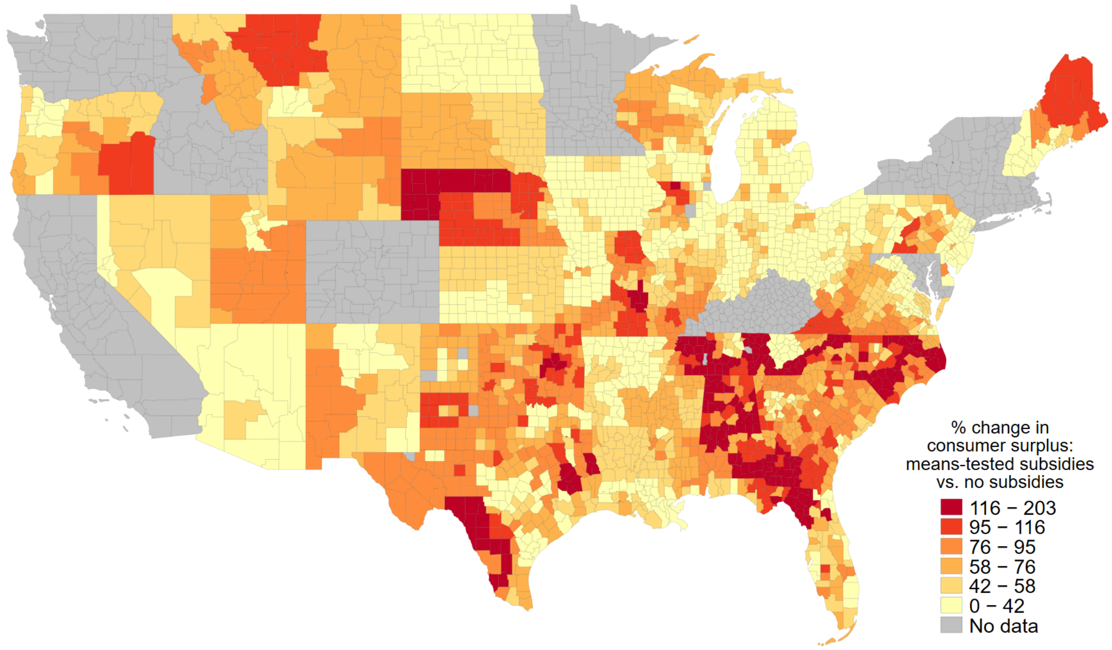

(b) Means-tested subsidies relative to flat subsidies

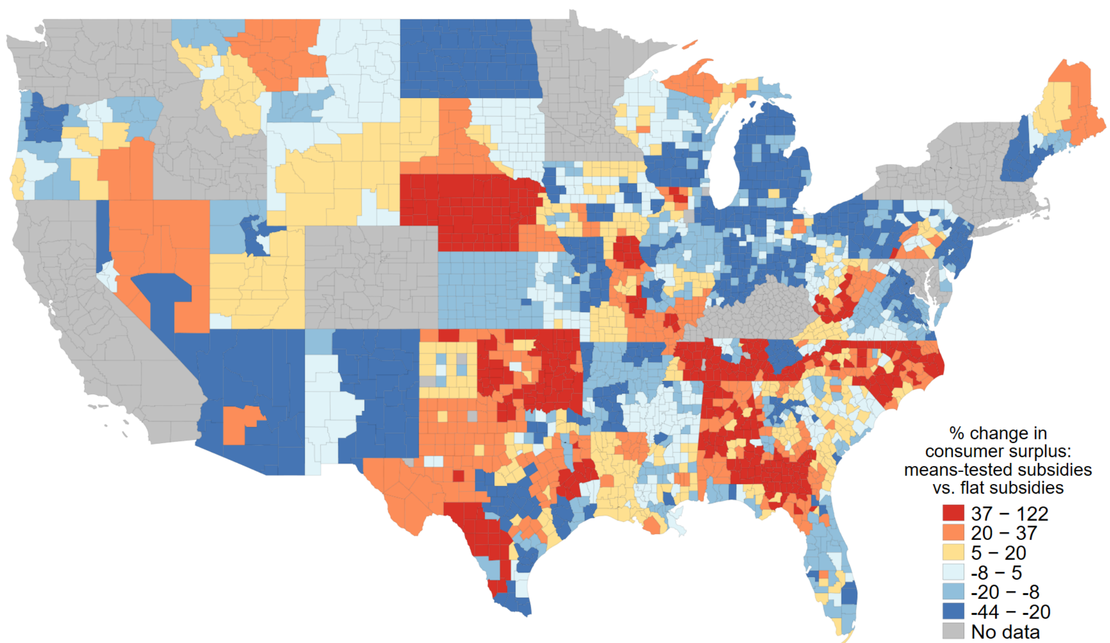

Notes: Panel (a) reports the percent change in average consumer surplus that results when we move consumers from an environment with no subsidies to means-tested subsidies at the level observed in the data. Insurers are allowed to retain market power and can adjust prices. Panel (b) reports the percent change in average consumer surplus from moving consumers from means-tested to flat (age-adjusted) subsidies, preserving total nominal government spending on subsidies. The baseline flat subsidy for a 20-year old that results in the same government spending as under means-tested subsidies is $1,105.

Figure 6: Flat subsidies vs. means-tested subsidies: efficiency effects

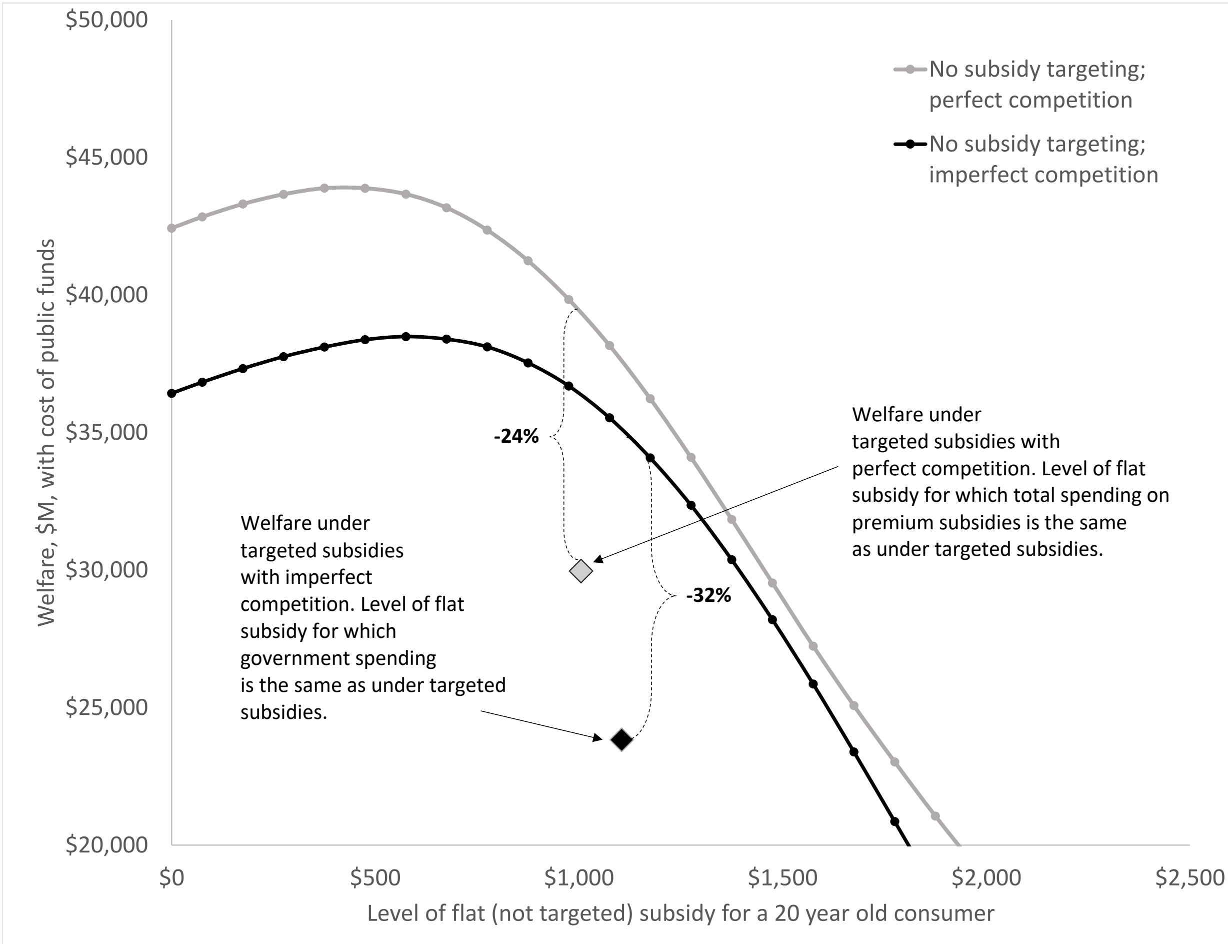

Notes: Figure reports estimated total welfare (y-axis), including the opportunity cost of government spending and the cost of raising public funds, under subsidies that do not vary with consumer income *within* an age group. The x-axis marks the value of the flat voucher that is offered to a 20-year old consumer. Consumers of all other ages receive vouchers that are scaled by the premium age curve as observed in the data. The grey line marks welfare estimates in counterfactuals that shut down market power in the market, setting baseline (for a 20 year olds) premiums to be equal to estimated marginal costs. The black line marks cases that preserve market power. Two diamonds mark the two levels of baseline (for a 20 year old) flat subsidies that lead to the same level of nominal government spending (premium and cost-sharing reduction subsidies combined) as under means-tested subsidies, in the perfectly-competitive (grey diamond) and the imperfectly-competitive (black diamond) environment. The y-axis value for both points record the level of total welfare achieved under means-tested subsidies. The values correspond to row (7), Columns (1) and (4) in Table 3. Curly brackets show the difference in welfare that is achieved - for the same level of government spending - between means-tested and flat (age-adjusted) subsidies. We observe that the welfare loss from subsidy means-testing is higher in the presence of market power.

Table 1: Summary statistics

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

‡ Across counties

‡‡ Based on Kaiser Family Foundation estimates

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

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

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

Table 2: Model Estimates

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

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

Table 3: Surplus under counterfactual subsidy mechanisms

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

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

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

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

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

Table 4: Demographic incidence of subsidies

Change in consumer surplus, %	Age<25 (1)	Age 25-40 (2)	Age >40 (3)	Age<25 (4)	Age 25-40 (5)	Age >40 (6)
	Panel A: Moving from no subsidies to means-tested subsidies					
	Under perfect competition			**Under imperfect competition**		
Income FPL <200	422.4	200.0	380.2	317.3	179.2	351.1
Income FPL > 200 and FPL < 400	148.6	52.5	203.6	58.6	30.6	126.3
Income FPL > 400	0.0	0.0	0.0	0.5	0.3	0.6
	Panel B: Moving from means-tested subsidies to flat subsidies^					
	Under perfect competition			**Under imperfect competition**		
Income FPL <200	-46.3	-37.6	-52.7	-48.4	-38.9	-57.7
Income FPL > 200 and FPL < 400	34.2	10.4	-26.4	59.2	16.2	-16.9
Income FPL > 400	108.1	33.9	52.1	127.0	35.9	58.4

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

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