NBER WORKING PAPER SERIES

SUBSIDY TARGETING WITH MARKET POWER

Maria Polyakova
Stephen P. Ryan

Working Paper 26367
http://www.nber.org/papers/w26367

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2019

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, University of Toronto, 2019 Penn State-Cornell Econometrics \& IO Conference, NBER Health Care Meetings, and NBER Public Economics Meetings for helpful comments. We thank Alan Jaske Lynn Hua, and Vinni Bhatia for excellent research assistance. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Subsidy Targeting with Market Power
Maria Polyakova and Stephen P. Ryan
NBER Working Paper No. 26367
October 2019
JEL No. H0,H2,H20,H21,H22,H23,H31,H41,H5,H51,I1,I10,I11,I13,I18,I38,L0

ABSTRACT

In-kind public transfers are commonly targeted based on observable characteristics of potential recipients. This paper argues that when the subsidized good is provided by imperfectly-competitive firms, targeting can give rise to a “demographic externality,” creating unintended redistribution of surplus and distorting efficiency. We illustrate this mechanism empirically in the context of means-tested subsidies for privately-provided health insurance plans under the Affordable Care Act (ACA). Using a structural model of supply and demand, we show that market power increases the welfare loss from subsidy targeting, vis-a-vis income-invariant subsidies, by 33 percent.

Maria Polyakova
Department of Health Research & Policy
Stanford University
Redwood Building T111
150 Governor's Lane
Stanford, CA 94305
and NBER
maria.polyakova@stanford.edu

Stephen P. Ryan
Olin Business School
Washington University in St. Louis
Campus Box 1133
One Brookings Drive
St. Louis, MO 63130
and NBER
stephen.p.ryan@wustl.edu

1 Introduction

Public welfare programs have a long history of linking their transfers to observable characteristics of potential recipients, such as income, age, employment, family, or disability 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 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.[1]

In practice, however, governments are increasingly turning to profit-maximizing firms to provide public benefits.[2] 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: just for privately-provided health insurance alone, the US government expects to pay $560 billion in subsidies in 2019, 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. The contribution of this paper to this debate is twofold. First, we combine insights from the literature in public economics and industrial organization to investigate how one common design choice—targeting of transfers—may interact with market power, leading to changes in the efficiency of allocations and their ability to redistribute surplus. Second, through our empirical application, we provide novel estimates of efficiency and incidence of subsidies, in a new market for non-group health insurance plans in the United States that insures about ten million individuals, making these estimates important in their own right.

The interaction of market power on the supply side and targeted taxes or subsidies on the demand side can generate distortions in equilibrium. There are two key, potentially countervailing, forces at work here. First, subsidy targeting introduces heterogeneity in

[1]See, for example, Currie and Gahvari (2008) and Hoynes and Schanzenbach (2015) for overviews of the literature on targeted and universal in-kind and in-cash transfers in the United States. Some recent examples of the literature documenting the demand-side distortions include Jacob and Ludwig (2012); Allcott et al. (2015); Lieber and Lockwood (2019); Basurto et al. (2017).

[2]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.

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. Second, targeted subsidies change the relative importance of different consumers in the firm's profit-maximization problem. If subsidized consumers are inherently 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*.[3]

We explore the equilibrium effects of subsidy targeting in the presence of market power empirically in the context of publicly-subsidized, privately-provided health insurance Marketplaces created under the Affordable Care Act (ACA). These relatively new markets insure about ten million Americans and present an important precedent for how a non-group health insurance market could work in the US. 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 resembling traditional means-tested transfers. 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 that leverages price variation that is induced by kinks in the subsidy formulas (the approach is similar in spirit to Tebaldi

[3]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.

et al. 2019). The regulatory variation in prices also allows us to capture the 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 our setting. We then move to our main analysis, by first quantifying the allocative efficiency of the observed subsidy regime. This exercise serves two purposes—it first delivers novel estimates of efficiency and incidence of public spending in this important market, and second, it establishes a quantitative baseline against which we interpret our counterfactual results. We next turn to a series of counterfactuals that modify the subsidy regime and 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 mechanism, we start with two stylized exercises that show the complex interplay between targeted subsidies and the distribution of consumer demographics within a market. In our first experiment we introduce generous subsidies for the highest-income consumers (who were previously not eligible for subsidies), keeping everything else fixed. This leads to a shift out in the aggregate demand curve and induces insurers to increase prices. The increase in prices implies that consumers not directly affected by the subsidy change experience a negative demographic externality: their consumer surplus falls because their (high-income) *neighbors* become subsidized. This simple case, however, does not capture the idea that the demographic externality depends both on the amount of the subsidy and also on who is eligible to receive it. To capture this nuance, in our second exercise we endow those newly-subsidized highest-income consumers with the demand parameters of the lowest-income consumers. Since in our empirical application lower-income consumers have more elastic demand, prices now move down in response to a more elastic demand even though subsidies have increased. In this second experiment, the change in aggregate elasticity dominates the change in subsidies, leading these newly-subsidized consumers to exert a positive demographic externality on their neighbors. How these two opposing forces play out in any given market in equilibrium is an empirical question, which we turn to next.

We start by computing two model simulations that allow us to measure the efficiency of the observed allocation in ACA Marketplaces and to also document the incidence of means-tested subsidies between consumers and producers under this allocation. We find that under the observed subsidy regime, the ACA Marketplaces generate \$42 billion in consumer and producer surplus for \$30 billion in nominal public spending. Net of \$16 billion, which we estimate the government saves on the same consumers (for example, by not paying for uncompensated care), the extra government outlays for the ACA program amount to \$14 billion. Accounting for the cost of public funds, we conclude that \$1 of extra public spending invested in this market generates \$2.33 of surplus in return. In other words, the ACA Marketplaces generate a substantial positive return on government spending. Contrasting this allocation to a simulated equilibrium that fully shuts down premium subsidies, we conclude that consumers capture the majority of subsidy payments, with 84 percent of surplus generated by subsidies accruing to consumers and 16 percent going 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 than higher-income consumers.

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 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 \$263 (24 percent), enrollment increases by 24 percent, consumer surplus goes up by nearly \$10 billion, from \$38 to \$47 billion, while government expenditures on premium subsidies 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.[4]

One of the major policy goals of using targeted subsidies is to achieve distributional objectives. To understand how subsidy design per se influences efficiency and equity of the market, we next compute an equilibrium where both subsidy targeting and market power are removed. We replace means-tested subsidies with an income-invariant voucher that we allow to vary be age, following the regulatory age curve for premiums. We focus on the voucher level that results in the same level of subsidy spending as under 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 strongly affects what type of consumers choose to enroll. Many more higher-income consumers enroll under flat vouchers. As a result, consumer surplus increases from \$47 billion to \$54 billion. 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.[5] The increase in surplus is not uniform among consumer types, however: consumers at the bottom of the income distribution experience about a 50 percent decline in average consumer surplus, while those with higher income experience a significant increase. These results are intuitive and highlight the equity-efficiency trade-off that is inherent in subsidy targeting. In our empirical setting, higher-income consumers have less elastic demand and a higher revealed valuation of the good. Thus, spending the same amount of public funds through non-targeted subsidies substantially increases the size of the pie, but it is not a Pareto-improving policy.

Finally, we contrast the efficiency implications of means-tested and income-invariant subsidies in the presence of market power. We again focus on the voucher level that holds government spending on subsidies constant between the two subsidy regimes. We estimate that moving from means-tested to universal subsidies increases enrollment, but, as in the case with perfect competition, there is a significant regressive re-distributional effect. 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 we

[4] A complementary comparison, in which we vary market power after dropping public subsidies altogether, gives us another way of measuring the effect of imperfect competition in this market. 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 28 percent lower.

[5] We note that in our setting, means-tested subsidies generate a welfare loss. If consumers that are targeted by subsidies had higher utility from the good, then targeting would generate a *gain* in welfare relative to an income-invariant subsidy at the same level of government spending.

observe under the means-tested subsidy regime. With vouchers, both consumer surplus and insurer profit increase by nearly \$9 billion to a combined \$50 billion. The 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 welfare loss generated by targeted subsidies is \$9.4 billion, or 24 percent of total welfare. In contrast, accounting for how targeted subsidies interact with market power changes that estimate dramatically; total welfare is \$6.1 billion lower and the welfare 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 welfare loss from targeted subsidies by 33 percent.[6]

Our analysis makes several contributions to the literature. First, starting with the specific context of the ACA markets that constitute our empirical application, our study relates to the rapidly growing literature that is examining the economics of these markets and their effects on a variety of socio-economic outcomes.[7,8] Several papers in this literature are related to our work. First, Ericson and Starc (2016); Abraham et al. (2017); Drake (2019); Saltzman (2019) among others, have estimated demand in state-based ACA Marketplaces, focusing on a handful of states, such as California, Colorado, Massachusetts, and Washington. Second, several studies have directly considered the question of optimal subsidy design in this market using stylized simulations[9] Finally, a handful of previous studies have consid-

[6]In general, a policy-maker that fails to account for the strategic responses of firms will have biased expectations about the welfare of the subsidy mechanism. In principle, this bias could go the opposite way than what we find in our specific empirical setting, leading the policy-market to overestimate welfare losses.

[7]Gruber and Sommers (2019) provide an extensive literature review of studies that have examined the impact of the ACA on the rate of insurance coverage, on healthcare utilization and health outcomes, on healthcare providers, on employment, and on state and federal budgets.

[8]A significant share of this literature has examined the role of adverse selection and policies designed to counteract selection in ACA markets. The results are somewhat mixed depending on the specific context, but several studies find that risk-equalization policies on these markets are effective at neutralizing a substantial amount of screening incentives (Geruso et al., 2019; Sacks et al., 2018). Other studies in this vein include Layton and McGuire (2017); Layton et al. (2017); Kowalski (2014); Panhans (2019); Diamond et al. (2018).

[9]For example, Einav et al. (2019) consider the difference in market outcomes between policies that channel

ered the supply side and the role of market power on ACA markets. For example, Dafny et al. (2015) show that premiums in ACA markets are lower when more insurers are competing in the market, while Cicala et al. (forthcoming) show that minimum loss ratio regulation constraints premiums. Further, using data from the California Exchange, Dickstein et al. (2015); Orsini and Tebaldi (2017) argue that the interaction of price-discrimination restrictions on plans and differences in the demographics across geographic markets can affects pricing incentives. A related idea is examined in Tebaldi (2017), which is the closest to our paper, both methodologically and in terms of the empirical context. Tebaldi examines the welfare implications of price-linked subsidies in California ACA Exchanges. He highlights a different type of externality in this market—an externality that arises because insurers pool risks across consumers of different ages, but are only allowed to partially price discriminate across consumers of different ages. The demographic externality that we study is different in spirit, as it is not specific to markets with selection. Risk-pooling is not needed—the demographic externality can arise in any market with subsidy targeting when the firms providing the good posses market power and have to set the same list price for subsidized and unsubsidized consumers.

By considering how subsidy targeting and market power interact in the context of a health insurance market, we also contribute to the literature that has examined the role of subsidy design in a broader health insurance context. This literature has investigated the efficiency and incidence of tax exclusions for employer-sponsored health insurance in the US (e.g., Gruber and Washington, 2005; Gruber, 2005; Gruber and Poterba, 1996); the design of employer contributions to employer-sponsored plans (Cutler and Reber, 1998; Ho and Lee, 2019); and the design of subsidies in several publicly-funded privately-provided health insurance programs in the US and other countries. Conceptually and methodologically, our paper is closely related to the latter group that includes Curto et al. (2015), Decarolis (2015), Decarolis et al. (forthcoming), Jaffe and Shepard (2018), Miller et al. (2019). These papers study the design of subsidies in health insurance markets such as Medicare Advantage, Medicare Part D, and Massachusetts Health Insurance Exchange, focusing on the idea that subsidies linked to prices of insurers may distort allocations. In many of these applications, subsidies are not universal and are frequently targeted on income; the contribution of our

public transfers through the demand-side (subsidies) or supply-side (through risk-adjustment) of the market; they test the theoretical model of these two mechanisms in the Covered California context. In a different vein, Saltzman et al. (2015); Taylor et al. (2015) used the RAND Corporations model of employer-sponsored insurance to simulate an extensive set of alternative subsidy designs for ACA Marketplaces in a stylized framework with no strategic firms.

work is to consider the efficiency and equity implications of this aspect of the subsidization mechanism.

Last, but not least, our 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, 2019 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 important empirical context. 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; and Collinson and Ganong (2018) examine how housing vouchers may change housing prices.

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 a standard result for why governments may target subsidies on the basis of observable information about consumers, such as income, we first begin with a textbook model of a benevolent government that aims to allocate in-kind subsidies in order to maximize social welfare subject to a fixed budget constraint. In this simple framework, We show that the optimal subsidy is differentiated across consumer types due to differences in the curvature of the utility function. Then, taking the idea that subsidies may be differentiated by observable consumer type, we model the incentives of a profit-maximizing firm that provisions the good to the consumers. We show that the combination of subsidy targeting and market power interacts in a complex fashion 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:

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

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

$$\eta U_L'(w_L) = (1 - \eta) U_H'(G - w_L). \tag{2}$$

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.[10]

Next, we introduce a profit-maximizing firm to act as an intermediary to provide the good.[11] 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

[10] 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 targeting.

[11] 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.

each type differentially. We are particularly interested in the interaction of the single-price 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.[12] 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, $\frac{\partial p_H(b)}{\partial b} = 1$, i.e. 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

[12] 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.

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

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 of 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, replacing subsidy targeting with a universal subsidy could lead to higher social welfare (albeit, importantly, not necessarily a Pareto improvement) 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 than the optimal subsidy level under public provision. Intuitively, targeted subsidies change the level of 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 by a public policy that differentially changes consumer demand function by providing targeted subsidies. This effect does not depend on the cost linkage and can be present in any type of market.

3 Economic Environment and Data

Institutions Our empirical application is the US market for non-group health insurance plans that was created by the Affordable Care Act and started its operation in 2014. The program (often refereed to as ACA "Marketplaces") allows consumers to purchase health insurance plans for themselves and their families. Enrollment is voluntary; however, individuals that do not have any health insurance face fines. Insurance plans offered on this market are 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 that is measured as a fraction of costs that a plan would cover for a standardized population. In addition to differences in cost-sharing and the scope of coverage, plans also vary in their provider networks and (possibly) customer service.

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 "rating areas." Insurers have to charge the same price in all counties they serve within the same rating area. However, insurers do not have to serve all counties in a rating area; it follows, as Fang and Ko (2018) argue, that it is reasonable to think about a county as the relevant market boundary in this setting.

Insurers in this market are not allowed to price-discriminate based on individual health risk, but they are allowed to charge different premiums based on age.[13] Age adjustment of prices is tightly regulated. Insurers have to follow a regulatory age curve that determines how much prices can differ across consumers of different ages.

The key institutional feature of ACA Marketplaces that is relevant for our analysis 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. The second type of subsidies, held constant in our analysis, are cost-sharing subsidies that reduce the out-of-pocket liability, such as deductibles, co-pays, etc. Both types of subsidies are means-tested: Families with lower household incomes receive higher

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

subsidies.

Premium subsidies are known as Advanced Premium Tax Credits (APTC). [14]. The APTC is calculated in several steps. First, the Modified Adjusted Gross Income (MAGI) for a tax family is converted to the percent of Federal Poverty Level (FPL).[15] The FPL varies with family composition and allows comparing incomes of families of different sizes using the same scale. The MAGI relative to FPL measure then determines the maximum dollar amount that the (tax) household "should be" paying for insurance premiums. Let us call this amount a "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. At 400 percent FPL, the CAP is equal to infinity: Households with income at or above 400 percent FPL are not subsidized.

Next, to determine the level of subsidies that a household is eligible to receive, the regulator considers the distribution of list premiums on ACA Marketplaces in the county of residence for this household. 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 amount is larger 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 SLSP and the CAP.

Data We combine several sources of data for our analysis. We focus our analysis on the 2017 edition of ACA Marketplaces that used the federal healthcare.gov platform. We use Marketplace Public Use Files that record detailed information on the choice sets that consumers faced in each geographic market. We also use two sets of data provided by The Center for Consumer Information & Insurance Oversight (CCIIO). The first, reports enrollment at county-metal level, at plan level, and at county-insurer level. The second set of a data are so-called rate review files that contain information on accounting costs as

[14]The APTC can be claimed concurrently with enrollment based on projected household income and then adjusted (if necessary) when consumers file taxes. Consumers can also choose to forgo receiving advanced credit and instead claim the subsidy as a regular tax credit in their tax return

[15]MAGI is reported on US tax form 1040. The AGI is the total income that includes wages, tips, self-employment income, etc., as well as taxable interest, dividends, taxable parts of the social security income, IRA, pension, and annuity distributions that is adjusted for a variety of deductions specific to the income source, such as, for example, student loan interest deduction. MAGI modifies the AGI by adding back certain deductions.

reported by insurance plans.[16] 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 2017 American Community Survey (ACS) to create a representative sample of potential consumers (individuals without employer-sponsored or public health insurance coverage) in each county, for whom we observe income, age, race, and gender.[17]

Table 1 summarizes the key data points on the choice sets, enrollment, and demographics. In 2017, on average consumers could choose among plans offered by 1 to 4 large national insurers and a number of smaller firms. The annual list premium 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. We observe 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. 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. Consumers pick plans that give them the highest utility at the lowest price, or choose not to participate in the market. In practice, demand operates at the family level. The ACA institutional design imposes that all consumers within a tax household that need health insurance coverage form a "coverage family" and enroll in the same health insurance

[16] All CCIIO data is public and can be accessed at https://www.cms.gov/CCIIO/Resources/Data-Resources/index.html.

[17] In ACS, we exclude individuals from our set of potential consumers if they report having no health insurance coverage, but would likely be eligible for Medicaid in expansion states. When applying for health insurance coverage on the Marketplaces, these individuals would in general be diverted to Medicaid.

plan. To operationalize such family-level choices in our empirical model, we assume that each consumer maximizes the average utility of her coverage family. We drop family-specific notation for simplicity in what follows, but each term that enters the utility function denotes an average within a coverage family.[18]

Formally, we posit that individual i in market t chooses plan j from a set of choices J. The set of choices that each consumer faces depends on the consumer's family geographic location, household income, and age distribution of family members that require insurance coverage.[19] The indirect utility function takes the following form:

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

Where p_{ijt} is individual-specific (average across members of the coverage family) price that consumer i (in family f), who lives in market t, faces for plan j. This price is indexed by i: consumers may face different prices for the same plan j in market t depending on their household income and age composition of their coverage families. ϕ_{ijt} is the amount of utility a consumer in market t gets from plan j. 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 may get more generous coverage.

We make several assumptions about α_i and ϕ_{ijt} to arrive at an empirically-tractable version of this utility function. 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), to capture the idea that older consumers may value insurance

[18] For example, price $p_{i(f)}$ is the average price that all consumers in coverage family f face. We account for the coverage family structure of each potential consumer in estimation.

[19] Recall that we define a county as the relevant geographic unit, following the evidence in Fang and Ko (2018). 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.

more, all else equal.[20] The second component captures the deviations in the generosity of a plan j that consumers may see if they have low enough income. For these consumers, cost-sharing reduction subsidies change the actuarial value (AV) of a plan. Finally, we include a plan-specific constant δ_j for each plan j that captures the average utility that a consumer gets from purchasing plan j. The remaining difference between ϕ_{ijt} and these three components is captured by the taste shock, ϵ_{ijt}, that is observed by the consumer but not the econometrician. We assume that this shock is distributed iid (at the family level) with Extreme Value Type 1 distribution, which leads to a logit discrete choice model.[21] 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 $\frac{1}{|f|}\sum\limits_{i \in f} u_{ijt} > \frac{1}{|f|}\sum\limits_{i \in f} u_{ikt}$ for all k in J such that k is not equal to j.

The variation in premiums p_{ijt} that we observe in the data does not stem from experimental assignment, implying that, in general, p_{ijt} could be correlated with ϵ_{ijt}. However, the regulatory design of this market generates variation in prices that is independent of demand shocks and helps 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.[22] 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

[20] 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.

[21] 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.

[22] We make a substantive assumption that all consumers that see plan j across different markets have the same δ_j. In the standard BLP setting, one estimates a separate δ_{jt} for each product in each market, which we cannot implement due to data limitations.

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 market-level enrollment in any plan by age and income intervals, this policy-induced 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 plan-level enrollments, 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 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 Marketplaces decide which geographic markets to enter, how to design their plans, and how to price them. In this paper we are interested in how the targeting of subsidies 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 numerous regulatory provisions. To fix ideas, 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.

For each consumer i of age a, plan j (im market t) collects revenue that consists of several pieces. First, plan j collects premium p_{ij} from the consumer. For consumers that are not eligible for premium subsidies, this premium is equal to the 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 and a subsidy from the federal government. Together, the consumer premium and the subsidy add up to b_j^a.[23] Second, the insurer collects additional revenue, if any, from three risk-equalization programs that we describe below.

[23] 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.

On the expenditure side, the insurer pays for consumers' healthcare bills and any administrative costs. Let the total *expected* healthcare spending of consumer i in plan j be h_{ij}. In general, this spending is a function of consumer's underlying health risk, r_i, and the plan's contract features, ϕ_j.[24] Plan j's expected cost for consumer i is usually not equal to h_{ij}. Instead, the plan expects to pay a portion of h_{ij}, net of consumer cost-sharing in the form of deductibles, co-pays, and co-insurance. 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 $c_{ij}(r_i, \phi_j, D_i) \leq h_{ij}$, where D_i denotes consumer i' s income.

Prior to any risk-equalization transfers, plan j's expected profit from enrolling consumer i of age a is generally:

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

Let c_j^a denote the baseline plan-specific cost of covering an average enrollee of age a. We can re-write c_{ij} as the sum of c_j^a and the 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, the expected profit of plan j from enrolling individual i of age a becomes:

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

where the individual-specific expected cost term $\tilde{c}$ allows for the presence of advantageous or adverse selection into plan j that is a function of plan characteristics ϕ_j.

Three risk-equalization programs existed on ACA Marketplaces within the time horizon we study. The programs aimed to equalize expected insurers' costs across all enrollees, thereby reducing the incentives for active cream-skimming by insurers and ameliorating the cost consequences of sicker consumers self-sorting 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.

[24] 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, so that for the same underlying risk r_i, total expenditures generated by the consumer in plan j are different from those in plan k.

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 was 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. For example, the reinsurance program effectively gave insurers additional individual-specific revenue for individuals with particularly high $\tilde{c}_{ij}(r_i, \phi_j, D_i)$, reducing the importance of this term in the profit function. Denote the difference between this additional revenue and the idiosyncratic cost-component with 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 $H_j(\phi_j) = \sum_{i \in j} \nu_{ij} - R_j$ denote any residual selection.[25] If the risk-equalization programs fully offset the ex ante net idiosyncratic shocks, this term would be zero. Without data on the accuracy of the reinsurance and risk-adjustment programs for the ACA Marketplaces, we can only estimate the profit function up to the constant $H_j(\phi_j)$.[26] We simplify the notation by assuming that $H_j(\phi_j) = 0$ in subsequent discussion. Importantly, the normalization of $H_j(\phi_j)$ to zero has no bearing on the first-order condition in prices.[27] Finally, 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'

[25] Note that H is a function of plan characteristics, but not plan prices. The assumption that $\tilde{c}_{ij}(r_i, \phi_j, D_i)$ primarily depends on contract characteristics, while screening of consumers on premiums is second-order follows from the following argument. 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 the prices of plans that these consumers face, eliminating the possibility of quantitatively meaningful risk-selection on premiums. Instead, 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.

[26] Geruso et al. (2019) find empirical evidence in support of risk-equalization programs performing well in this market. From other markets that employ risk-equalization policies, we know that while risk-equalization leaves some scope for residual selection, it goes a long way to reducing the differences in costs in expectation.

[27] In footnote 25 we argue that, *conditional on* ϕ_j, selection on prices is unlikely to be quantitatively important. In other words, the derivative of $H_j(\phi_j)$ with respect to prices is close to zero. 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). In the latter, we considered models that that H_j was and wasn't a function of prices and did not find that this assumption had first-order effects on the simulation of insurers' pricing decisions.

pricing incentives on the margin.

Under these assumptions about the structure of the risk-equalization programs, we can re-write 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 potential consumers of age a on the market, we get:

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

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.[28] 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:

[28] 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.

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

For an insurer that offers more than one plan in a market, the set of j first-order conditions accounts for own and cross-price elasticities of demand. In this case, the first-order condition takes a vector form: $S - \Omega(B - C) = \mathbf{0}$, 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. We observe age-specific prices charged by each plan and we want to recover costs c_j^a that are unknown up to the scaling factor τ. To accomplish this, 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.[29] 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-cheapest silver plans (SLSPs) and consumer's family income. SLSPs face a different set of incentives in the market. Under the observed subsidy regime, when these plans change their bids, it affects not only their own prices, but also the subsidies, and hence consumer premiums, of *all* plans.[30] 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.[31] Instead,

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

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

[31] Note that since marginal costs of each plan are separable in the inverted first-order condition, the

following (Decarolis et al., forthcoming), we take a hedonic regression approach and use other plans to impute the marginal costs of SLSPs. We project the estimated marginal costs (from the first-order condition inversion) for the non-SLSPs on a rich vector of plan characteristics and then use this projection to predict the marginal costs for the SLSPs.

4.3 Efficiency Metric

We define a welfare function (W) that consists of three pieces: consumer surplus (CS), insurer profits (Π), and government subsidy spending (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. 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, in return, is the average utility within a family) for plan j and is equal to utility net of the idiosyncratic ϵ term.[32] 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 and cost-sharing reduction subsidies. The former are computed either from the data or are adjusted following the simulation scenarios of Section 5. Cost-sharing

inversion of marginal costs for non-SLSPs is not affected by the SLSP.

[32]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.

reduction (CSR) spending is held at observed levels. Specifically, using CCIIO data reports we compute the average per capita spending on CSR subsidies by consumer type, based on income brackets.[33] In all counterfactual simulations, we then assign this average spending level to each consumer who falls into the respective income bracket and who enrolls in a plan where cost-sharing reduction is available. In the final step, we account for the fact that when a consumer enrolls into an ACA plan, the government likely saves some money on this consumer; for example, if a consumer enrolls in a formal insurance plan, this consumer is then unlikely to benefit from any public payments for uncompensated care. Following the Kaiser Family Foundation 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 on each consumer who buys an ACA plan.

4.4 Estimation results

4.4.1 Demand Parameters

We use non-linear least squares to estimate utility function parameters. 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 a larger impact on the utility of poorer consumers and younger consumers. The relationship between the overall value of insurance and age, as captured by age-specific intercepts, 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, all else equal. Figure 4a plots compensating variation (CV) by income and age, allowing us to summarize how much utility each demographic group gets from the observed plan offerings. When taking into account age-rated premiums and other plan features, we find a monotone relationship between consumer surplus and age - older consumers at any income point value insurance more than younger consumers. We find that the relationship between CV and income is, however, U-shaped. This follows intuitively from two countervailing forces: 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. Other plan characteristics are captured here by

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

the plan's δ that measure the relative attractiveness of each plan 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, ranging 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 our 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, we compare how enrollment moments predicted by the model compare to the moments observed in the data. Figure 1 illustrates one set of moments that was used for estimation and a related measure of 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 (across all metal levels) enrollment shares for each of 2,566 counties that are used in the estimation. The differences between the data and the model prediction are very close to zero, indicating that we closely match these aggregated enrollment shares and that the model is 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-SLSP plans. The estimates are reassuringly intuitive: 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 cost by approximately $500 ($300 metal-label effect plus the actuarial value adjustment of $192).[34]

[34] 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

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 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: Plans on the ACA Marketplaces are extremely heterogeneous, with some plans being offered by large national insurers and some by local co-operatives. Second, there are 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 at least 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, accounting costs do not include some ex post cost reconciliation, such as, for example, MLR payments. Nevertheless, the accounting cost data provide a valuable informational signal. They are likely to provide an accurate ordinal ranking of plans from the least to the most expensive, on average; they and also give a general sense of cost levels in the market. As we would expect given the existence of ex post cost reconciliation transfers, our estimates of marginal cost are on average lower than reported accounting cost, although they have the same general order of magnitude, as Figure 2b illustrates. In the same Figure, 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.[35]

fixed effects. The coefficients on these terms also have mostly intuitive signs.

[35] Related work in this areas has pursued a different approach - directly using accounting costs as inputs into the counterfactual exercises and 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 the two measures.

4.4.3 Demographic Externality

The combination of estimated demand and cost parameters allows us to empirically illustrate the demographic externality mechanism that we described conceptually in Section 2.

We entertain two thought experiments to illustrate this idea. In the first experiment, we 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. We then re-simulate the model and find the equilibrium. The results of this simulation are reported in Figure 3 and are marked as case "A." This figure shows how the average premiums and consumer surplus change for *directly unaffected* consumers—those with incomes between 150 percent and 400 percent FPL—when their neighbors with income above 400 percent start getting subsidies. Insurers take advantage of the fact that in the 400+ percent FPL market segment, consumers now face lower prices for any given list price, and raise list premiums. As a result, plans become more expensive for "unaffected" consumers with incomes under 400 percent FPL. As the light dashed line marked with "A" in the figure illustrates, the average annual effective premium 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 changes in prices and consumer surplus are relatively small on average, this masks a highly unequal distribution of changes across geographic areas. Some counties are completely unaffected by the simulated change, while 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 additionally endow higher-income consumers with the marginal utility of income parameter of 151 percent FPL consumers. In other words, we make 400+ percent 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+ percent FPL income to a county that had no 400+ percent 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-income 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 geographic markets. While some counties 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.

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 effects 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. These objects are also interesting in their own right, as they deliver novel estimates for the economic efficiency and incidence of public spending in ACA Marketplaces—an extensively debated new market that is setting an important policy precedent on providing publicly subsidized non-group individual health insurance in the United States. Having established this baseline, we then examine the economic consequences of subsidy means-testing and market power in isolation. Finally, we put the two together, in a set of counterfactual comparisons that allow us to estimate the extra welfare loss that is generated when a targeted subsidy is used as a policy instrument in an environment with market power.

Efficiency of ACA Marketplaces We begin our analysis by examining the economic costs and benefits of observed subsidy payments. To facilitate our analysis, we first solve for an equilibrium under the observed subsidy mechanism with one change. We remove the second-lowest silver plan incentives from the first-order conditions. This 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 keeps observed subsidy levels, but shuts down the SLSP mechanism, creating a baseline that can be compared to subsidy regimes in subsequent counterfactuals.[36] In practice, our model has a tight in-sample fit, so that the resulting simulated allocation is very close to the data and allows us to quantify the efficiency of the ACA market.

Column (1) of Table 3 reports the results. We estimate that consumer surplus in the ACA Marketplaces amounts to $38 billion. Producer surplus amounts to $3.9 billion (Column 1, row 2). In our simulation of the observed market equilibrium, the government is spending $30 billion in premium and cost-sharing subsidies, which is consistent with the subsidy spending reported by the Congressional Budget Office.[37]

Consumer surplus of $38 billion in total corresponds to $1,881 in annual surplus per capita, on average, among all potential consumers. Consumer surplus levels vary substantially across different socio-demographic groups and different areas of the country, as we already saw in Section 4.4. 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. This variation in surplus levels dovetails with enrollment patterns across income groups; 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

[36] 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. This implies that SLSPs end up being the only plans for which prices change relative to observed prices, and we find that the price changes are minor.

[37] 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)

rates still increase monotonically with age. Overall, the program attracts about 44 percent of potential enrollees (row 17).[38] Figure 4b illustrates the substantial geographic variation: consumer surplus ranges from $621 in Brazos County, Texas to $6,260 in McHenry County, Illinois.

Comparing consumer and producer surplus vis-a-vis government spending, we conclude that public subsidy expenditures generate $1.39 of surplus for $1 of spending: the total outlays on subsidies are $12 billion lower than the sum of consumer and producer surplus. This partial equilibrium efficiency computation, however, is incomplete. As is the case in many healthcare settings, the government is likely to incur costs for the same consumers even if they do not enroll in the ACA. To account for this opportunity cost of public spending, we report in row (5), the amount that the government would have spent on ACA enrollees in uncompensated care payments. 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 are approximately $14 rather than $30 billion. Taking into account this foregone spending, but also accounting for the cost of raising public funds, we estimate a total return of $2.33 on a dollar of public funds spent on the ACA Marketplaces (row 9).

Subsidy incidence in ACA Marketplaces Column 2 of Table 3 allows us to assess the incidence of observed subsidies between consumers and producers. In this column, we simulate an equilibrium without 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, $29 billion in premium subsidies generate $13.7 billion in additional surplus. The nominal within-ACA subsidy spending is partially mitigated by $9 billion of additional public savings, leading to a net increase in government spending of $20 billion. From this perspective, subsidizing premiums generates negative economic value, as this extensive extra spending generates only $13.7 billion in additional surplus, resulting in circa $6 billion of deadweight loss (or 68.5 cents in return to the extra dollar of public spending).

These subsidies, however, do significantly benefit consumers. Enrollment jumps from 19 percent without subsidies to 44 percent with subsidies. This is particularly true of the consumers below 400 percent FPL, who effectively exit the market without subsidy support.

[38]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.

Out of $13.7 billion in additional surplus, 84 percent accrues to consumers and 16 percent is captured by firms, which is a high pass-through for an imperfectly-competitive market.[39] Two margins are at work here: one, the MLR constrains profit margins, and two, targeted subsidies disproportionately attract lower-income consumers with substantially more elastic demand into the market. 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 $3,377, the average consumer surplus increases from $1,310 to $1,881—a substantially lower increase in surplus as compared to government expenditures.[40]

The incidence of subsidies on consumer surplus is 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 equilibrium 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. In these states, average consumer surplus increases by more than 100 percent in many counties when means-tested subsidies are introduced.

Isolating the efficiency cost of market power Before turning to the discussion of how market power interacts with subsidy design, we first estimate the 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 re-allocated from insurers to consumers, with consumer surplus increasing by nearly $10 billion (by ca. $460 on average) and consumer enrollment increasing from 44 to 55 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 20 percent of per capita

[39]These estimates suggest that the ACA market is on average more competitive—or, more likely, is effectively restricted by price caps imposed by the Minimum Loss Ratio—than Medicare Advantage, for which Cabral et al. (2018) find 45 (in premiums only) to 54 (in premiums and benefits) percent pass-through rate on average, and a range from 13 to 74 percent between least and most competitive markets.

[40]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 Lurie et al. (2019), and especially Finkelstein et al. (2019), who speculate about the role of uncompensated care and behavioral biases in accounting for the low revealed willingness to pay.

surplus due to the presence of market power in this market.

Isolating the efficiency cost of subsidy targeting To further understand how subsidy targeting per se affects this market, we shut down market power and compare equilibria under two different subsidy regimes. One with targeted subsidies as observed in the data, and one with age-adjusted, but income-invariant (and hence universal) subsidies. In the second regime, we pick a voucher level such that the total government spending remains the same as under the regime with targeted subsidies. Columns (4) and (6) of Table 3 report the characteristics of the respective allocations. We find that a voucher of $1,005 for a 20 year old consumer that scales with age according to the statutory age curve leads to the same amount of total government spending as under targeted subsidies.

Comparing these two allocations allows us to ask how much more (or less) consumer surplus would be achieved 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 delivers $6.6 billion more in total consumer surplus. This substantial gain in efficiency, however, comes at a large re-distributional cost. While overall enrollment in the market increases, enrollment shifts from the poorest to less poor consumers under universal subsidies. That is, means-tested subsidies work as intended—they lower prices for the lowest-income consumers and thus attract them to the market. When we keep government spending fixed, however, this means-testing is not efficient, since lowest-income consumers exhibit the lowest valuation of the good. In other words, the marginal consumers attracted by increasingly generous subsidies have an increasingly-declining willingness-to-pay for insurance, meaning that government spending grows faster than consumer surplus.

More generally, this exercise highlights the central role that subsidies play in ensuring high enrollment rates into the program. Without any subsidies and with marginal cost pricing, we simulate (in Column 5) that only 27 percent of potential consumers would enroll in this market (69 percent of these consumers coming from relatively higher-income households). With either means-tested subsidies or universal subsidies, enrollment increases to 55 or 60 percent, respectively, but comes from different points in the income distribution. Under means-tested subsidies more than half of the consumers are enrolling from households with income under 250 percent FPL, while under universal subsidies this share is less than 40 percent.

Interaction between targeting and market power We next consider the central question of the paper, which is whether the presence of market power can introduce additional distortions when subsidies to consumers are targeted. The thought experiment is the following. Consider a policy-maker who is deciding whether to allocate some amount of public spending through universal subsidies or whether to set means-tested subsidies that disproportionately benefit low income consumers. Without considering the effects of market power, the policy-maker will compare total welfare under the two subsidy regimes and conclude, as we saw above, that means-tested subsidies generate a welfare loss, but possibly have desirable re-distributional properties. Weighing the equity gains vis-a-vis efficiency losses, the policy-maker may be willing to accept efficiency losses. We argue that by not taking into account the presence of market power, the policy-maker is likely to significantly underestimate the extent of efficiency losses from means-tested (versus universal) subsidies, leading to a misguided policy choice. Our goal here is to measure whether the amount of "overlooked" efficiency loss that results from the interaction between subsidy targeting and market power can be quantitatively important.

We isolate this interaction effect through the lens of two comparative exercises. We first simulate the decision problem of the "naive" policy-maker, comparing the efficiency of two subsidy regimes (means-tested versus universal) under marginal cost pricing (Columns (4) versus (6) of Table 3). We compute how much efficiency is lost from means-testing in that scenario. We then simulate the decision problem of a "sophisticated" policy-maker, doing the same comparison, but allowing for the market power to exist in both subsidy regimes, and computing how much welfare loss is lost in that case (Columns (1) versus (3) of Table 3). Within each of the pairwise comparison, we focus on universal subsidies that keep the amount of public spending fixed in equilibrium. As the level of public spending differs across the two pairs, we focus on relative rather than absolute welfare loss as the key comparison metric.

Figure 6 summarizes our findings. The light grey curve illustrates the "naive" comparison case. Recall that we defined total welfare as the sum of consumer and producer surplus net of extra government spending (sans the savings on uncompensated care) adjusted for the cost of raising public funds. The light grey curve traces out the level of total welfare in the market in the absence of market power, under different levels of age-adjusted, but income-invariant, subsidies. The x-axis indexes subsidies using the voucher level for a 20-year old consumer as the baseline. We note that total welfare is concave with respect to the voucher level. Highest welfare is achieved at the baseline subsidy levels of $500, and starts falling at higher subsidy levels. The level of welfare under means-tested subsidies (at the level observed in the data), without market power, is marked with a light grey diamond. One the x-axis, the grey diamond is located at the baseline subsidy level of $1,005. As we see in Column (3) of Table 3, this is the level of universal subsidies that results in the same amount of public spending as under means-tested subsidies. We estimate that in this "naive" decision-making scenario—which ignores the presence of market power—the policy-maker would conclude that means-tested subsidies lead to a 24 percent efficiency loss.

We repeat the same analysis for the "sophisticated" decision-making situation that takes into account the presence of market power in this market and anticipates firms' responses to the subsidy regime. The black curve in Figure 6 now traces out the total welfare under different levels of (baseline) universal subsidy, while the black diamond marks that amount of welfare that is achieved under means-tested subsidies. The black diamond is located at the baseline subsidy level of $1,105, at which point we observe the same amount of public spending across the two subsidy regimes. we observe that in this case means-tested subsidies generates a loss of 32 percent in total surplus. We conclude that market power significantly exacerbates the efficiency losses from means-testing, leading to a 33 percent higher loss in welfare.

It is important to emphasize, however, that the allocation with income-invariant subsidies is not Pareto improving over the targeted subsidy mechanism, in either the "naive" or "sophisticated" worlds. In both cases, flat subsidies generate a re-allocation of surplus from highly subsidized consumers to less subsidized consumers. This is true both under perfect and imperfect competition, with the effects slightly 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.

The take-away from these two pairwise comparisons is that the "naive" policy analysis that does not consider the potential for interaction between subsidy design and actions of strategic firms with market power, may slightly underestimate the distributional gains, but also significantly underestimate, by a third, the welfare losses from means-testing.

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 in 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 vis-a-vis universal subsidies. 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 the evaluation of the equity-efficiency tradeoff in publicly-funded, but 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*, 2015, *105(5)*, 187–191.

Basurto, Pia M, Pascaline Dupas, and Jonathan Robinson, "Decentralization and Efficiency of Subsidy Targeting: Evidence from Chiefs in Rural Malawi," Working Paper 23383, National Bureau of Economic Research May 2017.

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.

Cabral, Marika, Michael Geruso, and Neale Mahoney, "Do Larger Health Insurance Subsidies Benefit Patients or Producers? Evidence from Medicare Advantage," *American Economic Review*, August 2018, *108* (8), 2048–87.

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.

Collinson, Robert and Peter Ganong, "How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality?," *American Economic Journal: Economic Policy*, May 2018, *10* (2), 62–89.

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

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," *NBER Working Paper No. 24668*, 2018.

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.

Drake, Coleman, "What are consumers willing to pay for a broad network health plan?: Evidence from covered California," *Journal of Health Economics*, 2019, *65*, 63 – 77.

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

Ericson, Keith M. Marzilli and Amanda Starc, "How product standardization affects choice: Evidence from the Massachusetts Health Insurance Exchange," *Journal of Health Economics*, 2016, *50*, 71 – 85.

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, Washington University in St. Louis*, 2019.

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

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

Geruso, Michael, 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, Stanford University, New York University, UCSD*, 2018.

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

_ **and Benjamin D. Sommers**, "The Affordable Care Act's Effects On Patients, Providers, And The Economy: What We've Learned So Far," *Journal of Policy Analysis and Management*, 2019, *38* (4), 1028–1052.

_ **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" in "Valuation of New Goods under Perfect and Imperfect Competition, eds. Bresnahan, T. and Gordon, R.," *University of Chicago Press*, 1996, pp. 207–248.

Ho, Kate and Robin Lee, "Health Insurance Menu Design: Managing the Spending-Coverage Tradeoff," *Princeton University and Harvard University, presentation slides*, 2019.

Hoynes, Hilary and Diane Whitmore Schanzenbach, *US Food and Nutrition Programs*, in Economics of Means-Tested Transfer Programs in the United States, Volume 1 by University of Chicago Press, October, 2015

Jacob, Brian A. and Jens Ludwig, "The Effects of Housing Assistance on Labor Supply: Evidence from a Voucher Lottery," *American Economic Review*, February 2012, *102* (1), 272–304.

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," *Brookings Papers on Economic Activity*, 2014, pp. 277–333.

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.

Lieber, Ethan M. J. and Lee M. Lockwood, "Targeting with In-Kind Transfers: Evidence from Medicaid Home Care," *American Economic Review*, April 2019, *109* (4), 1461–85.

Lurie, Ithai, Daniel J Sacks, and Bradley Heim, "Does the individual mandate affect insurance coverage? Evidence from tax returns," *Unpublished mimeo, Office of Tax Analysis and Indiana University*, 2019.

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?*, in Individual and Social Responsibility: Child Care, Education, Medical Care, and Long-Term Care in America by University of Chicago Press, January, 1996

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.

Saltzman, Evan, "Demand for health insurance: Evidence from the California and Washington ACA exchanges," *Journal of Health Economics*, 2019, *63*, 197 – 222.

Saltzman, Evan A., Christine Eibner, and Alain C. Enthoven, "Improving The Affordable Care Act: An Assessment Of Policy Options For Providing Subsidies," *Health Affairs*, 2015, *34* (12), 2095–2103.

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.

Taylor, Erin A., Evan Saltzman, Sebastian Bauhoff, Rosalie L. Pacula, and Christine Eibner, "More Choice In Health Insurance Marketplaces May Reduce The Value Of The Subsidies Available To Low-Income Enrollees," *Health Affairs*, 2015, *34* (1), 104–110.

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

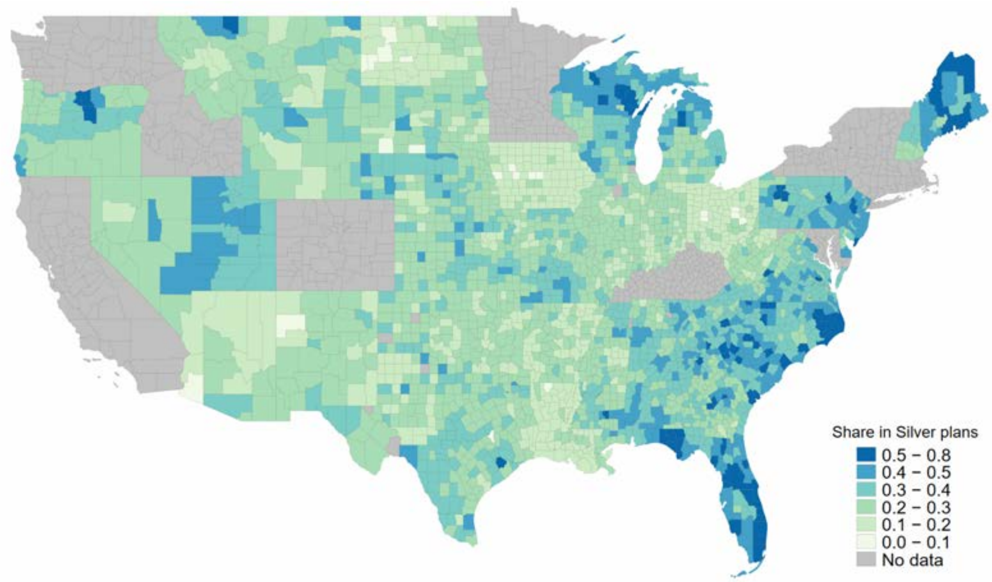

(b) Model in-sample prediction error of Silver plan market share

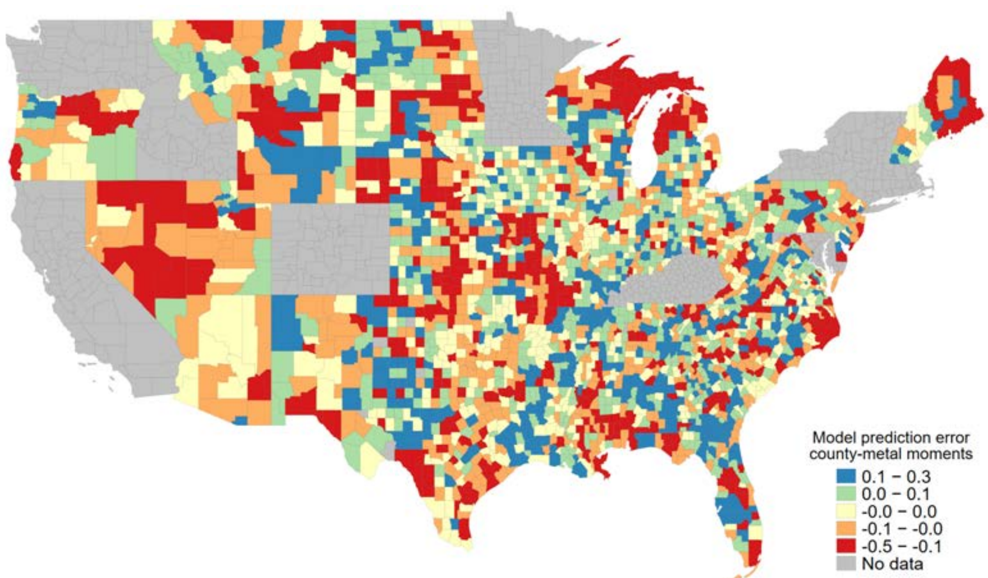

Notes: Map in panel (a) plots the share of potential consumers in each county that enrolled in a Silver plan on ACA Marketplaces. States that are marked with grey are not federally facilitated and do not enter our analysis. The counts of the pool of potential consumers (denominator) was provided by the 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

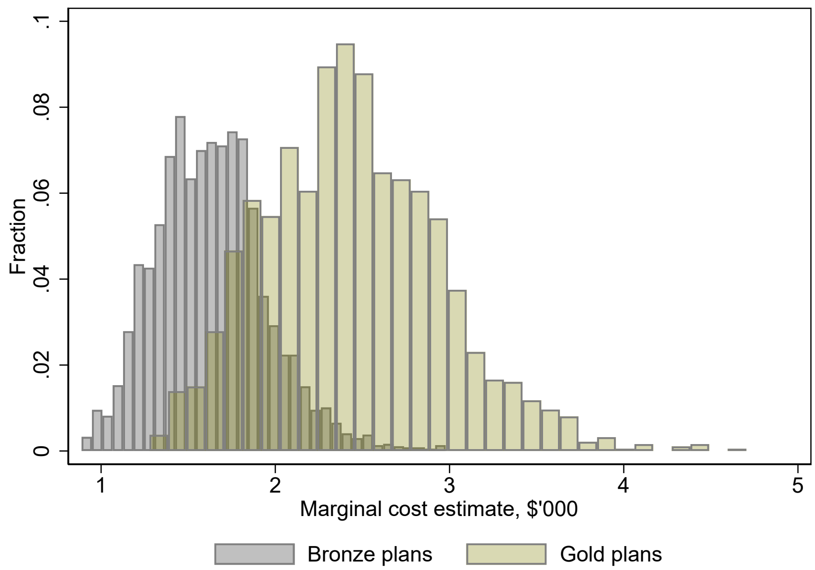

(b) Marginal cost estimates and observed accounting cost

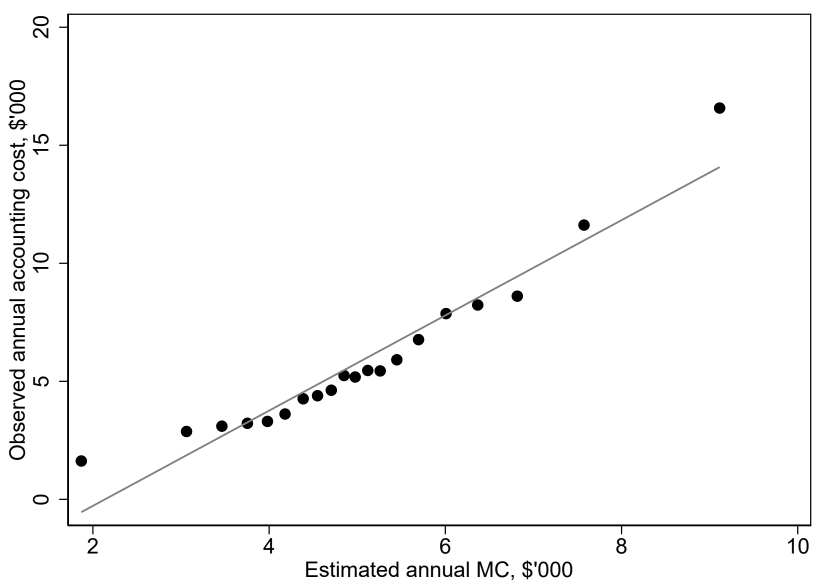

Notes: Panel (a) plots the distribution of plan-market level marginal costs as estimated in Section 4.2. The costs are plotted for a baseline, age 20, consumer. The costs are plotted separately for Bronze plans that provide the lowest amount of coverage and Gold plans that provide the highest amount 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

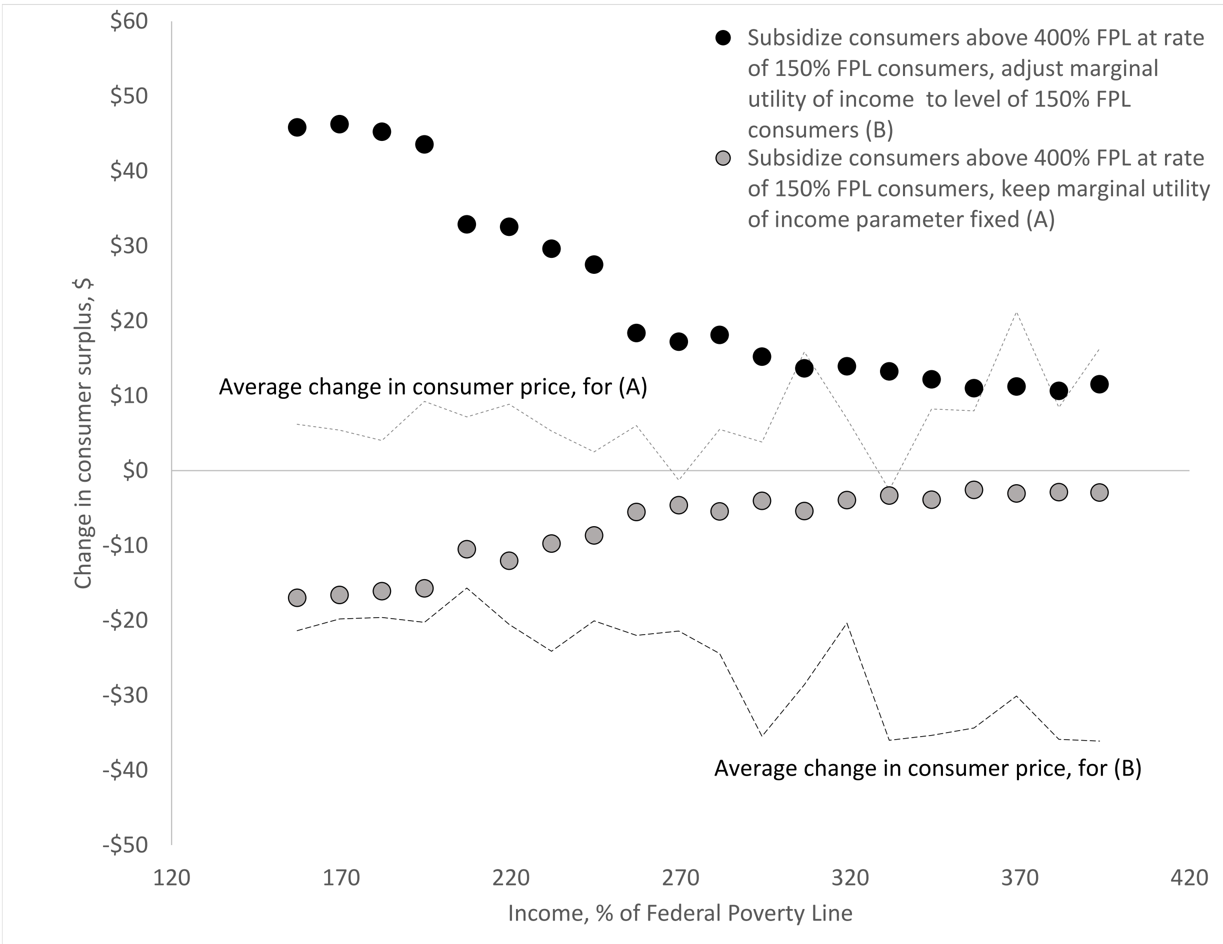

Notes: Figure reports estimated change in consumer surplus and consumer-facing prices (y-axis, in dollars) by income level (x-axis) in two counterfactual cases that capture the "demographic externality". The first counterfactual simulation (case A) changes income of consumers with true income of above 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

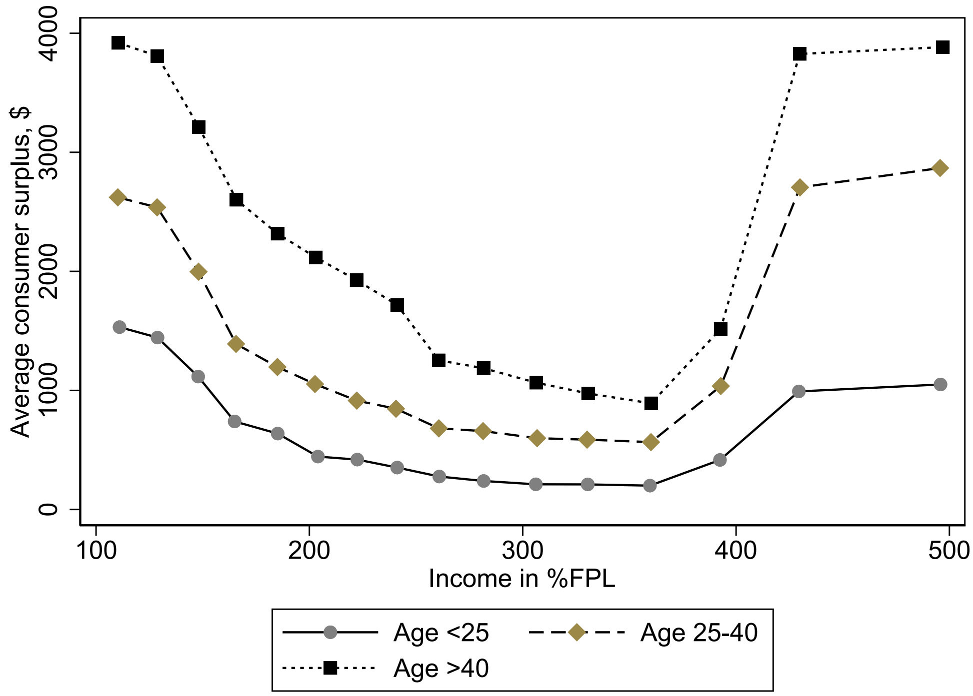

(b) Geographic variation in consumer surplus

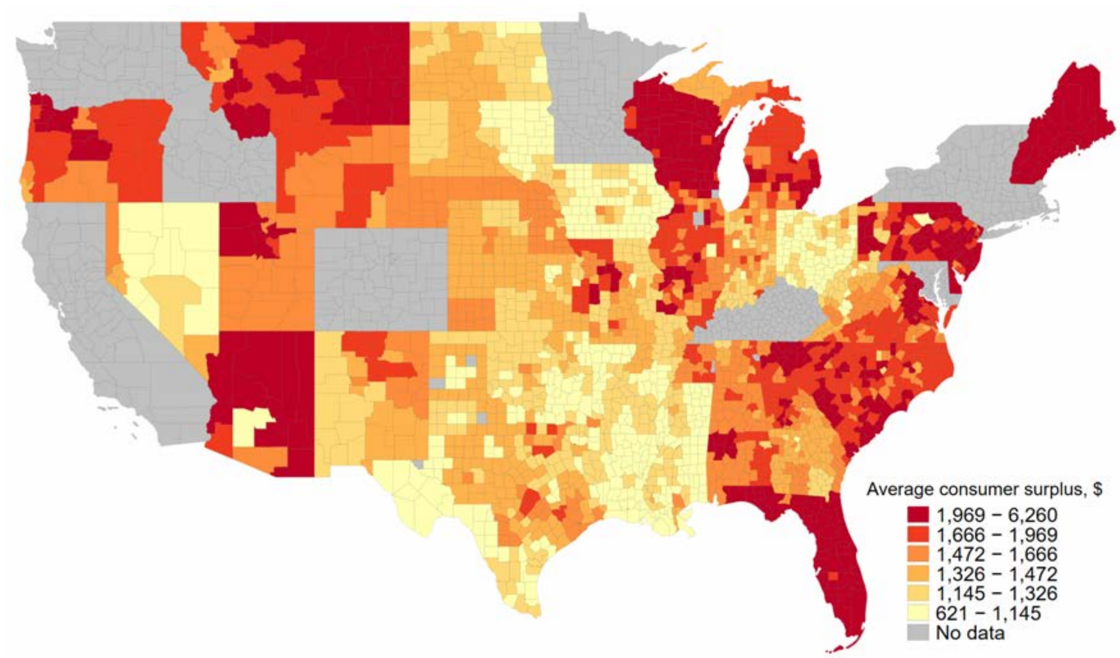

Notes: Panel (a) plots the average consumer surplus by income and age groups, as estimated under the observed subsidy regime. Map in Panel (b) plots average consumer surplus per county, also as estimated under the observed subsidy regime.

Figure 5: Geographic incidence of subsidies under imperfect competition

(a) Means-tested subsidies relative to no subsidies

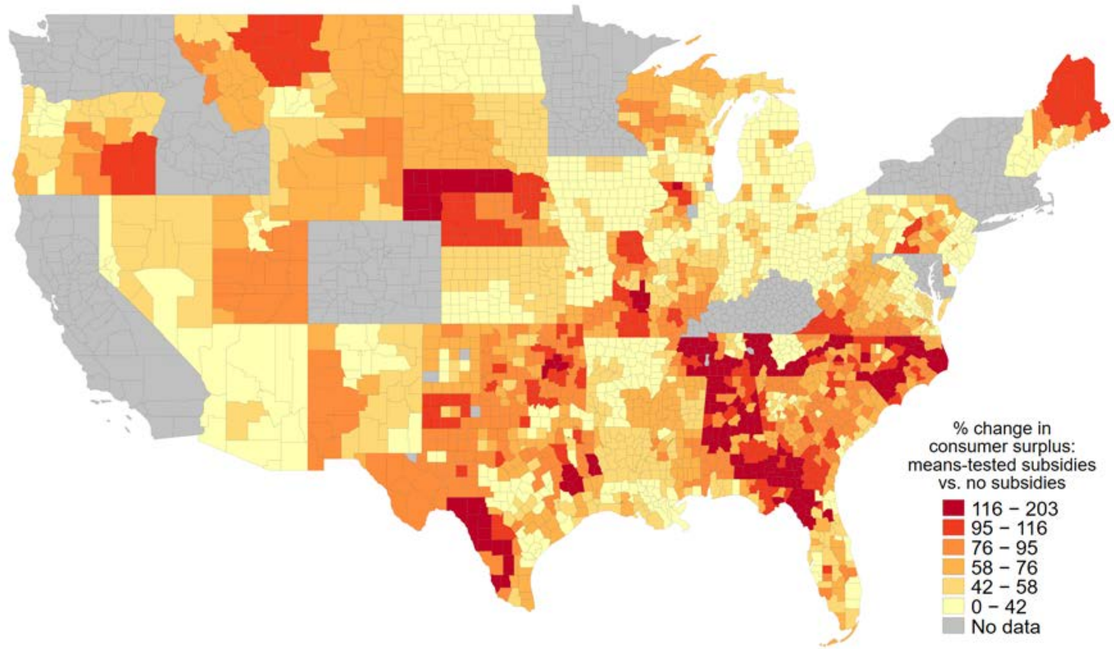

(b) Means-tested subsidies relative to flat subsidies

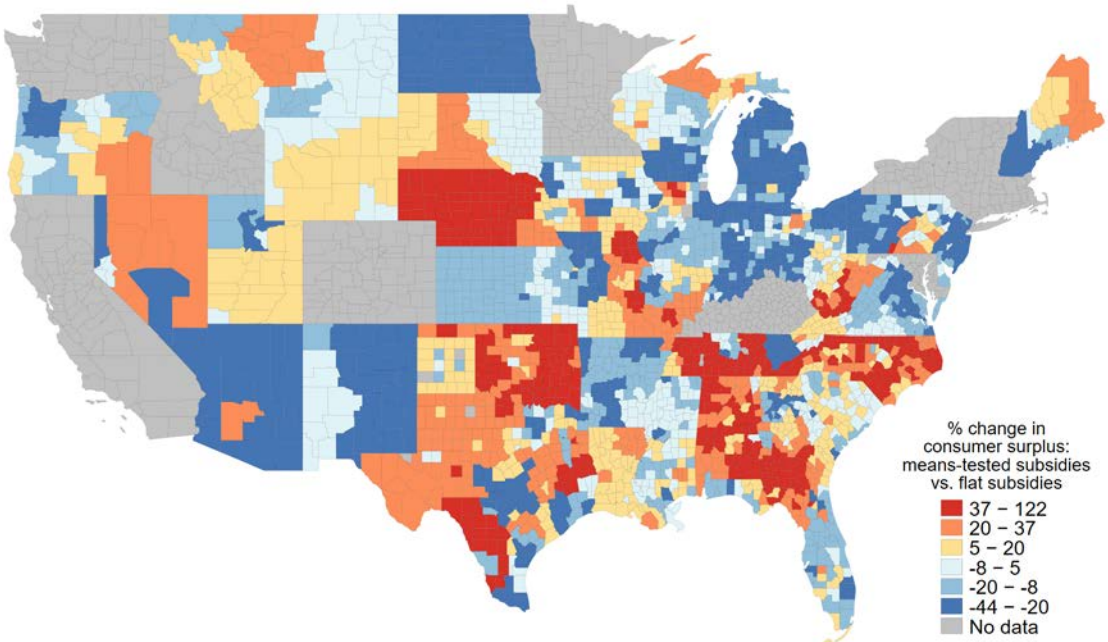

Notes: Panel (a) reports the percent change in average consumer surplus that results when we move consumers from an environment with no subsidies to means-tested subsidies at the level observed in the data. Insurers are allowed to retain market power and can adjust prices. Panel (b) reports the percent 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

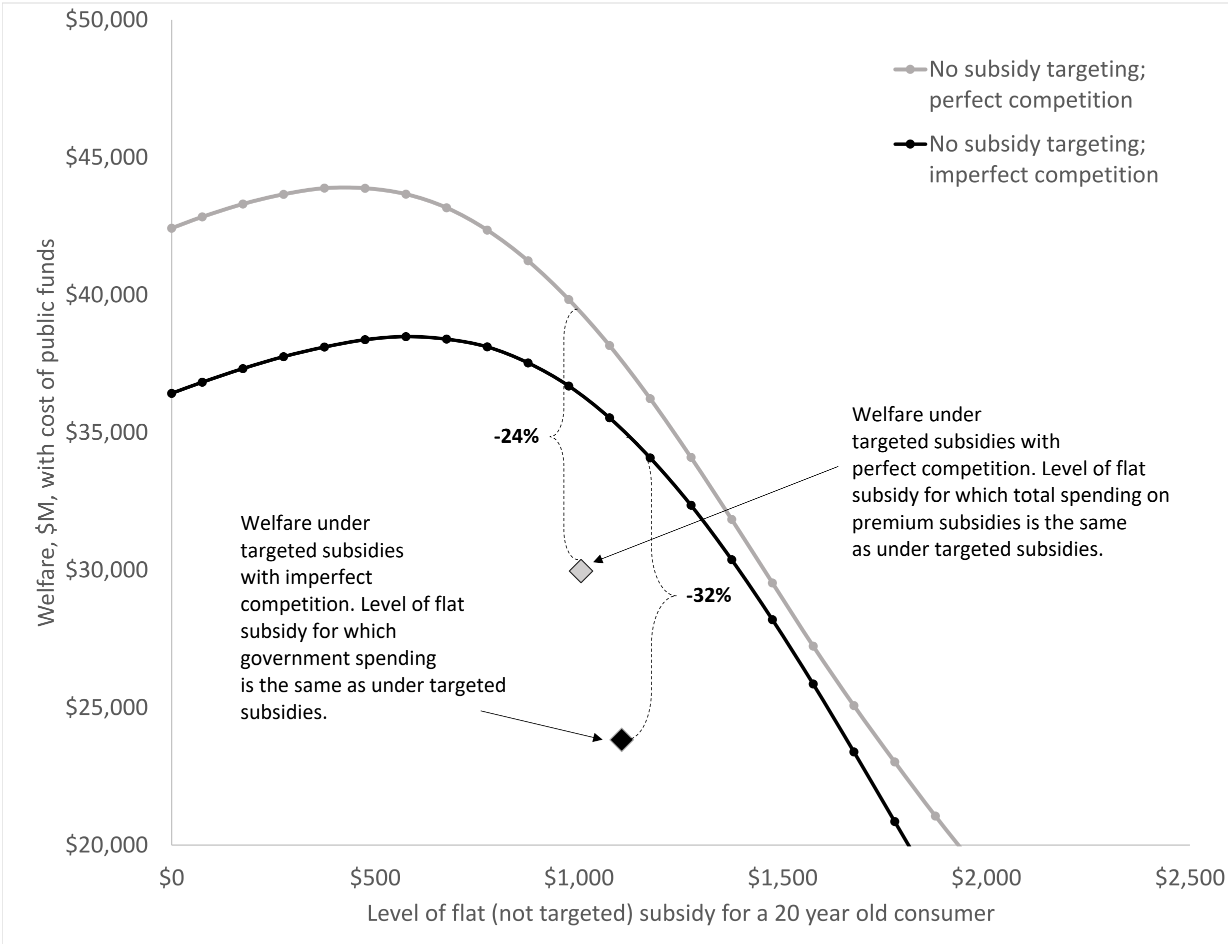

Notes: Figure reports estimated total welfare (y-axis), including the opportunity cost of government spending and the cost of raising public funds, under subsidies that do not vary with consumer income *within* an age group. The x-axis marks the value of the flat voucher that is offered to a 20-year 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

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

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

Table 3: 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)	Public 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	8,896	3,909	10,275	11,056	5,437	12,173	10,326	11,147
(12)	Inside option enrollment, percent of total market	44.2	19.4	51.0	54.9	27.0	60.4	51.3	55.3
(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, $	1,083	2,315	1,129	820	1,874	813	782	507
(29)	Average subsidy paid, $	3,377	186	2,935	3,021	253	2,741	4,294	4,782
(30)	Average consumer surplus, $	1,881	1,310	2,265	2,340	1,554	2,672	2,458	1,532

^ 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 (first by county, and then to total using county market size as weights). 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 consumers 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 per capita, 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.