

Subsidy Targeting with Market Power

Maria Polyakova* Stephen P. Ryan†

December 23, 2020

Abstract

Public programs frequently use observable characteristics of recipients, such as income, to target benefits. We show theoretically that when the provision of the subsidized good is decentralized to intermediaries with market power, targeting of subsidies induces a “demographic externality” that can distort the incidence and efficiency of public transfers. We examine this possibility empirically in the context of means-tested subsidies for privately-provided health insurance under the Affordable Care Act (ACA). We estimate that the overall incidence of premium subsidies on consumers in ACA Marketplaces is less than 50 percent, and a third of net government spending on premium subsidies is a deadweight loss. Market power in the presence of means-tested subsidies leads to regressive redistribution, lowering consumer surplus and rates of insurance in the poorer population targeted by subsidies. Under sufficiently high social preferences for redistribution, however, means-tested subsidies still dominate income-invariant transfers.

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

*Center for Health Policy, School of Medicine, Stanford University, Palo Alto, CA, 94305; NBER; and CESifo (maria.polyakova@stanford.edu).

†Corresponding Author: Olin Business School, Washington University in St. Louis, St. Louis, MO 63104; NBER; and CESifo (stephen.p.ryan@wustl.edu, 314-935-8221 (w), 314-935-6359 (f)).

1 Introduction

Many public programs, such as housing assistance, food stamps, health insurance, free school lunch, and energy assistance, are in-kind government transfers targeted to recipients based on observable characteristics such as income, age, employment, family, or disability status (Akerlof, 1978; Nichols and Zeckhauser, 1982; Smeeding, 1984; Currie and Gahvari, 2008; Hoynes and Schanzenbach, 2015). The literature studying the economic costs and benefits of such transfers has almost exclusively focused on potential distortions in recipients’ behavior, while assuming that the subsidized goods or services are provided by a benevolent government or perfectly competitive firms (recent examples include Jacob and Ludwig, 2012; Allcott et al., 2015; Lieber and Lockwood, 2019; Basurto et al., 2017; Hendren and Sprung-Keyser, 2020). In practice, however, governments are increasingly relying on imperfectly competitive intermediaries to deliver subsidized benefits.¹ The fiscal magnitude of this shift is staggering. In health insurance alone, for example, the US government spends \$0.6 trillion annually in subsidies for products provided by private firms (CBO, 2019a,b,c), with an expansion projected under the Biden administration for 2021 onward.² Yet, despite the importance of understanding the economics of such environments for policy-making, the existing literature provides little guidance on how the presence of profit-maximizing intermediaries may affect the productive efficiency and the allocative properties of targeted in-kind transfers.

In this paper, we seek to close this gap by extending the literature along two dimensions. First, we combine insights from both public economics and industrial organization to outline a set of fundamental economic forces that characterize the equilibrium outcomes of subsidy targeting under market power. Second, we apply these insights to the subsidized health insurance market that was launched under the Affordable Care Act in 2014 (hereafter: “ACA Marketplaces”), an important empirical setting that provides insurance to nearly ten million people in the United States and costs the federal government more than \$40 billion in annual outlays. Conceptually, we show that in addition to the standard imperfect pass-through of subsidies to consumers, a subsidy-induced “demographic externality” arises. Targeted transfers change the relative importance of different consumers in the firm’s profit-maximization problem and the equilibrium price a given consumer faces becomes a function

¹Private provision is frequently motivated by the desire to improve program efficiency through competition as well as to reduce fiscal uncertainty for federal and state program budgets.

²Contracting with private firms to provide public benefits is particularly common in health care. 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.

of her neighbors' eligibility for transfers.³ This externality can distort both which consumers benefit from subsidies and aggregate efficiency. In our ACA setting, we find that only 50% of surplus generated by subsidies is passed through to consumers, while the rest is captured by firms. Critically, we find that the interaction of market power with targeting generates regressive distributional effects in ACA markets, leading to disproportionately lower consumer surplus and insurance rates among lower income and older consumers relative to an environment with perfect competition, working counter to the policy-makers' targeting goals.

The subsidy-driven interdependency of consumers belongs to a broader class of mechanisms where equilibrium outcomes are a function of the distribution of heterogeneous consumers in the same market. An extensive literature has studied several distinct types of such mechanisms. The classic notion of adverse or advantageous selection, for example, is a *cost-driven* demographic externality in which consumers are linked through risk pooling (Rothschild and Stiglitz, 1976). Tebaldi (2017), which is the closest work to our paper, both methodologically and in terms of the empirical context, examines such cost-based externality in the California ACA market. He estimates efficiency distortions that arise because insurers pool risks across consumers of different ages that have different costs, but are constrained in how much they can price-discriminate based on age. *Demand-driven* demographic externalities can arise due to the pooling of consumers with heterogeneous preferences. For example, features of differentiated products available to a consumer may depend on the preferences of the consumer's neighbors (George and Waldfogel, 2003; Waldfogel, 2003). In general, when price discrimination is prohibited, the price a given group pays is a function of the distribution of other demand types (Stole, 2007). In markets with selection, differences in consumer demand types may correlate with differences in costs, combining preference-based and cost-based spillovers (Finkelstein and McGarry, 2006; Shepard, 2016). In this paper, we develop a modeling strategy that encompasses these disparate effects in a unified framework. This allows us to assess the equilibrium interactions of the cost, demand, and subsidy spillover channels with market power.

The aggregate direction of distortions attributable to the interaction of market power and targeted transfers is theoretically ambiguous and is thus an empirical question in any given setting. In this paper, we ask this empirical question in the context of a publicly subsidized health insurance market. We consider ACA Marketplaces facilitated through the federal online platform www.healthcare.gov that operates in 36 states. The ACA Marketplaces

³The externality does not arise if firms are able to perfectly price-discriminate and set a different price for each group of differentially targeted consumers. Such perfect price discrimination is rarely possible in practice. In our empirical context, for instance, insurers are regulated to set one price per market.

provide a fruitful empirical laboratory for studying the impact of market power on targeted transfers for several reasons. First, it is an important market with nearly 10 million enrolled consumers that is expected to grow in the future. Second, targeted transfers are a first-order feature of the market: the vast majority (85%) of consumers receive a means-tested tax credits to help defray premium costs. The generosity of the tax credit is a declining function of a consumer’s income, reaching zero for higher-income consumers. Third, the significant exercise of market power in this environment is likely, as many geographic markets are highly concentrated. Fourth, firm pricing is regulated and premiums cannot vary by income. Given these institutional features, we anticipate the interaction between means-tested public transfers and market power in this setting to be quantitatively important.

To quantify this interaction, we formulate and estimate a structural model of supply and demand for ACA Marketplace insurance plans. On the demand side, we utilize the institutional setting of the Marketplaces to implement a within-market identification strategy that leverages price variation generated across consumers by the regulatory design of subsidies. On the supply-side, we derive equilibrium first-order conditions that link prices and marginal costs, incorporating the demographic externality, risk-sorting, and regulatory constraints on prices.

With the estimates of demand and supply in hand, we perform several counterfactual simulations of the market equilibrium that allow us to measure the impact of market power on subsidies. We start by simulating a counterfactual environment without subsidies, holding all else equal. This allows us to measure the aggregate efficiency of existing subsidies and their incidence between consumers and producers. We report three major findings: subsidies make consumers and firms better off by an additional \$14.8 billion; firms capture \$7.9 billion of that surplus, or 53 percent; and the program generates a deadweight loss of \$5.7 billion, as it requires \$20.5 billion in additional (net) government expenditures.⁴

The distribution of surplus from targeted subsidies is highly heterogeneous across consumer types. While all subsidized consumers are better off with subsidies than without, the lowest-income consumers experience up to a four-fold higher increase in consumer surplus than partially-subsidized consumers, while unsubsidized consumers are modestly worse-off with subsidies than without. A similar pattern characterizes the effect of subsidies on the rate of insurance coverage. Subsidies increase insurance take-up among the lowest income consumers by up to 80 percentage points, while reducing coverage rates slightly among un-

⁴We compute government expenditures on subsidies net of savings on publicly paid uncompensated care. Further, if raising extra \$20.5 billion cost taxpayers 30 cents on a dollar, then the deadweight loss of the program accounting for the cost of public funds would be \$11.9 billion.

subsidized consumers. Dovetailing with our theoretical framework, subsidies exert countervailing forces on the equilibrium prices in this market, simultaneously attracting inherently more elastic, lower-income consumers with higher marginal costs while also directly lowering demand elasticity.

Next, we turn to the question of how much market power exerted by firms on ACA Marketplaces affects the distributional outcomes of means-tested transfers. As this is a central question of the paper, we pursue two complementary measurement approaches. First, we compare the level of consumer surplus and insurance rates among different consumer types under the observed regime with subsidies and imperfect competition and a counterfactual regime that holds subsidies fixed but removes market power. We find that in the presence of subsidies market power reduces overall consumer surplus by 21 percent and aggregate rate of insurance by nearly 15 percentage points (from 59 to 45 percent insured). The largest relative decline in consumer surplus from market power occurs among the poorest consumers: we estimate a decline in consumer surplus of 18 percent on average among unsubsidized consumers with income over 400% FPL and a decline of 46 percent among consumers with incomes between 150% and 200% FPL. Similarly, while consumers with incomes over 400% FPL experience a ten percentage point drop in insurance rate, consumers with incomes between 200% to 300% FPL see nearly a 25 percentage drop in the coverage rate.

Our second approach measures the difference in gains from subsidies with and without market power. We compare two scenarios. The first scenario assumes that means-tested subsidies are rolled out into a perfectly competitive market. We measure gains in consumer surplus and insurance coverage rates by consumer type. This variation in gains across consumer types reveals the distributional effects of transfers under perfect competition.⁵ We then contrast this to the gains in consumer surplus and insurance rates in the scenario where means-tested transfers are rolled out into an imperfectly competitive market. Comparing the gains from means-tested subsidies (relative to no subsidies) by different consumer types between the competitive and imperfectly competitive scenarios allows us to assess how much market market power distorts the distributional effects of subsidies. We find that market power lowers the average gain from subsidies across all consumers by \$270 per capita (from \$613 to \$343), or 44 percent, while disproportionately harming the lowest-income and oldest consumers. The gains in consumer surplus from subsidies are on average ca. \$1,000 (or 46 percent) lower among consumers age 60 or older with income under 300% FPL when markets

⁵If all consumer types have equal marginal costs, as is the case in most product markets, this incidence will be equal to the schedule of subsidies, subject to a zero-price lower bound. With risk sorting, even under perfect competition subsidies may change insurer costs and affect unsubsidized consumers.

are imperfectly competitive. Thus market power redistributes the marginal subsidy dollar away from the intended means-tested schedule.

The take-away from our results is that when imperfectly competitive intermediaries provide a subsidized good, the market power of these intermediaries can interfere with distributional objectives of the policy-maker. These results cast doubt on the effectiveness of the common approach of incorporating distributional policy instruments into environments where a publicly subsidized good is privately provided.

A natural question is whether there are alternative subsidy mechanisms that still provide incentives for consumers to buy insurance, but reduce efficiency distortions from market power. We consider this question in the last section of the paper by examining one commonly-proposed alternative mechanism: “flat” subsidies that remove means-testing completely but keep net government spending the same. We find that flat subsidies lead to substantially higher subsidy pass-through to consumers and reduce the deadweight loss. However, if the society has preferences for redistribution, then the choice between means-testing versus flat subsidies generates a stark equity-efficiency tradeoff. In the absence of preferences for redistribution, means-testing is a strictly dominated mechanism in this market. As we estimate higher marginal utility of income among lower-income households, our model implies that these households have low willingness to pay for health insurance, preferring cash transfers to in-kind subsidies for insurance.⁶ This results does not hold, however, if we allow for a welfare function that puts value on redistribution per se. We use the [Atkinson \(1970\)](#) welfare function with constant relative inequality aversion, which assigns higher welfare weights to consumers with lower income, to illustrate this point. We estimate that from the consumer surplus perspective, means-testing becomes preferred to flat vouchers once the preference for redistribution is relatively strong, with the inequality aversion parameter higher than 1.3, implying that the society values transferring \$1 to a household with income of \$17,820 (150% FPL in 2017 for a single person household) as much as transferring to \$3.6 to a household with income of \$47,520 (400% FPL in 2017 for a single person household).

Our work relates closely to the growing literature on the design of subsidies for health insurance in general ([Curto et al., 2015](#); [Jaffe and Shepard, 2018](#); [Decarolis, 2015](#); [Decarolis et al., 2020](#); [Miller et al., 2019](#); [Einav et al., 2019](#)) and in the ACA Marketplaces more specifically ([Aizawa, 2019](#); [Aizawa and Fang, 2020](#); [Aizawa and Fu, 2020](#); [Tebaldi, 2017](#);

⁶This finding is not unique to ACA Marketplaces, or even health insurance more generally, and has been documented in other health insurance settings ([Finkelstein et al., 2019](#)), as well as other markets such as subsidized housing ([Rosen, 1985](#)).

Tebaldi et al., 2019).⁷ Our results also speak to the optimal design of rating areas (a regulatory grouping of markets where prices must be set equally for a given plan), which has been investigated in Dickstein et al. (2015). Alternative rating areas would lead to different pooling of demographics and would thus generate different equilibrium outcomes under the mechanism that we investigate in this paper.

Outside of health insurance, our paper is conceptually similar to a set of recent papers that have examined the strategic motives of private firms delivering goods that are publicly subsidized for *some* consumers but not others, across a variety of domains. For example, Cellini and Goldin (2014) and Fillmore (2019) have examined the relationship between federal grants that only some students are eligible for and college tuition; Rothstein (2010) examined how firms set wages in the presence of the Earned-Income Tax Credit; Goldin et al. (2018); Meckel (2019); Meckel et al. (2020) consider the effects of food assistance programs on the pricing and entry of grocery stores; and Rosen (1985); Eriksen and Ross (2015); Collinson and Ganong (2018); Waldinger (2018) examine the efficiency of housing vouchers and how they may change housing prices. We seek to generalize the insights from this literature into a framework that is likely to be applicable in a variety of settings, including health insurance that accounts for by far the largest share of such transfers in the US.

The paper proceeds as follows. Section 2 discusses the theoretical framework. Section 3 gives a brief primer on the ACA Marketplaces and describes our data sources. Sections 4.1, 4.2, and 4.3 lay out the empirical models of demand and supply, and describe how we measure welfare. 4.4 reports estimation results. Section 5 then proceeds to simulate counterfactual equilibria to measure the efficiency and incidence of observed transfers and the role of market power. Section 6 briefly concludes.

2 Theory

We are interested in the interaction between targeted transfers and market power in general economic settings. Consider a market that is populated with consumers that have a distribution of types, D , which encapsulates salient consumer characteristics such as endowments,

⁷More broadly, our work also relates to the literature that has considered the incidence of tax exclusions that are effectively subsidies for employer-sponsored health insurance plans (e.g., Gruber and Washington, 2005; Gruber, 2005; Gruber and Poterba, 1996) as well as the design of *employer* rather than government subsidies for health insurance (Cutler and Reber, 1998; Ho and Lee, 2019). Further, 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.

preferences, or risk. On the supply side, firms offer a set of products to consumers and compete under a set of market rules, r , set by the policymaker. Such rules might include restrictions on the minimum set and types of products firms must offer or limits to price discrimination. Let $\theta \in \Theta$ represent the current market structure out of the set of all possible market structures; we normalize θ_0 to represent perfect competition. The government sets a subsidy schedule, $Z(D)$, which gives targeted transfers to consumers based on demographics for the sole purpose of purchasing a product in this market.

Within this environment, define the distribution of an equilibrium outcome—such as a social welfare function, some measure of the efficiency of government spending, or subsidy incidence—as:

$$H(Z(D), \theta, D, r), \tag{1}$$

which takes the subsidy schedule, market structure, demographic distribution, and market rules as arguments.

Various strands of the literature have examined components of Equation 1, but often in isolation from each other. For example, the extensive literature investigating the equilibrium effects of asymmetric information is often concerned with assessing the role that D plays in Equation 1. The vast majority of this literature has assumed that $\theta = \theta_0$. For example, [Rothschild and Stiglitz \(1976\)](#) and [Azevedo and Gottlieb \(2017\)](#) examine conditions under which an equilibrium with adverse selection will exist while assuming perfectly competitive insurance markets. [Hendren et al. \(2020\)](#) analyze the value of choice in insurance settings with adverse selection and moral hazard and derive optimal Pigouvian subsidies under the assumption of no market power. However, as highlighted by [Mahoney and Weyl \(2017\)](#), failing to account for the interaction of selection and market power can potentially lead to misleading policy conclusions.⁸

In contrast, the literature on optimal subsidy design has primarily focused on the role of $Z(D)$. A central object of interest in this literature is the effect on equilibrium outcomes of introducing a subsidy:

$$H(Z(D), \theta, D, r) - H(0, \theta, D, r), \tag{2}$$

where H would be, for example, the distribution of surplus across firms and consumers. The vast majority of this literature examines subsidy design in the environment of perfect competition or a benevolent government. A notable exception is a series of papers that have considered the optimal subsidy design problem in the context of health insurance markets

⁸A small but growing empirical literature has started examining the role of market power in markets with selection (e.g. [Crawford et al., 2018](#)).

(e.g., [Curto et al., 2015](#); [Jaffe and Shepard, 2018](#); [Decarolis, 2015](#); [Tebaldi, 2017](#); [Decarolis et al., 2020](#); [Miller et al., 2019](#); [Einav et al., 2019](#)). These papers keep the market structure environment fixed and focus on the changes in equilibrium outcomes with respect to the subsidy schedule given by the following functional derivative:

$$\frac{\partial H(Z(D), \theta, D, r)}{\partial Z(D)}. \quad (3)$$

In this paper, we extend this prior literature by investigating how the distribution of equilibrium outcomes H depends on the interaction between market structure θ and the subsidies' dependence on consumer type D . The change in outcomes going from the observed market structure to perfect competition that we are interested in is:

$$H(Z(D), \theta, D, r) - H(Z(D), \theta_0, D, r). \quad (4)$$

By interacting this difference with the calculation in [Equation 2](#), we can assess how market power, *ceteris paribus*, changes the distribution of equilibrium outcomes in the presence of targeted subsidies:

$$\underbrace{(H(Z(D), \theta, D, r) - H(0, \theta, D, r))}_{\text{Subsidy Effect with Observed Market Structure}} - \underbrace{(H(Z(D), \theta_0, D, r) - H(0, \theta_0, D, r))}_{\text{Subsidy Effect under Perfect Competition}}. \quad (5)$$

A primitive that drives the outcome of [5](#) is the elasticity of the objective function with respect to the distribution of demographics:

$$\frac{\partial H(Z(D), \theta, D, r)}{\partial D}. \quad (6)$$

In settings where perfect price discrimination is either not possible or prohibited by regulation, [Equation 6](#) will typically be non-zero on its support. We term the dependence of equilibrium outcomes on the distribution of types as a “demographic externality,” as generically a consumer’s outcome will depend on the distribution of its neighbors’ types.

We note that the effects captured by [Equations 4, 5, and 6](#) do not necessarily depend on any notion of asymmetric information in D ; a market setting can have non-trivial interactions between market power and subsidy targeting without suffering from either adverse selection or moral hazard. The fact that $Z(D)$ depends on D directly has the possibility of generating economically-relevant demographic externalities. To illustrate these forces in more detail, we turn to a simple worked example of a monopolist selling in an insurance market.

Monopoly Example Consider a market where a single good is sold at a uniform price. Suppose that there is a unit mass of consumers and each of them has the following utility for the good:

$$U_i = u(p; d_i, \theta_i, \epsilon_i), \quad (7)$$

where i indexes the consumer, p is the product's price, d_i are consumer characteristics, θ_i is a vector of utility parameters, and ϵ_i is a vector of preference shocks. Consumers purchase a unit amount of the good if and only if $U_i > 0$. Denoting the distribution of consumer characteristics by D and the joint set of preferences and shocks by F , market-level demand—share of consumers buying the good $s(p)$ —is formed by aggregating individual demands:

$$s(p) = \int 1(u(p; d, \theta, \epsilon) > 0) dF dD = \int s_d(p; d) dD, \quad (8)$$

where $s_d(p; d)$ is the share of consumers within group d who buy the good.

Now suppose the government introduces a schedule of targeted subsidies, $Z(D)$, where consumers of type d receive a transfer equal to z_d if they buy the good and do not receive any transfer if they do not. An extensive literature going back at least to [Hylland and Zeckhauser \(1981\)](#) and [Nichols and Zeckhauser \(1982\)](#) discusses when in-kind targeted transfers may improve target efficiency of distributional policies and be preferred by the government to unrestricted cash transfers. For the moment we sidestep the issue of the optimality of transfer design per se, and take the choice of the policy instrument $Z(D)$ as given. We come back to the question of subsidy design and the tradeoff between efficiency and redistribution in the context of our empirical application in [Section 5](#).

Given the subsidy schedule $Z(D)$, demand shifts outward:

$$s(p, Z(D)) = \int s_d(p - z_d; d) dD. \quad (9)$$

Next, consider the supply side of the market. Under perfect competition, prices are set equal to average marginal cost. Denoting the marginal cost of each consumer type by c_d , the competitive price solves the following equation:

$$p = \frac{1}{s(p, Z(D))} \int c_d \cdot s_d(p - z_d; d) dD. \quad (10)$$

If marginal costs are equal across all consumer types, as is the case in many product markets, [Equation 10](#) reduces to the familiar $p = c$ formulation.

When a private intermediary with market power is introduced to the model, there are

additional changes in equilibrium. Rather than setting price at average marginal cost, the firm sets prices to maximize profits:

$$\pi(p) = \int s_d(p - z_d)(p - c_d)dD. \quad (11)$$

The equilibrium price equalizes marginal revenue and marginal cost:

$$\int s_d(p - z_d) + p \cdot \frac{\partial s_d(p - z_d)}{\partial p} dD = \int c_d \cdot \frac{\partial s_d(p - z_d)}{\partial p} dD. \quad (12)$$

In general, the equilibrium price under both market structures depends on the distribution of demographics D —a relationship that we label a “demographic externality.” Formally, it is the functional derivative of (consumer-facing) price with respect to the distribution of demographics:

$$\frac{\partial(p - z)}{\partial D}. \quad (13)$$

While the demographic externality encompasses many mechanisms that generate interdependencies across consumers in a market, in our setting we focus on the demographic externality driven by means-tested public subsidies. To the first order, such subsidies change the demand curve for lower-income consumers, reducing demand elasticity of these consumers and increasing incentives for the firm to raise prices. In addition, subsidies change the relative importance of different consumer types in the firms’ profit function, with ambiguous effects on the equilibrium price. For example, if lower-income highly-subsidized consumers also have very elastic demand, attracting these consumers into the market could dampen incentives to raise prices.

We highlight the mechanism of the demographic externality in the context of means-tested public transfers because it may generate negative (and potentially ironic) equilibrium outcomes. For example, suppose that consumers are all identical except for income. If the government sets up a schedule of means-tested subsidies for the good that decline in income, consumers in the middle of that income distribution will end paying more out-of-pocket for the good when they are surrounded by relatively more poor consumers. This is a type of price discrimination effect that is inverted from the standard model, where being surrounded by high-income, low-elasticity neighbors can lead to an increase in prices. The demographic externality that stems from targeted subsidies generates inequality in consumption across markets solely due to demographic composition, which is an undesirable consequence of the regulatory design.

3 Institutional Primer and Data

Institutions Our empirical application is the US market for non-group health insurance plans that was launched in 2014 under the Affordable Care Act (ACA). The ACA “Marketplaces” platform allows consumers to purchase health insurance plans for themselves and their families. Insurance plans sold on this market are high dimensional products, offering a variety of cost-sharing levels and provider networks. The financial characteristics of plans are frequently summarized by its actuarial value that measures the fraction of costs that a plan would cover for a standardized population. Plans are grouped into “metal levels” on the basis of how much risk is borne by the insurer versus the consumer: Bronze (60% actuarial value), Silver (70%), Gold (80%), and Platinum (90%); there is also a “catastrophic” level, characterized by low premiums and high deductibles, that is only available to consumers under 30. 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 so-called “federally-facilitated” states in our analysis. About 9 million enrollees bought insurance on the Marketplaces in these states in 2017, which is the year of our data. Figure 1 provides an example of the user interface on the healthcare.gov platform.

Insurers selling on the ACA Marketplaces set premiums for their plans subject to several regulatory restrictions. First, insurers are not allowed to reject enrollees based on pre-existing health conditions and are not allowed to price-discriminate based on individual health risk. Insurers can collect different premiums from consumers based on age, but the age gradient in premiums has to follow a pre-specified regulatory age curve.⁹ Second, insurers have to charge the same premiums in all counties that belong to the same “rating area.” Rating areas are collection of counties pre-specified by each state. In practice, while insurers have to charge the same price if they offer their product in all counties in a rating area, they do not have to serve all counties in a rating area. Thus, following [Fang and Ko \(2018\)](#) we consider a county to be the relevant market boundary in this market. Finally, insurer premiums have to satisfy the medical loss ratio (MLR), which requires them to spend at least 80% of revenue on medical reimbursement, constraining markups to be at most 25 percent. The MLR has been documented to be binding for the majority of insurance contracts in this market ([Cicala et al., 2019](#)).

⁹Insurers 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.

Figure 1: Consumer interface on healthcare.gov

The screenshot displays the healthcare.gov website interface. At the top, there are navigation tabs for "Individuals & Families" (selected), "Small Businesses", "Log in", and "Español". A dark blue banner contains an important notice: "Important: The premiums below are only estimates. You'll need to fill out a Marketplace application to get actual plan prices. Some plans and details you see here may change." Below this, the main heading is "121 Health Plans".

The interface is divided into several sections:

- Navigation:** "BACK TO QUESTIONS" button.
- Filters:** "Viewing:" tabs for "HEALTH PLANS" (selected) and "DENTAL PLANS"; "Sort:" tabs for "BY MONTHLY PREMIUM" (selected) and "BY DEDUCTIBLE".
- Left Sidebar (NARROW YOUR RESULTS):**
 - Premium:** Options include "less than \$200 (17)", "less than \$300 (80)", "less than \$400 (119)", and "less than \$500 (121)". A "Get more details about premiums" link is present.
 - Coverage categories:** Options include "Bronze plans (33)", "Silver plans (42)", "Gold plans (33)", and "Platinum plans (13)". A "Get more details about categories" link is present.
 - Plan Types:** Options include "PPO (52)", "HMO (64)", and "POS (5)". A "Get more details about plan types" link is present.
 - Insurance companies:** Option is "Aetna (5)".
- Main Plan Details:**
 - Title:** "Health Choice Insurance Co. · Health Choice Value Bronze" with a "Compare" button.
 - Type:** "Bronze HMO" with Plan ID: "70239AZ0010043".
 - Key Metrics:**
 - ESTIMATED MONTHLY PREMIUM:** \$153
 - ESTIMATED DEDUCTIBLE:** \$5,000 (Estimated individual total)
 - ESTIMATED OUT-OF-POCKET MAXIMUM:** \$6,600 (Estimated individual total)
 - COPAYMENTS / COINSURANCE:**
 - Primary doctor: \$20 Copay after deductible
 - Specialist doctor: \$50 Copay after deductible
 - Emergency room care: \$500 Copay after deductible
 - Generic drugs: \$15 Copay after deductible
 - PEOPLE COVERED:** 1 (Age 40): Covered
 - MORE INFORMATION:** Links for "Summary of Benefits", "Plan brochure", "Provider directory", and "List of covered drugs".
 - Action:** "LEARN MORE ABOUT THIS PLAN" button.

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

The key institutional feature of ACA Marketplaces are means-tested subsidies that can be used by low income families to defray the cost of the insurance premium.¹⁰ Formally, premium subsidies are known as Advanced Premium Tax Credits (APTC).¹¹ The APTC is calculated in several steps. First, the Modified Adjusted Gross Income (MAGI) for a tax family is converted to the percent of the Federal Poverty Level (FPL).¹² The FPL varies with family composition and allows comparing incomes of families of different sizes using the same scale. The MAGI relative to the 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 specified by the IRS.¹³ For example, in 2017, if a household’s income was 200 percent of the FPL, then this household’s CAP was 6.34 percent of household income, while if income was 270 percent FPL, the CAP was 8.7 percent. Subsidies phase out at 400 percent FPL, so that households at or above this income threshold are not subsidized.

Once the CAP has been computed, the next step is to calculate the total premium that a household would owe for a benchmark plan in the absence of subsidies. The benchmark insurance plan is set to be the second-lowest cost plan with 70 percent actuarial value (“second-lowest cost silver plan,” or SLCSPP) in the household’s county of residence. “Premium owed” is the total list premium that the household would have to pay to enroll all family members who would like to be insured in the benchmark plan (known as “coverage family”). If the household’s CAP is lower than the list premium that the household would have had to pay for the benchmark plan, then the household gets an APTC that is equal to the difference between “premium owed” and the CAP. If the list price of the plan that the household actually buys is lower than the full APTC that the household is eligible for, APTC is reduced to the actual cost of the plan, so that the final consumer-facing premium can be zero, but cannot be negative. To summarize, the subsidy that household h receives

¹⁰In addition, low income families qualify for cost-sharing reduction subsidies. These subsidies reduce the out of pocket liability from deductibles, co-pays, and co-insurance. We take these subsidies as given and do not alter them in our counterfactual analyses.

¹¹The 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.

¹²MAGI 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.

¹³See IRS Form 8962 “Premium Tax Credit.”

as a function of its income Y_h is:

$$z_h = \max \left\{ 0, \sum_{i \in h} b_{i,SLCSP} - CAP(Y_h) \right\}, \quad (14)$$

where $\sum_{i \in h} b_{i,SLCSP}$ is the total list premium that the coverage household would have to pay for the benchmark plan. This mechanism both induces means-tested subsidies, as CAP varies with income, and also ties the level of generosity for that subsidy to equilibrium price through the price of the second-lowest cost silver plan. The out-of-pocket premium that that household pays for plan j is then:

$$p_{hj} = \max \left\{ 0, \sum_{i \in h} b_{ij} - z_h \right\}. \quad (15)$$

Data and summary statistics Our analysis relies on a combination of several data sources. First, we use 2017 Marketplace Public Use Files (PUF) that record detailed information on which plans were offered in each geographic market of the federally-facilitated Marketplaces, plan features, and list premiums. Second, we use data provided by the Center for Consumer Information and Insurance Oversight (CCIIO) that reports Marketplace enrollment counts.¹⁴ Enrollment is reported in several ways. First, we have enrollment reported at the plan level, where plans may be offered in several markets. Second, we also observe enrollment at the county level aggregated by different levels of plan actuarial values (i.e. metal levels). For example, we observe how many individuals purchased a plan with 70 percent actuarial value (“Silver” plans) in Cook County, IL. Finally, we also observe the market-level counts of consumers that enrolled in any plan by seven age and seven income buckets.¹⁵ We enrich enrollment data with a dataset, provided by Kaiser Family Foundation, that computes the number of potential ACA Marketplaces consumers for each geographic region.¹⁶

Next, we use the 2017 American Community Survey (ACS) to create a representative sample of potential ACA Marketplace consumers in each county. For these potential consumers

¹⁴CCIIO data is public and can be accessed at www.cms.gov/CCIIO/Resources/Data-Resources.

¹⁵The age categories are: under 18, 18-25, 26-34, 34-44, 45-54, 55-64, and over 65. The income categories are: below 100 FPL, 100-150 FPL, 150-200 FPL, 200-250 FPL, 250-300 FPL, 300-400 FPL, and over 400 FPL.

¹⁶This data was requested in 2017 and underlies computations of Marketplaces take-up that Kaiser Family Foundation (KFF) has been reporting continuously in its briefings “Marketplace Enrollment as a Share of the Potential Marketplace Population.” KFF analysis, in turn, is based on the counts of uninsured individuals from the American Community Survey.

we observe household composition, household income relative to FPL, and age.¹⁷ Finally, we use Optum database of claims and demographic information for the commercially insured under-65 population to estimate the gradient in average healthcare costs across different age and income groups. We use this database to create marginal cost multipliers as a function of age and income; these multipliers are then used in the marginal cost inversion below to obtain plan-type specific marginal costs, which is key for accounting for selection in this market.¹⁸ Appendix B provides details.

Table 1 summarizes the data. In 2017, consumers faced on average a choice of 21 plans. 2.2 large insurers operated in an average county. The annual list premium for a 40-year old consumer ranged from \$3,978 (10th percentile) to \$6,351 (90th percentile) with an unweighted average of \$5,160. The average number of potential enrollees per county 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, plans with 70 percent actuarial value 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, but plan sizes varied substantially from plans covering fewer than 50 consumers to plans with more than 6,000 enrollees within a county. Potential enrollees based on our ACS sample were on average 39 years old and had an average household income of 295 percent FPL. About a third of potential consumers had income under 200% FPL, making them eligible for the most generous subsidy levels, 37 percent had income over 200% FPL but under 400%FPL making them eligible for partial subsidies, and another third of potential consumers had household income over 400% FPL, making them ineligible for subsidies. On average, potential consumers qualified for \$2,349 in annual premium subsidies.¹⁹

¹⁷To construct the sample of potential consumers we keep individuals in ACS who do not report having employer-sponsored or public health insurance coverage. We further exclude individuals from our set of potential consumers if they report having no health insurance coverage, but would likely be eligible for Medicaid based on their household income in states that expanded Medicaid under the ACA.

¹⁸Tebaldi (2017) uses multiplier from the MEPS database in a similar fashion.

¹⁹We divide the household-level subsidy by the number of individuals that would need insurance in the household to compute average subsidy per “coverage family” member.

4 Empirical Model

4.1 Demand

Utility function We estimate demand for health insurance plans on ACA Marketplaces using a semi-parametric discrete choice random utility model. The unit of choice is a “coverage family,” f , consisting of $i = \{1, \dots, N_f\}$ members of the same tax household who are in need of insurance coverage. Within each family, we assume that each consumer i obtains the following utility of consuming insurance plan j :

$$u_{ij} = -\alpha_i p_{ij} + \phi_{ij}, \quad (16)$$

where ϕ_{ij} is the utility of plan j , p_{ij} is the plan’s out-of-pocket price, and α_i is the marginal disutility of price. We assume that the family f chooses a single insurance plan that maximizes the average utility across members of the family:

$$\epsilon_{fj} + \frac{1}{N_f} \sum_{i \in f} u_{ij} > \epsilon_{fk} + \frac{1}{N_f} \sum_{i \in f} u_{ik}, \forall k \in J \text{ s.t. } k \neq j, \quad (17)$$

where ϵ_{fj} is a family-level idiosyncratic taste shock for plan j , or chooses the outside option of no purchase, which has a normalized payoff of zero.

We make several assumptions about α_i and ϕ_{ij} to arrive at an empirically-tractable version of utility. First, we replace individual-specific α_i with a coarser schedule of marginal utilities of income. We allow α to vary across nine demographic groups, d . The demographic groups are defined as the cross product 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, and income above 400 percent FPL. Second, we decompose ϕ_{ijt} , the utility that a consumer gets from plan j , into several additively-separable components:

$$\phi_{ij} = \psi_{a(i)} + \gamma AV_{ij} + \delta_j. \quad (18)$$

The first component, $\psi_{a(i)}$, captures the average level of utility that consumers get from purchasing any insurance plan. We allow this intercept parameter to vary across the same three age groups as the price coefficient to capture the idea that the value of insurance may vary across ages, all else equal.²⁰ The second component, AV_{ij} , captures the deviations in

²⁰We 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

the generosity of plan j that consumers may face if they have sufficiently low income. For these consumers, cost-sharing reduction subsidies change the actuarial value (AV) of plans.²¹ Finally, we include a plan-specific constant δ_j for each plan j that captures the average utility that all consumers get from purchasing that plan.

Identification We briefly sketch out an intuitive identification argument for the parameters of our model given the type of data that we have. First, we note that [Berry et al. \(1995\)](#) established that, after integrating out individual-specific utility components, there exists a unique vector of mean product utilities, δ , for any vector of product shares within a market. A complication arises in our setting because CMS reports enrollments at the plan level, where plans are often offered in several markets. Therefore, we cannot construct plan-market level shares. However, as we know where each product is offered we are able to aggregate across individual markets to predict plan-level enrollments whenever a product is offered in more than one local market. Under the assumption that the common component of a plan’s utility is the same across all markets in which it is offered, there exists a unique δ_j that rationalizes observed plan-level enrollments reported in the data.²²

This establishes that δ is one-to-one in plan level enrollments conditional on individual-level variation in utility. To establish the identification of the parameters governing these individual-level utility components, we note that there are several dimensions of residual variation in our data that δ alone cannot account for. We consider each of these sets of parameters in turn.

Recalling that CMS reports metal-level enrollment by market, we can identify the marginal willingness to pay for actuarial value, γ , through the variation in the share of enrollment in Silver plans (the only plans that qualify for cost-sharing reductions) across markets with different demographic compositions but the same set of plan choices. For example, Alabama has a single set of plans offered to all consumers in the state. Demographic variation across markets makes Silver plans more attractive in places with more consumers that qualify for cost-sharing reductions; the degree to which Silver plans are more popular in those locations

eligible for Medicaid and are in theory not eligible for ACA subsidies, we observe some enrollees from this group in the data; a separate intercept for this group allows the model to rationalize a very low, but non-zero, inside share for this group.

²¹To be eligible for cost-sharing reduction subsidies, consumers need to enroll in “Silver” plans, i.e. plans with 70% actuarial value, and then the AV of the plan gets increased depending on consumers’ household income.

²²This follows from a simple modification of the definition of a share to span multiple markets in [Berry et al. \(1995\)](#)’s original proof, which shows that the difference equation defining δ is a contraction mapping, which is a sufficient condition for uniqueness.

identifies γ under the assumption that δ does not vary across markets.

Two additional sets of parameters vary across consumers: the disutility of price, $\alpha_{d(i)}$, and the baseline valuation of insurance plans, $\psi_{a(i)}$. In our data, we do not observe individual-level purchases; however, we do observe two moments of the data that are not usually present in [Berry et al. \(1995\)](#) style models: at the market level, we observe the share of consumers buying any product by intervals of age and income. While δ identifies the average level of enrollments, and therefore the aggregate share of the inside option, it cannot rationalize the within market variation in enrollments across ages or incomes. Within a given market, the rate at which enrollments decline as income increases across income bins gives a local estimate of α . Similarly, variation in the enrollment in the inside option across age bins identifies ψ , the baseline willingness-to-pay for any insurance plan. As with γ , one could also use variation across markets to identify those parameters; as the demographic distribution of consumers in two markets varies by age and income, the patterns in overall enrollments identify both sets of parameters.

Finally, we note that, in principle, both α and ψ can be flexible functions of income and age, respectively. Analogous to a local linear regression, one can estimate slopes in a neighborhood of any given age or income. Cross-market differences in the distribution of demographics creates rich variation allowing for the flexible recovering of local estimates of these two sets of coefficients.

Price Endogeneity A major concern in product markets is price endogeneity; the typical issue is that price is correlated with unobserved quality or demand shocks within a market, creating an inference problem. The usual solution to this problem in settings with aggregate cross-sectional data is to search for instruments that move prices across markets ([Hausman, 1996](#); [Berry and Haile, 2016](#)). Here, however, the regulatory design itself provides an innovative solution to this issue. The key observation is that the subsidy schedule generates variation in prices across consumers within a market that is, by regulatory fiat, independent of demand shocks or unobserved product quality. Subsidies vary across income levels according to a pre-specified administrative formula, generating a non-decreasing relationship between income and effective premiums, as discussed in [Section 3](#). The statutory age-adjustment curve does the same for consumers of different ages. These exogenous sources of variation stemming from the regulatory design allow us to estimate α and ψ using a within-market estimator while conditioning on δ . This approach is valid as long as consumers are not systematically sorting themselves across markets on the basis of these regulatory-induced

changes in prices.

4.2 Supply

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. To fix ideas, we first start with a brief accounting of payment flows in the market.

For each consumer i , plan j collects premium p_{ij} from the consumer. For consumers who are not eligible for premium subsidies, this premium is equal to the full list price, b_{ij} .²³ For consumers who are eligible for subsidies, the insurer collects $p_{ij} < b_{ij}$ from the consumer and a subsidy from the federal government. Together, the consumer premium and the subsidy add up to b_{ij} .²⁴ 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 a consumer’s underlying health risk, r_i , and the plan’s contract features, ϕ_j .

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. Let $c_{ij}(r_i, \phi_j) \leq h_{ij}$ denote the plan’s expected cost for offering insurance coverage to individual i . Plan j ’s expected profit from enrolling consumer i is then:

$$\pi_{ij} = b_{ij} - c_{ij}(r_i, \phi_j). \tag{19}$$

There are several features of the environment that allow us to simplify this very general formulation. On the revenue side, recall that plans are not allowed to perfectly price-discriminate across consumers. Insurers can collect higher premiums from older consumers;

²³Recall that the list price of plan j is the same in all markets t where plan j is offered and that fall into the same rating area. If a consumer is eligible for subsidies, then the price that consumers pay for j in two different counties may differ, since subsidies depend on the prices of other plans in the same market. Consumers not eligible for subsidies pay the same price for j in all markets.

²⁴If the subsidy is higher than the list price, the consumer pays zero and does not receive the cash value of the “unused” subsidy. In addition to premiums, insurers may receive revenue from three risk-equalization programs run by the federal government or the states that aim to reduce the ex post volatility in realized profits from unexpectedly high or low healthcare spending realizations.

however, this age-based underwriting is regulated—insurers set the base rate for a 20 year old consumer and then have to follow a pre-specified regulatory age curve. This regulatory restriction implies that revenue does not vary across consumers of age a and can be written as a product of the baseline list price for a 20 year old consumer b_j and a set of regulatory age-specific multipliers τ^a that do not vary across plans, i.e. $b_{ij} = \tau^{a(i)}b_j$.

On the cost side, we parametrize c_{ij} to make the problem empirically tractable. We preserve the idea that costs may vary across consumers, which is important for allowing risk sorting to exist in this insurance market, but we discretize this variation. Using a commercial claims database from Optum, we start by estimating how much average healthcare spending varies across individuals with age and income in the general population. We fit smooth functions of healthcare costs in age for three levels of income (see Section B in the Appendix for details). Three income levels correspond to the income brackets that were defined for demand estimation. These estimates recover a matrix of age-income cost multipliers κ^d . Now suppose that each plan has a plan-specific baseline cost c_j —the cost that plan j expects to incur for a 20 year old consumer in the lowest income bracket. This cost may vary across plans due to differences in benefit design and negotiated prices with providers. We make a simplifying assumption that income-age cost multipliers do not vary across plans. In other words, we assume that in expectation enrolling a 20 year old versus a 40 year old in the lowest income bracket is κ^{40} times more expensive for any plan. Under this assumption, c_{ij} can be written as a product of the baseline cost c_j and the age-income cost multiplier κ^d that does not vary across plans, i.e. $c_{ij} = \kappa^{d(i)}c_j$. Allowing demand and cost to both vary by age and income allows for adverse or advantageous selection in our model, as cost systematically varies with willingness to pay.

With this additional structure, we can re-write the profit equation for firm f offering plan portfolio F_f as:

$$\Pi_f(b) = \sum_{j \in F_f} \sum_{d \in D} [(b_j \tau^d - c_j \kappa^d) s_j^d(p(b)) M^d], \quad (20)$$

where $s_j^d(p(b))$ is the share of consumers in age-income group d that buys plan j among all consumers in this age-income group M^d , $p(b)$ is the link function between the out-of-pocket price that a consumer pays and the firm's list price (Equation 15 above), $\tau^d = \tau^a$ for any income level, b_j is the baseline list premium for a 20 year old that does not vary with income, and c_j is the baseline cost for a 20 year old consumer in the lowest-income bracket.²⁵

²⁵The share equation is defined over all markets for which the plan is offered (i.e. its rating area). For

First-order conditions Each insurer f chooses a vector of baseline list prices to maximize profits; subject to the MLR constraint, for each plan j , the optimal list price b_j satisfies the following first-order condition:

$$\frac{\partial \pi_f(b)}{\partial b_j} = \sum_{k \in F_f} \sum_{d \in D} \left[(b_k \tau^d - c_k \kappa^d) \frac{\partial s_k^d(p(b))}{\partial b_j} M^d + 1(j = k) \cdot \tau^d s_j^d(p(b)) M^d \right] = 0. \quad (21)$$

The first-order condition highlights how the theoretical demographic externality mechanism discussed in Section 2 enters our empirical model. Note that while the firm sets on baseline list price b_j , it is maximizing the sum of marginal profits with respect to b_j across all age-income groups d . This implies that if one consumer group d is affected by a targeted subsidy that changes its elasticity of demand, this will affect the list price not only for this group of consumers, but for all other demographic groups d . This interdependence of demands in the firms' profit function gives rise to the subsidy-induced demographic externality in our empirical model.

As in Nevo (2001), the first-order conditions can be written in the vector form:

$$S - \Omega(B - C) = 0, \quad (22)$$

where row j of vector S is given by $S_j = \sum_d \tau^d s_j^d(p(b)) M^d$ and row j of vector $(B - C)$ is given by $(B - C)_j = \sum_d (b_j \tau^d - c_j \kappa^d)$, while row k , column j of matrix Ω is:

$$\Omega_{kj} = - \sum_d \frac{\partial s_j^d(p(b))}{\partial b_k} M^d \quad (23)$$

for plans k and j offered by firm f . This expression is useful as one can invert Equation 22 to obtain 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. We constrain our estimates of marginal cost to conform to the MLR restriction described above.²⁶

Subsidies enter the first-order condition in two ways. One, they directly change the share equation $s(p(b))$ by changing out-of-pocket prices faced by a subset of consumers. Second, they also change the derivative of the market share with respect to price: $\frac{\partial s_j^d(p(b))}{\partial b_k^d}$. This

expositional clarity, we do not write out this dependence explicitly, but we do account for this in our empirical estimation and counterfactual simulations.

²⁶As marginal cost is additively separable in the first-order conditions, one can invert all marginal costs and impose the MLR without having to adjust the other non-binding marginal costs.

share derivative reflects how much the demand for plan j changes when plan k increases its baseline list price by a small amount. Unlike in a standard product-market setting, this term captures not only how consumers respond to price changes, but also the relationship between consumer-facing and list prices. Consumer premiums and list prices are linked via the subsidy mechanism captured by $\frac{\partial p_j}{\partial b_k}$, so that for any demographic group d ,

$$\frac{ds_j(p(b))}{db_k} = \frac{\partial s_j(p(b))}{\partial p_k} \cdot \frac{\partial p_k}{\partial b_k}. \quad (24)$$

In unsubsidized product markets, $p = b$, so $\frac{\partial p_k}{\partial b_k} = 1$. Here, the out-of-pocket price is determined by Equation 15 and that is no longer necessarily the case: for lower-income consumers where $b_k < z_h$, the derivative is zero; for all other consumers with $b_k > z_h$, the derivative is one.

Finally, when plan k is the second-lowest cost silver plan there are additional terms in Equation 24. This is due to b_k entering Equation 14, which determines subsidies for all consumers in the market. In this case, changing b_k not only shifts p_k (for a subset of consumers) but also shifts p_{-k} for all other products via a change in z_h . This effect generates a modified first-order condition which we require for both solving out marginal cost and computing counterfactual equilibria. However, two aspects of our setting allow us to avoid explicitly modeling these complex incentives. First, we note that firms setting the second-lowest cost silver plan have incentives to raise prices arbitrarily high, as the generosity of the subsidy the government gives consumers increases one-for-one in b_k . Therefore, we directly impose that the MLR constraint is binding for these plans. Second, our counterfactual simulations consider pricing environments that do not have the same subsidy mechanism where the second lowest-cost silver plan plays a role. As a result, we will not use the modified first-order condition in either estimation or counterfactual simulation.

4.3 Efficiency Metric

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

$$W = CS + \Pi - \lambda G, \quad (25)$$

where λ is the social cost of raising public revenues, which we assume to be 30 cents on a dollar. In our baseline analysis we define consumer surplus as a compensating variation that

puts equal social welfare weights on the utility of all consumers, following [Williams \(1977\)](#) and [Small and Rosen \(1981\)](#). Consumer surplus for consumer i with a vector of marginal utilities θ_i then takes the following form:

$$CS_i(\theta_i) = \frac{1}{\alpha_i} \left[\gamma + \ln \left[1 + \sum_{j=1}^J \exp(v_{ij}(\theta_i)) \right] \right], \quad (26)$$

where γ is Euler’s constant and v_{ij} is the deterministic component of utility for person i (recall that this is the average utility within a family) for plan j .²⁷ We integrate out over the empirical distribution (as observed in the ACS) of ages, income, and family composition to obtain average consumer surplus that is then scaled to the total market size:

$$CS = M \int CS(\theta) dF(\theta). \quad (27)$$

Producer surplus, Π , is computed as expected profits following equation 20. Government spending G includes three parts. The first component of G is the nominal spending on premium subsidies. These are computed either from the data or are adjusted following the simulation scenarios of Section 5. The second component is government outlays for cost-sharing reduction subsidies. We hold cost-sharing reduction (CSR) spending constant at the observed level across all counterfactual simulations. Specifically, using CCHIO data reports we compute the average per capita government spending on CSR subsidies by consumer type, based on income brackets.²⁸ 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 a cost-sharing reduction is available. The third component of government spending accounts for the fact that when a consumer enrolls into an ACA Marketplace 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 prior to the rollout of the ACA Marketplaces ([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.

²⁷Euler’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.

²⁸The data was accessed in June 2019 at [Health Insurance Marketplace Cost-Sharing Reduction Subsidies](#).

Consumer surplus with preferences for redistribution To analyze the alternative subsidy mechanism—income-invariant subsidies—in Section 5, we turn to an alternative definition of consumer surplus that directly allows for preferences for redistribution. We use the family of welfare functions with a constant relative inequality aversion (Atkinson, 1970).²⁹ Individual-level consumer surplus is recomputed as a function of individuals’ income y_i , as follows:

$$CS_i^\lambda = \begin{cases} \frac{1}{1-\lambda}[(y_i + CS_i)^{1-\lambda} - y_i^{1-\lambda}] & \text{if } \lambda \neq 1, \\ \log(y_i + CS_i) - \log(y_i) & \text{if } \lambda = 1. \end{cases} \quad (28)$$

Integrating out CS_i^λ over the empirical distribution of demographics then gives us average consumer surplus that is weighted by a function of consumers’ incomes for each level of λ . The parameter λ measures preferences for redistribution. At $\lambda = 0$, each consumer receives the same social welfare weight, recovering our baseline average consumer surplus. As λ increases, transfers to lower-income households become more valued by the society than equivalent transfers to higher-income households. The parameter λ captures how much welfare increases (in percentage points) when \$1 is transferred to a consumer with one percent lower income.

4.4 Estimation results

Demand We use non-linear least squares to estimate the utility function parameters.³⁰ Panel A of Table 2 reports the results. We find intuitive patterns for the variation in the marginal utility of income across demographic groups. A one dollar increase in price has a larger impact on the utility of poorer 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. We further 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. For example, consumers over age of 40 with income between 200% and 400% FPL who are eligible only for partial subsidies are willing to pay \$680 in premiums for each 10 percentage point increase in actuarial value of a plan. This estimate of the willingness

²⁹See also discussion in Waldinger (2018) for a recent empirical implementation in the context of subsidized housing.

³⁰Computational details are provided in Appendix A.

to pay is within the empirical support of differences in list premiums faced by a 40 year old consumer, who would need to pay \$730 on average in list prices across all markets to move from a 60 percent actuarial value (Bronze) to a 70 percent actuarial value (Silver) plan.

While the patterns of parameter values are intuitive, we are cautious about the interpretation of individual magnitudes, 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 the fit of the model, we compare how enrollment moments predicted by the model compare to the moments observed in the data. Appendix Figure C.4 illustrates one set of moments that was used for estimation and a related measure of model fit. In Panel C.4A we report observed market share of Silver plans in each county. In Panel C.4B we report the average in-sample difference between the data and the model’s prediction of county-level Silver enrollment shares. The in-sample fit of the model is tight. The average prediction error is -0.0007 relative to average empirical enrollment share across counties of 0.29, broadly distributed across counties, suggesting that the model is able to capture a substantial amount of variation in the data.

Supply Panel B of Table 2 reports the results of the supply side estimation. Inverting the first-order conditions results in an average (across all plans, averaging within plans across rating areas) estimate of the marginal cost for a 20 year old consumer with income under 200% FPL of \$1,561 (c_j), with a standard deviation of \$454. The least costly plan had an estimate baseline marginal cost of \$732, while the most costly plan having the baseline marginal cost of \$4,102. From commercial claims data we estimate three vectors of cost multipliers (κ^d) that allow us to compute the expected cost for each individual in each plan following our parametric assumption in Section 4.2 that $c_{ij} = \kappa^{d(i)}c_j$. We estimate that on average across all ages individuals in households with lowest incomes experience 40 percent higher healthcare costs than individuals in households with income over 400% FPL. These patterns are consistent with the widely documented socio-economic gradient in health (Marmot, 2015).

We observe two pronounced patterns in our estimates of marginal costs. First, there is

substantial heterogeneity in costs across plans with the same level of actuarial value. 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. This is also intuitive, as mechanically more generous plans cover a higher percent of consumers' healthcare expenditures on average. In addition, more generous plans may attract an overall sicker pool of consumers, leading to disproportionately higher costs than the mechanical difference in actuarial value would suggest. The ratio of (average) marginal cost between 60% actuarial value, 70% actuarial value, and 80% actuarial value plans that we estimate are consistent with mechanically higher costs among 70% AV plans and some adverse selection in 80% plans.

In Appendix Figure C.5, we compare our estimates of marginal costs from the first-order condition inversion to plan-level accounting costs as reported by plans to CMS for 705 plans for which accounting data was available. 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, as they likely generate an accurate ordinal ranking of plans from the least to the most expensive, on average. As we would expect given the existence of ex post cost reconciliation transfers, our estimates of marginal cost are on average 84% of the average reported accounting cost (\$4,010 versus \$4,768). We observe a strong ordinal correlation between accounting costs and marginal costs, which supports the idea that we are accurately able to differentiate more and less expensive plans.³¹

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

We entertain two thought experiments to illustrate this idea. In the first experiment, we increase the number of subsidized consumers without changing the distribution of marginal

³¹Related work in this area 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 for several reasons. First, accounting cost data were not available for all plans. Second, 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 be inconsequential for the subsequent analyses, given the strong correlation between the two measures.

utility of income or cost 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 2 and are marked as case “A.” This figure shows how the average premiums and consumer surplus change for consumers who are *not* directly affected by the change—those with incomes between 150 percent and 400 percent FPL—when their neighbors with income above 400 percent start getting subsidies, all else equal. Two forces are at play here. On one hand, insurers have an incentive to raise prices to take advantage of the fact that in the 400+ percent FPL market segment—which is relatively inelastic—consumers now face lower prices for any given list price and are more likely to buy insurance. On the other hand, however, the pool of insured individuals now includes more higher-income consumers who have lower costs. The net effect is that plans become less 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 premium for consumers with income under 400 percent FPL decreases by \$15 to \$20. Consumer surplus, marked with grey circles, in turn increases by up to \$35 for the poorest consumers. This counterfactual simulation cleanly illustrates the mechanism at the heart of our results: subsidizing one group of consumers in a market with market power, all else equal, changes prices and welfare for their neighbors. In this case, subsidized consumers exert a positive demographic externality on other consumers.

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 and the cost 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 and more costly consumers, which pushes prices up, the firms also face much more elastic consumers, which pushes prices down. Case B in Figure 2 illustrates that the first effect dominates in our empirical setting. In our context, moving to an environment with more poorer consumers increases list prices. As the dashed line marked with “B” in 2 illustrates, the annual average consumer-facing prices for consumers that are not directly affected by our simulation go up by ca. \$90.

This increase in premiums leads to a decrease 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. On average, however, in this simulation the lowest-income subsidized consumers exert a negative demographic externality on other consumers in the market.

The average changes in prices and consumer surplus depicted in Figure 2 mask a highly unequal distribution of changes across geographic areas. Some counties are nearly completely unaffected by the simulated change, while others experience a substantial gain or loss in average consumer surplus among consumers not directly affected by the counterfactual policy change. An important factor contributing to the geographic variation is the degree of competition and market power. As would be predicted by our theoretical model in Section 2, the magnitude of the demographic externality is amplified in markets with fewer insurers, and especially in monopoly markets. Figure 3 illustrates the relationship between the average change in consumer surplus (by county) among only indirectly affected consumers by the number of insurers in the county, controlling for demographic characteristics of the county such as age, sex, income, and racial composition, as well as state fixed effects. All else equal, the effect of the demographic externality on consumer surplus decreases by \$11 for each additional insurer on the market in the first scenarios where the demographic externality is positive, and increases by \$26 for each additional insurer in the case where the demographic externality is negative.

We next examine how these economic forces affect the efficiency and distributional properties of the more realistic targeted transfer structure on the ACA Marketplaces.

5 Counterfactual simulations

With the model of demand and supply of ACA plans we conduct several counterfactual simulations that allow us to examine how market power of intermediaries on this market is affecting the efficiency and incidence of targeted subsidies. We start by characterizing the aggregate efficiency and the economic incidence of the observed market allocation. In addition to providing the baseline for our subsequent analyses, these objects are of central policy interest in their own right. We are able to deliver novel estimates of economic efficiency and incidence of public spending in ACA Marketplaces—an extensively debated market that is setting an important policy precedent for how to privately provide publicly funded health insurance coverage in the United States. Having established this important baseline, we then

examine the central question of the paper, which is whether and by how much market power distorts the efficiency and the distributional effects of targeted subsidies. Finally, we consider an alternative budget-neutral subsidy mechanism that could reduce these distortions, but at the cost of an equity-efficiency tradeoff.

5.1 Efficiency and incidence of observed ACA subsidies

Efficiency We begin our analysis by examining the economic costs and benefits of observed means-tested subsidy payments. To facilitate our analysis, we first simulate the equilibrium under the observed level of subsidies within our model. This allows us to establish a baseline that differences out any demand model simulation error.³² In practice, our demand model has a tight in-sample fit, so that the resulting simulated allocation is very close to the data. Column (1) of Table 3 reports the characteristics of this baseline allocation. We estimate that consumer surplus in the ACA Marketplaces amounts to \$50 billion. Producer surplus amounts to \$15 billion. The government is spending \$29 billion in premium and cost-sharing reduction subsidies.³³

Consumer surplus of \$50 billion in total corresponds to \$2,495 in annual surplus per capita, on average, among 20 million potential ACA consumers. Consumer surplus levels vary substantially across different socio-demographic groups and different areas of the country. While lower income consumers receive the highest level of transfers, they experience the lowest level of consumer surplus given the high disutility from a dollar in premiums that we estimate for these consumers. At the same time, we estimate that older workers experience higher consumer surplus than younger workers for all income levels. The variation in surplus and subsidy levels dovetails with differential enrollment patterns across income and age groups: 40 percent of ACA Marketplace enrollees have incomes under 200 percent FPL, while these consumers account for 32 percent of the potential market as we saw in Table 1. Similarly, consumers over age 40 account for 56 percentage points of inside option enrollment, while their share in the population eligible for ACA insurance is 48 percent. Overall, the program attracts 45 percent of potential enrollees.³⁴ We further observe stark variation in average

³²The supply-side returns back observed list prices from the first order conditions and has no simulation component.

³³This level of government subsidies as we estimate in ACS data and our model is consistent with the subsidy spending reported by the Congressional Budget Office. CBO reports \$39 billion in net premium and cost-sharing subsidy spending for 2016, which includes spending in non-federally facilitated states that are not part of our computation (<https://www.cbo.gov/sites/default/files/114th-congress-2015-2016/reports/51385-healthinsurancebaselineonecol.pdf>).

³⁴Surplus monotonically increasing with age is consistent with the observation in Tebaldi (2017), who

consumer surplus across US counties (Appendix Figure C.6). Average consumer surplus ranges from ca. \$750 in many counties in Texas to over \$2,500 on average in parts of Arizona, Illinois, Florida, and Maine. The spatial variation in surplus stems both from differences in income and age composition of potential consumers across local geographies (Appendix Figure C.2), but also from pronounced differences in local market structure (Appendix Figure C.3); jointly, the variation in age, income, and the count of insurers and plans can explain 58% of the geographic variation in average consumer surplus.

Comparing consumer and producer surplus vis-a-vis government spending, we conclude that public subsidy expenditures generate \$2.25 of surplus for \$1 of spending: the outlays on subsidies are \$36 billion lower than the sum of consumer and producer surplus. This partial equilibrium computation of subsidy efficiency, 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 compute the amount that the government would have spent on ACA enrollees in uncompensated care payments (row 5 Table 3). We estimate that the savings of public funds on the same set of consumers amounts to \$16.5 billion, which implies that the *net* additional government outlays for premium and cost-sharing subsidies on the Marketplaces are approximately \$12 billion rather than \$29 billion. Taking into account this foregone spending, we estimate a total return of \$5.25 on a dollar of public funds spent on the ACA Marketplaces without accounting for the cost of raising public funds, and \$4.04 with the cost of public funds included at 30 cents on a dollar.

Incidence 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 \$50 billion to \$43 billion, while producer surplus drops from \$14.7 billion to \$6.8 billion. Overall, \$28 billion in premium subsidies generates \$15 billion in additional surplus. The extra subsidy spending is partially mitigated by \$8 billion of additional public savings, leading to a net increase in government spending of \$20.5 billion between a scenario with observed premium subsidies and a scenario with no premium subsidies. From this perspective, subsidizing premiums generates negative economic value, as this extensive extra spending generates only \$14.8 billion in additional surplus, resulting in circa \$6 billion of deadweight loss, even when

argues that the structure of ACA subsidies discourages enrollment of the youngest consumers.

we do not take into account the cost of raising public funds.

These subsidies, however, significantly benefit consumers. In fact, we find that subsidies play the key role in stimulating enrollment in ACA Marketplaces. In the absence of any subsidies, we simulate (in Column 2) that 4.7 million, or 23 percent, of potential consumers would enroll in this market. With subsidies, enrollment increases to 9 million (45 percent), or nearly doubles. The increase in enrollment is particularly pronounced among consumers below 400 percent FPL, who effectively exit the market without subsidy support.³⁵ At the same time, insurance firms also enjoy a substantial benefit from subsidies. Out of \$14.8 billion in additional surplus, only 47 percent accrues to consumers and 53 percent is captured by firms. The less than 50 percent pass-through of subsidies to consumers suggest an important role of market power in this market (Pless and van Benthem, 2019). These estimates suggest that the competitiveness of the ACA market is similar to Medicare Advantage, for which Cabral et al. (2018) find 45 (in premiums only) to 54 (in premiums and benefits) percent pass-through rate to consumers.

The incidence of subsidies on consumer surplus is highly heterogeneous across demographic groups and locations. Four panels in Figure 4 illustrate the demographic and geographic distribution of the estimated gains in consumer surplus and insurance coverage under the observed means-tested subsidies relative to an equilibrium with no subsidies. As would be intended by the subsidy schedule, consumers with lowest incomes benefit the most, both in terms of gains in consumer surplus (Panel A) as well as in terms of gains in insurance coverage rates (Panel C), which increases by up to 80 percentage points for the lowest income consumers. Geographically, the highest changes in surplus (Panel B) and insurance rates (Panel D) occur in the Southeastern United States, including parts of Tennessee, North Carolina, Alabama, Georgia, and Florida. In many counties in these states, average consumer surplus increases by more than \$500 per capita, while insurance rates in ACA eligible population increase by more than 30 percentage points on average due to means-tested subsidies.

³⁵The 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.

5.2 Distortionary effects of market power

Before considering how market power may distort which consumers benefit from targeted transfers, we estimate the aggregate efficiency cost of market power in this market. Counterfactual simulations in Columns (1) and (4) of Table 3 compare equilibria with and without market power, keeping everything else about the environment, including means-tested transfers, fixed. In column (4), the allocation without market power is simulated by setting prices equal to average marginal cost. Removing market power increases consumer surplus by \$13 billion from the baseline of \$50 billion. Total surplus remains almost the same, implying that in the absence of market power \$13 billion of surplus is re-allocated from insurers to consumers. Average consumer surplus increases by \$652 or 26 percent. Baseline list premiums in the absence of market power are \$658 lower on average, which attracts more mid-income consumers into the market. Insurance coverage rate in the absence of market power is nearly 15 percentage points higher (45 versus 59 percent), implying that overall market power plays a central role in determining insurance coverage rates in ACA Marketplaces.

Another way of looking at the aggregate efficiency cost of market power is to consider how the presence of market power changes how much consumers gain on average when means-tested transfers are introduced. We start by comparing perfectly competitive allocations with no transfers (Column 5) and with means-tested transfers fixed at the observed level for each individual in our data (Column 4).³⁶ Introducing means-tested subsidies in the perfectly competitive environment increases average consumer surplus by \$613 (24 percent) and increases insurance enrollment rate by 25 percentage points (from 34 to 59). We now consider the gains from subsidies in the environment that is characterized by market power. We compute the gain in average consumer surplus and insurance coverage rates between Columns (3) and (1) in Table 3. In the presence of market power, keeping individual level maximum subsidies the same, average consumer surplus increases by a much smaller \$343 (16 percent). Aggregate insurance coverage rates under subsidies goes up by 21 percentage point (from 23 to 45). In sum, we find that market power substantially reduces the aggregate consumer gains from means-tested transfers.

We next consider whether market power differentially distorts which consumers benefit from targeted transfers due to the presence of demographic externality that we examined in Section 4.4. We again entertain two comparisons. We first compare the *levels* of con-

³⁶We keep the maximum possible transfer fixed in the data at the individual level. If consumers buy cheaper plans in the equilibrium with perfect competition, realized transfers could be lower, as we impose a zero lower bound on consumer-facing premiums following the ACA statues.

sumer surplus by consumer demographic groups between allocations with and without market power, keeping subsidies fixed. We then consider how market power changes how much consumers *gain* from subsidies in a world with and without market power.

We start by comparing the level of consumer surplus between: (i) the observed equilibrium with means-tested subsidies and market power, and (ii) a counterfactual equilibrium that holds the level of means-tested subsidies fixed, but shuts down the market power channel. The results are illustrated in Figure 5A. We find that market power leads to the largest relative decline in consumer surplus among poorest consumers: we estimate a decline in consumer surplus of 18 percent on average among consumers with income over 400% FPL and a decline of 46 percent among consumers with incomes between 150% and 200% FPL.

Our second approach measures the difference in the gains from subsidies in the presence and absence of market power. The first consider a scenario that assumes that means-tested subsidies are rolled out into a perfectly competitive market. Under this scenario, we measure the gains in consumer surplus and insurance rates across different consumer types relative to an environment without any subsidies. This variation reveals the distributional effects of transfers under perfect competition.³⁷ We then contrast this to the scenario where means-tested transfers are rolled out into an imperfectly competitive market. Comparing the gains from means-tested subsidies (relative to no subsidies) by different demographic groups between the competitive and imperfectly competitive scenarios allows us to assess how much market market power distorts the targeting of subsidies. Figure 5B reports the results of this “difference in gains” computation. We find that market power disproportionately harms the lowest-income and oldest consumers, with the gains in consumer surplus declining by circa \$1,000 per capita (or 46 percent) among consumers age 60 or older with income under 300% FPL. In other words, market power redistributes the marginal subsidy dollar away from the intended means-tested schedule.

Appendix Figure C.7 correlates the market-level average of this distortionary effect of market power with the number of insurers in the market, conditioning on the distribution of demographics in each market. Consistent with our theoretical argument in Section 2, we find that the distortionary effect of market power on the gains in consumer surplus from means-tested subsidies is more pronounced in markets with fewer insurers, and is especially large in monopoly markets.

³⁷If all consumer types had equal marginal costs, as is the case in most product markets, this incidence will be equal to the schedule of subsidies, subject to a zero-price lower bound.

5.3 Equity, Efficiency, and Redistributive Preferences

To further understand how subsidy targeting per se affects this market and whether alternative mechanisms could reduce its impact, we compute an equilibrium to a counterfactual environment with age-adjusted, but income-invariant (and hence universal) subsidies. We set the voucher level such that the total government spending remains the same as in the observed data. We perform this exercise under both market power and perfect competition.

Columns (4) and (6) of Table 3 report the characteristics of the respective allocations in the case without market power. 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 allowing them to vary by age). We find that a voucher of \$1,055 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.

Income-invariant subsidies are substantially more efficient. The same amount of total public spending delivers \$13 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 (35 percent of consumers in income bracket under 200% FPL versus 26 percent). That is, means-tested subsidies work as intended—they lower prices for the lowest-income consumers and thus attract them to the market. In the absence of explicit preferences for redistribution, however, this is not efficient, since lowest-income consumers exhibit the lowest valuation of the good. Thus, 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.

In Figure 6 we explicitly consider how high preferences for redistribution need to be for the means-tested subsidies to be considered more efficient than flat vouchers. We plot the ratio of average consumer surplus between the case with means-tested versus universal subsidies point by point for different values of λ that measure the preference for redistribution. The horizontal dashed line at 1 marks the value of λ at which the society is indifferent between the two subsidy regimes. The indifference λ is equal to 1.3. This implies that if the society values transferring \$1 to a household with income of \$17,820 (150% FPL in 2017 for a single person household) as much as or more than transferring to \$3.6 to a household with income of \$47,520 (400% FPL in 2017 for a single person household), then the preference for redistribution is high enough to prefer means-tested subsidies.

We repeat the same analysis for the situation that retains market power. We observe that in this case, means-tested subsidies generate a higher loss in total surplus. As Figure 6, this conclusion is true for any levels of preference for redistribution. In other words, market power exacerbates the efficiency losses from means-testing. Thus, in the presence of market power, the society needs to have a higher preference for redistribution to prefer means-tested transfers in the presence of imperfectly competitive intermediaries.

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 purchases of goods or services from private intermediaries that contract with the government. 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 observable characteristics of consumers has the potential to change the productive efficiency and the distributional effects of these transfers. The intuition is simple. Targeted transfers differentially alter demand across different types of consumers. Yet, since firms typically cannot price discriminate based on the same consumer characteristics, the price faced by one consumer type will depend on the composition of other consumer types in the market and the targeting schedule. We call this economic force a subsidy-induced “demographic externality.”

We examined this theory in the empirical context of targeted subsidies on ACA Marketplaces. These relatively new markets provide health insurance coverage for millions of individuals in the United States and may expand significantly in 2021 onward.³⁸ Our estimates suggest that market power leads to substantial efficiency distortions in this market. On aggregate, imperfectly competitive firms capture more than 50 percent of the surplus generated from public transfers. The impact of market power is differential across consumer types. Market power makes it harder for the policy-makers to achieve the distributional objectives, since in the presence of targeted transfers, market power redistributes marginal subsidy dollars away from the intended beneficiaries. We show that switching to subsidies

³⁸For example, the Biden administration has proposed significantly extending subsidy coverage to a broader group of consumers, which could lead to a doubling of federal expenditures (see, e.g., <https://www.kff.org/health-reform/issue-brief/affordability-in-the-aca-marketplace-under-a-proposal-like-joe-bidens-health-plan/>).

that do not vary with income can reduce the efficiency losses from market power. However, the choice between means-tested and universal subsidies faces a stark equity-efficiency tradeoff.

Overall, our results suggest that re-distributional policy tools that have a long history in direct public provision and have been frequently adopted one for one into environments with private provision should be used with caution in such environments. Market power of private intermediaries that contract with the government to provide publicly subsidized goods or services is likely to distort the ability to achieve the policy-makers' distributional objectives.

While it is infeasible to directly change the distribution of demographics within any given market, we note that regulating market boundaries serves much of the same purpose. In ACA Marketplaces, for instance, grouping together different markets into uniform rating areas, the regulator effectively changes the composition of a consumer's neighbors. This paper provides a mechanism to understand the equilibrium effects of different groupings; we leave a more detailed investigation of optimal rating area design to future work.

Acknowledgements

We are grateful to David Card, Keith Ericson, Mark Shepard, and Pietro Tebaldi for their discussions of the paper, and seminar participants at Daniel McFadden 80th Birthday Conference, Junior Health Economics Summit, Ohlstadt Workshop, Stanford University, London School of Economics, American Society of Health Economists, Yale University, University of Connecticut, Columbia University, 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, MIT, 2019 Penn State-Cornell Econometrics & IO Conference, Stanford University, University of Wisconsin Madison, 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.

References

Aizawa, Naoki, "Labor market sorting and health insurance system design," *Quantitative Economics*, 2019, *10* (4), 1401–1451.

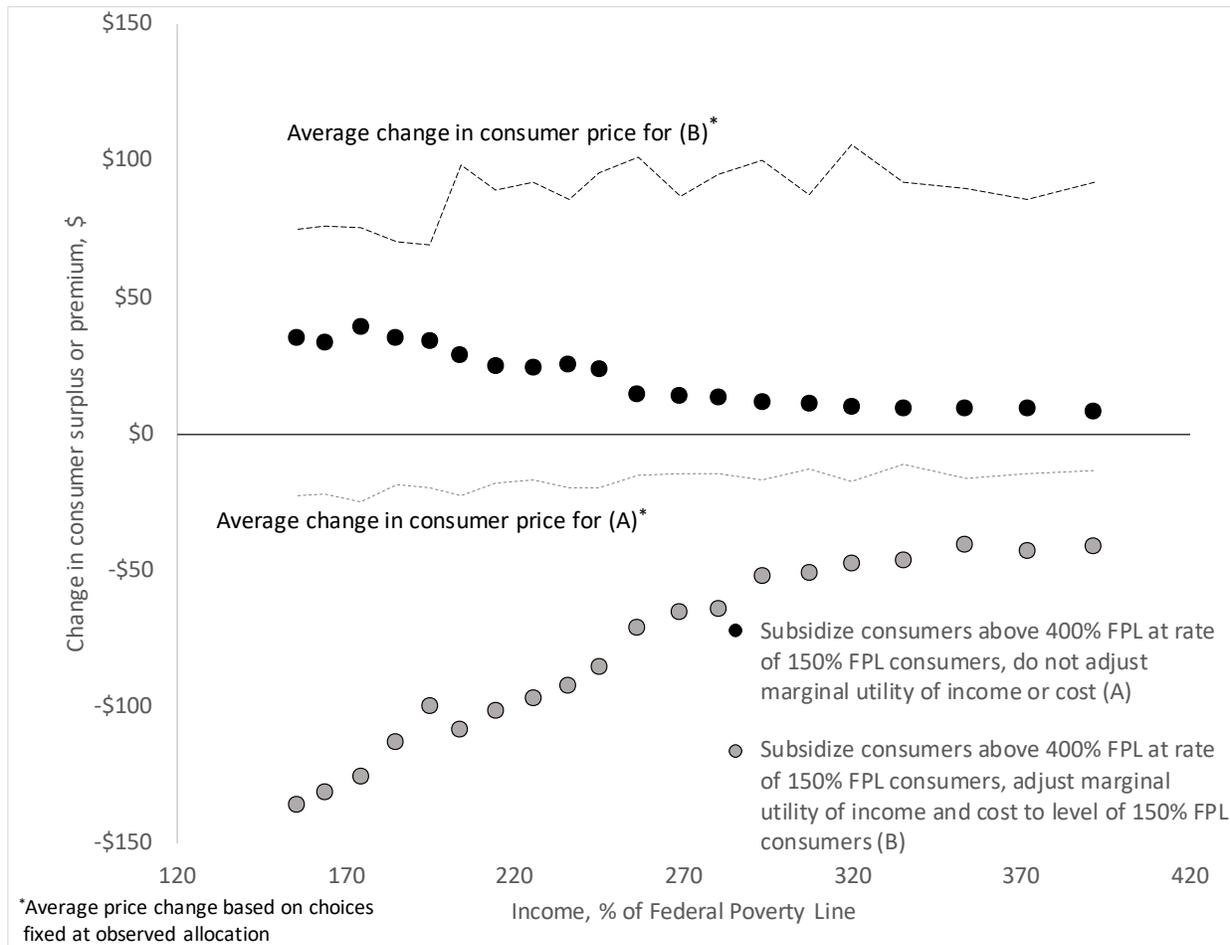
- **and Chao Fu**, “Interaction of the Labor Market and the Health Insurance System: Employer-Sponsored, Individual, and Public Insurance,” Technical Report, National Bureau of Economic Research 2020.
- **and Hanming Fang**, “Equilibrium labor market search and health insurance reform,” *Journal of Political Economy*, 2020, *128* (11), 000–000.
- Akerlof, G.**, “The Economics of “Tagging” as Applied to the Optimal Income Tax, Welfare Programs, and Manpower Planning,” *American Economic Review*, 1978, *68*(1), 8–19.
- Allcott, Hunt, Christopher Knittle, and Dmitry Taubinsky**, “Tagging and Targeting of Energy Efficient Subsidies,” *American Economic Review*, 2015, *105*(5), 187–191.
- Atkinson, Anthony B.**, “On the measurement of inequality,” *Journal of Economic Theory*, 1970, *2* (3), 244–263.
- Azevedo, Eduardo M and Daniel Gottlieb**, “Perfect competition in markets with adverse selection,” *Econometrica*, 2017, *85* (1), 67–105.
- 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 Student Aid Raise Tuition? New Evidence on For-Profit Colleges,” *American Economic Journal: Economic Policy*, 2014, *6*(4), 174–206.
- Chernozhukov, Victor and Han Hong**, “An MCMC approach to classical estimation,” *Journal of Econometrics*, 2003, *115* (2), 293–346.
- Cicala, Steve, Ethan MJ Lieber, and Victoria Marone**, “Regulating markups in us health insurance,” *American Economic Journal: Applied Economics*, 2019, *11* (4), 71–104.
- 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.

- Crawford, Gregory S, Nicola Pavanini, and Fabiano Schivardi**, “Asymmetric information and imperfect competition in lending markets,” *American Economic Review*, 2018, 108 (7), 1659–1701.
- 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.
- 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 P Ryan**, “Subsidy Design in Privately Provided Social Insurance: Lessons from Medicare Part D,” *Journal of Political Economy*, 2020, 128 (5), 1712–1752.
- 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*, 2015, 105 (5), 120–25.
- Einav, Liran, Amy Finkelstein, and Pietro Tebaldi**, “Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies,” *Unpublished mimeo, Stanford University, MIT, and University of Chicago*, 2019.
- Eriksen, Michael D. and Amanda Ross**, “Housing Vouchers and the Price of Rental Housing,” *American Economic Journal: Economic Policy*, August 2015, 7 (3), 154–76.
- 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, A. and K. McGarry**, “Multiple Dimensions of Private Information: Evidence from the Long-Term Care Insurance Market,” *American Economic Review*, 2006, 96(4), 938–958.
- 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.
- 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 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.
- Hendren, Nathaniel and Ben Sprung-Keyser**, “A Unified Welfare Analysis of Government Policies,” *Quarterly Journal of Economics*, 2020, 135 (3), 1209–1318. View the estimates online at www.policyinsights.org Watch the Econimate Video.
- , **Camille Landais, and Johannes Spinnewijn**, “Choice in Insurance Markets: A Pigouvian Approach to Social Insurance Design,” Technical Report, National Bureau of Economic Research 2020.
- 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
- Hylland, Aanund and Richard Zeckhauser**, “Distributional objectives should affect taxes but not program choice or design,” in “Measurement in Public Choice,” Springer, 1981, pp. 123–143.
- 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.
- 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.
- Mahoney, Neale and E Glen Weyl**, “Imperfect competition in selection markets,” *Review of Economics and Statistics*, 2017, 99 (4), 637–651.
- Marmot, Michael**, “The health gap: the challenge of an unequal world,” *The Lancet*, 2015, 386 (10011), 2442–2444.
- Meckel, Katherine**, “Is the Cure Worse than the Disease? Unintended Effects of Payment Reform in a Quantity-based Transfer Program,” *Unpublished mimeo*, 2019.
- , **Maya Rossin-Slater, and Lindsey Uniat**, “Efficiency Versus Equity in the Provision of In-Kind Benefits: Evidence from Cost Containment in the California WIC Program,” *NBER Working Paper*, 2020, (w26718).
- Miller, Keaton S, Amil Petrin, Robert Town, and Michael Chernew**, “Optimal Managed Competition Subsidies,” Working Paper 25616, National Bureau of Economic Research February 2019.
- Nevo, A.**, “Measuring Market Power in the Ready-to-Eat Cereal Industry,” *Econometrica*, 2001, 69, 307–342.

- Nichols, Albert L and Richard J Zeckhauser**, “Targeting transfers through restrictions on recipients,” *The American Economic Review*, 1982, 72 (2), 372–377.
- Pless, Jacquelyn and Arthur A. van Benthem**, “Pass-Through as a Test for Market Power: An Application to Solar Subsidies,” *American Economic Journal: Applied Economics*, October 2019, 11 (4), 367–401.
- 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
- Rosen, Harvey S**, “Housing subsidies: Effects on housing decisions, efficiency, and equity,” in “Handbook of Public Economics,” Vol. 1, Elsevier, 1985, pp. 375–420.
- Rothschild, Michael and Joseph Stiglitz**, “Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information,” *The Quarterly Journal of Economics*, 1976, 90 (4), 629–649.
- Rothstein, Jesse**, “Is the EITC as Good as a NIT? Conditional Cash Transfers and Tax Incidence,” *American Economic Journal: Economic Policy*, 2010, 2(1), 177–208.
- 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.
- Shepard, Mark**, “Hospital network competition and adverse selection: evidence from the Massachusetts health insurance exchange,” Technical Report, National Bureau of Economic Research 2016.
- Small, Kenneth A. and Harvey S. Rosen**, “Applied Welfare Economics with Discrete Choice Models,” *Econometrica*, 1981, 49(1), 105–130.
- Smeeding, Timothy**, “Approaches to measuring and valuing in-kind subsidies and the distribution of their benefits,” in “Economic transfers in the United States,” University of Chicago Press, 1984, pp. 139–176.
- Stole, Lars A**, “Price discrimination and competition,” *Handbook of industrial organization*, 2007, 3, 2221–2299.
- 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.
- Waldinger, Daniel**, “Targeting in-kind transfers through market design: A revealed preference analysis of public housing allocation,” Technical Report, Working paper 2018.
- 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.

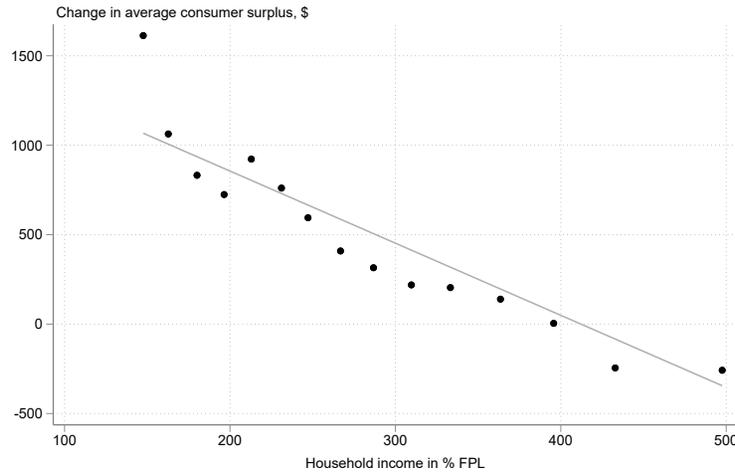
Figure 2: Demographic externality with targeted subsidies and market power



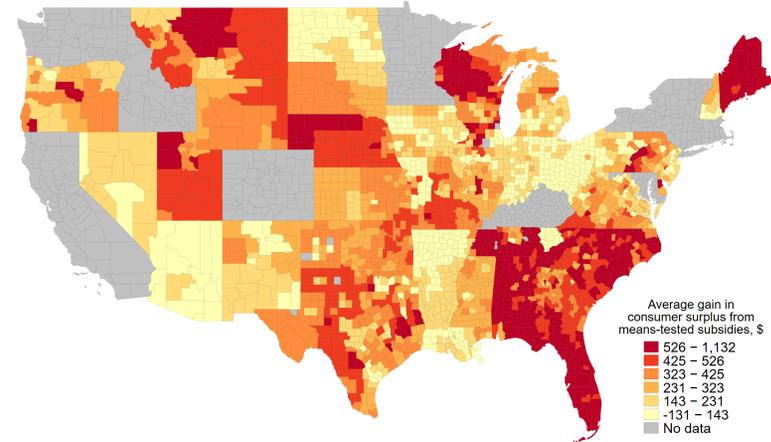
Notes: Figure reports estimated change in consumer surplus and consumer-facing prices (y-axis, in dollars) by income level (x-axis) in two counterfactual cases that capture the “demographic externality”. The first counterfactual simulation (case A) changes income of consumers with true income of above 400 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: Effect of ACA Marketplace subsidies on welfare and insurance rates

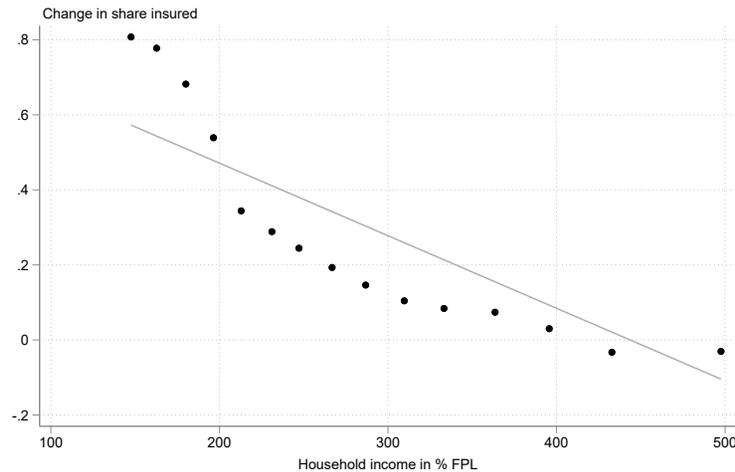
(A) Gain in consumer surplus, by income



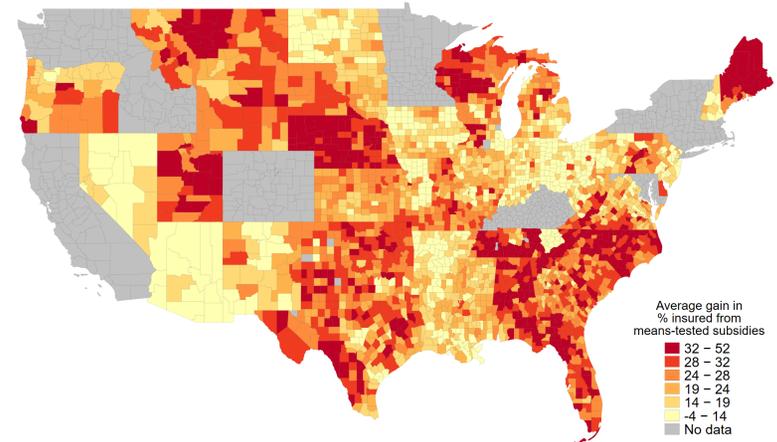
(B) Gain in consumer surplus, by geography



(C) Gain in share insured, by income



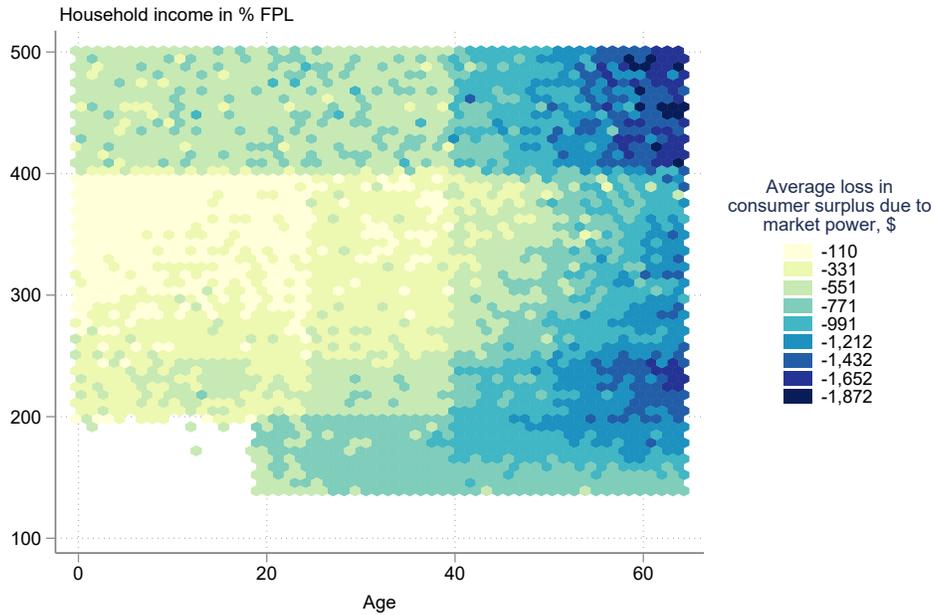
(D) Gain in share insured, by geography



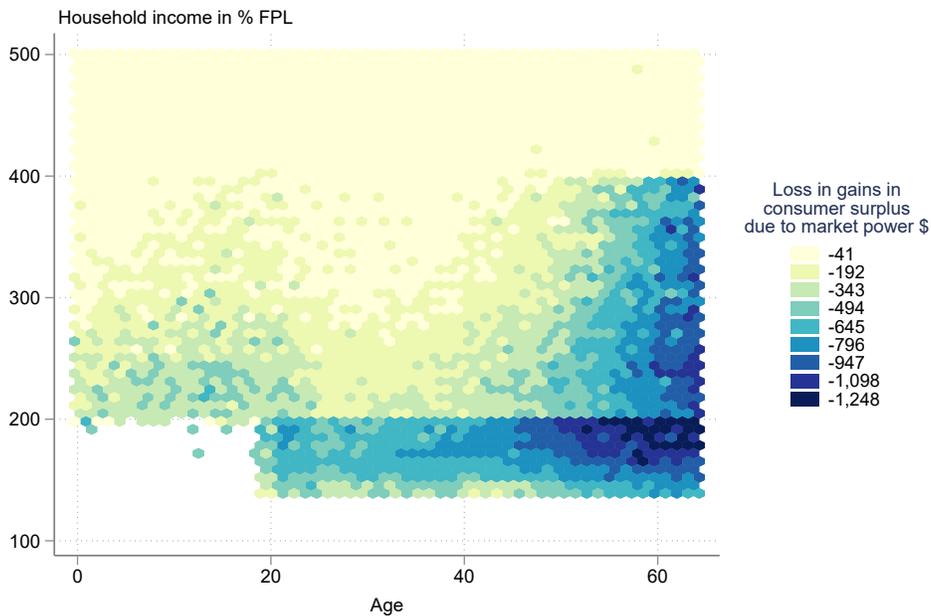
Notes: The figure reports the change in average consumer surplus by income (Panel A) or county (Panel B) and the change in the rate of insurance coverage by income (Panel C) or county (Panel D) between an equilibrium with observed means-tested subsidies and an equilibrium without subsidies, all else equal. Market power is kept constant. For each income cell the average is computed as a weighted average of compensating variation in ACS consumer sample in that income cell; each individual within the corresponding income cell receives a weight that is equal to the potential ACA market size in the individual's county.

Figure 5: ACA Marketplace subsidies and market power

(A) Loss in consumer surplus levels, perfect vs. imperfect competition

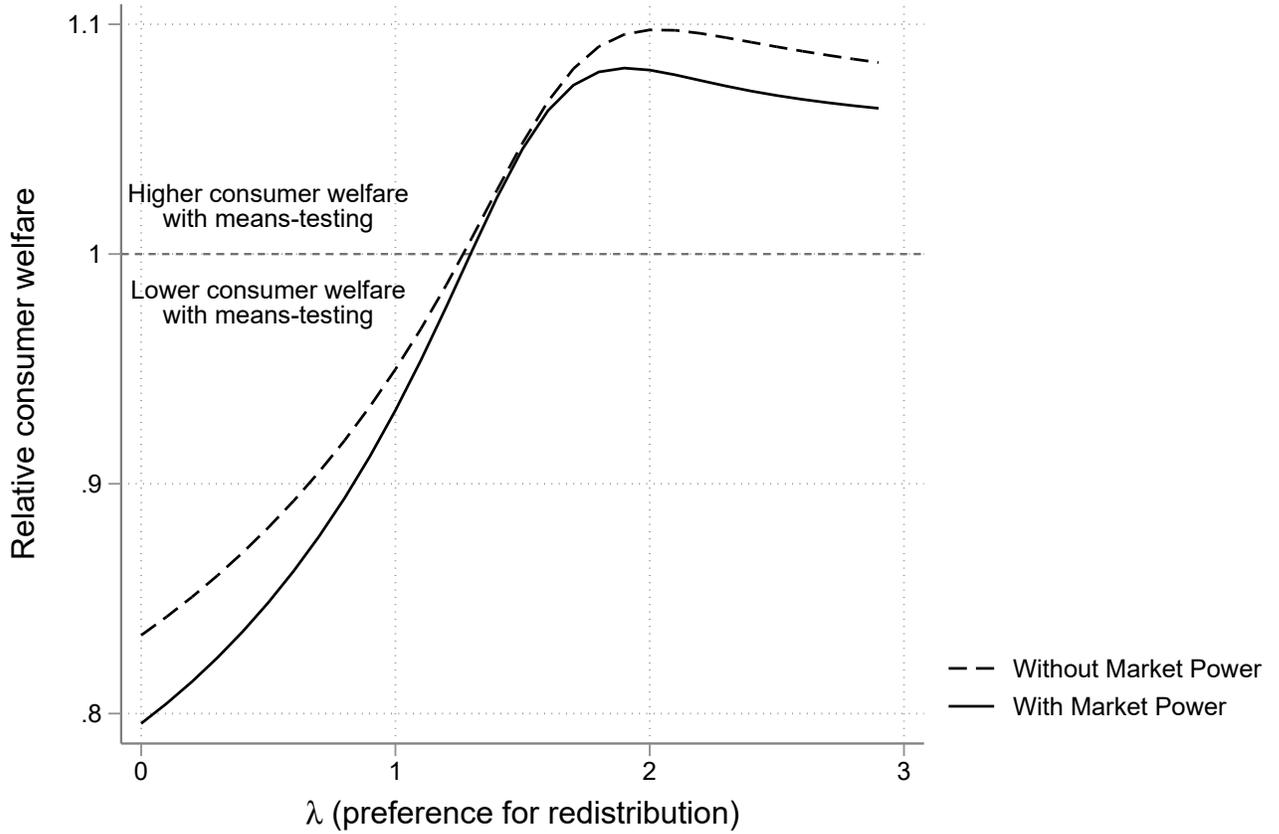


(B) Loss in consumer surplus gains, perfect vs. imperfect competition



Notes: The figure reports the simulated effect of market power on average consumer surplus on ACA Marketplaces by age and income. Panel A shows the difference in the levels of average consumer surplus between the observed allocation and a counterfactual that removes market power, but keeps subsidies fixed. Panel B shows the difference in gains in consumer surplus between a counterfactual with no subsidies versus means-tested subsidies, computed first under perfect and then, separately, under imperfect competition. The average within a demographic cell is computed as a weighted average of compensating variation in ACS consumer sample in that demographic cell; each individual within the demographic cells receives a weight that is equal to the potential ACA market size in the individual's county.

Figure 6: Equity-efficiency trade-off in subsidy design



Notes: Figure reports the point-wise (for each value of parameter λ) ratio of average consumer surplus between an equilibrium with means-tested subsidies and flat subsidies, keeping total net government spending fixed. The blue line measures this ratio for the case with perfect competition. The red line measures this ratio for the case with imperfect competition. Average consumer surplus for each level of λ is computed as specified in Equation 28 in Section 4.3. At $\lambda=0$ (social welfare weights are the same for every household), the ratio of average consumer surplus equals to the ratio of entries in line (26) of Table 3. At higher levels of λ , social welfare weights increase for consumers with lower income.

Table 1: Summary statistics

	Mean [‡]	Std.Dev.	10th pctile	90th pctile
	(1)	(2)	(3)	(4)
A. Choice set				
(1) Number of plans	21	13	8	37
(2) Number of large 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 in 60% actuarial value plans	0.09	0.05	0.04	0.2
(4) Share in 70% actuarial value plans	0.3	0.1	0.2	0.4
(5) Share in 80% actuarial value 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	39	2.3	36	42
(2) Income in % FPL	295	52	231	365
(3) Annual premium subsidy, \$ ^{^^}	2,349	1,244	919	4,226
(4) Share income under 200% FPL	0.32	0.14	0.13	0.50
(5) Share income 200-400% FPL	0.37	0.07	0.28	0.47
(6) Share income 400% FPL and above	0.31	0.11	0.19	0.46
(7) Share age under 25	0.22	0.05	0.16	0.28
(8) Share age 25-40	0.30	0.05	0.24	0.36
(9) Share age 40 and above	0.48	0.07	0.41	0.57

[‡] Across counties. One market is defined as one county. There are 2,561 counties in the data. For Panel C all statistics are weighted by total market size

^{‡‡} 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 (**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 report having employer-sponsored insurance, or any type of public health insurance coverage, and were not eligible for insurance under Medicaid expansion in those states that expanded Medicaid based on their income. Annual premium subsidies were computed using the ACS records of income relative to the Federal Poverty Line and tax family composition following instructions for 2017 IRS Form 8962 (Premium Tax Credit).

Table 2: Model estimates

	Mean (1)	Age<25 (2)	Age 25-40 (3)	Age >40 (4)
A. Demand: parameters of utility function				
Coefficient on premium, \$000 (α)				
Income <200% FPL	-	-5.17 (0.33)	-2.47 (0.16)	-2.21 (0.14)
Income > 200% FPL and < 400% FPL	-	-4.32 (0.27)	-0.64 (0.04)	-3.94 (0.26)
Income > 400% FPL	-	-1.13 (0.07)	-0.20 (0.01)	-0.46 (0.04)
Age-specific intercepts	-	1.52 (0.10)	-1.72 (0.11)	base
Actuarial Value	26.83 (1.69)			
B. Supply: inversion of first-order conditions				
	Mean (5)	Std. dev. (6)	Max (7)	Min (8)
Marginal cost for a 20 year old with income <200% FPL, \$	1,561 [^]	457 [^]	732 [^]	4,102 [^]
60% actuarial value plans	1,332	265	747	2,710
70% actuarial value plans	1,506	368	732	3,268
80% actuarial value plans	2,137	467	1173	4,102
Estimated cost multipliers [‡]				
Income <200% FPL	2.77 ^{**}			
Income > 200% FPL and < 400% FPL	2.15 ^{**}			
Income > 400% FPL	1.96 ^{**}			

[^] Average across all plans after averaging within plans across rating areas

[‡] Income categories in OptumTM demographic data do not fully correspond to our FPL classifications, as we do not observe continuous income records and do not observe the family structure. We assign individuals with income reported to be <\$40,000 to "<200% FPL" category; individuals with income over \$40,000 but under \$75,000 to "> 200% FPL and < 400% FPL" category; and individuals with income over \$75,000 to "> 400% FPL" category. All multipliers are computed relative to 20 years olds in the lowest income category.

^{**} Average across 65 age categories, age 0 to 64

Notes: Panel A reports non-linear least squares parameter estimates for demand model described in Section 4.1. The NLLS objective function minimizes the squared distance between estimated and observed age- and income-specific enrollment shares in each market (county). The model includes, but we do not report, intercepts for each insurance plan, as well as an intercept for consumers with income below 100% FPL. Bootstrapped standard errors are reported in parentheses. Panel B reports inversion of the first-order conditions as well as estimates of cost multipliers by income-age from commercial claims data.

Table 3: Equilibria under counterfactual scenarios

	Preserve market power			No market power			Demographic externality	
	Baseline [^]	No premium subsidies	Flat subsidies with net G ≈ net G in (1) (\$1,125 subsidy for a 20 year-old) [~]	Keep premium subsidies as in (1) - targeted subsidies	No premium subsidies	Flat subsidies with net G ≈ net G in (4) (\$1,055 subsidy for a 20 year-old) [~]	Subsidize consumers above 400% FPL at rate of 150% FPL consumers	(7) + change α_{400} to α_{150} + change cost multiplier to k^{200FPL}
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
(1) Consumer surplus (CS), \$M	50,268	43,355	63,233	63,403	51,059	76,081	62,847	22,611
(2) Insurer profit (Π), \$M	14,683	6,802	15,355	-	-	-	17,208	22,437
(3) Consumer and producer surplus, \$M	64,951	50,157	78,588	63,403	51,059	76,081	80,055	45,049
(4) Public spending on premium and cost-sharing subsidies (G), \$M	28,882	454	32,424	35,754	1,391	39,280	44,095	51,959
(5) Public savings on uncompensated care for un(under-)insured, \$M	16,510	8,625	20,134	21,701	12,437	25,158	18,918	19,372
(6) Net government spending, \$M	12,372	(8,171)	12,289	14,053	(11,046)	14,122	25,177	32,587
(7) Total welfare, including the cost of public funds, \$M ^{^^}	48,867	60,779	62,611	45,133	65,419	57,723	47,325	2,685
(8) CS+ Π per dollar of nominal public spending, \$	2.25	-	2.42	1.77	-	1.94	1.82	0.87
(9) CS+ Π on a dollar of net public spending \$	5.25	-	6.39	4.51	-	5.39	3.18	1.38
(10) Characteristics of the allocation								
(11) Inside option enrollment, in 1,000s	9,037	4,721	11,020	11,878	6,807	13,770	10,355	10,603
(12) Inside option enrollment, %	45%	23%	55%	59%	34%	68%	51%	53%
(13) Composition of inside option enrollment by consumer type								
(14) Income under 200% FPL	40%	6%	22%	35%	14%	26%	35%	31%
(15) Income 200-400% FPL	25%	17%	33%	32%	23%	37%	22%	19%
(16) Income 400% FPL and above	35%	77%	44%	33%	63%	36%	43%	50%
(17) Age under 25	18%	21%	25%	20%	21%	25%	19%	19%
(18) Age 25-40	26%	25%	26%	26%	27%	26%	25%	26%
(19) Age 40 and above	56%	54%	49%	54%	52%	49%	56%	55%
(20) Average income of consumers in inside option, % FPL	296	427	343	299	388	323	319	338
(21) Average age of consumers in inside option	41	39	38	40	39	38	41	41
(22) Average list premium among consumers in inside option, \$	5,618	4,788	4,676	4,044	3,426	3,599	5,667	5,760
(23) Average list premium for a 20 year old consumer, \$	2,401	2,239	2,301	1,743	1,592	1,653	2,382	2,492
(24) Average subsidy paid out to consumers in inside option, \$	3,196	96	2,942	3,010	204	2,853	4,258	4,900
(25) Average cost of coverage among consumers in inside option, \$	3,993	3,348	3,283	4,045	3,425	3,599	4,005	3,644
(26) Average consumer surplus, \$	2,495	2,152	3,139	3,147	2,534	3,776	3,120	1,122

[^] Observed equilibrium is recomputed in the model

[~] Subsidies are constant within an age group across income levels, but not across ages. Subsidy for other ages is scaled following the same regulatory age curve as the premiums

^{^^} The cost of public funds λ is assumed to be 30 cents for a dollar of public spending

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 that allow simulating 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 nationally using county market size as weights). Consumer surplus is computed as a compensating variation (see Section 4). Firm profits reported in row (2) indirectly account for risk-equalization programs, as marginal cost estimates are 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 200% 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 prior to ACA Marketplace rollout. Negative quantities throughout are reported in parentheses.

ONLINE APPENDIX

A Computational details

To estimate demand, we follow a multi-step procedure:

1. Start with an initial guess, $\theta = 0$.
2. Given θ , solve for δ such that plan-level enrollments match simulated enrollments. This requires simulating the plan choice of each family in the ACS sample and aggregating together market-level enrollments up to the plan level.
2. At θ and δ , compute enrollments by age and income and aggregate to the seven income intervals (e.g. 300 to 400 FPL) and seven age intervals (e.g. age 35 to 44) reported by CMS in the enrollment data. Calculate enrollments at the market level by metal level tier.
3. Compute squared error of predicted inside share of enrollment by age and income interval for each market and sum over all markets. Compute squared error of predicted plan metal level tier by market and sum over all markets and add to prior sum.
4. Update θ according to nonlinear optimizer. We first use a derivative-free Laplace-type estimator (LTE) from [Chernozhukov and Hong \(2003\)](#) before turning to a Newton's method style optimizer (KNITRO) when the LTE has converged to the neighborhood of the solution.
5. Repeat from point 2 until convergence criterion is met. The LTE is run 500 times, which we found sufficient to locate the local neighborhood of the parameter vector. KNITRO uses more sophisticated measures for determining convergence. Repeated runs of this approach produced identical answers.

B Construction of cost multipliers

We construct a matrix of cost-multipliers κ^d using commercial claims data from the Optum database. These data cover individuals enrolled in employer-sponsored insurance administered by a large national US insurer. We compute total claim amounts (amounts charged) at the individual level across inpatient, outpatient, and drug claims for year 2016, which

precedes the year of our analysis. For each individual, Optum reports basic socio-economic variables. We use the year of birth variable that gives us age and household income bracket. Income brackets in Optum do not exactly correspond to the income brackets that are used by CMS and that we use in demand estimation. We make the following approximation. We map Optum income bracket “under \$40,000” annual income to “<200% FPL” category; Optum brackets \$40K-\$49K, \$50K-\$59K, and \$60K-\$74K, to “over 200% FPL and under 400% FPL” category; and finally, \$75K and over Optum bracket into “400% FPL and above” category.

We collapse the individual data to the age-income bracket level, computing the mean medical cost for each age in three income brackets. We then normalize each age-income bracket estimate of the average medical cost to that among 20 year old individuals in the lowest income bracket. Denote the relative cost for each demographic group with RC . RC is a matrix with three columns for each income bracket and 65 rows for each age from 0 to 64. For each of the three income brackets separately we fit the following regression:

$$\ln(RC_a) = \alpha + \beta a + \gamma 1\{a = 0\} + \epsilon_a \quad (29)$$

The regression fits smooth exponential cost curves in age a , allowing for a separate intercept for newborn children, who typically have higher costs relative to what the data would predict for their age from charges during birth. We then use the coefficient estimates from Equation 29 for each income bracket to predict age-specific total costs. Re-normalizing these predictions to age 20 lowest income demographic bracket, gives us the matrix of κ^d that are shown in Figure B.1.

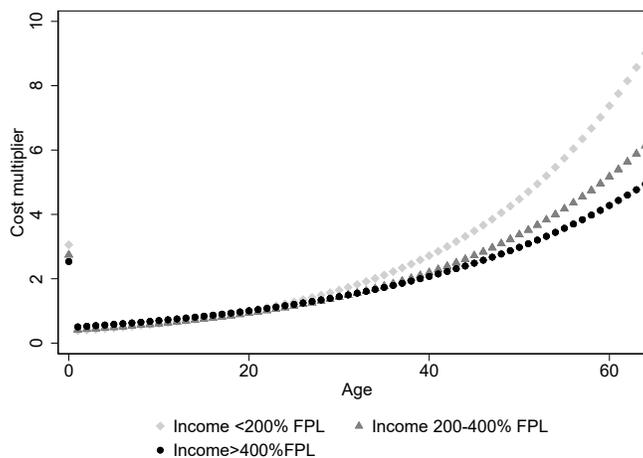
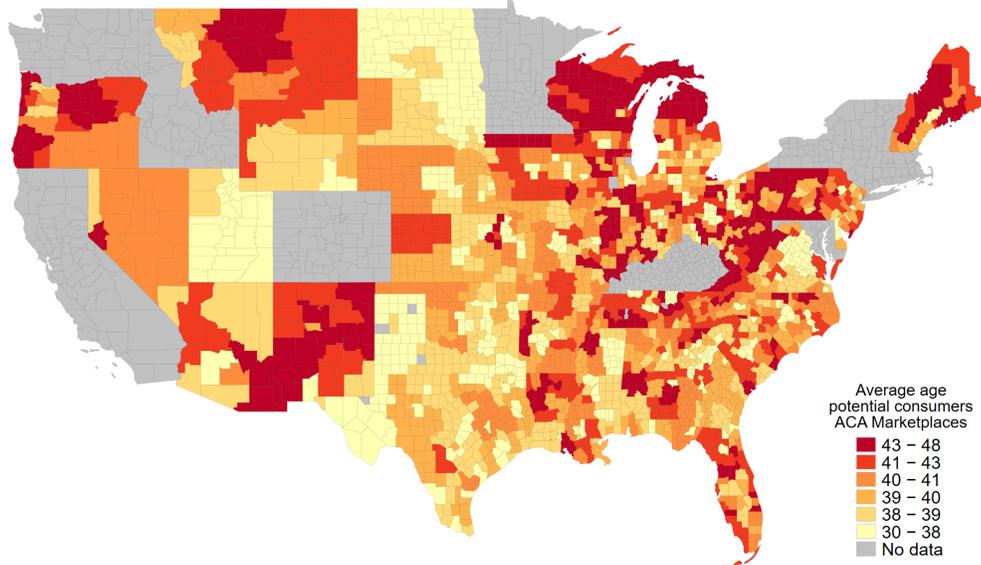


Figure B.1: Matrix of cost multipliers κ^d

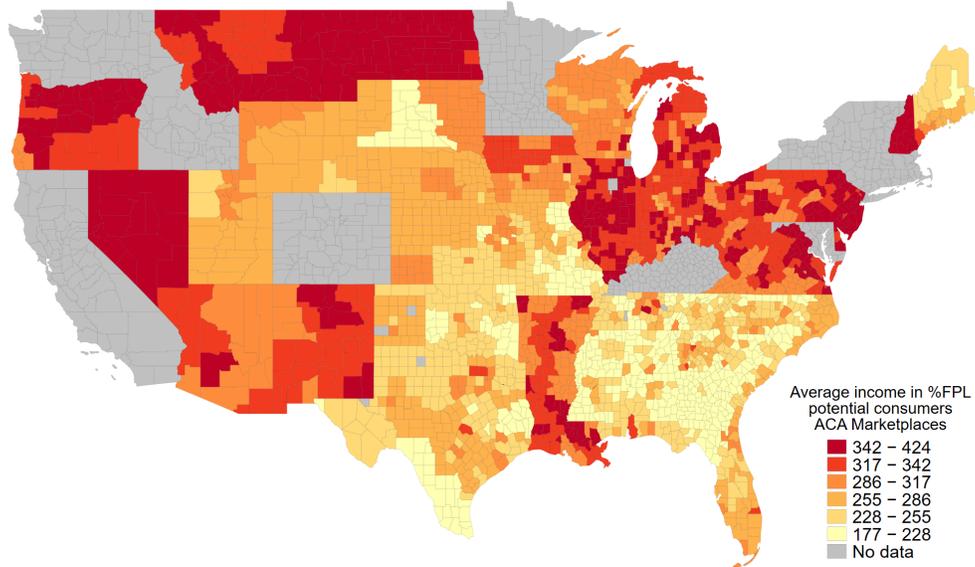
C Additional figures and tables

Figure C.2: Demographics of potential ACA Marketplace consumers

(A) Average age of potential consumers



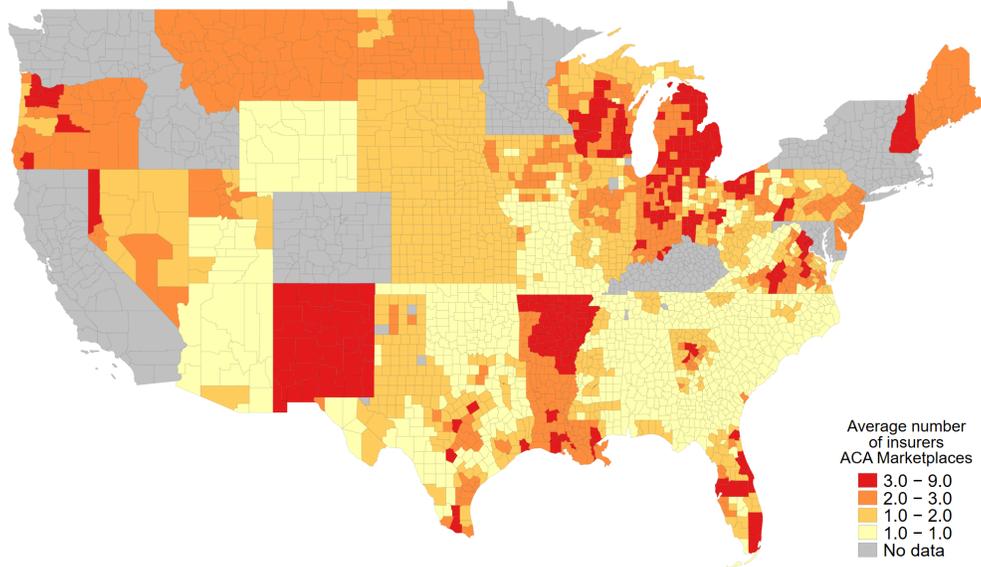
(B) Average income of potential consumers



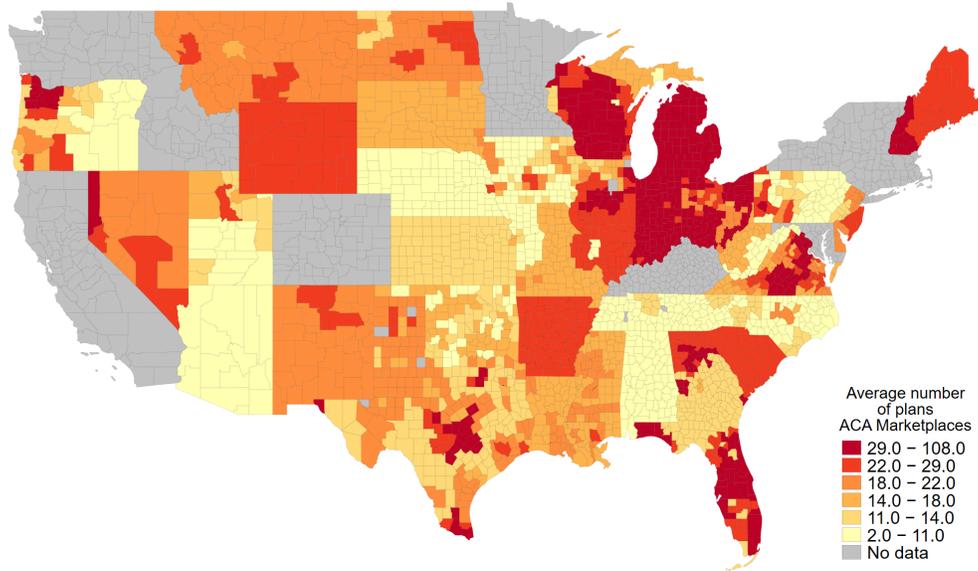
Notes: Maps plot the average age and income in % of Federal Poverty Line of potential ACA Marketplace consumers based on ACS 2017 sample as described in Section 3. Potential consumers in ACS are identified as consumers who have no employer-sponsored or public health insurance coverage and were not eligible for Medicaid expansions in states that expanded Medicaid.

Figure C.3: ACA Marketplace market structure

(A) Average number of insurers



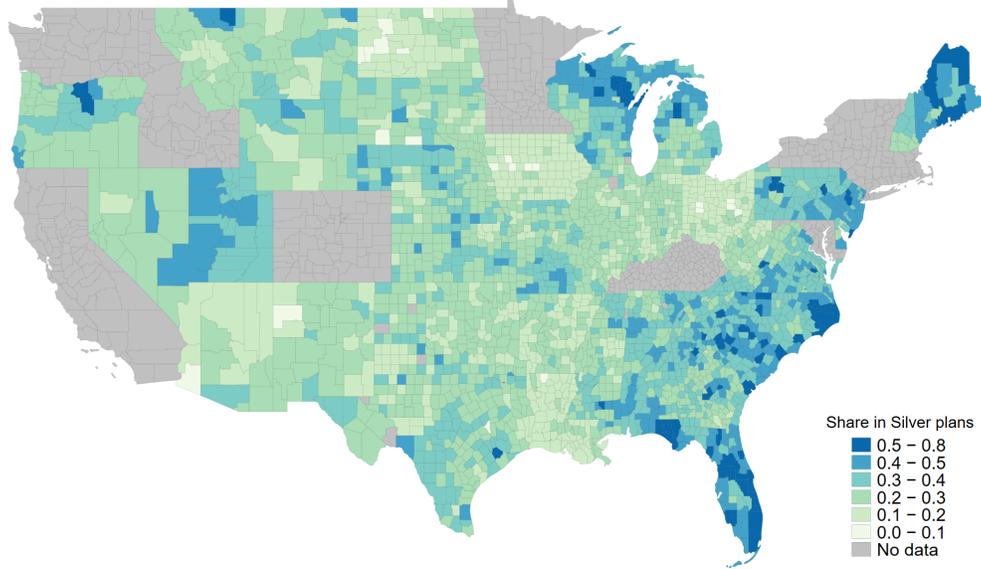
(B) Average number of plans



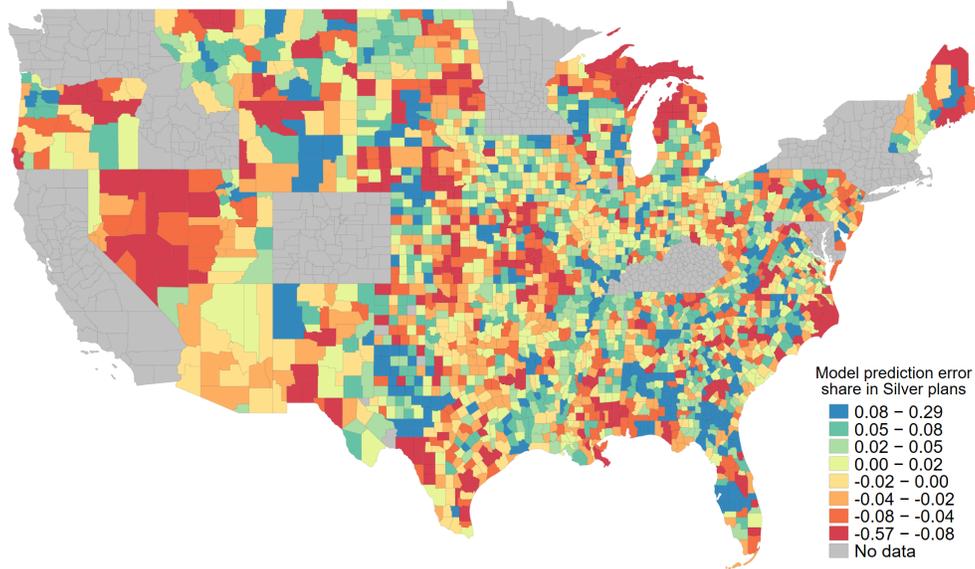
Notes: Maps plot the average number of insurers and plans offered in each county on federally facilitated ACA Marketplaces in 2017. An insurer is defined as a unique “issuer id.” A plan is defined using unique plan ids.

Figure C.4: Empirical moments and demand model fit for Silver plans

(A) Observed Silver plan market share

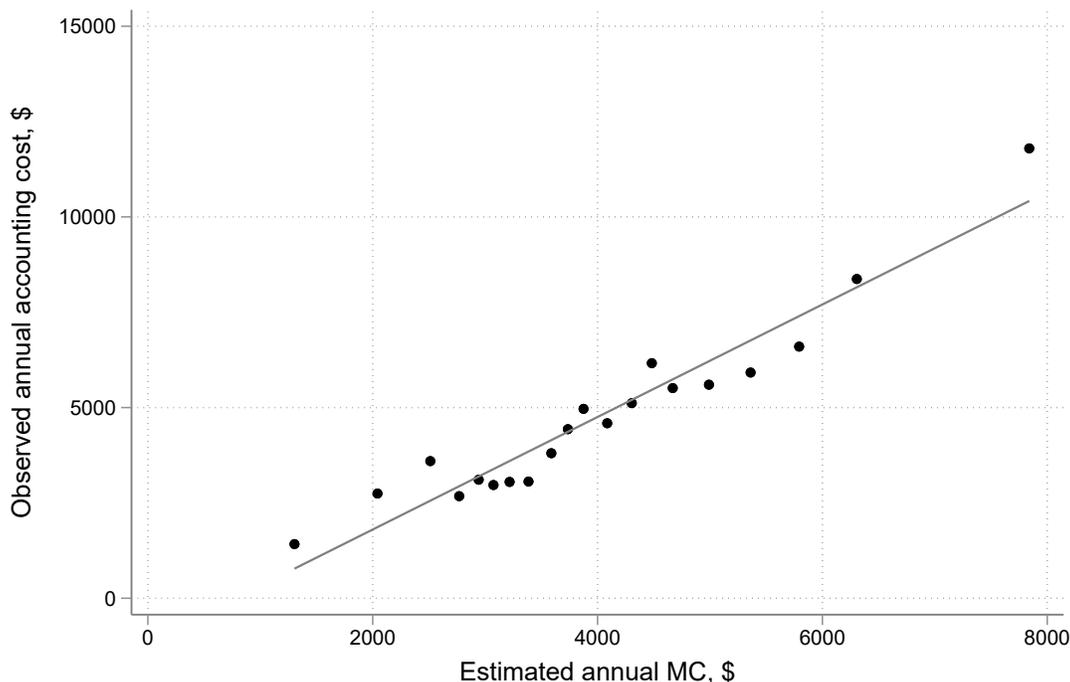


(B) Prediction error of Silver plan market share



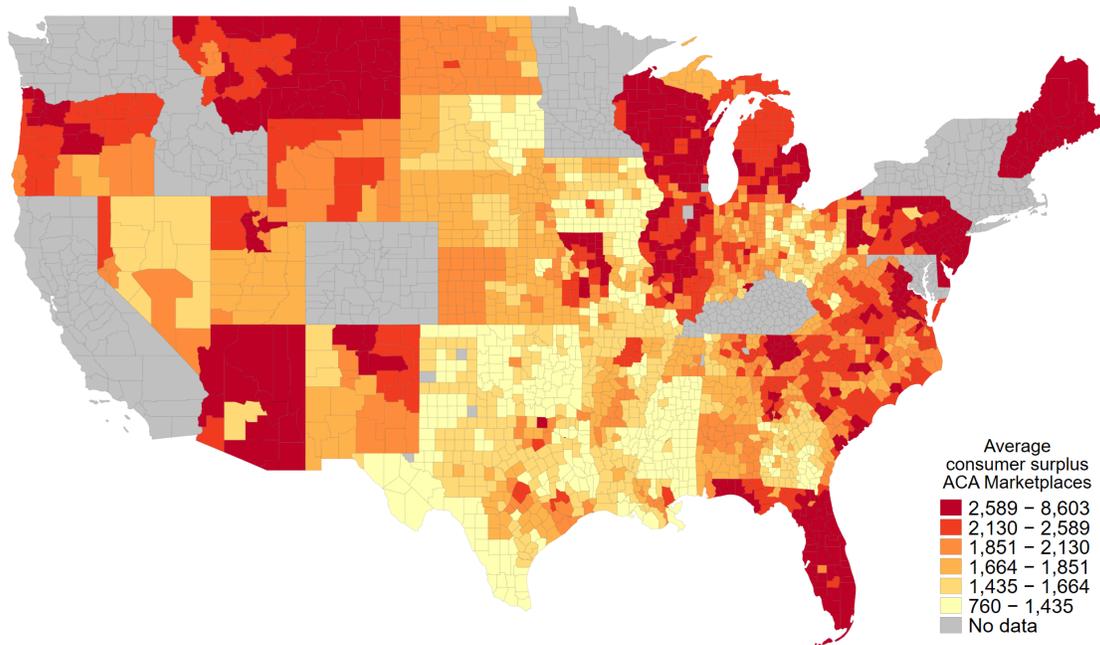
Notes: Map in Panel C.4A plots the share of potential consumers in each county that enrolled in a silver plan on ACA Marketplaces. States that are marked with grey are not federally facilitated and do not enter our analysis. The counts of the pool of potential consumers (denominator) were provided by the Kaiser Family Foundation and are 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 C.4B plots the difference between the observed share of enrollees in Silver plans and the share of enrollment in Silver plans as predicted by demand model of Section 4.1.

Figure C.5: Marginal cost estimates and accounting costs



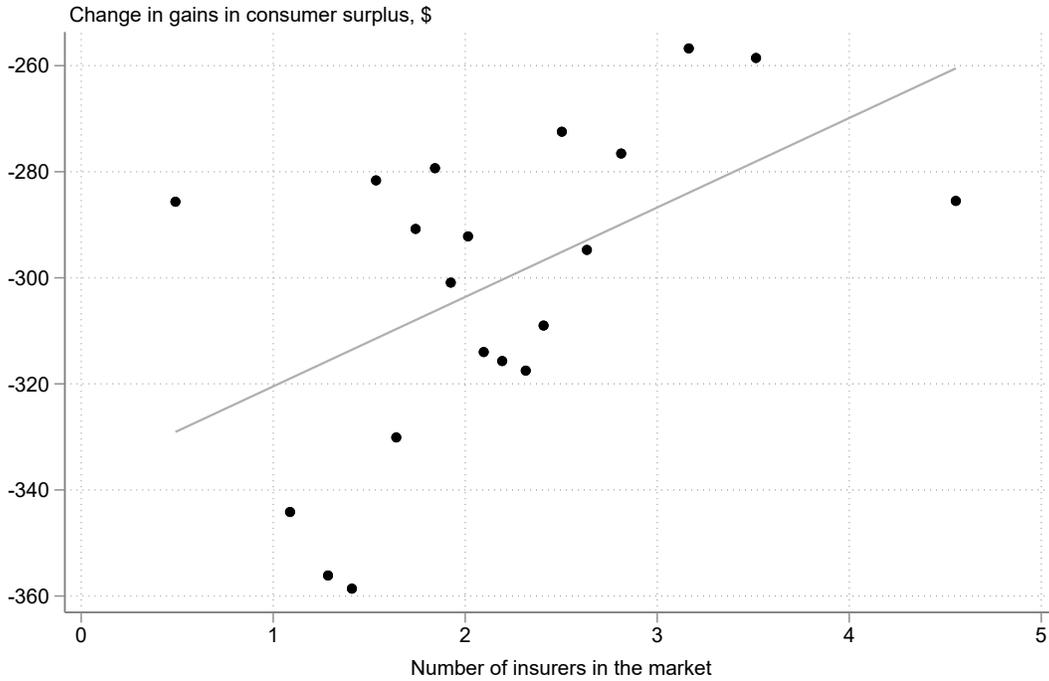
Notes: Figure plots the average accounting cost reported by plans by ventiles of plans’ estimated marginal costs. The underlying data includes only 705 plans for which accounting cost reports were available. Estimated marginal cost at the plan level was computed in several steps. First, we invert the firms’ first order condition as described in Section 4.2 to estimate the baseline marginal cost for a 20 year old, lowest-income category consumer. Second, for each consumer in the ACS sample, we compute plan-specific margin cost using income and age information together with income-age specific cost multipliers κ^d . Third, plan-level marginal cost is then computed by averaging across all consumer in the ACS sample who are predicted to enroll in each plan of interest in our demand model simulation. Accounting cost is computed based on public use rate review files released annually by CMS and CCIO. We pool information from rate review files released in 2017, 2018, and 2019, as all of them contain information about realized accounting costs for 2017. We use “experience incurred claims per member per month” as a measure of accounting cost. We multiple this cost by 12 to arrive at the average annual incurred cost for each plan. Incurred cost are defined by CMS as “cost of service paid by insurer,” and thus excludes patient cost-sharing payments.

Figure C.6: Consumer surplus under observed ACA Marketplace subsidies



Notes: Figure plots the average consumer surplus (in \$) from ACA Marketplace program that we estimate among potential consumers in each county in federally-facilitated ACA states. Subsidies and resulting consumer-facing prices are kept at levels implied by ACA Marketplace rules as implemented in 2017 (see Section 3 for details of subsidy computation. The distribution of potential consumer age and incomes is taken as observed in 2017 ACS. Average consumer surplus is computed using Equation 27.

Figure C.7: Gains in consumer surplus with and without market power, by number of insurers in the market



Notes: Figure plots the market-level average difference in consumer surplus gains from means-tested subsidies between an environment with perfect competition and an environment that preserve the observed degree of market power on ACA Marketplaces. The computation follows the same algorithm as in Figure 5B in the main text, but rather than reporting the national average difference in gain by income and age, we compute the average at the market (i.e. county) level. We then estimate a county-level regression of average change in consumer surplus gain by county on the number of insurers in the county. The x and y axes are both residualized to control for average income, average income, share of income and age in each of three income and age categories as specified in demand estimation, average share female, and average share of white individuals in each county.