Market Power and Redistribution: Evidence from the Affordable Care Act

Maria Polyakova* Stephen P. Ryan†

August 26, 2022

Abstract

Regulatory oversight of market power has traditionally focused on losses in aggregate consumer surplus and not on the distribution of losses across consumers. This paper quantifies both the average and the distribution of efficiency distortions due to market power in the context of the Affordable Care Act's federal health insurance Marketplaces. Using an equilibrium model of supply and demand for insurance plans, we estimate that market power leads to 21% lower surplus and 15pp lower rate of insurance for the average consumer. 47% of the subsidy-generate surplus accrues to consumers, while 53% is captured by firms. The loss in consumer surplus and in insurance coverage is not uniform, disproportionately affecting lower-income consumers. We show that the means-tested design of product subsidies *exacerbates* these distortions, as it reduces elasticity of demand on the margin among more price-elastic, higher marginal cost consumers. Our results caution for the use of means-tested transfers as policy tools in environments with imperfect competition and heterogeneous consumers. We conclude with a brief discussion of what our findings imply for antitrust regulation with redistributional considerations.

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

*Department of Health Policy, Stanford University, Stanford, CA, 94305; NBER; and CESifo (maria.polyakova@stanford.edu).

†Olin Business School, Washington University in St. Louis, St. Louis, MO 63130; NBER; and CESifo (stephen.p.ryan@wustl.edu).

1 Introduction

Governments utilize a variety of policy tools targeted at specific segments of their population in pursuit of redistributional objectives. While redistribution is possible via transfers in the tax system, targeted in-kind subsidies are also a popular policy choice.[1] Housing assistance, food stamps, health insurance, free school lunch, and energy assistance are all examples of in-kind government transfers targeted to recipients based on observable characteristics such as income, age, employment, family, or disability status (Akerlof, 1978; Smeeding, 1984; Currie and Gahvari, 2008; Hoynes and Schanzenbach, 2016). The literature studying the efficiency and distributional implications of such transfers has focused almost exclusively 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.[2] In this paper, we highlight the importance of accounting for a non-competitive supply side in such settings: strategic firms, aware of both the targeted subsidy schedule and distribution of consumer types, may set prices that result in an equilibrium incidence of benefits distorted away from the "intended" distribution embedded in the subsidy schedule. As a consequence, the efficiency benefits of using private intermediaries may be in direct conflict with the government's redistributional objectives. To help advance our understanding of the role that market power plays in settings where subsidies are set with distributional objectives in mind, we address three key interrelated questions: how does imperfect competition change the distributional outcomes of targeted subsidies? What are the efficiency implications of the exercise of market power? And finally, combining the answers to the prior two questions, how do governmental preferences for the trade-off between equity and efficiency inform the choice of whether subsidies should be targeted?

We take a two-pronged approach to answering these questions. First, we combine insights from both public economics and industrial organization to outline a set of fundamental eco-

[1] The literature has offered several explanations for why in-kind transfers may be preferred to cash transfers, such as restricting increased consumption to the desired good or creating better self-selection incentives (Nichols and Zeckhauser, 1982; Besley et al., 1990). Recently, Cunha et al. (2019) highlight that in-kind transfers can increase local supply, leading to relatively lower prices compared to cash transfers.

[2] Private 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. 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).

nomic forces that characterize the equilibrium outcomes of subsidy targeting under market power. There are three key ingredients to our model: heterogeneous consumers that vary in their willingness-to-pay and costs; a uniform pricing rule that prohibits price discrimination; and firms with market power. When subsidies are targeted to consumers under these conditions, we show that a subsidy-induced "demographic externality" arises in addition to the standard imperfect pass-through caused by the exercise of market power. The essential intuition is that targeted transfers change the relative importance of different consumers in the firm's profit-maximization problem and, under the restriction that firms cannot perfectly price discriminate, the equilibrium price a given consumer faces becomes a function of her neighbors' eligibility for transfers. The covariance between demand, costs, and the subsidy schedule can lead to complex equilibrium effects depending on which consumers the subsidies attract at the margin. In general, however, our key result is that equilibrium prices before and after the introduction of targeted subsidies differ. This then changes the equilibrium distribution of subsidy incidence and aggregate efficiency, although we show the direction (and, by extension, magnitude) of changes attributable to the interaction of market power and targeted transfers is theoretically ambiguous and is thus an empirical question in any given setting. The important policy takeaway is that the equilibrium benefit of being targeted will generically not equal the naive estimate of the pre-subsidy equilibrium price minus the targeted subsidy.

Second, we empirically evaluate these conceptual insights in the context of subsidized health insurance markets launched under the Affordable Care Act in 2014 (hereafter: "ACA Marketplaces"). The ACA Marketplaces 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, with an associated federal outlay of more than $40 billion, 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 (even to monopoly). Fourth, firm pricing is highly regulated, with premiums restricted to not vary with income. Given these institutional features, we anticipate the interaction between means-tested public transfers and market power in this setting to be quantitatively important.

We combine detailed data on plan offerings, prices, costs, and enrollment from 2017 for

Figure 1: Diagram of Counterfactuals

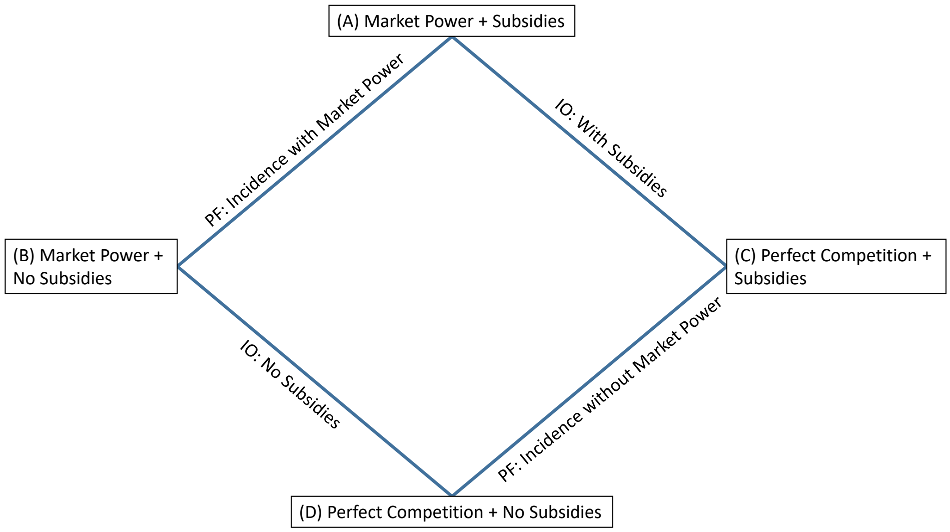

federally-facilitated exchanges covering most of the United States to estimate a structural model of supply and demand for insurance plans on ACA Marketplaces. On the demand side, we allow for the marginal utility of income (and hence price elasticity) to vary across income levels. We utilize the institutional 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 allow for adverse selection and incorporate regulatory constraints on prices.

With the estimates of demand and supply in hand, we are able to simulate several key counterfactuals, summarized in Figure 1, by removing market power, subsidies, or both from the economic environment. These three counterfactuals, along with the observed equilibrium, allow us to understand the effects of market power on equity and efficiency.

First, we simplify the economic environment to highlight our fundamental economic insight: market power has distributional effects. To show this, we remove all subsidies from the market, simulate equilibria with market power (point B) and under perfect competition (point D). In this environment without subsidies, market power leads consumer surplus to fall by $380 dollars per capita, or 15 percent, with an associated decline in enrollment from 34 percent to 23 percent, while producer profit increases by $340 per capita. These aggregates mask significant disparities across income groups, however: the share of consumers with income below 200 percent of the Federal Poverty Level (FPL) buying an insurance plan drops from 4.8 percent to 1.4 percent, while the share of consumers with incomes above 400 percent FPL decreases from 21 percent to 17 percent.

Our second set of counterfactual experiments layer the targeted in-kind subsidies back into the economic environment. Subsidies have the effect of differentially changing the demand curves across consumer types, with ambiguous equilibrium effects depending on which consumers are marginal and how much they cost. We perform two related calculations to help illustrate the effects in our empirical setting. The first calculation addresses the question of how market power directly changes equilibrium outcomes with subsidies. Holding the subsidy schedule constant, we compare the equilibrium with market power (point A) to that under perfect competition with subsidies (point C). We find that per-capita consumer surplus without market power is \$652 higher, implying that market power leads to a 21% loss in average consumer surplus. In terms of insurance coverage, market power leads to 15 percentage point (or 25%) lower rate of coverage—59% of consumers buy insurance without market power, relative to 45% of insurance coverage in the observed environment with market power.

Our last calculation addresses a related, albeit more subtle, follow-up question: how does market power change the distribution of benefits from subsidies across consumers? If marginal costs were equal across all consumers, pass-through in a perfectly competitive world would be 100 percent and the equilibrium incidence of subsidies would exactly reflect the subsidy schedule. However, when marginal costs vary across consumer types, that is no longer the case as average cost pricing now depends on which types of consumers purchase the good. Therefore, to isolate the effect of market power, ceteris paribus, on the incidence of the subsidies, we first compute the incidence under perfect competition (point C minus point D) before then comparing it to the same calculation under market power (point A minus point B); the resulting difference ((A-B)-(C-D)) is the isolated view of market power's effect on the distribution of consumer gains accruing from targeted subsidies.[3] 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 consumers in relative terms.

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

[3]We note that one could have arrived at this same calculation by a conceptually different but algebraically equivalent path by computing the effect of market power on equilibrium outcomes with (A-C) and without (B-D) subsidies, and then taking the difference ((A-C)-(B-D)) to quantify how mean-tested subsidies alter the effects of market power.

the common approach of incorporating distributional policy instruments into environments where a publicly subsidized good is privately provided. We also note that this has implications for antitrust enforcement. Typically, regulatory authorities consider the costs and benefits of a proposed merger on the basis of aggregate measures such as total consumer surplus. When mergers are set in economic environments where redistributional objectives are a first-order concern, antitrust regulators may want to additional consider the additional distortions in the equilibrium incidence of subsidy benefits that may result from the reduction of competition.

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.[4] This results does not hold, however, if we allow for a welfare function that puts value on redistribution per se. As in Waldinger (2018), 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 \$3.60 to a household with income of \$47,520 (400% FPL in 2017 for a single person household).

[4]This 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).

1.1 Related Literature

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; Handel et al., 2019; 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.

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; Tebaldi et al., 2019).[5] 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

[5]More 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 CorporationâĂŹs 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.

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 equity and efficiency effects of of market power. Section 6 briefly concludes.

2 Theory

To deepen our understanding of how market power interacts with targeted subsidies, and how efficiency and equity objectives are affected in equilibrium, we pose a simple theoretical model illustrating several key forces. Our model has three primary ingredients: heterogeneous consumers, who may differ in demand, cost, and subsidies; a uniform pricing requirement that comes from a regulatory restriction on price discrimination; and firms that may have market power.

We show three major results in this setting: first, the uniform pricing requirement itself generates both efficiency and equity effects. The intuition for this is that uniform pricing pools together consumers of many types into a single market; as a result, the equilibrium price is a function of the demographic composition of all types in the market. This pooling is the source of the demographic externality that we described above. Second, the use of targeted subsidies generates an additional dimension of consumer heterogeneity; in combination with the first result, this also implies that equilibrium outcomes are going to depend on who a consumers' neighbors are. Third, we show that the use of strategic intermediaries with market power will have a further effect on the pass-through and incidence of benefits from targeted subsidies. The central point is that firms with market power equate marginal revenues and costs, while perfect competition equates average revenue and cost. Layering market power onto the dual effects of uniform pricing with targeted subsidies can result in equilibrium levels of benefits to targeted consumers that are very different than what the nominal schedule of subsidies might otherwise imply.

Our theory has policy implications for both the provision of goods using private firms and antitrust enforcement. Our results highlight the potential problems of using strategic intermediaries in environments where targeted subsidies reflect the government's preferences for redistributional objectives—while private firms may provide goods and services at lower cost, this may come at the expense of broader societal goals regarding redistribution. This has direct implications for merger policy, as well, as those same government preferences should be reflected not only in efficiency measurements, but also equity concerns in terms of who loses the most from the exercise of market power. This suggests that antitrust authorities may want to broaden their view of the negative effects of a potential merger beyond aggregate measures of efficiency to also encompass the distribution of losses across consumer types.

Uniform Pricing Rule We begin with an analysis of the uniform pricing rule under perfect competition in an environment without any subsidies. Consider a market with a unit mass of consumers. Each consumer faces a menu of possible options, $j = 1, \dots, J$, with an associated utility for each product j:

$$U_{ij} = u_j(p_j, w_i, \theta_i, \epsilon_{ij}), \tag{1}$$

where i indexes the consumer, p_j is the product's price, w_i are consumer characteristics, θ_i is a vector of utility parameters, and ϵ_i is a vector of preference shocks. Normalizing the utility of the outside option (purchasing nothing) to zero, consumers will purchase a single unit of product j if and only if $U_{ij} > U_{ik}, \forall k$ and $U_{ij} > 0$. If utility of all inside options are negative, the consumer will not purchase anything. Denoting the distribution of consumer characteristics by $G(w)$, with an associated density $g(w)$, and the joint set of preferences and shocks by F, market-level demand—the share of consumers buying good j, $s_j(p)$—is formed by aggregating demands:

$$s_j(p) = \int s_{jd}(p_j, w)g(w)dw, \tag{2}$$

where $s_{jd}(p; d)$ is the share of consumers within group d who buy good j.

Under perfect competition, prices are set equal to average marginal cost. Denoting the marginal cost of each consumer type for product j by c_{jw}, the competitive price solves the following equation:

$$\bar{p}_j = \frac{1}{s_j(\bar{p})} \int c_{dw} \cdot s_{jw}(\bar{p}, w)g(w)dw. \tag{3}$$

If marginal costs are equal across all consumer types, as is the case in many product markets, pricing simplifies to the familiar $\bar{p}_j = c_j$ formulation.

Inspection of Equation 3 yields our first observation: the regulatory prohibition on price discrimination has distributional implications. In an unrestricted market, competitive pricing sets price to marginal cost for each consumer type. In contrast, the uniform pricing requirement pools together consumers of different types in the pricing equation, leading to a competitive price that is the sum of marginal costs weighted by each consumer type's share of market demand. Even without any market power considerations, the equilibrium price paid by a consumer of type w depends on the demographic composition of their market via a pooling mechanism in the vein of Rothschild and Stiglitz (1976). We label this economic relationship a "demographic externality."

Introduction of Targeted Subsidies Pooling has additional implications when targeted subsidies are introduced to the market. Let the schedule of targeted subsidies, $Z(w)$, represent the mapping where consumers of type w receive a transfer equal to z_w that can only be spent on goods in this market. Demand shifts outward:

$$s_j(p, Z(w)) = \int s_{jw}(p, z_w)g(w)dw, \tag{4}$$

with the competitive price now determined by:

$$\hat{p}_j = \frac{1}{s_j(\hat{p}, Z(w))} \int c_{jw} \cdot s_{jw}(\hat{p}, z_w)g(w)dw. \tag{5}$$

Our second primary observation follows directly from Equation 5, which is simply Equation 3 with targeted subsidies: the pass-through of the subsidy, which is the out-of-pocket reduction in expenditures for consumers of type w, will generally not equal z_w since $\bar{p} \neq \hat{p}$.

To see this precisely, consider the change in price in response to a marginal change in the subsidy to only type a:

$$\frac{d\hat{p}_j}{dz_a} = -\frac{(\hat{p}_j - c_{ja})\frac{\partial s_{ja}(\hat{p},z_a)}{\partial z_a}g(a)}{\int s_{jw}(\hat{p},z_w) + (\hat{p}_j - c_{dw})\frac{\partial s_{jw}(\hat{p},z_w)}{\partial \hat{p}_j}g(w)dw}. \tag{6}$$

Equation 6 illustrates several important equilibrium relationships, most important of which is that the price change is generically non-zero. The numerator is the change in type-a-specific profit due to an increase in their targeted subsidy; this is the margin for that type multiplied by the change in demand, weighted by the density of that type in the population. Unless demand is invariant to the subsidy or price is exactly equal to that type's marginal cost, this term is non-zero (although it may be positive or negative, depending on the price-cost margin for that type). Likewise, the denominator is the change in profit with an increase in price; it is the usual oligopoly first-order condition for each type integrated over the distribution of types. This term is positive (an increase in price from the zero profit condition will increase profits). The upshot is that subsidizing relatively low marginal cost types will lead to a decrease in prices, with the converse true for relatively high cost types. In either case, subsidizing a single type changes the price faced by all consumers.

That linkage between the targeted subsidy for a specific type and the price faced by all consumers via the uniform price restriction has important policy implications. For one, who a consumers' neighbors are is critical to how high of a price she faces in the marketplace. This can generate potentially ironic outcomes; if the government targets substantial subsidies to high-cost, low-demand types, it can lead to a situation where equilibrium prices are higher the more poor consumers are present in the market. While out-of-pocket prices for the targeted consumers will fall due, prices for everyone else increase. This is a type of price effect that is inverted from the standard model, where being surrounded by high-income, low-elasticity neighbors can lead to an increase in price.

More generally, the uniform price requirement generates a very complex balancing act between the schedule of subsidies, which reflect the government's preferences for redistribution, and the actual distribution of benefits across consumers. As Equation 6 illustrates, the equilibrium price after the introduction of subsidies is a complex function of the complete schedule of subsidies, the demand curve and costs associated with each type, and the distribution of those types in a market. The interplay between all of these economic forces may result in an equilibrium distribution of benefits across consumers that departs substantially from a naive estimate of pre-subsidy prices minus the targeted subsidy, $\bar{p} - z(w)$.

Firms with Market Power Finally, our third observation is to note that intermediaries with market power will further distort the equilibrium distribution of benefits from the targeted subsidy. The key point is this: private intermediaries with market power no longer set prices to equate *average* revenues and costs, but rather maximize profits by equating *marginal* revenue and cost for each product j:

$$\int s_{jd}(\tilde{p}, z_d) + \tilde{p}_j \cdot \frac{\partial s_{jd}(\tilde{p}, z_d)}{\partial \tilde{p}_j} dD = \int c_{jd} \cdot \frac{\partial s_{jd}(\tilde{p}, z_d)}{\partial \tilde{p}_j} dD. \tag{7}$$

The introduction of market power changes equilibrium outcomes in several ways. The principal change is that price is higher under market power than perfect competition, as firms seek to increase margins to generate non-zero profits. This increase both leads to a lower quantity of the good consumed in equilibrium (an efficiency effect) and a change in the relative consumption of the good across demographic types (a distributional effect). In addition, for any increase in subsidies, the firm captures some of the subsidy, leading to an additional channel influencing pass-through.

Equation 6 under market power becomes:

$$\frac{d\tilde{p}_j}{dz_a} = -\frac{\frac{\partial s_{ja}(\tilde{p}, z_a)}{\partial z_a} + (\tilde{p}_j - c_{ja}) \cdot \frac{\partial^2 s_{ja}(\tilde{p}, z_a)}{\partial \tilde{p}_j \partial z_a}}{\int 2\frac{\partial s_{jd}(\tilde{p}, z_d)}{\partial \tilde{p}_j} + (\tilde{p}_j - c_{jd}) \cdot \frac{\partial^2 s_{jd}(\tilde{p}, z_d)}{\partial \tilde{p}_j^2} dD}. \tag{8}$$

This equation is the higher-order derivative analogue of its perfect competition counterpart. As before, this expression again highlights the demographic externality that a type's neighbors exert on equilibrium outcomes. On the other hand, signing this equation is complicated by both reasons that some of the terms may change sign depending on the level of the function (e.g. the second derivative of the share equation if the underlying demand function is the logit) and that the two cross-partials may have opposite signs. The broader point is that firms with market power may react very differently to the introduce of subsidies than perfectly competitive firms; it is even possible that prices move in opposite directions between the two regimes. Given this ambiguity, assessing the effects of market power on markets with targeted subsidies is an empirical exercise.

Figure 2: **Consumer interface on healthcare.gov**

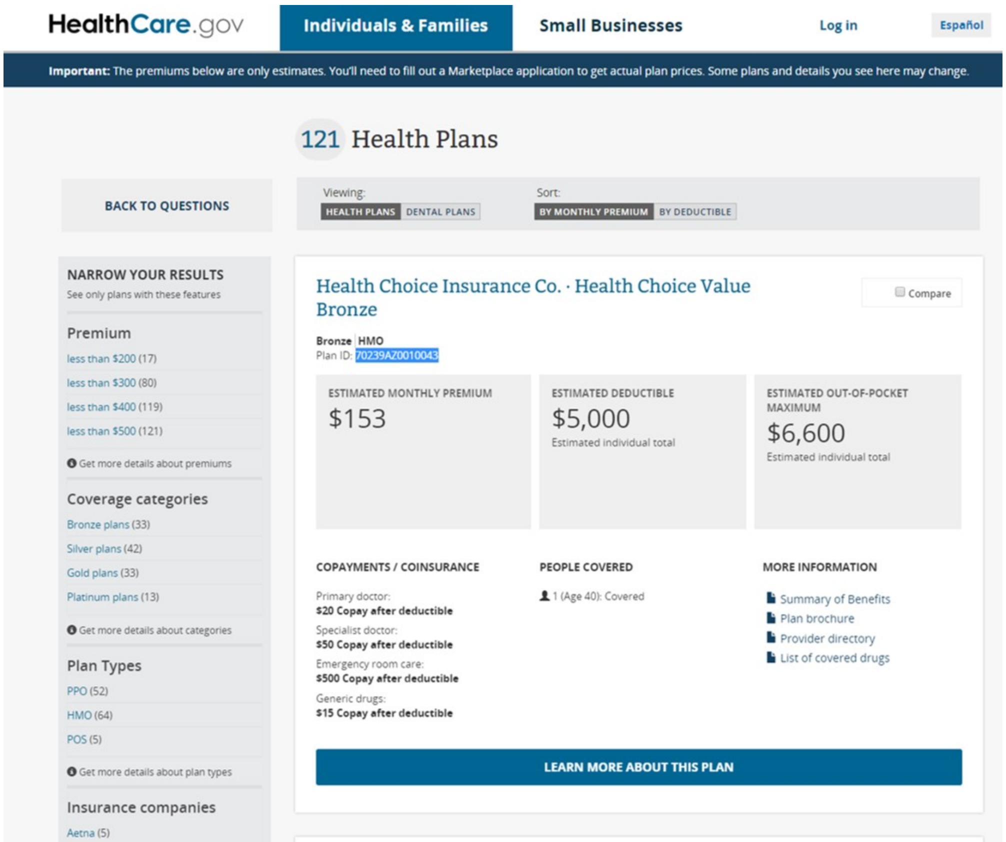

Notes: Screenshot of healthcare.gov for a plan that as offered to a 40-year old individual in Cook County, IL. The premium that individuals see displayed incorporates the premium subsidy if individuals enter their income information during the selection process.

3 Institutional Primer and Data

3.1 Institutions

Our empirical setting is the US market for non-group health insurance plans — Health Insurance Marketplace — launched in 2014 under the Affordable Care Act (ACA). The Marketplace allows anyone living in the US to purchase a health insurance plan for their families.[6] In most states consumers are directed to a federal online platform, www.healthcare.gov, to purchase coverage. We focus on these 37 so-called "federally-facilitated' states in our analysis. Figure 2 provides an example of the user interface on the healthcare.gov platform.

[6]Formally, eligibility is restricted to individuals living in the US, who are lawfully present, not incarcerated, and do not have Medicare coverage.

Insurance firms selling on the Marketplace set list premiums for their plans subject to several regulatory restrictions. First, insurers are not allowed to reject coverage and are not allowed to price-discriminate based on individual health risk. The regulator allows an insurer to collect different premium based on age, but the age gradient is regulatorily set (Tebaldi, 2017).[7] Second, insurers have to set the same list premium in all counties that belong to the same "rating area." A rating area is a collection of counties within a state pre-specified by the regulator. While insurers have to set the same list price for a plan if they offer the plan in all counties in a rating are, 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 an insurer to spend at least 80% of revenue on medical reimbursement and quality improvement (Cicala et al., 2019).

The key institutional feature of the Marketplace are means-tested subsidies that can be used by lower income families to defray the cost of the insurance list premium. Formally, premium subsidies are known as Advanced Premium Tax Credits (APTC).[8] 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).[9] The FPL varies with family composition and allows comparing incomes of families of different sizes using the same scale. The MAGI as percent of the FPL then determines the maximum dollar amount that the (tax) household "can be expected" to pay for health insurance coverage. Let us call this amount a CAP. The CAP is based on a non-linear sliding schedule specified by the IRS.[10] For example, in 2017, if a household's MAGI was 200 percent of the FPL, then this household's CAP was 6.34 percent of the household income, while if income was 270 percent FPL, the CAP was 8.7 percent. Subsidies normally phase out at 400 percent FPL.[11]

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

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

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

[10] See IRS Form 8962 "Premium Tax Credit."

[11] American Rescue Plan Act of 2021 for COVID-19 Relief temporarily extended subsidies to households with income over 400 percent FPL.

The CAP is then compared to the list premium (b) that the household (or specifically the "coverage family"—everyone in the household who is buying insurance) would have to pay for a benchmark plan (described below) in the household's county of residence. If the list premium for the benchmark plan is above the CAP, then the household is eligible to get an APTC that is at most equal to the difference. To summarize, the maximum subsidy that a household h can receive as a function of household income Y_h is:

$$z_h = \max\left\{0, \sum_{i \in h} b_{i,benchmark} - CAP(Y_h)\right\}, \tag{9}$$

where $\sum_{i \in h} b_{i,benchmark}$ is the total list premium that the coverage family would have to pay for the benchmark plan. If the list price of the plan that the household actually buys is lower than the maximum APTC, APTC is reduced accordingly, so that the final consumer-facing premium can be zero, but cannot be negative. The out of pocket premium that household h pays for plan j is then:

$$p_{hj} = \max\left\{0, \sum_{i \in h} b_{ij} - z_h\right\}. \tag{10}$$

Insurance plans sold on this market are high dimensional products, offering a variety of cost-sharing rules for different services and different provider networks. The financial characteristics of a plan are summarized by the plan's actuarial value that measures the fraction of costs that a plan would pay for a standardized population. Plans are grouped into "metal levels" based on their actuarial value: 60% actuarial value (bronze plans), 70% actuarial value (silver plans), 80% actuarial value (gold plans), and 90% actuarial value (platinum plans).[12] Some families will qualify for cost-sharing reduction (CSR) subsidies that increase the actuarial value of a plan by reducing deductibles, co-pays, and co-insurance. We take these subsidies as given in our analysis. The second-cheapest 70% actuarial value (silver) plan in each county is designated as the benchmark plan for that county.

3.2 Data

Choice set data Our data on the set of Marketplace plans that was available to consumers in each county comes from 2017 Marketplace Public Use Files (PUF) provided by the Center for Medicare and Medicaid Services (CMS) that oversees the Marketplace. The data record

[12] Insurers also offer "catastrophic" plans with high deductibles that do not follow any specified actuarial value and are only available to consumers under 30.

detailed information on which plans were offered in each geographic market of the federally-facilitated Marketplaces, plan features, and list premiums. We restrict our analysis to year 2017, as (arguably) the year in which the market most resembled the long-run equilibrium. Market stabilization programs (reinsurance and risk corridors) were set for the initial program years 2014-2016, while starting with 2018 the individual mandate and cost-sharing reduction subsidies were repealed and COVID-19-related changes followed closely after.

Enrollment data We use enrollment data for year 2017 provided by the Center for Consumer Information and Insurance Oversight (CCIIO). Enrollment is reported in several ways. First, we observe enrollment reported at the plan level (which can span several counties). Second, we observe enrollment at the county level aggregated by different metal levels: for example, how many individuals purchased a Silver plan in Cook County, IL. Third, we observe county-level counts of consumers who enrolled in any plan, by seven age and seven income groups.[13] We enrich CCIIO enrollment data with a dataset from the Kaiser Family Foundation (KFF) that provides an estimate (based on year 2015) of the number of potential Marketplace consumers for each Public Use Microdata Area.

Demographic data We use 2017 American Community Survey (ACS) to create a representative sample of potential Marketplace consumers in each county. We flag an ACS respondent as a potential consumer if the individual did not report having employer-sponsored or public health insurance coverage, and likely would not have been eligible for Medicaid (based on household income) in states that expanded Medicaid under the ACA. For the resulting sample of potential consumers we observe household size and composition, household income relative to FPL, and age. To allow for adverse (or advantageous) selection, i.e. non-constant marginal costs, in this market, we assign cost multipliers to each consumer in our ACS sample based on age and income. We estimate the gradient in average healthcare costs across different age and income groups using the Optum database for the commercially insured under-65 population. Appendix B provides details.

Summary statistics Table 1 summarizes the data. An average consumer faced a choice of 21 plans. About 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

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

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 who did purchase a Marketplace plan, 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 percent FPL, making them eligible for the most generous subsidies. 37 percent had income over 200 percent FPL but under 400 percent FPL, making them eligible for partial subsidies, and another third of potential consumers had household income over 400 percent FPL, making them ineligible for subsidies. On average, potential consumers qualified for $2,349 in annual premium subsidies.[14]

[14] We 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.

Table 1: **Summary statistics**

	Mean‡	Std. Dev.	10th pctile	90th pctile
A. Choice set				
Number of plans	21	13	8	37
Number of large insurers	2.16	1.13	1	4
Average annual premium (age 40), $	5,106	902	3,978	6,351
B. Enrollment				
Market size‡‡	7,867	25,756	479	15,671
Share outside option	0.60	0.17	0.43	0.76
Share in 60% actuarial value plans	0.09	0.05	0.04	0.16
Share in 70% actuarial value plans	0.29	0.15	0.17	0.43
Share in 80% actuarial value plans	0.01	0.02	0.00	0.03
Market-level enrollment	3,536	13,798	168	6,411
Plan-level enrollment^	3,165	12,040	39	6,353
C. ACS Sample of Potential Consumers				
Age	39	2	36	42
Income in % FPL	295	52	231	365
Annual premium subsidy, $^^	2,349	1,244	919	4,226
Share income under 200% FPL	0.32	0.14	0.13	0.50
Share income 200–400% FPL	0.37	0.07	0.28	0.47
Share income 400% FPL and above 0	0.31	0.11	0.19	0.46
Share age under 25	0.22	0.05	0.16	0.28
Share age 25–40	0.30	0.05	0.24	0.36
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).

4 Empirical Model

4.1 Demand

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

$$u_{ij} = -\alpha_i p_{ij} + \phi_{ij}, \tag{11}$$

where ϕ_{ij} is the utility of plan j, p_{ij} is the price that consumer i would have to pay for plan j, including any subsidies, and α_i is the marginal disutility of price. We assume that family f chooses a single insurance plan j or the outside option to maximizes the average utility across members of the family:

$$\epsilon_{fj} + \frac{1}{N_f} \sum_{i \in f} u_{ij} > \epsilon_{fk} + \frac{1}{N_f} \sum_{i \in f} u_{ik}, \forall k \in J \text{ s.t. } k \neq j, \tag{12}$$

where ϵ_{fj} is a family-level idiosyncratic taste shock for plan j. We normalize the payoff from the outside option of not purchasing a Marketplace plan to zero.[15]

We make several assumptions about α_i and ϕ_{ij} to arrive at an empirically-tractable utility function. 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 three additively-separable

[15] outside option can constitute either choosing to have no formal health insurance coverage or purchasing a health insurance plan outside the Marketplace. Other sources of health insurance coverage can be plans purchased directly from health insurance companies or brokers, typically known as "off-exchange" plans. These can either be pre-ACA grandfathered individual plans, or new plans that are sold with no subsidies. We assume that individuals who have the option to purchase employer-sponsored plans or a eligible for public insurance program, always choose those options and thus are not participating in the Marketplace market. Under this assumption, the outside option does not vary across consumers and hence the normalization of the outside option payoff to zero for all consumers is without loss of generality.

components:

$$\phi_{ij} = \psi_{a(i)} + \gamma AV_{ij} + \delta_j. \tag{13}$$

The first component, $\psi_{a(i)}$, captures the average level of utility that consumers get from purchasing any Marketplace plan. We allow this intercept parameter to vary across three age groups to capture the idea that the value of insurance may vary across ages, all else equal.[16] The second component, AV_{ij}, captures the deviations in the level of cost-sharing of plan j among consumers eligible for cost-sharing reduction subsidies based on their income. Finally, we include a plan-specific constant δ_j for each plan j that non-parametrically captures the average utility that all consumers get from purchasing plan j.

Identification 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 we observe 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 the geographic markets in which each plan j 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.[17]

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 the 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 we observe 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

[16] We also allow for a separate intercept for the group of consumers with income under 100 percent FPL. While this group of consumers should not be participating in the Marketplace, as they are commonly eligible for Medicaid, we observe some enrollees from this group in the data.

[17] This follows from a 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.

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 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 age or income groups. 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 Marketplace 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 models of 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). In our empirical context, 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 do not systematically sort across markets on the basis of these regulatory-induced changes in prices.

4.2 Supply

Profit function For each consumer i purchasing plan j, the plan collects a premium p_{ij}. For consumers who are not eligible for premium subsidies, p_{ij} is equal to the list price, b_{ij}. For consumers who are eligible for subsidies, the insurer collects $p_{ij} < b_{ij}$ from the consumer, and a subsidy equal to the difference between b_{ij} and p_{ij} from the government. On the expenditure side, the insurer pays for consumers' healthcare bills and any administrative costs. The total *expected* cost of healthcare of consumer i in plan j is a function of the consumer's underlying health risk, r_i, and the plan's contract features, ϕ_j. Contract design can change how much care i consumes for any given r_i. Plan j expects to pay a portion of the total healthcare cost, net of consumer cost-sharing. Let c_{ij} denote plan's j expected cost for offering insurance coverage to individual i. Plan j's expected profit from enrolling consumer i is then:

$$\pi_{ij} = b_{ij} - c(r_i, \phi_j). \tag{14}$$

Regulatory restrictions on pricing imply that revenue does not vary within age groups 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 multipliers τ^a that do not vary across plans:

$$b_{ij} = \tau^{a(i)} b_j. \tag{15}$$

As we do not observe individual-level health status, we parameterize c_{ij}. We preserve the idea that costs may vary across consumers, but we discretize this variation. Using a commercial claims database, we estimate a matrix of cost multipliers κ^d for d demographic groups defined as a cross product of 1 year age intervals and three income intervals (see Appendix B). 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 (which in turn lead to different quantity of healthcare consumed—the moral hazard effect, as well as different plan liability for any given amount of healthcare consumed) 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 that varies across plans and the age-income cost multiplier κ^d that does not vary across plans:

$$c_{ij} = \kappa^{d(i)} c_j. \tag{16}$$

Allowing both demand and cost to 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 J_f as:

$$\Pi_f(\mathbf{b}) = \sum_{j \in J_f} \sum_{d \in D} \left[(b_j \tau^d - c_j \kappa^d) s_j^d(p(b)) M^d \right], \tag{17}$$

where $s_j^d(p(b))$ is the share of consumers in age-income group d that buys plan j and M^d is the number of potential consumers in age-income group d. $p(b)$ is the link function between the firm's list price b and the out-of-pocket price p that a consumer pays (Equation 10 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. The share equation is defined over all markets in which the plan is offered (i.e. its rating area). For expositional clarity, we do not write out this dependence explicitly, but we do account for this in our empirical estimation and policy simulations.

First-order conditions Each insurer f chooses a vector of baseline list prices $\mathbf{b}$ to maximize profits. Subject to the MLR constraint, the optimal list price b_j for each plan $j \in J_f$ has to satisfy the following first-order condition:

$$\frac{\partial \Pi_f(\mathbf{b})}{\partial b_j} = \sum_{k \in J_f} \sum_{d \in D} \left[(b_k \tau^d - c_k \kappa^d) \frac{\partial s_k^d(p(b))}{\partial b_j} M^d + 1(j = k) \cdot \tau^d s_j^d(p(b)) M^d \right] = 0. \tag{18}$$

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

$$S - \Omega(B - C) = 0, \tag{19}$$

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:

$$\Omega_{kj} = -\sum_d \frac{\partial s_j^d(p(b))}{\partial b_k} M^d \tag{20}$$

for plans k and j offered by firm f. This expression is useful as one can invert Equation 19 to obtain the baseline marginal cost c_j for each plan as a function of observed equilibrium prices and the elasticity of demand for each demographic group d as given by the demand parameters from Section 4.1. We constrain our estimates of marginal cost to conform to the MLR regulation.[18]

In a standard product market, $p = b$, so that

$$\frac{\partial s_j(p(b))}{\partial b_k} = \frac{\partial s_j(p)}{\partial p_k}. \tag{21}$$

In our empirical setting, many consumers receive means-tested subsidises, which introduces an additional relationship between list prices b and consumer-facing prices p. The share of consumers choosing to buy plan j depends on consumer-facing prices p, which are in turn related to the list price b through Equation 9. The derivative of the market share with respect to list price b then becomes:

$$\frac{ds_j(p(b))}{db_k} = \frac{\partial s_j(p(b))}{\partial p_k} \cdot \frac{\partial p_k}{\partial b_k}. \tag{22}$$

For lower-income consumers eligible for subsidies, if $b_k < z_h$, the derivative of the out of pocket price with respect to a plan's bid is zero; for all other consumers $b_k > z_h$ and the derivative is one.[19] In addition, when plan k is the second-lowest cost silver plan there are additional terms in Equation 22. This is due to b_k entering Equation 9 as the benchmark price, 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 the 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. Note that firms aiming to become the second-lowest cost silver plan have incentives to raise prices arbitrarily high, as the generosity of the subsidy increases one-for-one in b_k. This incentive will be bound by

[18] As 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.

[19] The primary distinction is between subsidized and unsubsidized consumers. If subsidizes consumers, however, buy relatively expensive plans, then their subsidy does not cover the full cost and they also face the full marginal increase in the list price.

the MLR constraint. We assume that in equilibrium a plan knows whether or not it will be the second-lowest Silver plan and use the MLR constraint to recover marginal costs for these plans.

The first-order condition illustrates that while the firm sets one 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 interdependence of demands in the firms' profit function gives rise to a *demographic externality* and has implications for any public policy that alters demand only for a subset of consumers. For example, if one consumer group d is affected by a means-tested subsidy, as is the case in our empirical application, this will affect the list price not only for this group of consumers, but for *all* other demographic groups d. We will come back to this conceptual observation when we evaluate how the subsidy structure (whether means-tested or income-invariant) can change the effect of market power on the market.

4.3 Efficiency Metrics

Consumer surplus 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:

$$CS_i(\theta_i) = \frac{1}{\alpha_i}\left[\gamma + \ln\left[1 + \sum_{j=1}^{J} \exp(u_{ij}(\theta_i))\right]\right], \tag{23}$$

where γ is Euler's constant and u_{ij} is the deterministic component of utility for person i (recall that this is the average utility within a family) for plan j.[20] 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:

$$CS = M \int CS(\theta) dF(\theta). \tag{24}$$

In some of our policy simulations we use an alternative definition of consumer surplus that allows for differential welfare weights and, hence, preferences for redistribution. We use the family of welfare functions with a constant relative inequality aversion (Atkinson,

[20] Euler's constant is the mean value of the Type I Extreme Value idiosyncratic shock under the standard normalizations in the logit model, and is approximately equal to 0.577.

1970).[21] Individual-level consumer surplus is recomputed as a function of individuals' income y_i, as follows:

$$CS_i^{\lambda} = \begin{cases} \frac{1}{1-\lambda}[(y_i + CS_i)^{1-\lambda} - y^{1-\lambda} & if\ \lambda \neq 1, \\ log(y_i + CS_i) - log(y_i) & if\ \lambda = 1. \end{cases} \tag{25}$$

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.

Producer surplus Producer surplus, Π, is computed as expected profits following equation 17.

Government spending 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 policy simulations 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 policy simulations. Specifically, using CCIIO reports we compute the average per capita government spending on CSR subsidies by consumer type, based on income brackets.[22] We then assign the average CSR subsidy 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 a Marketplace plan, the government saves money on this consumer in other healthcare domains. 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 Marketplace (Coughlin et al., 2014), we assume that the government saves \$1,827 per capita in public funds on each consumer who buys a Marketplace plan.

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

[22]The data was accessed in June 2019 at Health Insurance Marketplace Cost-Sharing Reduction Subsidies.

4.4 Parameter estimates

Demand We use non-linear least squares to estimate the parameters of the function.[23] 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 premiums 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. Consumers get positive utility from purchasing plans with a higher level of coverage (as measured by the actuarial value). 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 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 to a 70 percent actuarial value 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, all marginal utility parameters capture family level preferences, and as such reflect, for example, parental preferences for health insurance for their children.

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.2 illustrates one set of moments that was used for estimation and a related measure of model fit. In Panel C.2A we plot the observed market share of 70 percent actuarial value (Silver) plans in each county. In Panel C.2B we plot the average difference between the data and the model's prediction of county-level Silver enrollment shares. The average prediction error is -0.0007 relative to the average empirical enrollment share of 0.29, and is broadly distributed across counties, suggesting that the model is able to capture a substantial amount of variation in the data.

[23] Computational details are provided in Appendix A.

Table 2: **Model estimates**

A. Demand: parameters of utility function	Mean	Age <25	Age 25–40	Age >40
Coefficient on premium, $000 ($\boldsymbol{\alpha}$)				
Income <200% FPL		−5.17	−2.47	−2.21
		(0.33)	(0.16)	(0.14)
Income >200% FPL and <400% FPL		−4.32	−0.64	−3.94
		(0.27)	(0.04)	(0.26)
Income >400% FPL		−1.13	−0.20	−0.46
		(0.07)	(0.01)	(0.04)
Age-specific intercepts		1.52	−1.72	base
		(0.10)	(0.11)	
Actuarial Value	26.83			
	(1.69)			

B. Supply: inversion of first-order conditions	Mean	Std. dev.	Min	Max
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	1,173	4,102
Estimated cost multipliers‡				
Income <200% FPL	2.77‡‡			
Income >200% FPL and <400% FPL	2.15‡‡			
Income >400% FPL	1.97‡‡			

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

‡ Income categories in Optum™ 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 year 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.

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 lowest cost plan has an estimated baseline marginal cost of $732, while the highest cost plan has an estimated baseline marginal cost of $4,102. In general, there is substantial heterogeneity in estimated marginal costs across plans both across and within actuarial value levels.

In Appendix Figure C.3, we compare our estimates of marginal costs from the first-order condition inversion to plan-level accounting costs. The accounting data is available from CMS rate review files for 705 plans. The accounting costs are measured with error, as insurers are allowed to report their costs equally split across their plans rather than providing a 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 accounting cost ($4,010 versus $4,768). We observe a strong ordinal correlation between accounting costs and inverted marginal costs, which lends credence to our estimates.[24]

The table also reports our estimates of cost multipliers (κ^d). We find that on average across all ages individuals in households with income under 200 percent FPL experience 40 percent higher healthcare costs than individuals in households with income over 400 percent FPL. This difference is consistent with the widely documented socio-economic gradient in health (Marmot, 2015).

[24]An alternative empirical approach, also frequently used in the literature, would be to directly use accounting costs as inputs into our analyses and rather than inverting 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 is second-order for the comparative statics of the subsequent policy simulation, given the strong correlation between the two measures.

5 Policy simulations

5.1 Aggregate effects of market power

We start by computing the effect of market power on an average consumer in the Marketplace before turning to the distributional analysis. To facilitate this computation, we first simulate the observed allocation within our model to establish a baseline that differences out any demand model simulation error.[25] Next, we re-simulate the model, shutting down the market power channel—which we implement by forcing firms to price at average cost—keeping all else equal.

The first column of Table 3 reports the characteristics of the (simulated) observed equilibrium. We estimate that consumers enjoy on average $2,495 in annual consumer surplus in this market. 45% of potential consumers buy an insurance plan on the Marketplace. Both government spending and insurer profits are substantial. Insurers collect $5,618 in revenue for the average buyer. The average cost of a buyer, however, is $3,993, leaving insurers with $1,625 (or 28%) in economic profits. Subsidies, which we estimate at $3,196 per buyer, amount to 57% of the firm revenue and 80% of the cost of coverage.[26]

The high profit margin points to a substantial degree of market power in this market. To quantify the distortions created by market power, we turn to the third column of Table 3 that reports the simulation of the same environment, but sets prices equal to average cost. List 20-year-old price drops by 27%, as firms are forced to have a zero profit margin. The drop in price is not equal to the full profit margin in the baseline, as the marginal consumers attracted into the market when prices drop are higher cost—the average cost of the buyer population is about about 1% higher in this scenario. Consumer surplus without market power is $652 higher, implying that market power leads to a 21% loss in average consumer surplus. In terms of insurance coverage, market power leads to 15 percentage point (or 25%) lower rate of coverage—59% of consumers buy insurance without market power. Average taxpayer spending on premium subsidies for the buyers is nearly $200 (or 6%) higher in the

[25] The supply-side returns back observed list prices from the first order conditions and has no simulation component. In practice, our demand model has a tight in-sample fit, so that the resulting simulated allocation is very close to the data.

[26] The level of government subsidies that we estimate in the model is close to the subsidy spending reported by the Congressional Budget Office. For all Marketplace markets (including non-federally facilitated states that are not part of our analysis) CBO reports $53 billion in net premium and cost-sharing subsidy spending for 2017 for 15 million enrollees, amounting to an average spending of $3,467 per buyer. (https://www.cbo.gov/sites/default/files/114th-congress-2015-2016/reports/51385-healthinsurancebaselineonecol.pdf).

presence of market power (the average subsidy among all potential buyers is lower, as fewer people buy insurance).

Subsidy pass-through The pass-through of subsidy spending to consumers is another key way of assessing the efficiency implications of market power. Computing subsidy pass-through requires knowing what the prices and the economic surplus (of consumers and producers) would have been in the absence of premium subsidies. In columns (2) and (4) of Table 3 we report the results of simulations that set premium subsidies to zero, keeping all else equal relative to columns (1) and (3), respectively. In other words, in column (2) we set subsidies to zero, but keep the price-setting power of the firms; while in column (4) we also force firms to price at average cost.

We find that in the presence of market power, subsidy spending of about $1,400 per market participant generates an extra $734 in (consumer and producer) surplus. This implies the rate of return of 52 cents on a dollar of subsidy spending—or a deadweight loss of 48 cents. Consumers capture 47% of the surplus generated by subsidies, or 24 out of 52 cents, while 53% is captured by firms.[27] In the absence of market power, the pass-through of extra surplus to consumers is 100% (by construction). The deadweight loss is still substantial. $1,706 in premium subsidy spending on an average potential consumer generates $613 in additional consumer surplus, implying a welfare return of only 36 cents on a dollar of extra taxpayer spending. Our results suggest that first, consumers have a low willingness to pay for insurance and value a dollar of cash at 36 cents when it is given through insurance, i.e. consumers would have substantially preferred to receive these transfers in cash.[28] Second, in the presence of market power firms capture about half of the subsidy surplus and the incidence of a dollar in taxpayer spending on consumer surplus is ultimately a third lower in the presence of market power than would have been the case with average cost pricing.

Subsidies play a crucial role in consumers' decisions to buy plans—we observe that enrollment in column (2) drops by half when we remove subsidies. However, our calculations also suggest a high degree of deadweight loss, as consumers have a relatively low willigness

[27] A pass-through to consumers of less than 50 percent is consistent with the importance 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.

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

to pay for this (effectively) in-kind benefit. Our computation so far, however, is incomplete. The institutional environment in healthcare suggests that the government is likely to incurs some costs for the same consumers even if they do not enroll in the ACA Marketplace. To account for this, we re-compute subsidy spending, taking into account how much the government "saves" on each consumer who enrolls in a Marketplace plans. We use the estimates of public spending on uncompensated care prior to the rollout of ACA as a proxy for government savings. As spending on uncompensated care was high, incorporating these savings reduces extra taxpayer spending needed to finance Marketplace premium subsidies by more than a half; indeed, in the scenarios where we remove subsidies, the government saves money on the Marketplaces. Taking net taxpayer cost as a basis for the incidence calculations, we get that in the presence of market power, the extra $734 in consumer and producer surplus is now generated by $1,020 in extra government spending (foregone savings plus the extra spending). In other words, the rate of return increases to 72 cents on a dollar of public spending, in aggregate, and to 34 cents on a dollar for consumers. The analogous number in the absence of market power is 49 cents of extra surplus for a dollar of public spending.

Table 3: **Policy Experiments: Aggregate Results**

	With market power		Perfect competition	
	Baseline – observed	Remove (premium) subsidies	Keep subsidies; firms set $p = AC$	Remove subsidies; firms set $p = AC$
Average across potential consumers ($)				
Consumer surplus	2,495	2,152	3,147	2,534
Insurer profit	729	338		
Taxpayer cost of subsidies	1,434	23	1,775	69
Taxpayer cost net of savings on uncompensated care	614	−406	698	−548
Insurance rate	0.45	0.23	0.59	0.34
Average 20 year old list premium (unweighted), $	2,401	2,239	1,743	1,592
Among consumers buying insurance ($)				
Average cost of covering a buyer	3,993	3,348	4,045	3,425
Average list premium among buyers	5,618	4,788	4,044	3,426
Insurer profit per buyer	1,625	1,441		
Taxpayer cost of subsidies per buyer	3,196	96	3,010	204

5.2 Distributional effects of market power

We next move to examining the distributional effects of market power. We use the same simulations of the observed equilibrium and three counterfactual equilibria, as reported in Table 3, to ask how much the implied effects of market power on consumer surplus, the rate of insurance coverage, and subsidy incidence vary across consumers with different family incomes. Our results illustrate how two economic primitives are needed to be able to assess the distributional effects of market power: the variation in elasticity of demand by income and the distribution of incomes in the population (in contrast to the standard analysis of market power that requires only the aggregate elasticity of demand).

No Subsidies We start our analysis of the distributional effects of market power in a simulated environment that sets all premium subsidies to zero. Columns 2 and 4 of Table 3 show that market power leads to a 41% increase in (unweighted) list prices[29], 15% decline in average consumer surplus (and a 2% decline in consumer and producer surplus jointly), an 11 percentage points (32%) decline in insurance coverage. Market power screens out higher cost consumers on the extensive margin of the market, with the average cost of covering a buyer dropping by 2%.[30]

The light grey lines in Figure 3 show how the impact on consumer surplus and insurance coverage rates vary across consumers with different levels of income. We observe that (proportional) losses in consumer surplus get higher as income gets lower. While market power leads to a 23% loss in consumer surplus among consumers with lower than average incomes, the loss among consumers with higher than average incomes is 14%. Lower income consumers with more elastic demand that creates more price pressure on insurers exert a positive demographic externality on higher income consumers, but are more affected by any dollar increase in prices than higher-income consumers. This is also reflected in the pattern of insurance coverage. Consumers with incomes below average in this population already have limited willingness to pay for insurance—in the absence of subsidies, only 24% buy a plan. The share of those covered drops to 11% or by 13 percentage points (or by 54%) when market power is present and leads to higher prices. Among consumers with incomes above the average, market power leads to a 9 percentage point reduction in coverage from a much higher base, from 49% to 40%, implying an 18% proportional drop in coverage. Using the estimates of average (not enrollment-weighted) list prices for a 20 year old in Table 3, we get

[29] And a 40% increase in enrollment-weighted list prices.

[30] The risk sorting pattern is different from a standard Akerlof (1978)-style model, because high risk consumers are also lower-income consumers, who are more price sensitive.

that prices increase by 41% due to market power (in the absence of subsidies). This implies an extensive margin elasticity of 1.3 among consumers with lower incomes and of 0.4, or three times lower, among consumers with higher incomes.

In the absence of market power, the traditional cost-based demographic externality is the key characteristic of the market. Higher-income consumers exert a positive externality on lower-income consumers, as higher-income consumers cost less on average. In a perfectly competitive environment, lower-income consumers thus prefer to pool with higher-income consumers. In the presence of market power, the nature of demographic externality changes. Now higher-income consumers also exert a *negative* externality on lower income consumers, as higher-income consumers are substantially less elastic and thus put much less downward pressure on prices. Pooling of consumers of different incomes becomes less attractive to lower-income consumers, who nearly exit the market. Market power thus exacerbates the gaps in the rate of insurance coverage, moving the gradient from 25 to 29 percentage point difference in coverage rates between the top half and the bottom half of the income distribution in our sample.

With Targeted Subsidies We next consider the effects of market power on consumers with different incomes in an environment with means-tested subsidies as observed on ACA Marketplaces. Conceptually, means-tested price subsidies introduce another income-correlated dimension of heterogeneity among consumers. Lower-income consumers are still higher cost and more inherently price elastic, but they now face lower out of pocket price for any list price. This tends to reduce the elasticity of demand among these consumers from the firm's perspective. Indeed, for some combinations of income and list prices, lowest-income consumers may always face a fixed price (equal to a percent of their income), so that their effective elasticity of demand is zero. The presence of means-tested subsidies may thus lead to a flip in the sign of the price-related demographic externality that lower-income consumers exert on higher-income consumers—they no longer put no pressure on insurers to reduce prices.

The black lines in Figure 3 show how surplus and insurance enrollment changes with market power for consumers of different incomes, when means-tested subsidies are present. Market power again has a larger (proportional) effect on consumer surplus among lower-income consumers (-34% vs -17%) and a larger effect on the absolute change in the rate of insurance enrollment (20pp drop from 77% to 57% among consumers with lower than average incomes and 15pp drop from 56% to 41% among potential consumers with higher than average incomes). The proportional effect on insurance enrollment is similar, with rate of coverage falling by 26% in both above mean and below mean income groups. In the presence of means-tested subsidies, lower-income consumers have higher enrollment rates than higher-income consumers, thus the "gaps" in coverage rates are negative. Market power shrinks the negative gap, reducing the effectiveness of a dollar of subsidies in increasing enrollment in the lower-income group. Overall, the level of consumer surplus and insurance coverage rates are substantially higher among lower-income consumers than they are when there are no subsidies, but the relative losses in surplus induced by market power are still larger for lower-income consumers.

5.3 Mechanisms

The central economic mechanism that is reflected in our results is the "demographic externality." Heterogeneous (on potentially many dimensions) potential consumers are buying a good from a firm with market power and this firm is required to set one list price. Each consumer would then prefer to be in a market with low-cost, price elastic neighbors as the firm will have incentives to set the lowest prices in those markets. In this environment, each consumer is exposed to the "demographic externality" from her neighbours.

Consumers in insurance markets are inherently heterogeneous on cost and on demand elasticity.[31] In our application, the heterogeneity in cost and elasticity varies systematically with income—lower-income consumers are higher cost, but more elastic. Under uniform pricing, lower-income consumers exert negative demographic externalities on neighbors through the cost channel (as in a classic Akerlof, 1978) model but positive demographic externalities by putting pressure on prices.

Means-tested transfers that target lower-income consumers introduces another dimension of heterogeneity that breaks the inverse relationship between cost and elasticity. Means-tested subsidies reduce the elasticity of demand among lower-income consumers. This reduces or even flips sign of the positive demographic externality that lower-income consumers were exerting on their unsubsidized higher-income neighbors.

In this section, we simulate two thought experiments to illustrate the subsidy-driven demographic externality channel. Relative to other sources of demographic externalities, this mechanism, to the best of our knowledge, has received no attention in the prior literature.

[31] The heterogeneity in demand elasticity can be driven by many primitives, such as differences in risk aversion, interest in one or the other type of providers, or informational/behavioral differences. While the exact source is important to understand on its own, the resulting differences in the elasticity of demand are a "sufficient statistic" from the perspective of the firm's pricing problem.

Figure 3: **Distributional effects of market power**

(A) Loss in consumer surplus

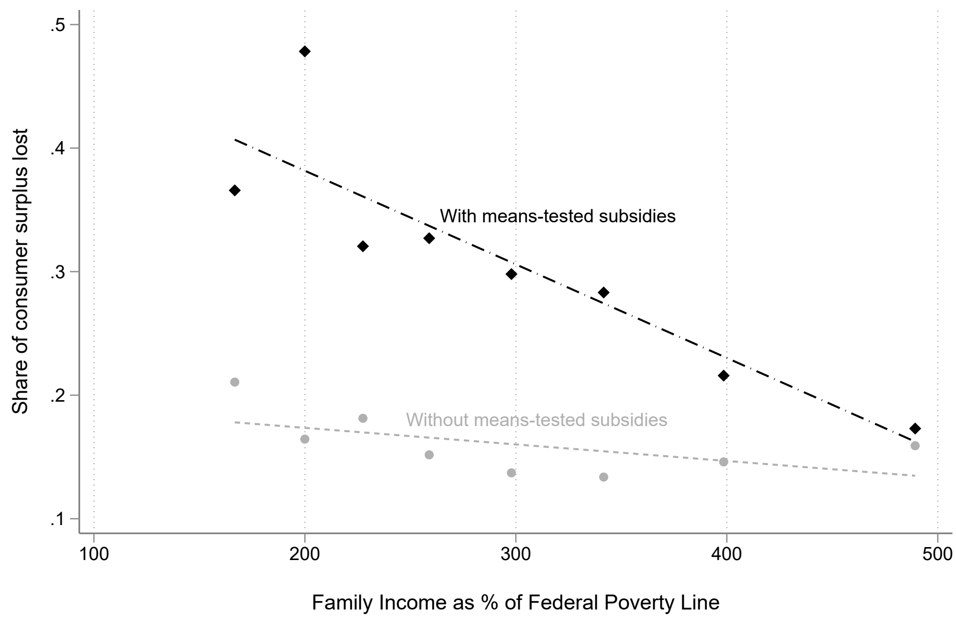

(B) Loss in insurance coverage

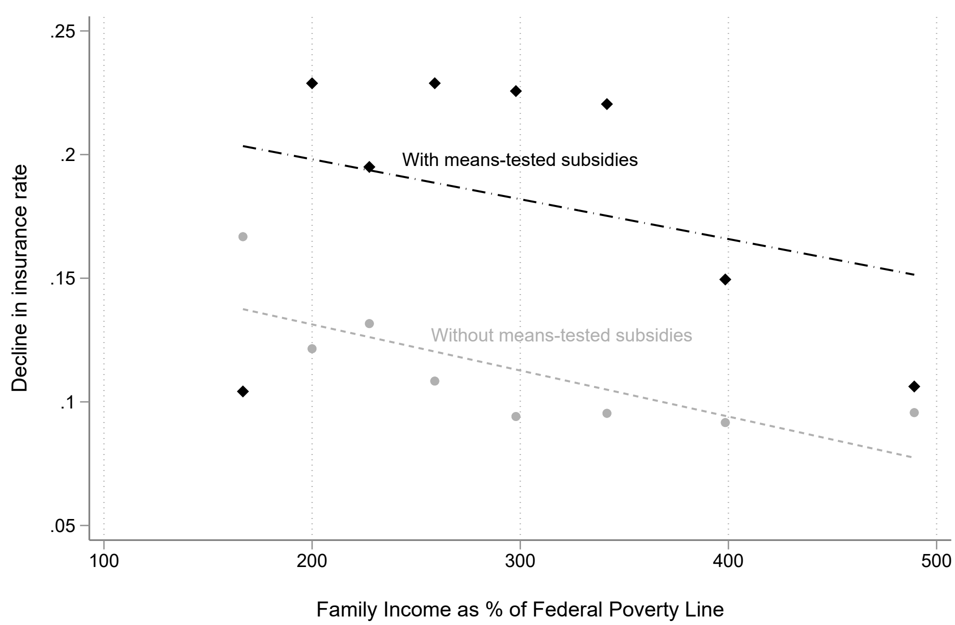

Notes: TBD.

In the first experiment, we increase the number of subsidized consumers in each geographic market without changing the distribution of marginal 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 4 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. One 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 cost effect dominates and 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. In this case, (newly) subsidized consumers exert a positive demographic externality on other consumers.

In the second exercise 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 unsubsidized consumers. Relative to the previous scenario, the effects are more nuanced. While the firms now face more costly consumers, which pushes prices up, the firms also face much more inherently elastic consumers, which pushes prices down unless the subsidies are large enough to mute, or even flip, this channel and make these consumers less price elastic. Case B in Figure 4 illustrates that the forces pushing prices up dominate in our empirical setting. Moving to an environment with more poorer consumers increases list prices. As the dashed line marked with "B" in 4 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

Figure 4: **Mechanism: demographic externality**

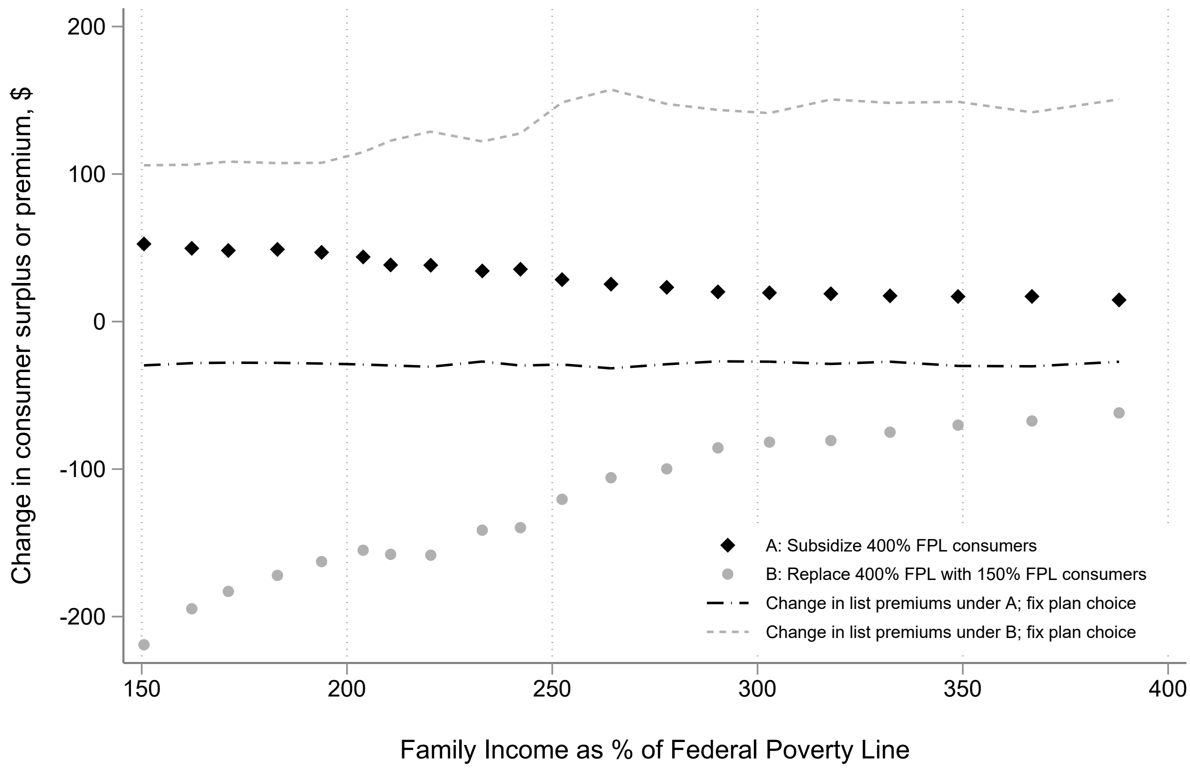

Notes: tbd

400 percent FPL, whose subsidies or utility functions are not directly manipulated in the simulation. In this case, the lower-income subsidized consumers exert a negative demographic externality on other consumers in their local market.

5.4 Role of Subsidy Design

To further understand whether means-testing of subsidies may exacerbate the aggregate and distributional losses from market power, and whether alternative mechanisms could result in both more equitable and more efficient allocations, we re-simulate the model with income-invariant subsidies.[32] We require the change in subsidy structure to be budget-neutral, so that the total (net) spending of the government remains the same as under means-tested subsidies.

[32] We still allow subsidies to vary by age following the same regulatory age curve as the list premiums.

Columns (2) and (3) of Table 4 report the characteristics of the allocation with income-invariant, budget-neutral, subsidies in the case with market power (column 2) and in the case of average cost pricing (column 3). We find that when we retain the firms' pricing powers, the budget-neutral level of income-invariant subsidy would be a voucher of $1,125 for a 20 year old consumer (allowing the voucher to scale with age according to the observed regulatory age curve).

Income-invariant subsidies are substantially more efficient than means-tested subsidies. This is true in our empirical application both in the absence and in the presence of market power. In the absence of firm pricing decisions, the same amount of public spending delivers $629 more in average consumer surplus and 9 percentage points more insurance coverage relative to the allocation with means-tested subsidies. We get similar gains in surplus and insurance when we retain market power. Under market power, most of the extra surplus accrues to consumers. List premiums and firms' profits per buyer decline.

This substantial gain in efficiency, however, comes at a substantial re-distributional cost. While overall enrollment in the market increases, enrollment shifts from the poorest to less poor consumers under universal subsidies. Means-tested subsidies work as intended—they lower prices for the lowest-income consumers and thus attract more of them into 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.

In Figure 5 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, separately for the case with market power (solid line) and with average cost pricing (dashed line). 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 when there is no distortion due to market power. 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.

For any level of preference for redistribution, surplus losses from means-testing are higher when market power is present (because means-testing exacerbates the distortionary effects of market power as we saw before). Thus, in the presence of intermediaries market power,

Figure 5: **Equity-efficiency trade-off in subsidy design**

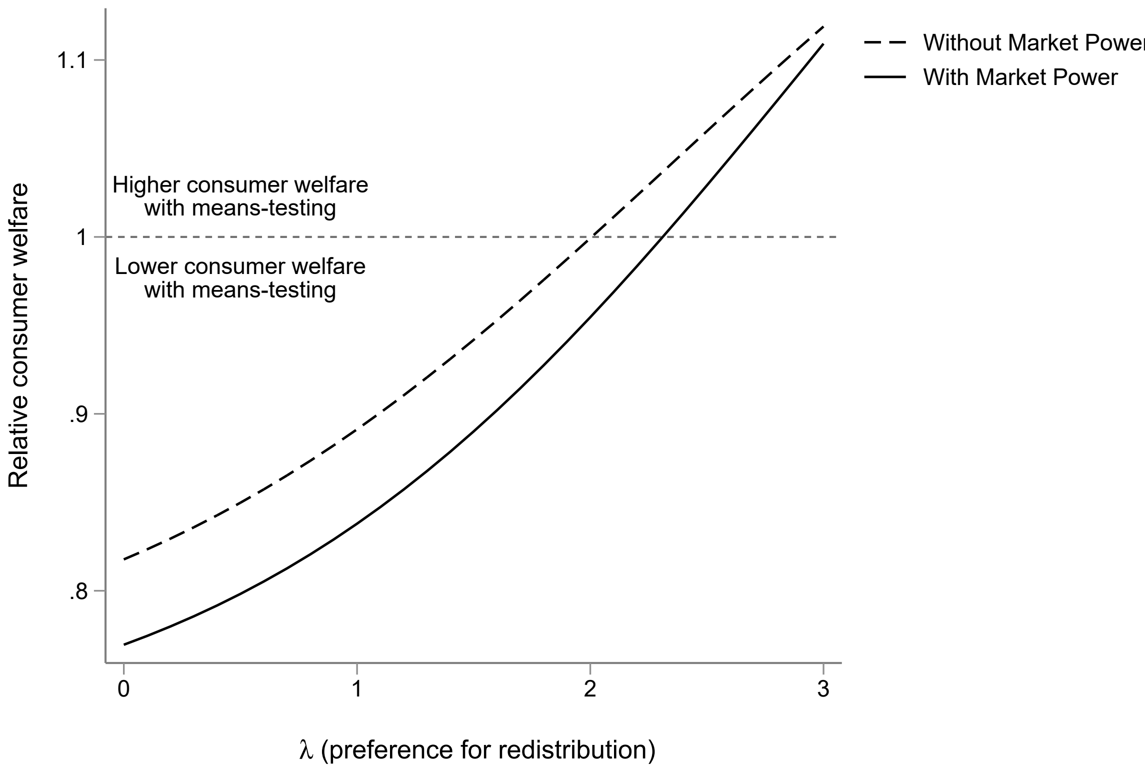

Notes: Figure reports the point-wise (for each value of parameter λ) ratio of average consumer surplus between an equilibrium with means-tested subsidies and flat subsidies, keeping total net government spending fixed. The blue line measures this ratio for the case with perfect competition. The red line measures this ratio for the case with imperfect competition. Average consumer surplus for each level of λ is computed as specified in Equation 25 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 5. At higher levels of λ, social welfare weights increase for consumers with lower income.

the society needs to have a higher preference for redistribution to prefer means-tested rather than income-invariant transfers.

Table 4: **Demographic Externality and Preferences for Redistribution**

	Baseline^	Subsidize consumers above 400% FPL	Replace 400% FPL with 150% FPL consumers	Flat subsidies, budget-neutral ($1,125 subsidy for a 20 year-old)~	Flat subsidies, budget-neutral; $p = AC$ ($1,055 subsidy for a 20 year-old)~
Average across potential consumers ($)					
Consumer surplus	2,495	1,122	3,120	3,139	3,776
Insurer profit	729	1,114	854	762	
Taxpayer cost of subsidies	1,434	2,579	2,189	1,609	1,950
Taxpayer cost net of savings on uncompensated care	614	1,618	1,250	610	701
Insurance rate	0.45	0.53	0.51	0.55	0.68
Average 20 year old list premium (unweighted), $	2,401	2,492	2,382	2,301	1,653
Among consumers buying insurance ($)					
Average cost of covering a buyer	3,993	3,644	4,005	3,283	3,599
Average list premium among buyers	5,618	5,760	5,667	4,676	3,599
Insurer profit per buyer	1,625	2,116	1,662	1,393	
Taxpayer cost of subsidies per buyer	3,196	4,900	4,258	2,942	2,853

^ ~

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.[33] 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 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. Market power of private intermediaries that contract with the government to provide publicly subsidized goods or services is likely

[33] For 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/).

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.

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.

Besley, Timothy, SM Ravi Kanbur, and Banco Mundial, *The principles of targeting*, Vol. 85, Office of the Vice President, Development Economics, World Bank, 1990.

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

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

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

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

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

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.

Cunha, Jesse M, Giacomo De Giorgi, and Seema Jayachandran, "The price effects of cash versus in-kind transfers," *The Review of Economic Studies*, 2019, *86* (1), 240–281.

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," in Tax Policy and the Economy, Volume 19 by MIT Press, September 2005, pp. 39–64.

_ 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," in Empirical Foundations of Household Taxation by University of Chicago Press, January 1996, pp. 135–168.

Handel, Benjamin R., Jonathan T. Kolstad, and Johannes Spinnewijn, "Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets," *The Review of Economics and Statistics*, 05 2019, *101* (2), 326–340.

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.orgWatch the Econimate Video.

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 Robert A. Moffitt, ed., *Economics of Means-Tested Transfer Programs in the United States, Volume I*, University of Chicago Press, 2016, pp. 219–302.

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.

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," Technical Report w26718, NBER Working Paper 2020.

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," in Individual and Social Responsibility: Child Care, Education, Medical Care, and Long-Term Care in America by University of Chicago Press, January 1996, pp. 277–308.

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 Goog as an 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.

ONLINE APPENDIX

A Computational details

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

1. Start with an initial guess of parameter vector, $\theta = 0$.

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.

B Construction of cost multipliers

We construct a matrix of cost-multipliers κ^d using a large commercial health insurance claims data (Optum). 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:

$$ln(RC_a) = \alpha + \beta a + \gamma 1\{a = 0\} + \epsilon_a \tag{26}$$

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 26 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.

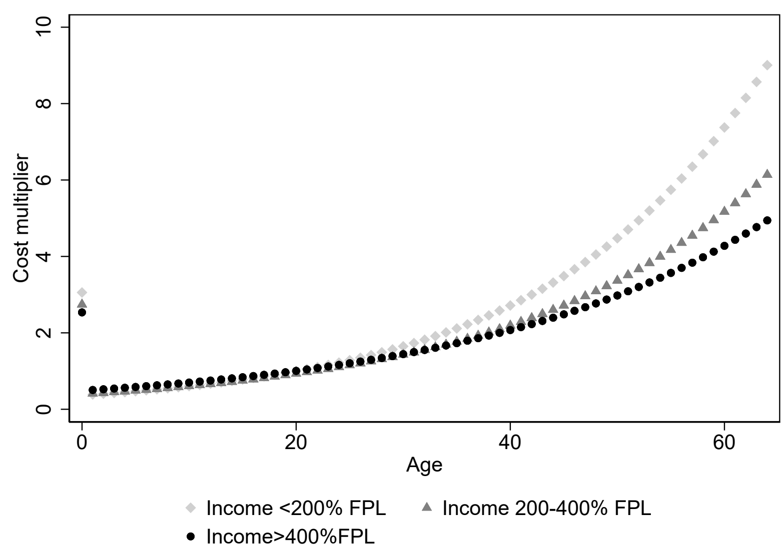

Figure B.1: Matrix of cost multipliers κ^d

C Additional figures and tables

Figure C.1: **Geography of income among potential consumers**

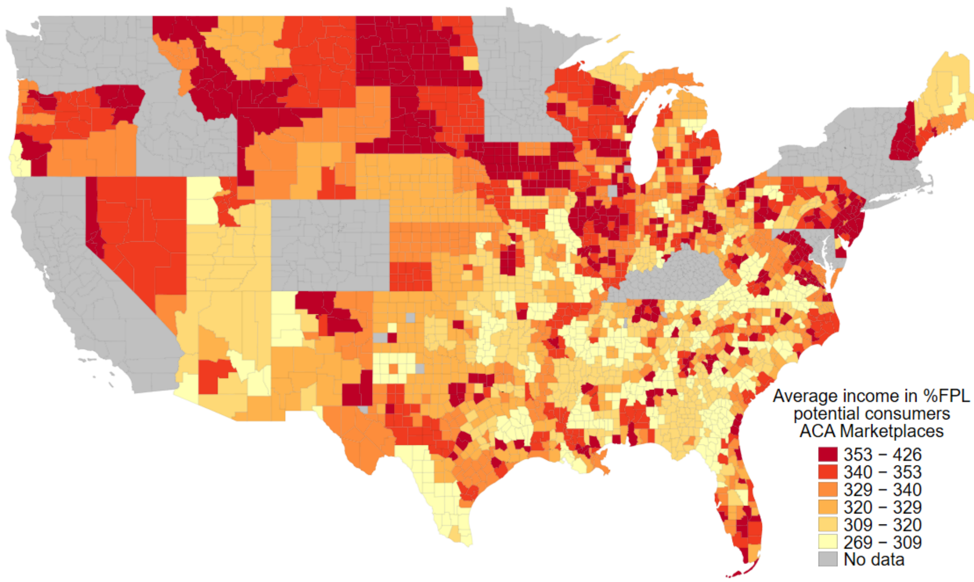

Notes: Map plots the average income in % relative to the Federal Poverty Line of potential ACA Marketplace consumers based on ACS 2017 sample as described in Section 3. Potential consumers in ACS are identified as consumers who have no employer-sponsored or public health insurance coverage and were not eligible for Medicaid expansions in states that expanded Medicaid. States marked with grey are not federally facilitated and do not enter our analysis.

Figure C.2: **Empirical moments and demand model fit for Silver plans**

(A) Observed Silver plan market share

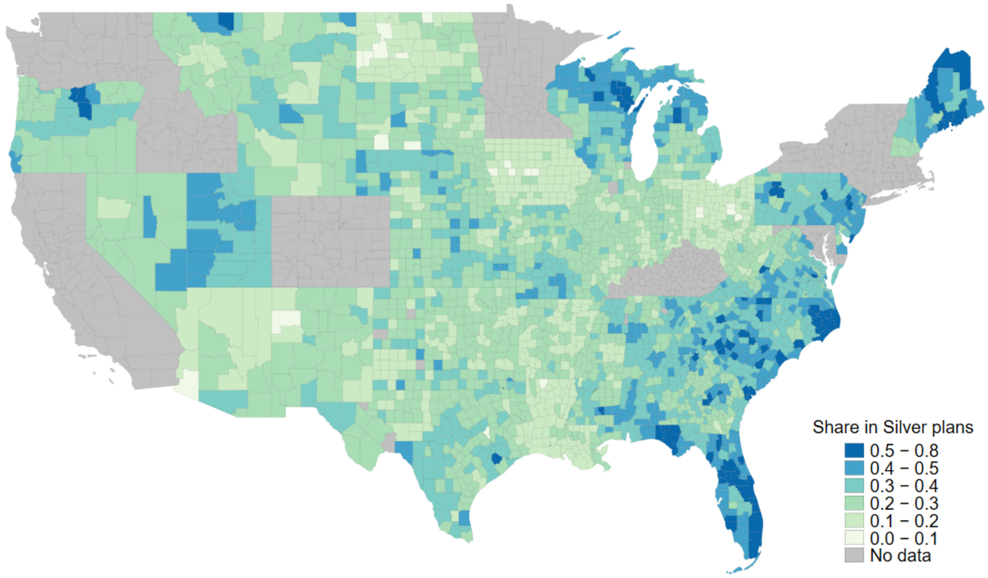

(B) Prediction error of Silver plan market share

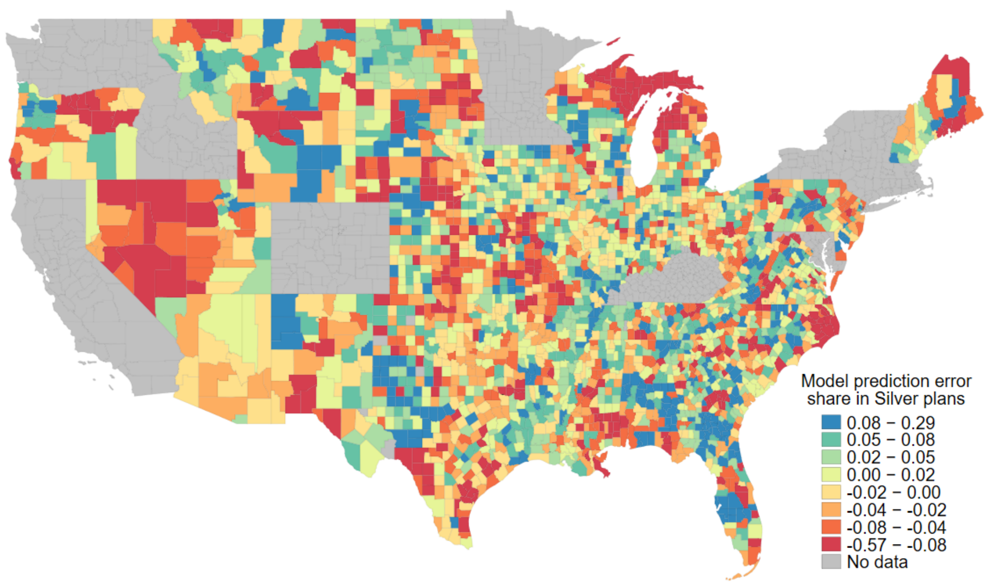

Notes: Map in Panel C.2A plots the share of potential consumers in each county who enrolled in a Silver plan on ACA Marketplaces. The counts of the pool of potential consumers (denominator) were provided by the Kaiser Family Foundation and are based on estimates from national surveys of how many people 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.2B 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. States marked with grey are not federally facilitated and do not enter our analysis.

Figure C.3: **Marginal cost estimates and accounting costs**

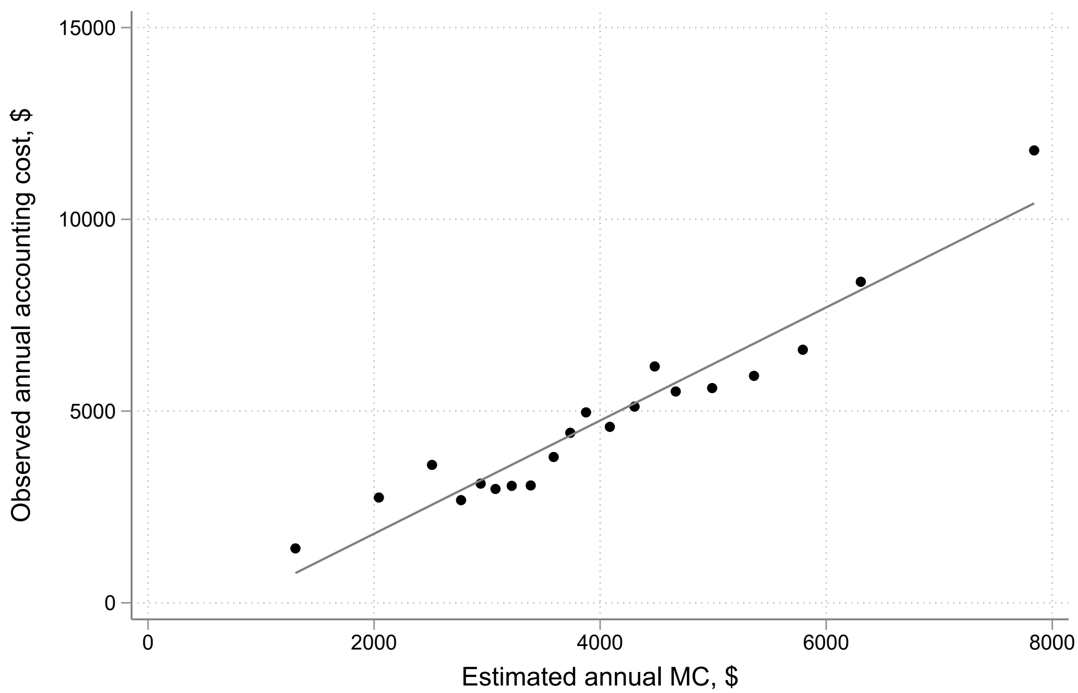

Notes: Figure plots the average accounting cost reported by plans by ventiles of plans' estimated marginal costs. The underlying data includes only 705 plans for which accounting cost reports were available. Estimated marginal cost at the plan level was computed in several steps. First, we invert the 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.4: **Geography of average consumer surplus in ACA Marketplaces**

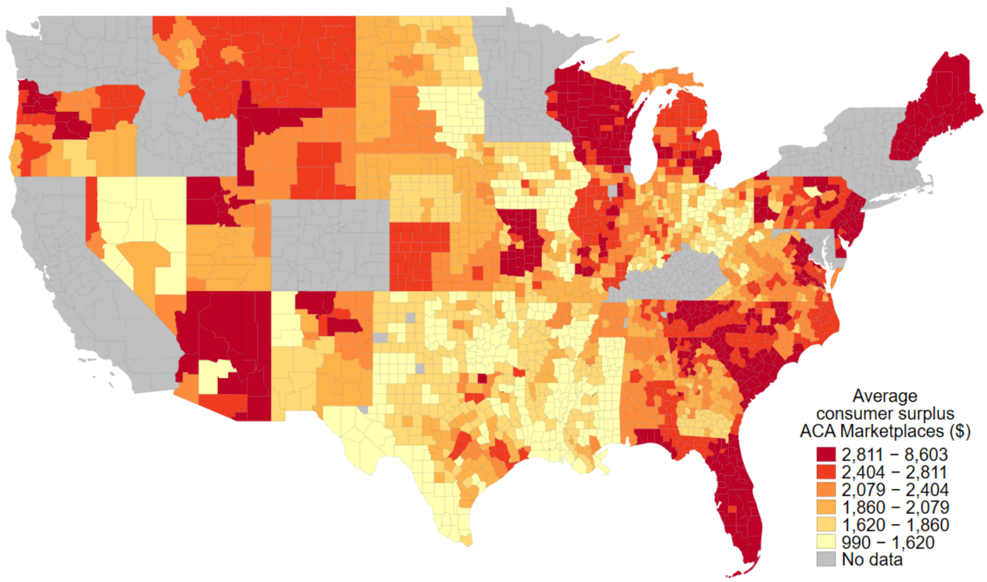

Notes: Map plots the average consumer surplus (in $) from ACA Marketplace program by county. Subsidies and prices are as observed in 2017. The geographic distribution of potential consumer age and incomes is taken as observed in 2017 ACS. Average consumer surplus is computed using Equation 24.

Table 5: **Policy Experiments: by Income Level of Consumers**

	No subsidies		With subsidies	
	Perfect competition	With market power	Perfect competition	With market power – baseline
Average across potential consumers ($)				
Consumers with income $<$ average	909	701	2,294	1,511
Consumers with income $>$ average	4,636	3,992	4,673	3,891
Insurance rate				
Consumers with income $<$ average	0.24	0.11	0.77	0.57
Consumers with income $>$ average	0.49	0.40	0.56	0.41