How Do Supply-Side Regulations in the ACA Impact Market Outcomes? Evidence from California

JOB MARKET PAPER

Richard Domurat*

November 12, 2017

For latest draft, visit: `http://www.richarddomurat.com`

Abstract

This paper examines how two supply-side regulations–modified community rating and risk adjustment–impact the Affordable Care Act (ACA) exchange in California. Using data on individual-level choices and networking providers, I estimate a model of health insurance demand and supply incorporating heterogeneity in preferences, plan characteristics, and costs. The results indicate consumers in this market are highly price-sensitive, and margins are modest (between 3% and 8%). Without risk adjustment, modified community rating in the ACA would lead to a significant reduction in enrollment in desirable plans and in take-up overall. Risk adjustment under the ACA roughly restores relative shares across plans to what they would be without community rating. The reduction in overall take-up from community rating is not impacted by risk adjustment. An alternative risk adjustment method can increase enrollment by 2.6% and would have little impact on government spending. Other policies besides risk adjustment would be needed to address low take-up among price-sensitive, low-cost consumers under community rating.

*University of California, Los Angeles. Contact: Richard.Domurat@ucla.edu. I am grateful to my PhD committee John Asker, Adriana Lleras-Muney, Martin Hackmann, and Wesley Yin for their helpful advice. I also thank the attendees from UCLA Industrial Organization student seminar, as well as John Bertko and Isaac Menashe for their helpful discussions on data and policies.

1 Introduction

The Patient Protection and Affordable Care Act (ACA) was passed in 2010 with the goal of expanding health insurance coverage in the United States. One way of accomplishing that goal was establishing subsidized and regulated insurance marketplaces ("exchanges") for consumers without other insurance options. The regulations in the exchanges aim to balance equity with efficiency, i.e. granting access to those most in need while maximizing consumer surplus. However, despite generous subsidies and guaranteeing access to those willing to pay for insurance, roughly 9% of the U.S. was still uninsured in 2016.[1] Along the same lines, nationwide enrollment was disproportionately lower in plans with either generous coverage or with wide provider networks (Polsky and Weiner, 2015). Some have used these facts to suggest that the exchanges are experiencing adverse selection–that high-cost enrollees are driving other consumers out of quality plans or out of the market entirely. At the same time, low-cost plans that garner high enrollment also face financial difficulties (Holahan et al., 2016). This paper aims to reconcile these facts by examining the impact of two main supply-side regulations in the ACA: *community rating* and *risk adjustment*. How do these policies affect enrollment across plans in the market? And do they in fact lead to adverse selection, or is there another force that is driving these market outcomes? In this paper, I address these questions using a structural supply and demand model of ACA exchanges.[2]

I examine these two specific policies because they largely dictate how enrollment composition impacts prices and thus market outcomes in equilibrium. Yet they have received relatively little academic attention. First consider community rating, which limits what consumer characteristics plans can use to set premiums.[3] In the ACA, premiums are based only on age, and cannot exceed a ratio of 3:1 across ages. By eliminating the volatility in premiums due to preexisting health conditions, community rating increases welfare (Handel et al., 2015). However, as long as the preference for insurance is positively correlated with expected costs, community rating without other correcting policies can create adverse selection. This arises from two effects: 1) a *direct effect*: marginal enrollees face higher premiums as a result of being pooled with higher cost enrollees and opt out, and 2) an *equilibrium effect*: as low-cost enrollees opt out, the price perpetually increases, possibly even unraveling the entire market (Akerlof, 1970). These combined effects can reduce coverage along both the *extensive margin* of getting any insurance (Rabin and Abelson, 2013), and the *intensive margin* of the quality of the plan chosen (Cutler and Reber, 1998; Buchmueller and DiNardo, 2002).

One of the key mechanisms to limit adverse selection across different plan types in the ACA

[1] See: `http://www.kff.org/other/state-indicator/total-population/`

[2] The regulations in the ACA exchanges also apply to the unsubsidized or "off-exchange" individual insurance market. I do not study that market in this study, but the findings of this paper would generally apply there as well.

[3] Guaranteed Issue, which prohibits denials of coverage, is another major part of the ACA reforms. It can be thought of as a part of community rating since it prohibits setting "infinite" prices to certain consumers based on health conditions.

is risk adjustment, the other focus of this study. Risk adjustment prevents the unraveling effect described above by having plans with healthier enrollees make transfers to plans with sicker enrollees. By doing so, the effective marginal costs faced by plans do not depend on enrollee risk composition. Hence, prices do not spiral upward in equilibrium, and high quality plans that attract sick enrollees can remain in the market. In this way, risk adjustment is effectively a Pigovian-style tax that corrects for some of the welfare losses from adverse selection.

This paper quantifies the impact of community rating and risk adjustment in the ACA using data from Covered California, California's ACA exchange. To do this, I estimate a structural model of demand and supply for health insurance, which includes specific details from the regulations. The demand side of the model is built from the common discrete choice framework with preference heterogeneity. Since plan heterogeneity is key to selection in this study, I include the utility generated from hospital and physician networks as a determinant of plan demand. The model also accounts for premium subsidies and the mandate to buy insurance. I estimate the demand model using individual-level plan choice data and provider directories from Covered CA in 2015. The supply side of the model allows for imperfect competition and assumes static profit maximization of all insurers–i.e. a static Nash-in-prices equilibrium. As in the regulations, I include community rating and risk adjustment in firm profit functions. Finally, I use the estimated model to simulate prices and enrollment under counterfactual scenarios where either a) risk adjustment is eliminated, or b) community rating is eliminated. I also simulate an alternative method of risk adjustment which has different properties from the method currently employed. This alternative somewhat resembles risk adjustment in Medicare Advantage, a similar market in the U.S. for the elderly.

There are a number of new findings from this study. First, consumers in this market are highly price-sensitive. For a typical plan, a premium increase of $1 is associated with a decrease in shares between 3.5% and 4.5%. While these are quite high relative to most health insurance products, they are close to those found in similar markets (Shepard, 2016; Finkelstein et al., 2017a).[4] This implies profit margins on these plans are slim and range between 3% and 8% of total revenues, excluding any fixed costs. Second, using counterfactual simulations, I find that ACA *community rating* without risk adjustment would reduce the share of enrollment in generous plans and enrollment in the market overall, consistent with adverse selection. For example, the market share of non-HMO plans would decrease by 15%. Similarly, shares for plans with the most generous cost sharing would decrease by 85%. Community rating would also lead to a decline in the total number of insured households by 14.2%.[5] Adding *risk adjustment* as implemented in the ACA has little effect on

[4]See footnote 56 in Ho (2006) for a good discussion of estimated elasticities in health insurance markets. My estimates are roughly between 3 and 6 times the elasticities in those markets with group insurance. The individual market in the ACA would be expected to have higher elasticities for a number of reasons (lower incomes, larger choice sets, transparency and plan standardization, etc.).

[5]This decline in coverage should not be interpreted as welfare reducing. The main motive of community rating is to reduce reclassification risk from unpredictable health shocks. This study doesn't capture the true variation in underlying costs and so cannot capture the true benefit of eliminating this form of risk.

total enrollment but does restore over half of the effect on shares from community rating. Under ACA risk adjustment, there is a small degree of persistent adverse selection. For a typical Gold or Platinum non-HMO plan, the marginal consumer has expected health expenditures \$2 less than the inframarginal average, net of risk adjustment. Finally, I simulate an alternative form of risk adjustment that can eliminate this persistent adverse selection. For each plan, this method flattens the perceived marginal cost curve and sets it equal to the plan's average cost in the population. In this case, I find that total enrollment would increase by 2.6%, with a disproportionate increase in plans that attract sicker enrollees. Net government spending would be roughly unchanged under this method. Premiums would decline and therefore so would premium subsidies, offsetting increases in government payments for risk adjustment.

The policy implications of these findings are threefold. First, the price elasticities in Covered CA suggest exchanges can be an effective mechanism to deliver care at competitive prices to consumers. Despite relatively high measures of market concentration, the high price-sensitivity in the market limits markups, and premiums are set close to costs. These low margins might partly explain why there are so few plans in other ACA markets around the country. Second, risk adjustment in the ACA plays an important role in mitigating adverse selection along the intensive margin. Without risk adjustment, generous plans and those with well-known providers would face significant cost pressure from high-risk enrollees that would drive down their enrollment. Risk adjustment does not, however, substantially reduce the decline in overall coverage caused by community rating. Any policies aimed at universal coverage will need to consider the high price-sensitivity of low-cost, low-demand consumers, as has been suggested in other work (Tebaldi, 2016; Panhans, 2017; Finkelstein et al., 2017a). Finally, the alternative risk adjustment method described above can increase the total rate of coverage with a negligible impact on government spending. However, the welfare consequences of this change, especially in the long run, are not clear.

On a broader level, this study can reconcile the nationwide observations described earlier: 1) a continuing high rate of uninsured, 2) the popularity of low generosity and narrow network plans, and 3) financial difficulties for low-cost plans despite popularity. Enrollment is high among low-cost plans not because of adverse selection but because consumers prefer low-price options. This high price-sensitivity puts downward pressure on prices and limits profit margins for all plans. And finally, despite the mandate and subsides, community rating can decrease total enrollment in ways that cannot be corrected by risk adjustment. The use of community rating and risk adjustment extends well beyond the individual market established by the ACA. Medicare Advantage covers roughly 16 million elderly Americans in a similar market that uses both community rating and risk adjustment. Countries such as Germany, Switzerland, and the Netherlands also use these policies. The U.S. and other countries are increasingly relying on private markets to deliver public health care benefits. This paper can help policymakers better understand how these common regulations shape insurance markets and ultimately benefit consumers.

Contributions

This paper makes contributions to three bodies of academic literature. The first is related to existing work on ACA regulations. A number of studies have examined the demand-side regulations like the mandate and subsidies using similar approaches to this study (Hackmann et al., 2015; Tebaldi, 2016; Jaffe and Shepard, 2017). There has also been some work on the community rating and risk adjustment policies in the ACA, but it is either based on data from other markets (Handel et al., 2015; Layton, Forthcoming; Ericson and Starc, 2015) or discussed in a theoretical sense (Geruso and Layton, 2017; Layton et al., 2017). The paper by Handel et al. (2015) is particularly relevant to this study since they examine the welfare implications of community rating and different corrective policies. Importantly, they find community rating would cause some plans to unravel but the benefits dominate welfare losses since community rating eliminates *reclassification risk* (exposure to high premiums given health shocks). That paper also addresses risk adjustment, but in less detail. While these papers have all informed about how policies in the ACA impact welfare and other outcomes, they have not studied ACA risk adjustment in detail or they have used data from other markets. Hence, this paper adds insight to how these policies shape the market in practice.

Relatedly, this paper fits into a long history of work looking at risk adjustment in general (Newhouse et al., 1989; Glazer and McGuire, 2000; Brown et al., 2014; Einav et al., 2016). That literature has largely focused on the information that should be included when quantifying enrollee risk–i.e. risk scores. This paper, conversely, assumes risk scores are perfect and observable, and instead focuses on *how* they are used in the risk adjustment process. As I show in this paper, this has implications on market outcomes that are generally ignored in this strand of literature.

The second contribution is to the extensive body of work on adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976; Cutler and Reber, 1998; Buchmueller and DiNardo, 2002; Einav et al., 2010; Shepard, 2016). These papers have theoretically or empirically shown how community rating, or hidden information more generally, can lead to a reduction in coverage in high quality plans. Notably, Einav et al. (2010) identify sufficient conditions for adverse selection using the slope of the marginal cost curve. Their application is in a competitive environment with just two options in the choice set, and reduced form cost and demand functions. To add to that work, I introduce a metric of adverse selection in the presence of modified community rating and risk adjustment as in the ACA. My measure is derived from a structural model and hence allows for imperfect competition and many options in the choice set. I use this measure to explicitly quantify adverse selection in the market using the estimated model.[6] Additionally, this is the first paper to look at adverse selection at the plan level for ACA markets.

Finally, this paper contributes broadly to work examining health plan competition (Town and

[6]The primary motivation of the adverse selection test in Einav et al. (2010) is to subsequently measure welfare losses. I make simplifying assumptions in my model that would underestimate welfare measures and hence I do not make welfare comparisons. Although, it could be added in the future with a refined choice model. Layton (Forthcoming) measures the welfare implications of risk adjustment when risk scores are imperfect.

Liu, 2003; Ho, 2006; Curto et al., 2014). This paper is the first to measure the degree of competition in an ACA market that incorporates provider network differentiation. I include networking hospitals and physicians in the demand for insurance, and find both are important. Broader network plans attract higher cost enrollees, which is relevant for measuring selection. The estimates are used to quantify demand elasticities across all plans in the market, which I find to be quite high. Additionally, the estimation makes two methodological contributions that can be applied for other demand models in ACA markets: 1) I specify a method of estimation that accounts for the *within-plan* price heterogeneity created by the ACA pricing policies, and 2) I exploit exogenous variation in insurance choice sets to identify the preference for getting coverage.

2 Background

This section discusses the theoretical motivation for regulation in insurance markets, and the details of those regulations in the ACA.

2.1 Theory of Community Rating and Risk Adjustment

Before discussing the details of the ACA, it is useful to think about how community rating creates adverse selection, and how risk adjustment can offset it. For this exposition, I use the graphical model presented by Einav et al. (2010). While it assumes perfect competition unlike in this paper, it still conveys the main economic forces.[7] I present the model with a single homogeneous insurance product relative to being uninsured, but it can also loosely be thought of as between plans of different desirabilities–though there are slight differences (Layton, Forthcoming).

Consider a market of consumers indexed by health status θ. For each θ, consumers have an expected health expenditure $c(\theta)$ and willingness-to-pay for insurance $v(\theta)$, both of which are increasing in θ. The implied demand and marginal cost curves are given in panel (a) of Figure 1. Note that the marginal cost curve is decreasing since $v(\theta)$ and $c(\theta)$ are positively correlated.

In a competitive market with full information (no community rating), prices would be set for each θ and equal to the cost of coverage $c(\theta)$. All consumers with valuations greater than the price would buy coverage. Since consumers purchase coverage if and only if their valuations exceed the cost, this is the efficient allocation. This is represented by Q_{Eff} in panel (a) of Figure 1, where the marginal cost curve intersects the demand curve.

Next consider the effect of community rating (or "hidden information" more generally). In this case, firms can only charge a single price for the entire population of θs. Since markets are assumed to be competitive, firms make no profits and prices are set equal to the average cost of the insured pool. The point at which the average cost curve equals the demand curve gives the equilibrium

[7] See Ericson and Starc (2015) for a discussion of how imperfect competition changes the results from community rating.

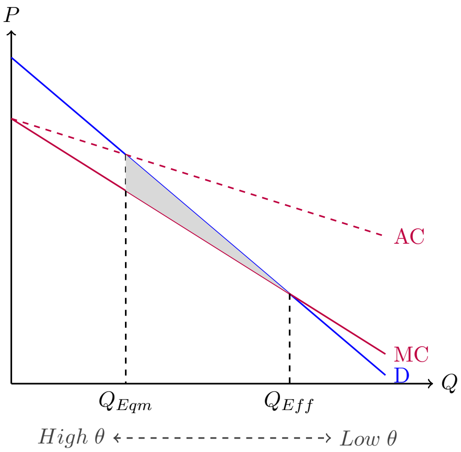

(a) Total Effect

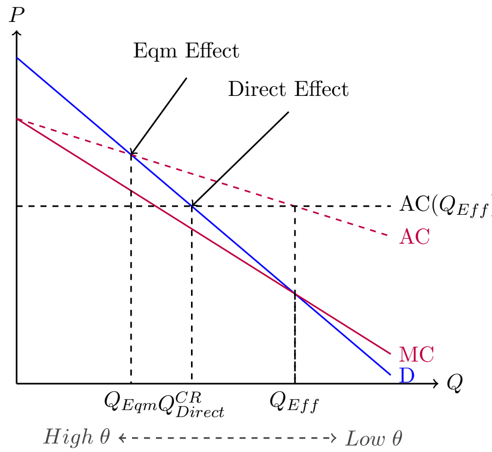

(b) Direct and Equilibrium Effects

Figure 1: Adverse Selection from Community Rating

allocation under community rating. This is given by Q_{Eqm} in the same figure. The shaded region between these two quantity levels gives the welfare loss from community rating.

In panel (b) of Figure 1, I decompose the effect of community rating on quantities into the two effects described in the introduction: a direct effect and an equilibrium effect. The direct effect gives the reduction in quantities as a response to pooling enrolled consumers, holding the underlying composition constant. This is where the demand curve intersects the average cost of the efficient quantity Q_{Eff}. I label this point as Q^{CR}_{Direct}. The equilibrium effect is that under this new allocation, the average cost has also increased and hence prices exceed the willingness-to-pay for marginal consumers at Q^{CR}_{Direct}. In equilibrium, quantities reduce until average costs equal demand at Q_{Eqm}. This decomposition is useful for two reasons: 1) it is related to how risk adjustment can correct for adverse selection, as will soon be clear, and 2) it is the second effect that can lead to the type of unraveling "death spirals" many policymakers are concerned about. For example, adding convexity to these curves could lead to a large unraveling effect.

In Figure 2, I present how risk adjustment impacts prices and quantities in this simple model. I use a simple "textbook" version of risk adjustment that has been used in other Economics literature (Mahoney and Weyl, 2017; Geruso and Layton, 2017). It need not be superior to alternatives, such as the method in the ACA (Layton et al., 2017), but it allows for an intuitive exposition.[8] Risk adjustment in this textbook form aims to match each firm's perceived marginal cost to their average cost in the population. As noted in the introduction, it can be thought of as a Pigovian tax to roughly correct for the equilibrium effect created by community rating. For each enrollee that is higher than average risk, plans receive payments equal to the difference between the marginal cost and the average cost of that enrollee. For each enrollee that is lower than average risk, plans make

[8]Mahoney and Weyl (2017) discuss the impact of risk adjustment under imperfect competition.

payments equal to the difference. The marginal cost curve perceived by firms after risk adjustment is depicted in the figure as the horizontal line equal to $AC(Pop)$. Under perfect competition, the equilibrium quantity will be where $AC(Pop)$ intersects demand, and is given by Q_{RA}.[9]

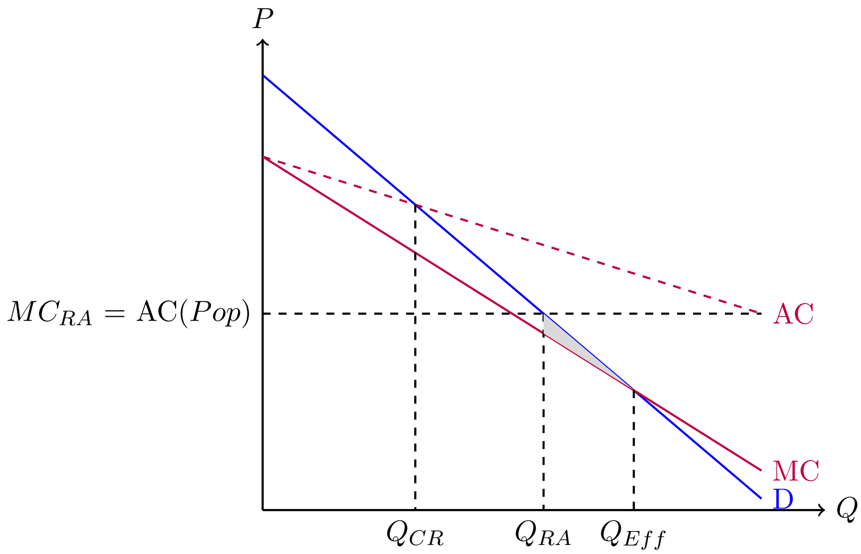

Figure 2: Adverse Selection with Risk Adjustment

There are two observations in this figure worth highlighting. First, risk adjustment in this model reduces the welfare loss from community rating and increases the equilibrium quantity to be closer to Q_{Eff}. However, this need not always be the case. Second, this form of risk adjustment creates a constant marginal cost curve that inhibits the equilibrium effect–i.e. unraveling–described in the prior figures.

In reality, the market is much more complicated by having many multi-product firms, imperfect competition, complex heterogeneity in costs and preferences, and subsidized demand. Hence, to properly quantify the effects of these policies, I use a much richer model with all of these features. Motivated by this simple environment, I estimate the same objects (demand and marginal/average costs) in the market to quantify the degree of adverse selection.

2.2 ACA and Covered CA Regulations

This section gives an overview of the ACA policies most relevant to this study. A more detailed discussion can be found in Layton et al. (2017).

[9] Layton (Forthcoming) discusses how this differs when considering two plans instead of insured vs uninsured. The main difference comes from the fact that the "other" option is also getting risk adjustment. As a result, the perceived marginal cost curve is still flat but will be shifted down. This is because the average cost curve at the right-side intercept is really the average cost of the population in the high plan relative to the *healthiest enrollee* in the low cost plan (not average cost in the population for the low cost plan).

2.2.1 Exchange Motivation

The individual health insurance market in the U.S. is for consumers who do not have coverage through other sources (e.g employer or other public programs). Before 2014, the individual market was characterized by a lack of affordability, poor coverage, and even outright denials of coverage (Claxton et al., 2016). This was especially true for anyone with preexisting health conditions. Hence, a major goal of the ACA was to fix any potential market failures in the individual market through a regulated environment. The most commonly discussed of these regulations are: guaranteed issue, "modified" community rating, the mandate, and premium subsidies. The two former policies target reclassification risk and equity motives but create the type of adverse selection described above. The mandate and the subsidies, joint with risk adjustment, aim to mitigate that adverse selection.

Modified community rating in the ACA is a variant of the simple form described above. Premiums are allowed to vary across consumers but only based on age. Hence, community rating operates mainly across health conditions and genders. Additionally, premiums across ages cannot exceed a ratio of 3:1. To the extent that underlying costs across ages vary by more, this would imply there is community rating across ages as well. Guaranteed issue is a large part of the reform, but for the purposes of this paper, it can be considered part of community rating (where denials were effectively infinite prices). For reasons discussed in the prior section, these provisions on their own would lead to adverse selection along both extensive and intensive margins.

2.2.2 Extensive Margin Selection and Remedies

To address adverse selection along the extensive margin, the ACA requires that all individuals have health insurance–i.e., the individual mandate. Anyone that does not have health insurance pays a penalty, which increases with income. For lower income households for which premiums would be too high, the policy provides subsidies through an Annual Premium Tax Credit (APTC, see Section 2.2.5). Despite these policy remedies, the extent of adverse selection along this margin is unclear. Nationwide, only 40% of those eligible chose to enroll in a plan in 2015.[10]. It's theoretically possible that some of this was from high risk enrollees driving up prices, and is empirically true in some states (Panhans, 2017).

2.2.3 Intensive Margin Selection and Risk Adjustment

Community rating creates adverse selection along the intensive margin in a similar manner. Plans that have more generous coverage and those with better known providers attract the sickest enrollees, and would experience disproportionate increases in premiums in response. To address this

[10]See: `http://www.kff.org/health-reform/state-indicator/marketplace-enrollment-as-a-share-of-the-potential-marketplace-population-2015`

margin, the ACA uses risk adjustment.[11] The goal of risk adjustment in the ACA is for "plan premiums [to] reflect differences in scope of coverage and other plan factors, but not differences in health status" (Kautter et al., 2014). Hence, allowable differences would be due to cost sharing, provider payment rates, care management strategies, or administrative costs.

Risk adjustment in practice is generally calculated with two pieces of information: 1) some benchmark level of expected spending for a plan, and 2) some measure of the relative *expected* cost of a person that is invariant to the plan, called the "risk score."[12] For each enrollee that is relatively sicker, plans receive a risk adjustment transfer in addition to the premium paid by be consumer, and vice versa for healthy enrollees. Risk adjustment payments for person i in plan j usually take the form:

$$RiskAdj_{ij} = RiskScore_i * Benchmark_j$$

In the simplified "textbook" model in the prior section, the benchmark is the average cost among the population if all consumers were enrolled in plan j. The risk score in that model is the (plan invariant) expected cost of consumer i divided by the expected cost in the population.

In practice, the ACA uses slight variants of both of these objects. For a benchmark, the ACA uses the average premium in the state, standardized for the local region and generosity (i.e. metal tier). This has a number of advantages for policymakers. First, it ensures budget neutrality since it is the same across plans. Second, it incentivizes cost reductions in the long run. For plans that receive risk adjustment transfers, those with low costs will receive proportionately higher risk adjustment payments. Similarly, high cost plans do not receive the full wedge between the average and marginal cost, so face a degree of residual adverse selection. However, for plans that attract healthy enrollees and make transfer payments, the opposite is true. This mechanism requires larger payments from low-cost plans than the textbook model described above. In the short run, this puts upward price pressure on low-cost plans relative to the other model.[13]

The risk score also differs from the model described above. Instead of being costs relative to the whole population, the ACA uses costs relative to the insured population. First, there is a practical matter which is the uninsured have no medical claims data from which to measure their risk. Second, this is also partly because the intention of ACA risk adjustment is to correct intensive margin selection, leaving the other policies to address the extensive margin.[14]

[11] While risk adjustment can also partially address adverse selection along the extensive margin as in the prior section, this is not the motive under ACA risk adjustment.

[12] In practice, this is usually calculated by projecting expenditures on demographic and health information, and using that model for prediction. There is a long literature and constant debate on what should be included in that model, even in the ACA. Moreover, there have been studies that show health information is insufficient to eliminate adverse selection due to other unobservable factors (Glazer and McGuire, 2000; Einav et al., 2013; Shepard, 2016).

[13] This could also have long-run social benefits since it incentivizes efficient low-cost plans to attract sicker enrollees. These long-run incentives are important but beyond the scope of this study.

[14] Risk adjustment in the ACA combines enrollment from the on-exchange (i.e. subsidized) and off-exchange (unsubsidized) markets. In this paper, I only consider risk adjustment within the exchange, given data availability. Enrollment and premiums in these two groups likely differ due to income differences. However, evidence from MEPS

Finally, risk adjustment in the ACA also discounts for revenue from age-based community rating. Since plans receive different premiums for different ages, this variation is subtracted from risk adjustment payments. The exact formula and other details are in Pope et al. (2014). In this paper, I simulate an alternative method similar to the one described in the prior section. These different methods have different implications on the market–especially in the long run–and the optimal choice ultimately depends on the objective function of regulators.

2.2.4 Other Standardizations in the ACA and Covered CA

The ACA has additional regulations to enhance competition in the event that markets are dominated by a few firms or if there is any residual adverse selection. Plan generosities are standardized to 4 "metal tiers" which indicate the share of total expenditures that is covered by the plan, known as the actuarial value (AV). Standardizing plans in this way reduces differentiation and increases transparency, both of which should decrease markups. Covered CA adds other regulations to their exchange beyond the federal requirements, which further limit adverse selection and reduce price volatility. First, plan benefit designs are largely standardized, even beyond the AV requirements. For example, copays for a primary care visit are the same for all plans within the same tier. Similarly, Covered CA ensures that all plans meet high provider network adequacy standards. Finally, firms must offer their plans in all metal tiers to prevent segmenting the market.

2.2.5 Plan Competition and Pricing in Covered CA

In Covered CA, premiums are set at the region level. There are 19 rating regions in the state made up of counties or aggregates of counties.[15] While prices within a rating region are all the same, whether or not the plan is actually offered is determined at the zip code level. While most plans are offered in the entire region, partial coverage areas occur when plans have geographic service constraints given networks.

In each rating region, an insurer can offer one or more of three networks types: a Preferred Provider Organization (PPO), an Exclusive Provider Organization (EPO), or a Health Maintenance Organization (HMO). These different levels of managed care, sorted from least to most, put constraints on what health services an enrollee can access.[16] All else equal, less constraints are associated with higher utility but also higher costs.

About halfway through the year, plans set their premiums for the following calendar year. In California, the premium variation by age is determined by an exogenous scaling profile set by

shows the impact of income on utilization (i.e. cost) is small relative to the demographics used in this study.

[15]A map of the rating regions can be found on page 26 of the Covered CA Rate Book for 2017: `https://www.coveredca.com/news/PDFs/CoveredCA-2017-rate-booklet.pdf`

[16]PPOs have a preferred set of providers in-network where services are offered at a lower negotiated price. For EPOs, an enrollee is not allowed to see any providers outside of this network. Finally for HMOs, enrollees have additional restrictions on the types of providers (or services) an enrollee can access, even if they are in-network.

regulators which preserves the 3:1 age requirement. This rating profile is presented in Figure 3. Plans submit their premiums for a 21 year-old enrollee only, and the rest of the rates are determined by the rating profile.

Figure 3: Pricing Profile in Covered CA

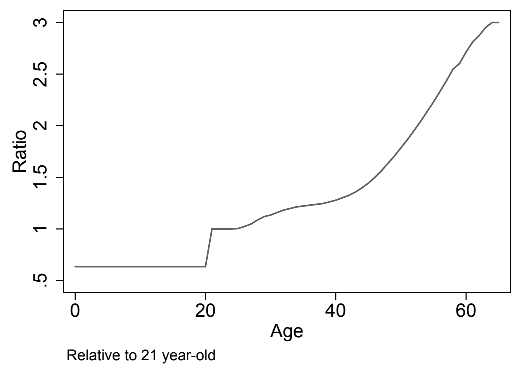

Note: This plot shows how premiums for a plan vary by age. Plans set a price for a 21 year-old and the premiums for other ages are scaled by this function. The function represents the expected health care expenditures by age, but compressed to preserve the 3:1 ratio requirement in the ACA.

Premiums for the same plan also vary across households due to the aforementioned subsidies (APTC). Households with an income less than 400% of the federal poverty line (FPL) are eligible for a subsidy for enrolling in any plan. In 2015, this income threshold was $47,080 for singles and $97,000 for a household of 4. The subsidy is the amount of money that ensures the second lowest available silver (SLS) plan costs exactly a certain share of the household's income. That share increases continuously with income from 2% to 9.5%. Note that this implies two subsidized individuals with the same income will face the same premium for the SLS, regardless of age. Since the APTC is household-specific, it does not affect the relative difference in prices between any two plans.[17] It does, however, affect the price of all plans relative to remaining uninsured, in effect increasing the effective penalty of the mandate.

In summary, the difference in prices between any two given plans is higher for older consumers. The relative prices of *all* plans compared to being uninsured varies with income, such that the high ATPCs from lower incomes makes getting any plan more attractive. Since the APTC sets the SLS premium regardless of age, older households get larger APTCs as well.

[17]The small exception is when APTC is maxed out and exceeds a plans premium, as with some low-cost Bronze plans. In this case the premium is $1 per enrollee and the APTC is partially unused.

3 A Model of ACA Exchanges

In this section, I describe a full model of demand and supply for insurance in the ACA exchanges that I use for the empirical analysis. The timing of events in the model is as follows. Plans are endowed with demographic-specific reputations and networks of physicians and hospitals. Before the start of the year, plans simultaneously choose premiums in a static Nash-in-prices equilibrium. Finally, at the start of the year, all potential enrollees choose the plan that gives them the highest utility given expectations about medical needs.

3.1 Demand for Health Insurance

The insurance choice model is similar to other discrete choice models in the Industrial Organization and Health Care Economics literatures (Berry et al., 2004; Goolsbee and Petrin, 2004; Town and Liu, 2003; Ho, 2006). The key features of this model are that plan choice depends on *observable* characteristics of the plans and enrollees, while there is unobserved heterogeneity in the preference to remain uninsured–as in the "nested logit" model (Berry, 1994). I treat all plans as part of one "nest," and the option to remain uninsured as its own separate nest. In the context of random coefficient models, this can analogously be thought of as a random coefficient on the outside option of remaining uninsured. The reason for using this structure is it allows for flexible substitution patterns along the extensive margin, and relaxes the *independence of irrelevant alternatives* (IIA) property of logit models.[18]

Potential enrollees are offered a menu of plan options and choose that which yields the highest idiosyncratic utility. The indirect utility of household i for plan j in market t is given as follows:

$$u_{ijt} = \underbrace{\alpha_i p_{ijt}}_{\text{Prem Disutility}} + \underbrace{EV^H_{ijt} + EV^P_{ijt}}_{\text{Network Utility}} + \underbrace{x'_{ijt}\beta_i}_{\text{Other Obs. Utility}} + \underbrace{\xi_{jt} + \zeta^{Ins}_{it} + (1-\sigma)\varepsilon_{ijt}}_{\text{Unobservable Utility}} \tag{1}$$

where p_{ijt} is the net-of-subsidy premium and x_{ijt} is a vector of plan attributes (e.g. carrier and tier). EV^H_{ijt} and EV^P_{ijt} are the expected utilities derived from hospitals and PCPs respectively and are explained below in Section 3.2.[19] ξ_{jt} is an unobservable component of utility for plan j which is common to all households in market t, and can take on any distribution. ε_{ijt} is the idiosyncratic

[18]Intuitively, IIA means that within observable demographic groups, as prices increase, marginal consumers substitute to other options proportional to the group's market share. Adding in unobserved heterogeneity conversely allows for disproportionate substitution to certain options in the choice set. In the context of this model, it allows for consumers to disproportionately substitute to other insurance options rather than being uninsured. In Covered CA, roughly half of the market remains uninsured. Removing unobserved heterogeneity would imply roughly half of the marginal enrollees would choose to be uninsured rather than another insurance option when facing price increases, which is unlikely to be true.

[19]Note the functional form assumes separability of utility from hospitals and doctors. It is possible these interact in a non-additive way. For example, in 2014 Anthem Blue Cross contracted with Stanford Health Care hospital but *not* the physicians group. One could argue the hospital being in-network generated relatively little utility without the complimenting physicians that worked at the hospital. These cases are rare and are beyond the score of this paper.

utility for each plan, and is i.i.d. Type I extreme value. ζ_{it}^{Ins} is the household-specific utility of having insurance and is common across all plans. As in the nested logit model, the distribution of this variable is characterized by $\sigma \in [0, 1]$, and satisfies the property that $\zeta_{it}^{Ins} + (1 - \sigma)\varepsilon_{ijt}$ is itself generalized extreme value (Berry, 1994; Train, 2009). σ can be thought of as the degree of correlation in the unobserved utility of all plans relative to the outside option. If it is 0, then there is no correlation in the unobserved utilities and it becomes a standard logit model. In the context of random coefficients, σ can be thought of as the relative variance of the random coefficient on the insurance indicator variable. As this variance approaches its maximum, there is no substitution along the extensive margin.

Preference parameters are indexed by i to indicate heterogeneity across households. I assume this heterogeneity is purely based on observable household characteristics so $\beta_i = \mathbf{B}z_i$ where z_i is a vector of household demographics and $\mathbf{B}$ is a matrix of parameters to be estimated. Since the ACA exchange population excludes the elderly, many chronically disabled, and those with the lowest incomes, this assumption is more plausible than in other government programs.

As discussed in section 2.2.5, net premiums are a function of income and age and can be written succinctly as:

$$p_{ijt} = \theta_i^P p_{jt} - \tau(FPL_i, p_{(2)t}) \tag{2}$$

where p_{jt} is the premium set by plan j for a 21 year-old. θ_i^P is the rating profile set by Covered CA ranging from 1 to 3 based on age. Since i is the household, this term is the sum of rating profiles for all members in the household. $\tau(.)$ is the premium subsidy (APTC), which is bounded above so that p_{ijt} is constrained to be at least \$1 per enrollee.

The outside option is the choice to not enroll in any of the available insurance plans and remain uninsured. The associated utility depends on individual characteristics z_i and is given by:

$$u_{i0t} = z_i'\omega + \varepsilon_{i0t} \tag{3}$$

where ε_{i0t} is also distributed i.i.d. Type I extreme value. z_i includes a constant and income so this is the reduced-form utility of remaining uninsured, which includes the penalty from the individual mandate. Note that this is a degenerate "nest" since it only contains a single option, so there's no nest-level heterogeneous coefficient.

Given that this model follows a nested structure, it is helpful to think of the choice probabilities as two different levels (Train, 2009)–though to be clear, this doesn't reflect the timing of the decision process. The different levels can be thought of as an "upper model"–the choice of which nest–and a "lower model"–which choices within the nests. In the context of this study, the upper model is the decision whether or not to buy *any* plan. The lower model is *which* plan to purchase conditional on buying a plan. In the nested logit framework, "nest 0" should be thought of as remaining uninsured, and the other nest contains all plans in the market. Hereafter, I refer to the upper model as the "take-up" model and the lower model as the "plan choice" model.

Working backwards, first consider the plan choice model. Household i chooses plan j if it generates the highest utility–i.e. if $u_{ijt} \geq u_{ikt}\ \forall k \neq j$. Let $s_{ijt|Ins}$ be the probability that household i choses plan j, conditional on choosing to buy insurance. Given the distributional assumptions, the probability of making such a choice is:

$$s_{ijt|Ins} = \frac{\exp(v_{ijt})}{\sum_{k \in J_i} \exp(v_{ikt})} \tag{4}$$

where v_{ijt} is u_{ijt} less the idiosyncratic unobservable utility $(\zeta_{it}^{Ins} + (1-\sigma)\varepsilon_{ijt})$ and is normalized by $(1-\sigma)$. Note from Section 2.2.5 that in Covered CA, different households within a market have different choice sets, represented by J_i.

Next consider the take-up decision. Let s_{it}^{Ins} be the probability of choosing to buy any plan, i.e. $s_{it}^{Ins} \equiv \Pr(\max_{j \neq 0}\{u_{ijt}\} > u_{i0t})$. Then the assumptions of the model imply:

$$s_{it}^{Ins} = \frac{\exp((1-\sigma)I_{it})}{\exp(z_i'\omega) + \exp((1-\sigma)I_{it})} \tag{5}$$

Where $I_{it} = \ln(\sum_{k \in J_i} \exp(v_{ikt}))$. I_{it} is often referred to as the "inclusive value" since it represents the expected utility of the entire set of plans (Train, 2009). Combining the above formulas, the unconditional probability of choosing any particular plan is $s_{ijt} = s_{it}^{Ins} s_{ijt|Ins}$.

Market shares s_{jt} are defined by summing the individual probabilities over the population in a market:

$$s_{jt} = \frac{1}{M_{jt}} \sum_i s_{ijt}$$

where M_{jt} is the number of households that are offered plan j in market t.

Given this framework, the effect on plan j's market share of increasing the premium by \$1 is given by:

$$\begin{aligned} \partial s_{jt}/\partial p_{ijt} &= \frac{1}{M_{jt}} \sum_i \partial s_{ijt}/\partial p_{ijt} \\ &= \frac{1}{M_{jt}} \sum_i \alpha_i s_{ijt} \Big(\frac{1 - s_{ijt|Ins}}{1-\sigma} + s_{ijt|Ins}(1 - s_{it}^{Ins}) \Big) \end{aligned}$$

and similarly, the cross-price effect of a price change for plan k on enrollment for plan j is:

$$\begin{aligned} \partial s_{jt}/\partial p_{ikt} &= \frac{1}{M_{jt}} \sum_i \partial s_{ijt}/\partial p_{ikt} \\ &= \frac{1}{M_{jt}} \sum_i -\alpha_i s_{ijt} s_{ikt|Ins} \big(\frac{1}{1-\sigma} - (1 - s_{it}^{Ins}) \big) \end{aligned}$$

To get semi-elasticities, I normalize these derivatives by dividing by market shares s_{jt}.

3.2 Provider Models

An important part of plan choice is access to providers, particularly for the purposes of studying adverse selection (Shepard, 2016). Provider network utilities, represented by EV^H and EV^P in (1), are derived from 2 features: 1) beliefs about likelihood of needing certain medical providers, and 2) access available given the plan networks. In this section, I present a general overview of the hospital and PCP models that underly the values for EV^H and EV^P. In the Online Appendix, I provide a more detailed explanation of each.

The hospital model is similar to those used in other papers on health plan competition (Ho, 2006; Gowrisankaran et al., 2014; Shepard, 2016). Enrollees have beliefs about their likelihood of needing different types of hospitalizations given their age and gender. Conditional on needing a hospitalization of a certain type, each plan's network generates a certain *expected* utility for each enrollee. The expectation is over idiosyncratic provider-patient utility which is unknown to even the enrollee until the hospitalization is needed. The deterministic portion of network utility is based on the "quality" of the hospitals in the network, discounted by their distance from the enrollee. Similar to the plan choice model, the hospital quality is inferred from revealed preferences, based on which hospitals patients choose given the possible alternatives. EV_{ij}^H is the expected utility of a network given the utilities of networking hospitals, and weighted by the probability of needing each type of hospitalization:

$$EV_{ijt}^H = \sum_{\kappa} \lambda_i^{\kappa} \ln\Big(\sum_{h \in \mathfrak{N}_j^H} \exp\big(\phi_h^{\kappa} - d_{\kappa}^H(dist_{ih})\big)\Big) \tag{6}$$

where κ is the hospitalization type ({Labor, Other inpatient, Outpatient} $\times$ {Childrens', Adults'}), and λ_i^{κ} is the probability of needed a κ hospitalization (which can sum to more or less than 1). $\mathfrak{N}_j^H$ is plan j's hospital network. ϕ_h^{κ} is the utility of hospital h for hospitalization κ. And $d_{\kappa}^H(dist_{ih})$ is the disutility of travel.

The PCP model follows a similar structure, though I make a number of simplifying assumptions to handle the thousands of doctors that could be in plan networks. Most crucially, I assume the mean utility of each physician is the same. I also discretize the distance metric to bins. This reduces EV_{ij}^P to be a function of the *number* of plan j PCPs ($N_{ij,b}^P$) within different distance bins (b) from the enrollee:

$$EV_{ijt}^P = \lambda_i^P \ln\Big(\sum_{b} N_{ij,b}^P \exp\big(-d^b\big)\Big) \tag{7}$$

where λ_i^P is the probability of needing to visit a PCP. d^b is the disutility of traveling to distance b.

A complete discussion of these models and the assumptions can be found in the Online Appendix.

3.3 Market Supply

Firm entry, network types, and networks are exogenous in this model. Firms know their costs and those of competitors. They also know the distribution of preferences and demographics in the market. Importantly, firms also anticipate risk adjustment when setting premiums.[20] Given this information, firms set a price for each plan for a 21 year-old, and the premiums for other ages are scaled exogenously according to the regulated 3:1 age factor.

For firm f in market t, profits are given as follows:

$$\pi_{ft}(\mathbf{p_{ft}};\mathbf{p_{-ft}}) = \sum_{j\in f}\sum_{i} s_{ijt}(\mathbf{p_{ft}};\mathbf{p_{-ft}})\Big(\theta_i^p p_{jt} - \theta_i^c c_{jt} + T_{ijt}\Big) \tag{8}$$

where p_{jt} is the premium for plan j for a 21 year old as in (2) and c_{jt} is the expected cost of a 21 year-old *male*. $\mathbf{p_{ft}}$ is the vector of the firm's premiums and $\mathbf{p_{-ft}}$ is the vector of all competitor premiums. θ_i^p is the regulated 3:1 rating factor established by Covered CA. θ_i^c is the actual expected cost factor relative to a 21 year old male–this can be thought of as the sum of the risk scores for household i. T_{ijt} is the risk adjustment and is the (possibly negative) payment firms receive from the program.

The precise formula I use for risk adjustment is based on that of the ACA (Pope et al., 2014) and is provided in the Online Appendix. The difference is that I use θ_i^c as my measure of risk, since I do not observe risk scores. As described in Section 2.2.3, the conceptual idea of ACA risk adjustment can loosely be represented by the following form:

$$T_{ij}^{ACA} \approx \Big(\frac{\theta_i^c}{\bar{\theta}_{CC}^c} - \frac{\theta_i^p}{\bar{\theta}_{CC}^p}\Big)\bar{P}_{t,tier}$$

Where $\bar{\theta}_{CC}$ represent averages among the Covered CA enrollees, and $\bar{P}_{t,tier}$ is the average premium in the region for the tier. This is just a stylized representation, however, and does not equal the formula in the ACA that I use in the model.

There are two crucial assumptions in this model regarding the structure imposed on cost variation across enrollees and plans. First, I assume multiplicative separability in the form of $c_{ijt} = \theta_i^c c_{jt}$, i.e. costs are risk scores times a baseline plan cost. This assumption rules out different plans having relative cost advantages covering different risk types. Variation in baseline plan costs (c_{jt}) is driven by prices for providers, care management strategies, and administrative costs. Within an insurer, c_{jt} would vary across tiers due to higher actuarial value and the associated moral hazard. Second, I assume θ_i^c is only a function of age and gender. Even though these two variables explain a small portion of total spending (Layton, Forthcoming), they explain a much higher share of the portion of expected spending that is correlated with preferences. Since the latter is the focus of this

[20] Bindman et al. (2016) note that Covered CA's active purchasing role helped plans to anticipate how risk adjustment would impact costs, and were ultimately passed through to premiums.

study, this formulation still can provide insight to the research question. It is, however, a strong assumption and I discuss the implications in greater detail in Section 8.

The degree of community rating in the model is given by the deviation between θ_i^c and θ_i^p. The source of adverse selection is the degree to which this deviation is correlated with plan preference parameters in (1).

Firms simultaneously set premiums in a static full information Nash equilibrium. While there is evidence of dynamic pricing in other insurance markets (Ericson, 2014; Ho et al., 2015), the high rate of churn and the uncertainty of the ACA make dynamics less relevant in this market. Assuming full information is plausible as long as I exclude 2014, the first year of the market.

First order conditions of the above profit functions with respect to prices (and dropping t subscripts) can be written as:

$$\frac{\partial \pi_f}{\partial p_j} = \sum_i \left\{ \theta_i^p s_{ij} + \sum_{k \in f} \frac{\partial s_{ik}}{\partial p_j} \left(\theta_i^p p_k - \theta_i^c c_k + T_{ik} \right) \right\} = 0 \tag{9}$$

This is similar to the standard formulation with discrete choice demand and heterogeneous preferences (Berry et al., 1995). The two main differences come from the policies in this study, risk adjustment and community rating. The derivative $\frac{\partial s_{ik}}{\partial p_j}$ is with respect to the 21 year-old price, given the modified community rating of the ACA.[21]

These first order conditions can be re-written in a more familiar form. Consider a single-product firm for convenience. Then (9) implies the profit-maximizing 21 year-old price for plan j is:

$$p_j = \underbrace{\frac{1}{\sum_i s'_{ij}\theta_i^p} \sum_i s'_{ij}(\theta_i^c c_j - T_{ij})}_{\text{“Standardized” Marginal Cost}} \underbrace{- \frac{\sum_i s_{ij}\theta_i^p}{\sum_i s'_{ij}\theta_i^p}}_{\text{Markup}} \tag{10}$$

where $s'_{ij} = \frac{\partial s_{ij}}{\partial p_j}$. Hence, as usual, the optimal price is equal to some measure of marginal cost plus a markup which is the inverse of the semi-elasticity. I label the first term as the “Standardized” Marginal Cost (SMC) to reflect that it is net of compensation from the premium scaling θ_i^p. It effectively down-weights the true marginal cost to account for the fact that older enrollees are generating more revenue though a higher value of θ_i^p. Note that the cost of the marginal consumer is a combination of both direct costs and the risk adjustment transfer (T_{ij}) associated with that enrollee–receiving positive transfers reduces effective costs and hence lowers premiums.[22]

Note that I have ignored premium subsides (APTC) in the context of firm optimization which

[21] There are two issues regarding risk adjustment that are being ignored in this setup. Premiums enter risk adjustment in two ways. First the benchmark is the average premium, so changing the premium changes the benchmark. Second, premiums affect the market composition which affects the “relative” risk in the risk score. I assume these effects are negligible when setting premiums.

[22] This formulation is the same idea as the model presented in Ericson and Starc (2015), except slightly more general and adding risk adjustment. They also describe the impact of multi-product firms on first order conditions, which affects both markups and “marginal costs.”

is important because APTCs makes prices enter demand in a non-traditional way. Recall that the APTC ensures that the second lowest silver (SLS) premium exactly matches some income amount. Hence, if a plan is the SLS, for every dollar increase in the (scaled) premium, the APTC increases dollar-for-dollar in kind. This effectively shuts down competition with the outside option for the SLS. However, when the SLS premium increases by \$1, there is still a \$1 impact on the price difference between plans. Therefore, for the SLS only, APTC should theoretically enter the FOC (ds/dp specifically) by shutting down competition with the outside option only. In most ACA markets, this is likely a non-trivial force (Jaffe and Shepard, 2017). In Covered CA, however, this force will be much smaller.[23] For this reason and for simplification, I omit this endogeneity from the model.

3.4 Identifying Adverse Selection

The form in (10) is particularly useful because it ties pricing back to adverse selection. Recall from Section 2.1 that a sufficient condition to identify adverse selection is if the marginal cost curve is decreasing in quantity or increasing in price (Einav et al., 2010).

In this model with risk adjustment and community rating, the sufficient condition to identify adverse selection is:

$$\frac{\partial SMC_j(p_j)}{\partial p_j} > 0 \tag{11}$$

where SMC_j is the standardized marginal cost defined above. From (10), it is easy to see how this property could lead to upward spiraling prices.

In theory, $\frac{\partial SMC_j(p_j)}{\partial p_j}$ could be calculated for each plan using the model. In this study, I use an alternative metric for adverse selection: the marginal cost relative to the average cost. Let the "Standardized" Average Cost be defined in a similar way:

$$SAC_j(p_j) = \frac{1}{\sum_i s_{ij}(p_j)\theta_i^p} \sum_i s_{ij}(p_j)\Big(\theta_i^c c_j - T_{ij}\Big) \tag{12}$$

where the average cost is discounted by revenue from age-based community rating (θ_i^p).

Under reasonable monotonicity assumptions on SMC (as are often assumed by parametric cost functions), $SMC < SAC \iff \frac{\partial SMC_j(p_j)}{\partial p_j} > 0$. Hence, $SMC < SAC$ is also a sufficient condition for adverse selection, in addition to (11). I use this measure because it is straightforward to calculate, and has an intuitive interpretation: it implies the marginal consumer costs less than the inframarginal consumers who more highly demand the plan. In competitive markets, SAC would

[23] Premiums within a carrier have the same ratio across tiers for all regions–for example, Kaiser's premium ratio of Silver to Bronze in Los Angeles is the same as in San Francisco. Since carriers generally enter in many markets, being the SLS in any particular market cannot be fully exploited for this purpose. This brings up another issue which is that it is inconsistent with the model presented in (9) where I've assumed pricing across markets is independent. Since the purpose of this paper is not estimating the precise costs by tier or exact profits, this is less relevant. Moreover, for the purposes of measuring adverse selection, it should have little impact at all.

be the amount that dictates the price faced by the marginal consumer, who costs SMC. In the analysis, my measure of adverse selection is $SMC - SAC$. I discuss the monotonicity assumptions more in the Online Appendix.

Under what conditions can the policies in this model eliminate adverse selection (at least the equilibrium effect) by flattening the SMC curve? The first is if there is no community rating, and the pricing profile reflects the exact expected costs (i.e. $\theta_i^p = \theta_i^c$). Then $SMC_j = c_j$ (and p_j would reflect the cost of a 21 year old *male*). The other case is if risk adjustment (T_{ij}) was designed in a way to make SMC constant. There are infinite possibilities for this since it could be set to any arbitrary level. However, motivated by the earlier discussion on "textbook" risk adjustment in Section 2.1, I consider how risk adjustment can set SMC equal to SAC in the population (SAC^{Pop}).

Specifically, let $SAC_j^{Pop} = \frac{\sum_i \theta_i^c c_j}{\sum_i \theta_i^p} = \frac{\bar{\theta}^c}{\bar{\theta}^p} c_j$, where bars represent population averages. Setting $SMC_j = SAC_j^{Pop}$ implies that risk adjustment is:

$$T_{ij}^{PlanPop} = \left(\frac{\theta_i^c}{\bar{\theta}^c} - \frac{\theta_i^p}{\bar{\theta}^p}\right)\bar{\theta}^c c_j \tag{13}$$

I call this "PlanPop" risk adjustment to denote that it has the properties of the textbook case defined in Section 2.1–it is benchmarked to *plan* costs ($\bar{\theta}^c c_j$) and scales relative to risk in the total *population* ($\theta_i^c/\bar{\theta}^c - \theta_i^p/\bar{\theta}^p$). One difference is that plans already receive revenue for age-based community rating, and so that compensation is subtracted.[24] Note that it need not generate greater welfare than the ACA method, and as seen earlier, it will not deliver the efficient allocations on it's own. It is only "textbook" in the sense that it perfectly flattens the marginal cost curve and sets it exactly equal to the average cost in the population. I use this risk adjustment in the main analysis to see how it alters the market relative to the current method. It is similar in spirit to risk adjustment used in Medicare Advantage and hence has additional policy relevance. The main difference for Medicare is that it is benchmarked to premiums as opposed to costs.

4 Fitting the Model to Data

I estimate the parameters of the above model using the data and methods discussed in the following section. Each of the following three subsections describes the data, identification, and estimation respectively.

[24] Risk in this model is relative to 21 year-old males, but it could just as easily be relative to the average cost person in the population as is normally done for risk scores. In that case, risk adjustment would take the very familiar form: $T_{ij} = (\theta_i^c - \theta_i^p/\bar{\theta}^p)\bar{c}_j$.

4.1 Data Sources

4.1.1 Covered CA Data

The primary data sources for this study come from Covered CA and include information on enrollment and provider networks. The enrollment data include the unidentified set of all households enrolled in any plan in 2015. The associated variables are each household's plan choice, income, zip code, age, gender, and race. Because carriers might offer their plans to only a subset of a rating region, I use data on coverage areas at the region-county-zip level to construct choice sets for each household. I apply a number of exclusions to simplify the analysis, which have a small impact on the overall sample. Specifically, I exclude households with any of the following characteristics: more than 4 members, an income less than 138% of the FPL, any members greater than 64 years old, chose a plan that wasn't offered in the household's zip code, or enrolled in a highly uncommon plan.[25] The size exclusion simplifies the demand estimation, while the income and age exclusions omit enrollees that are eligible for other government insurance programs. Combined, all exclusions account for 8.3% of total enrollment, virtually all coming from the income and household size exclusions in equal parts. Other than plans which are explicitly omitted, these exclusions do not affect plan shares. While the data include the universe of enrollees, I only use a 20,000 household sample, weighted evenly across rating regions. This significantly speeds up the analysis but otherwise has little impact on the estimation. More details on these exclusions and formatting of the enrollment data are included in the Online Appendix.[26]

This study also uses provider directories collected by Covered CA, which are used to ensure plans meet network adequacy standards outlined in the ACA. The directories include listings of all hospitals and physicians considered to be in-network for each plan. Provider networks are identified by a carrier and network type (HMO, PPO, EPO), and are assumed to be the same across tiers. For this study, I use hospitals and Primary Care Physicians (PCPs) for a given plan. From the hospital data, I use provider identifiers which are later merged with external data listed in the next section. From the PCP data, I use the office zip code.

4.1.2 Supplemental Data

While the data from Covered CA is very detailed, it does not include all the necessary information for this study. First, it doesn't include data on households that could have enrolled in a plan

[25] Specifically, I exclude households that enrolled in any of the following rare plans: "catastrophic" plans, Kaiser in Regions 1 and 13, Anthem's HMO in 3, 11, 19, Health Net's PPO in 5, and Molina in 19.

[26] There are households that enroll in individual plans outside of Covered CA. The data in this study comes specifically from the Covered CA enrollment website. Hence, it includes all subsidized enrollment (i.e. less than 400 FPL *and* electing to receive APTCs), but it need not include enrollment among those that do not receive subsidies. In fact, among those with incomes above 400% FPL, the data generating process that makes a household enroll through Covered CA is unclear. Among this income group without group insurance, there is actually a significant share that buys individual coverage. See footnote 27 for how this impacts the analysis.

but chose not to (i.e. taking the "outside option"). Additionally, while the data indicates which providers are in different networks, it does not indicate the cost of those providers or their relative value to consumers. Hence, I bring in external data sources to fill these gaps.

For this study, I consider households choosing the outside option as those that were uninsured in 2015.[27] To measure uninsured households, I use the individual-level 2015 American Community Survey (ACS) from IPUMS (Ruggles et al., 2017). The key variables are health insurance status, location of residence (as a PUMA), income, and household composition (i.e. age, gender, race of each household member). More details on how the ACS data is supplemented with the Covered CA data are provided in the Online Appendix. I make the same exclusions in the ACS as described above for enrollees when possible.

Another data source is needed for exogenous determinants of plan premiums (i.e. "cost shifters") to identify how enrollment responds to prices. For this, I use the costs of each plan's networking hospitals. Data on hospital costs can be found in the Healthcare Cost Report Information System (HCRIS) collected by the Centers for Medicare and Medicaid Services (CMS).[28] From these reports, I create a measure of the risk-adjusted "price" of a hospitalization at every hospital using methods from Dafny (2009).

To estimate the utility in the provider models, I use external hospital and PCP utilization data. For hospitals, I use data from the Office of Statewide Health Planning and Development (OSHPD). I use the 2014 Patient Origin and Market Share Reports[29] which give the number of discharges from a given patient zip code at each hospital by age category and type of hospitalization. To measure value of PCPs, I use estimates derived from an external model on Medicare Part B claims in 2012.[30] While Medicare patients are greater than 65 years old and hence quite different from the Covered CA population, I only use their PCP travel distances for this context. All distance measures are based on zip code centroids and come from NBER.

Finally, I use the Medical Expenditure Panel Survey[31] (MEPS) Household Component to estimate the how expected spending and utilization vary by age and gender.

[27]This could also include those in the "off-exchange" individual market (as in Tebaldi (2016)). However, since most of the Covered CA market is receiving premium subsidies, it's not unreasonable to think there is little substitution between on and off-exchange plans for the majority of Covered CA enrollees. Furthermore many of the plans available off the exchange are the same as those on the exchange, and have the same premiums. Omitting this group simplifies the problem of the outside option utility changing with price changes in the market.

[28]Source data is found here: `https://www.cms.gov/Research-Statistics-Data-and-Systems/Downloadable-Public-Use-Files/Cost-Reports/`. I use the formatted data created by Jean Roth at the National Bureau of Economic Research (NBER): `http://www.nber.org/data/hcris.html`

[29]See: `https://www.oshpd.ca.gov/HID/POMS-Report.asp`

[30]This model and the associated estimated parameters come from unpublished work examining plan competition in Medicare Advantage in which I examine similar market mechanisms.

[31]See: `https://meps.ahrq.gov/mepsweb/`

4.1.3 Descriptive Summaries

Table 1 gives a summary of the 19 rating regions in California. For each region, the table displays the number of firms, share of enrollment in each network type, and the market Herfindahl-Hirschman Index (HHI). The southern regions (region 12 and above), tend to have more firms, be larger, and have a much higher HMO share. In Northern CA, non-Kaiser HMO enrollment makes up less than 10% in each region, with the exception of San Francisco. Kaiser is a key player throughout the state but its role varies regionally. In the Bay Area counties for example, Kaiser makes up around a half of all enrollment. In the southern counties, Kaiser's enrollment ranges from 14% to 30%. The level of competitiveness ranges across the state with regions in Southern California generally being the most competitive. Unsurprisingly, the rural counties in the Central Valley, the Eastern Region, and the far North (Regions 10, 13, and 1) are the least competitive with HHIs that exceed 5,000, well above what's considered a concentrated market.[32]

The patterns in enrollment are not surprising given premiums across the state. Figure 4 plots 21 year-old Silver premiums across all regions, indexed by network type (Non-HMO, Kaiser, or Other HMO). The first pattern is that all premiums are significantly lower in Southern California (regions 12+). While insurance markets were identified to be more competitive in this region, the main results of this paper show this difference is more a function of the provider costs than plan markups. Moreover, the relative price of Kaiser plans is higher in Southern California than in the rest of the state which explains part of the Kaiser enrollment patterns. in Figure 5, I plot the enrollment in each plan against the relative premium. There is a clear negative correlation which suggests consumers are very price sensitive.

Figure 6 plots the distributions of ages and incomes for the Covered CA and uninsured populations, those who could have enrolled but chose not to. Those taking up insurance are older and have lower incomes than their uninsured counterparts. Given that subsidies are decreasing in income, this latter result is not surprising, though the difference is relatively large.

Last, Table 2 presents statewide characteristics of the individuals enrolled in each plan type, by either network type or metal tier. Examining the first three rows, Non-HMO plans (known to be "preferred" but more expensive) have intermediate ages and incomes, but have disproportionately more female and white enrollees. Kaiser enrollees are disproportionately younger with higher incomes. The second set of rows presents the same for metal tier selection. Gold and Platinum enrollees are actually younger. This is likely because these plans are relatively lower priced for the young and there are many younger households with relatively higher incomes. High generosity enrollees are also higher income and female, which is less surprising.[33]

[32]Markets with HHIs greater than 2,500 are considered highly concentrated and those with HHIs between 1,500 and 2,500 are moderately concentrated. See U.S. Department of Justice & FTC, *Horizontal Merger Guidelines* §5.3 (2010).

[33]Note that Silver plans are associated with large cost sharing reductions (CSRs) for over half of the population, so tiers are not necessarily increasing in coverage. However, we can compare Bronze to Gold/Platinum to get a sense

Table 1: Market Summary

Rating Region	Total Firms	Offering HMOs	Num HHs (1,000s)	NonHMO Share	Kaiser Share	Other HMO Share	HHI
1. Northern Counties	2	0	39	1.00	-	-	8,508
2. N. Bay Counties	5	2	40	0.41	0.51	0.08	3,520
3. Sacramento	4	2	53	0.54	0.43	0.02	3,458
4. San Francisco	5	2	33	0.40	0.38	0.22	2,744
5. Contra Costa	4	1	29	0.42	0.58	-	4,683
6. Alameda	3	1	50	0.47	0.53	-	3,922
7. Santa Clara	5	3	47	0.62	0.32	0.06	3,819
8. San Mateo	5	2	20	0.38	0.54	0.08	3,672
9. Central Coast I	3	0	23	1.00	-	-	5,222
10. Central Valley I	4	1	43	0.79	0.21	-	5,427
11. Central Valley II	3	1	20	0.72	0.28	-	3,723
12. Central Coast II	3	1	45	0.88	0.12	-	4,121
13. Eastern Region	2	0	5	1.00	-	-	5,802
14. Central Valley III	4	1	12	0.85	0.15	-	4,061
15. LA-East	6	5	118	0.39	0.15	0.46	2,773
16. LA-West	6	5	161	0.34	0.18	0.48	2,174
17. Inland Empire	5	4	86	0.37	0.23	0.40	2,221
18. Orange County	4	3	94	0.58	0.14	0.28	2,866
19. San Diego	5	3	91	0.34	0.30	0.36	2,202

Note: This table gives summary statistics for each of the 19 rating regions in Covered CA. The first column is the number of firms operating in each region. The second is the number of HMOs offered in each region. The third column is the number of households enrolled in a Covered CA plan (in thousands). The following three columns are the shares of enrollment in Non-HMOs, Kaiser, and HMOs respectively. The last column is the Herfindahl-Hirschman index (HHI). All statistics are net of sample exclusions described above.

Figure 4: Premiums for all Rating Region and Network Type

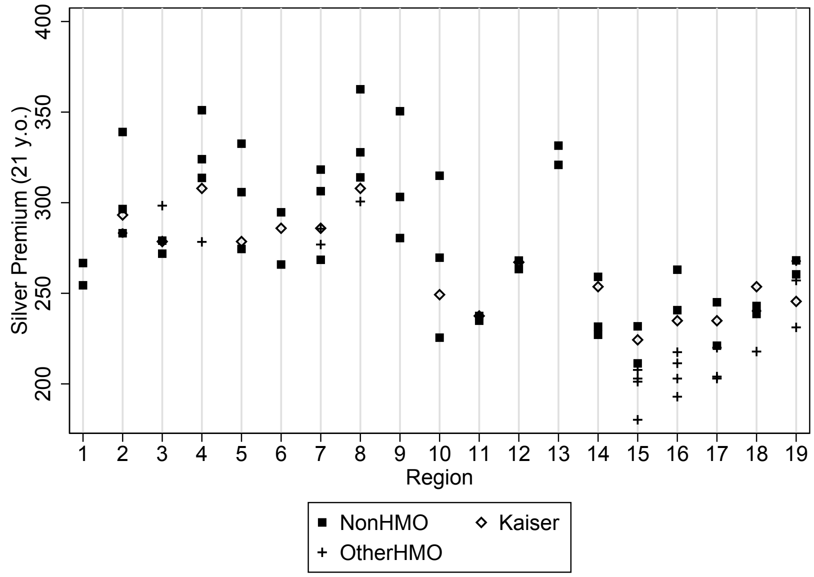

Note: This figure gives Silver plan premiums in 2015 for each rating region. Premiums based on 21 year-olds. Markers differ by network type. Only among plans that meet exclusion criteria. Regions 1-11 are Northern California.

Figure 5: Enrollment Shares by Premium and Network Type

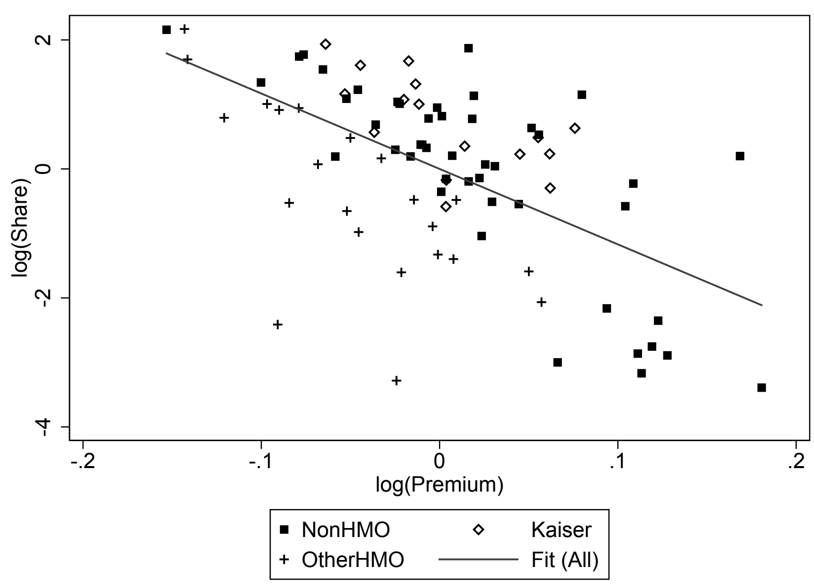

Note: This figure plots log of market shares by the log of the plan premium. Both measures are de-meaned at the region level. Only among Silver plans, which represent the majority of enrollment. Markers differ by network type. Fitted line based on all points in the figure. All data based on 2015.

for selection along the generosity margin.

Figure 6: Demographics by Insurance Coverage

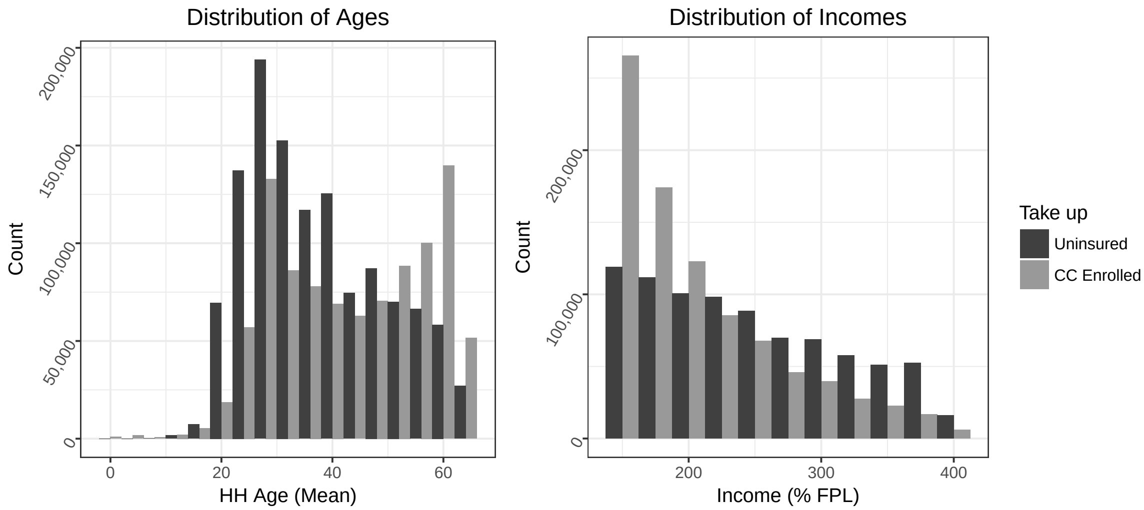

Note: These figures plot the distributions of ages and incomes. The unit of analysis is household–hence, "age" is the mean of the ages of all members in the household. The distributions are stratified by whether or not the household chose to buy insurance from Covered CA, or remain uninsured. The income distribution is censored at 400% FPL.

These descriptive results are suggestive that preferences are positively correlated with costs, the key feature of selection markets. Enrollees with higher expected spending–older or females–have higher rates of coverage, and are generally getting better plans. Since these consumers are facing different prices and different choice sets, the more detailed model discussed above is needed to draw more robust conclusions.

4.2 Identification

In this subsection, I describe the moments in the data that identify each parameter. For most parameters, identification is relatively straightforward. Two parameters for which this is not the case are the mean price sensitivity ("α_0") and the nesting parameter σ, which therefore receive disproportionate attention in this section.

Identification for all parameters in the plan choice model (1), come from Covered CA data combined with HCRIS cost reports. Parameters are identified from a combination of "micro" and "macro" moments. The micro moments identify how α_i and β_i deviate from some baseline–i.e. heterogeneity across individuals–and are given by correlations between individual characteristics and the chosen plan's characteristics (Train, 2009). Macro moments identify baseline values of these preference parameters based on how market shares correlate with plan characteristics (e.g. metal tier or insurer).

In the plan choice model, the only parameter for which identification is not straightforward is the baseline price sensitivity (the deviation based on individual characteristics is identified as above).

Table 2: Mean Demographics by Plan Type

	Members (1,000s)	Age	FPL	Female	Asian	Nonwhite
Network Type						
1. NonHMO	843	41.76	223	0.53	0.16	0.48
2. Kaiser	394	40.71	231	0.52	0.15	0.50
3. OtherHMO	394	42.73	202	0.51	0.20	0.60
Tier						
1. Minimum Coverage	12	24.22	326	0.45	0.18	0.44
2. Bronze	344	40.47	249	0.49	0.17	0.52
3. Bronze Hsa	72	41.04	260	0.49	0.14	0.46
4. Silver	1,038	42.97	200	0.54	0.18	0.53
5. Gold	89	38.39	276	0.52	0.12	0.47
6. Platinum	77	38.17	268	0.52	0.10	0.45

Note: This table presents characteristics of the enrollees choosing each plan type. The first column is the number of individuals enrolled in each plan type. The rest of the columns represent the mean of the individual characteristic among enrollees in the given plan type.

First, as is always the case, premiums are endogenously set in response to anticipated demand, creating bias which underestimates the true magnitude of the parameter. The second problem is that in this environment, market prices vary across consumers so the standard "inversions" to separate out a market price are not valid (Berry et al., 2004; Goolsbee and Petrin, 2004).[34] The logic I use to address this problem of within-market heterogeneous pricing is fairly intuitive and involves a combination of the standard endogeneity methods, plus the micro moments outlined in the prior paragraph.

More formally, notice that (1) can loosely be re-written as follows:

$$\begin{aligned} u_{ijt} &= \alpha_i p_{ijt} + \overbrace{\nu_{ijt}}^{\text{All other utililty}} + \xi_{jt} + \varepsilon_{ijt} \\ &= (\widetilde{\alpha}_i + \alpha_0)(\widetilde{p}_{ijt} + p_{jt}) + \nu_{ijt} + \xi_{jt} + \varepsilon_{ijt} \\ &= \underbrace{\alpha_0 p_{jt} + \bar{\nu}_{jt} + \xi_{jt}}_{\delta_{jt}} + \alpha_0 \widetilde{p}_{ijt} + \widetilde{\alpha}_i p_{ijt} + \widetilde{\nu}_{ijt} + \varepsilon_{ijt} \end{aligned}$$

where p_{jt} is the premium for a 21 year-old as discussed above, and tilde ("~") variables represent deviations from baseline values. α_0 is the parameter of interest. I isolate "δ_{jt}" to highlight the usual plan-market mean utility that is often used with instrumental variables to estimate α_0 in an unbiased way (Berry et al., 2004; Goolsbee and Petrin, 2004). Written this way, it's clear that the usual method will not work because α_0 appears outside of the δ_{jt} due to within-product variation in prices.

[34] The usual logic is to isolate the mean plan utility, including the market price and unobservable ξ_{jt}, using a market-plan fixed effect. Then form moments given exogenous instruments and ξ_{jt} to identify the price coefficient. Either of the aforementioned references explain this in great detail.

However, this formulation also suggests there are two potential moments that can identify α_0. The first is the usual macro moment $E(\xi_{jt}Z_{jt})$ where Z_{jt} is some instrumental variable that exogenously changes premiums. The instrument I use is "prices" of networking hospitals as cost shifters, as defined in Dafny (2009). To be a valid instrument, this measure need only be correlated with the rates paid by Covered CA plans. Even though this price index differs from the rates Covered CA, I assume that the prices in this index are roughly proportional to negotiated rates in Covered CA. To construct a plan-level price, I take the average price across the hospitals in each plan's network, weighted by the number of annual discharges.[35] The identifying assumption is that plan utility doesn't depend on hospital prices, *conditional on other observables*. Given that I have specified hospital network utility explicitly in the model (EV^H), this assumption seems reasonable.[36] Using this as an instrument for *plan* prices, α_0 will be identified based on how market shares correlate with networking hospital prices. See the Online Appendix for additional robustness checks on the validity of this instrument.

The second moment to identify α_0 is more simple and is the same type of micro moment used for other heterogeneity parameters. Since there is variation in premiums *within plans*, it can be used to see how enrollment probabilities within a market-plan also vary in kind. More formally, this moment is $E\big((e_{ijt} - s_{ijt|Ins})p_{ijt}\big)$ where $s_{ijt|Ins}$ is the model probability of enrollment into plan j and e_{ijt} is a dummy for whether the plan was chosen (Train, 2009; Berry et al., 2004). The validity of this moment relies on the fact that I condition on ξ_{jt}, which is correlated with premiums.

Since α_0 is over-identified, it is reasonable to consider using just one moment (particularly the micro moment) for simplicity. While in theory this is a valid strategy, in this environment it would likely be inadequate. Using only the micro moment would be using within-plan price variation which comes only from the regulated age rating profile in (2). Price differences between any two plans are larger for older consumers. This could create bias if α_i varies with age as well, which is a reasonable assumption. One might want to "control for age" and rely on the functional form difference in the rating profile, but this is a strong assumption since this profile represents expected spending. Hence, adding the macro moment uses data variation that is plausibly exogenous and will yield an unbiased estimate. Not surprisingly, adding this moment substantially increases the estimated price sensitivity.[37] This strategy differs from other studies on the pre-ACA Massachusetts exchange where pricing discontinuities can be used (Shepard, 2016; Ericson and Starc, 2015; Finkelstein et al.,

[35] Kaiser Permanente is a vertically integrated health system, so they have no hospital "prices" based on how I have defined them. Therefore, I set hospital prices to 0 for all Kaiser plans. Since there are brand fixed effects in the model, identification is based only on networks of non-Kaiser plans.

[36] A hospital's prices will be higher if they are a more demanded hospital (Ho, 2009), but this variation will not be used to identify α_0. Instead, identification comes from hospital price variation due to hospital market power (Gowrisankaran et al., 2014) or underlying operational costs. Conditional on the explicit network utility, these should not directly effect plan utility.

[37] Alternatively, using only the macro moment would be relying on only JT sample observations. While the estimate would be unbiased, it would also be much less precise. Hence adding the micro moment with the large sample significantly reduces variance of the estimate.

2017a).

Next, consider the take-up model in (5). In most settings, the nesting parameter $(1-\sigma)$ can be identified from the correlation between take-up probabilities and variation in the value of all plans in the choice set I_{it} (which is "data" once the plan choice model has been estimated). Unfortunately, most of the variation in I_{it} in this setting comes from the APTC which is a function of income, among other things. Since income impacts I_{it} in a highly nonlinear way that interacts with other demographic variables (all of which enter the take-up decision), it is likely that $E(\varepsilon_{i0t} I_{it}) \neq 0$, which biases the estimates. Since lower-income consumers likely have a lower preference for insurance, large values of I_{it} from high APTCs will be among those with relatively lower values for insurance, and $(1-\sigma)$ will be biased down.[38]

Since parametric models with non-linearities and interactions does not appear to fix this problem, I rely on an exogenous shifter of I_{it}. Recall that within rating regions, some plans are only offered to a subset of zip codes. Hence, I use variation in choice sets *within rating regions* to identify the nesting parameter $(1-\sigma)$. The primary sources of within-region choice set variation are Blue Shield in Marin, Alameda, and Santa Cruz, and Chinese Community in San Mateo. Importantly, this variation is correlated with I_{it} but not with the unobserved errors ε_{i0t} (conditioning on income and age).

4.3 Estimation

For tractability, the estimation is conducted sequentially and in the reverse order of the model timing.[39] The sequence of estimation is as follows: the provider models, the plan choice model, the take-up model, then finally inferring costs given demand and market structure.[40]

4.3.1 Estimating Provider Models

The full details on the provider model estimation can be found in the Online Appendix. In this section, I present a general overview.

[38] In the model, this bias would imply very little substitution along the extensive margin–hence underestimating (in magnitudes) elasticities.

[39] This implies that inference in later steps of the sequence is not valid unless standard errors have been corrected for estimation error in earlier steps (see Ho (2006) for a discussion of this). The standard errors presented are not corrected for this sequence and hence are *underestimates* of the true standard errors. However, in the context of this paper, this isn't too problematic. First, the sample sizes are sufficiently large so that most parameters are estimated with an extremely high degree of precision in all stages (e.g. the ratio of point estimates to standard errors is often far greater than 10). Second, as is often the case when relying on counterfactual policies for drawing conclusions, inference on the point estimates is less important than the robustness of the results to the model specification. To this end, I dedicate extensive discussion to robustness in the Online Appendix.

[40] The majority of the analysis is performed in R. The estimation routines and simulations were developed for this project, but rely heavily on the R development community which deserves a great deal of credit (R Core Team, 2017). I also use the following R packages: *nleqslv* (Hasselman, 2017) for solving non-linear systems of equations. *SQUAREM* (Varadhan, 2016) for increasing the speed of the contraction mapping (Reynaerts et al., 2012). *survey* (Lumley, 2017) for conducting a weighted logit for the take-up model. *ggplot2* (Wickham, 2009) for figures and *stargazer* (Hlavac, 2015) for outputting all results to LaTeX.

I estimate network utilities separately for Kaiser and Non-Kaiser plans, since Kaiser is generally a closed system. I estimate the hospital demand model for each hospitalization type using maximum likelihood. Choice sets are all hospitals visited from a particular patient zip code in the OSHPD data.[41] I estimate the probability of needing each hospitalization type non-parametrically based on age and gender using MEPS. Hospital network utility can then be calculated for each plan using (6).

For PCP demand, I use estimates from an external model using Medicare Part B claims data.[42] That model measures the disutility of traveling discrete distance bins to visit a PCP. By applying these parameters to the size and location of plan networks in Covered CA, I can estimate the conditional utility associated with PCP networks for each consumer. I estimate utilization probabilities from MEPS in the same manner as for hospitals. Given the external parameter estimates and utilization probabilities, I calculate unconditional PCP network utilities (EV_{ijt}^{P}) using (7).

4.3.2 Estimating Plan Demand

The next step is to estimate the utility function in (1). I estimate the model in two stages: The plan choice ("lower") model includes all parameters in (1) less σ. The take-up ("upper") model is the decision whether or not to be insured, and identifies σ and the parameters in the outside utility in (3). [43]

Working backwards through the nested model, I first estimate the plan choice model–conditional on take-up. Restating the main indirect utility function with a slight abuse of notation, let the plan utility conditional on buying insurance be:

$$\begin{aligned} u_{ijt} =& \alpha_i p_{ijt} + \gamma^H EV_{ijt}^H + \gamma^P EV_{ijt}^P + x_{ijt}'\beta_i + \xi_{jt} + \varepsilon_{ijt} \\ =& \underbrace{x_{jt}'\beta + \xi_{jt}}_{\delta_{jt}} + \alpha_i p_{ijt} + \gamma^H EV_{ijt}^H + \gamma^P EV_{ijt}^P + \widetilde{x_{ijt}'\beta}_i + \varepsilon_{ijt} \end{aligned}$$

Notice that the ζ_{it}^{Ins} has dropped out and the $(1-\sigma)$ coefficient has been removed, which means

[41] This assumes patients in OSHPD data have unrestricted provider networks, different from the plans in this study. Since the majority of patients are covered by employer coverage which tends to have large networks, this assumption is less problematic in this context. Any violation in this assumption will lead to underestimated utilities for highly demanded hospitals. Capacity constraints are also ignored and would have a similar effect. This assumption is not expected to have a large impact on the conclusions in this paper since I also allow insurer utilities to vary by health status, which would absorb anything not measured by network variables.

[42] This model and the associated parameter estimates come from unpublished work examining plan competition in Medicare Advantage in which I examine similar market mechanisms.

[43] It's possible to estimate these parameters simultaneously and is often preferred (Train, 2009; Hensher, 1986). By subsetting the data, I am throwing away information and increasing standard errors (or need to correct the standard errors in the take-up model which uses estimated data). Also, if there are any parameters in multiple nests, they cannot be estimated at all (Heiss, 2002). In this setting, sequential estimation is preferred for 2 reasons: 1) each stage is identified from two distinct data sets (Covered CA vs. ACS); and 2) separately estimating the take-up model allows for a concise way of using instrumental variables to identify σ. Additionally, because of the large enrollment sample and there being only 1 "nest," the usual reasons not to use sequential estimation are not of particular importance.

all other parameters have implicitly been rescaled by $1/(1-\sigma)$ (Heiss, 2002). Also, coefficients have been added to network EV terms to convert units from each provider utility model to the units in the demand model.

A full description of the estimation routine is given in the Online Appendix, but the idea is very similar to "Micro BLP" Berry et al. (2004). The two exceptions are that I use two moments to identify price sensitivity and I remove unobserved heterogeneity. The routine is as follows: I minimize a GMM objective function made up of micro and macro moments over the parameters related to plan heterogeneity (including "α_0")–call these Θ_2. Baseline utility parameters (Θ_1, or β above) can be determined analytically, given market shares and the current guess for Θ_2. The macro moments are characterized by the vector $\sum_{jt} \xi_{jt}(\Theta_2)Z_{jt}$, where $\xi_{jt}(\Theta_2)$ matches model shares with actual shares[44] and Z_{jt} is a vector of plan-market instruments: brand, tier, network type, and networking hospital prices. The micro moments are $\sum_{ijt}(s_{ijt|Ins}(\Theta_2, \delta_{jt}) - e_{ijt})z_i' x_{ijt}$, where z_i are individual characteristics that dictate heterogeneity in α_i and β_i, and e_{ijt} is an indicator for whether the plan was chosen. The model is over-identified because of the two price moments so I construct a GMM weighting matrix as the inverse of the covariance matrix of instruments as in two-stage least squares.

Once the parameters in the plan choice model are estimated, I estimate the take-up model of whether or not to buy insurance. Recall the probability of take-up in the model from (5):

$$\begin{aligned} s_{it}^{Ins} &= \Pr(\max_j u_{ijt} > u_{i0t}) \\ &= \frac{\exp((1-\sigma)\hat{I}_{it})}{\exp(z_i'\omega) + \exp((1-\sigma)\hat{I}_{it})} \end{aligned}$$

where $\hat{I}_{it}$ is implied by estimates in the plan choice model: $\hat{I}_{it} = \ln(\sum_{k \in J_i} \exp(\hat{v}_{ikt}))$.[45] z_i is a vector of individual attributes including income, age, gender, region, race, and population density. This formulation implicitly includes the mandate in a reduced-form way.

Recall the problem identifying $(1-\sigma)$ discussed in the above Section 4.2, and the proposal to use exogenous variation in the choice sets alone (i.e. J_i). Since this is a nonlinear model, standard two-stage least squares cannot be applied. My preferred method is using a control function approach due to its relative simplicity. This method separates the error ε_{i0t} into two components, only one of which is allowed to be correlated with the endogenous I_{it}. This correlation is explicitly estimated using residuals from a first stage based on exogenous instruments (See Train (2009), Section 13.4.1). The method is as follows: I regress I_{it} on the instruments (including J_i). I add the residuals from that first-stage regression as a "control variable" along with I_{it} in the take-up model. The control

[44] I use a "contraction mapping" on δ_{jt} with the SQUAREM (Varadhan, 2016) algorithm as recommended by Reynaerts et al. (2012).

[45] Since I_{it} is estimated rather than data, the standard errors are biased down and should be interpreted as such. Since I have a large sample, there is sufficient power to estimate these parameters with precision either way.

function is treated as an exogenous variable, and I estimate the model with MLE.[46]

Without controlling for the endogeneity of $\hat{I}_{it}$, standard MLE generates a value of $(1-\sigma)$ close to 0 and a value of insurance that is decreasing in income. This is expected given how subsidies decrease with income, but is inconsistent with theory. Using the control function approach outlined above, I get estimates that are theoretically plausible. These comparisons can be found in the Online Appendix.

Finishing the estimation in the take-up model, I now have all estimates to characterize demand across the state.

4.3.3 Estimating Costs

I estimate plan costs using a standard procedure in the IO literature based on the first order conditions in (9) (Nevo, 2001; Berry et al., 1995). In these models, once demand is known, equilibrium assumptions provide a system of JT equations and JT unknown costs, which can therefore be inferred. In my model, there are two complications from the standard implementation. First, this is a "selection market" where the cost of a product depends on which consumer is buying the product (Einav et al., 2010). Second, profits rely on risk adjustments which have not yet been calculated in the model. Omitting risk adjustment in the first order condition would overestimate the amount of price variation that is attributed to enrollee health.[47]

With respect to selection, I impose the cost structure $\theta_i^c c_{jt}$. I estimate θ_i^c non-parametrically from health care spending by age and gender in MEPS (See Online Appendix for details). This reduces each FOC to a single standardized plan cost c_{jt}. Since I assume expected cost variation comes only from age and gender, this underestimates the true variation. I discuss the implications and remedies in Section 8.

With respect to risk adjustment transfers, I use the estimated demand model combined with publicly available risk adjustment data to estimate transfer amounts for each potential enrollee.[48] I calculate risk adjustment transfers using the same formula as given by the ACA, where θ_i^c is my measure of risk scores. Since I do not observe enrollment off the exchange, I use "relative risk" $(\theta_i^c/\bar{\theta}^c)$ based on the on-exchange enrollment implied by the demand model.

Given θ_i^c and risk adjustment transfers, costs can be computed using the standard method. Given demand estimated in the prior section and observed prices assumed to satisfy the JT first order conditions in (9), I solve for JT costs for each plan.

[46] This application is not precisely correct since I am assuming away the remaining variance in the unobservable component that is correlated with I. However, I estimated the correct specification and found this variance to be precisely near 0. Hence I proceed without it for simplicity.

[47] Including the risk adjustment transfers yields cost estimates which much more closely parallel premiums. By including risk adjustment, high generosity plans and those with generous networks (non-HMOs) have much higher estimated costs than if risk adjustment was ignored.

[48] Specifically, from the public data, I use geographic cost factors and the state average premium. See: `https://www.cms.gov/CCIIO/Programs-and-Initiatives/Premium-Stabilization-Programs/`

5 Estimated Model

This section discusses the estimated parameters from the model outlined above. Details on the estimates from the provider models can be found in the Online Appendix. I proceed straight to the estimation of plan demand taking the provider models as given.

5.1 Demand Estimates

I begin with the plan choice model, conditional on choosing to buy a plan. The parameter estimates for baseline utilities (Θ_1) and heterogeneity (Θ_2) are given in Tables 3 and 4 respectively. Given the results in Table 3, I find that there is a strong increasing utility in the level of coverage, with higher tiers yielding higher utilities. Also as expected, EPOs and PPOs generate higher utility than HMOs conditional on network quality and prices.

In Table 4, I find evidence of preference heterogeneity for different plan characteristics. I find a higher price sensitivity among those with lower expected health spending (θ_i^c), Asians and other nonwhites, as well as those living in denser areas. This latter result is consistent with two theories: 1) high costs of living in less rural areas (Weinberg and Kallerman, 2017), and 2) possibly greater access to uncompensated care (Finkelstein et al., 2017a).[49] While non-HMOs yield higher utility in general, the difference is even larger among households with high incomes and those with higher expected health care spending.

The impact of providers on plan choice is also notable. Hospital networks matter, but not for Kaiser plans. This isn't surprising since hospital qualities do not vary in the Kaiser system as they do otherwise. For a typical enrollee, the increase in *hospital network utility* from the UCLA hospital, for example, is about 0.1. This implies an average valuation of having UCLA in-network of only about \$2.70. This seems relatively low compared to other studies (Shepard, 2016; Ho, 2006). It could be due to actually low valuations in this market, especially given the high price elasticity, or bias for reasons described in Section 4.3.1. It could also be that those studies have upward biased estimates as a result of omitting physicians, which are included in my model. Either way, any bias is not expected to change the results of this paper, since omitted heterogeneity is absorbed into carrier and network type preferences. Regarding PCP networks, they also matter to consumers, and among both Kaiser and non-Kaiser plans. For non-Kaiser plans, given the PCP utility estimate of 0.663 combined with that of hospital networks (1.297), this implies a hospital like UCLA would generate the same utility as a 19.6% increase in the number of nearby PCPs.

[49] Unlike in many other papers, I do not allow price sensitivity to vary with income. I consistently find that high income households are not less price-sensitive in this population. I attribute this to the APTC and the fact that high income households are facing higher prices overall. Hence my results are consistent with disutility in price that is convex rather than linear. Any conclusions should be interpreted within the context of the ACA rather than the broader individual insurance market. Also note that while income doesn't interact with price, it does interact with other product attributes which are associated with higher prices (e.g. metal tiers, network type, brand), as well as the preference for insurance in general.

Table 3: Select Demand Estimates: Baseline Utilities (Θ_1)

	Coef	SE	p-value
Constant	-2.672	0.400	***
Bronze HSA	-2.161	0.193	***
Silver	3.906	0.210	***
Gold	5.638	0.214	***
Platinum	8.062	0.236	***
PPO	2.578	0.259	***
EPO	2.488	0.270	***

***p<0.001; **p<0.01; *p<0.05; .p<0.1

Note: This table presents the baseline estimated coefficients from the plan choice model. "Baseline" represents the fact that this is common to all households in the market. Brand and region dummies omitted. Standard errors are not corrected for estimated provider models.

Given the size of many medical groups, this implies large physicians groups have similar bargaining abilities with plans to many hospitals (Ho, 2009). Finally, I also find that higher AV plans are more attractive for households with higher incomes and higher expected health expenditures.

Note that these results imply the potential for adverse selection is large in this market–higher cost patients have stronger preferences for all of the following plan characteristics which are likely to have higher costs: better networks,[50] a non-HMO network type, and more generous cost sharing.

Table 5 provides the estimates for the take-up model. Since this is a binary decision, I've moved all variables to the take-up side of the inequality, which implies these coefficients are the relative utilities of having insurance. The key parameter from this model is the coefficient on the "Inclusive Value," estimated at 0.549. Recall that this value $(1-\sigma)$ is the degree of *independence* between the unobservable utilities across insurance plans (hence in a logit model it is assumed to be 1). That it is significantly lower than 1 implies a correlation between the plan unobservable utilities. In other words, conditional on age, gender and race, consumers still disproportionately substitute between plans relative to choosing to be uninsured.

The coefficient on the control function is also important. It loosely represents the correlation between I_{it} and ε_{i0t}. That it is negative implies that the APTC (and hence I_{it}) is highest among households with low willingnesses-to-pay for insurance in unobservable ways. Regarding other coefficients, we can see that the preference for insurance is increasing in income and lower for minorities and urban consumers.[51] A detailed comparison of alternative estimation procedures for

[50] Recall that age and gender (both of which determine costs) enter the measures of EV directly through higher expectations of utilization.

[51] Notice that expected spending (as a function of age and gender) is associated with a negative coefficient in this

Table 4: Select Demand Estimates: Heterogeneity Utilities (Θ_2)

Plan Char	HH Char	Coef	SE	p-val
p_{ijt}		-4.731	0.163	$***$
p_{ijt}	$E(MedSpending)\,(\theta_i^c)$	0.967	0.036	$***$
p_{ijt}	$HHSize$ =2	2.009	0.078	$***$
p_{ijt}	$HHSize$ =3	3.185	0.117	$***$
p_{ijt}	$HHSize$ =4	3.915	0.139	$***$
p_{ijt}	Asian	-0.278	0.036	$***$
p_{ijt}	Nonwhite	-0.154	0.024	$***$
p_{ijt}	Density	-0.071	0.019	$***$
Silver AV94		2.405	0.099	$***$
Silver AV87		1.472	0.071	$***$
Silver AV73		0.480	0.068	$***$
HMO (non-Kaiser)	FPL	-0.217	0.042	$***$
HMO (non-Kaiser)	$E(MedSpending)\,(\theta_i^c)$	-0.122	0.050	$*$
$EV_{ijt}^{H,nonKais}$		1.297	0.440	$**$
$EV_{ijt}^{H,Kais}$		0.117	0.221	
$EV_{ijt}^{P,nonKais}$		0.663	0.057	$***$
$EV_{ijt}^{P,Kais}$		0.721	0.057	$***$
AV	FPL	0.112	0.009	$***$
AV	$E(MedSpending)\,(\theta_i^c)$	0.167	0.015	$***$
AV	$HHSize$ =2	0.384	0.036	$***$
AV	$HHSize$ =3	0.461	0.079	$***$
AV	$HHSize$ =4	0.221	0.100	$*$

***p<0.001; **p<0.01; *p<0.05; .p<0.1

Note: This table presents heterogeneity coefficients from the plan choice model. These are the coefficients that dictate deviations from baseline utilities in the prior table. Preferences for plan charateristics are based only on the observable characteristics as indicated. Brand-demographic (income, spending, race) interactions omitted. Expected spending is the average θ_i^c within the household. Density in Logs. FPL is censored at 400% FPL, and effect of additional income assumed negligible beyond that point. FPL, expected spending, and density are standardized. Prices in units of $100.

this part of the model can be found in the Online Appendix.

Table 5: Demand Estimates: Upper Model (Θ_3)

HH Char	Coef	SE	p-val
Constant	1.026	0.182	$***$
I_{it}	0.549	0.078	$***$
Control Function	-0.537	0.079	$***$
FPL	1.226	0.241	$***$
$E(MedSpending)$	-0.241	0.113	$*$
Asian	-0.019	0.076	
Nonwhite	-0.446	0.049	$***$
Density	-0.307	0.041	$***$

***p<0.001; **p<0.01; *p<0.05; .p<0.1

Note: This table presents coefficients from the take-up model. Utility is associated with taking up insurance. FPL is censored at 400% FPL, and effect of additional income assumed negligible beyond that point. FPL, expected spending, and density are standardized. Region dummies omitted.

Figure 7 plots the distribution of elasticities by network type (non-HMO, Kaiser, Other HMO), tier, and geographic region (North vs. South). Price-sensitivity in the market is quite high, with market shares generally decreasing between 3% and 5% for a price increase of $1. These are much higher then seen in the unsubsidized individual market (Cutler and Reber, 1998; Ho, 2006; Ericson and Starc, 2015), but close to those in other subsidized individual markets (Shepard, 2016; Finkelstein et al., 2017a). The fact that these are even slightly higher than those papers isn't surprising. Relative to the pre-ACA Massachusetts exchange or the ACA markets in other states, Covered CA has a number of standardizing regulations that intensify price competition (see Section 2.2).

The elasticities vary somewhat by region, network type, and tier. Silver plans appear to be the least price elastic, which makes sense given the large cost sharing subsidies tied to silver plans for a majority of enrollees. Generally, Gold plans are more elastic, especially for HMOs. Elasticities by network type depend on geography. Non-HMOs are generally less elastic in the South, while Kaiser plans are less elastic in the North. Interestingly, despite the relative concentration in Northern California (see Table 1), elasticities are not markedly lower.

model. Since the price sensitivity is lower among those with higher expectations of health spending, the value of insurance is already increasing in this variation through I_{jt}. Hence, this is a "reduced-form" effect, net of the already modeled preference for insurance.

Figure 7: Semi-Elasticities by Geography and Network Type

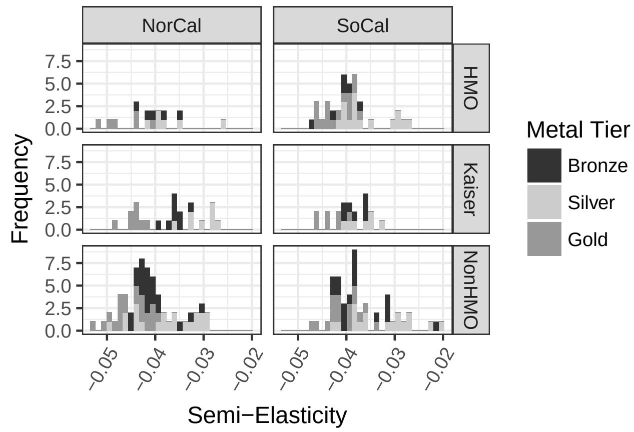

Note: This figure plots the distribution of own-price semi-elasticies in the market. These give the percent change in plan enrollment associated with a \$1 increase in plan premium. Each "point" in the distribution is a single plan (jt). Regions 1, 9, and 10 omitted.

5.2 Supply Estimates

Figure 8 plots the expected spending by age and gender relative to that of a 21 year-old male. The figure indicates that women generally have higher expected spending than men before the age of 45, at which time they roughly converge. For men, the ratio of costs between 65 year olds and 21 year-olds is approximately 4:1, as opposed to the 3:1 ratio allowed in the regulations. The figure also plots the regulated rating profile by age as specified by Covered CA. Hence the two dashed curves and the solid curve represent θ_i^c and θ_i^p respectively.

Figure 8: Estimated Spending by Age and Gender

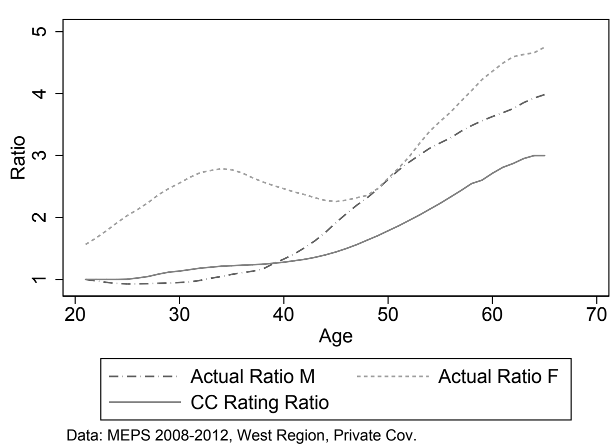

Note: This figure plots the non-parametrically estimated spending levels by age and gender used in the model, θ_i^c. All levels are relative to a 21 year-old male. The solid line gives the pricing profile set by Covered CA.

Using estimates for θ_i^c and first order conditions, I get implied plan costs and markups. I plot implied plan profit margins in Figure 9. The results are consistent with the high price elasticities described above. Namely, price markups on each plan are very low with variable profits making up only about 3%-8% of revenues (the main exception being Kaiser in the North). Interestingly, profit margins in the North for Non-HMO plans are especially low.

Figure 9: Profit Margins at Baseline

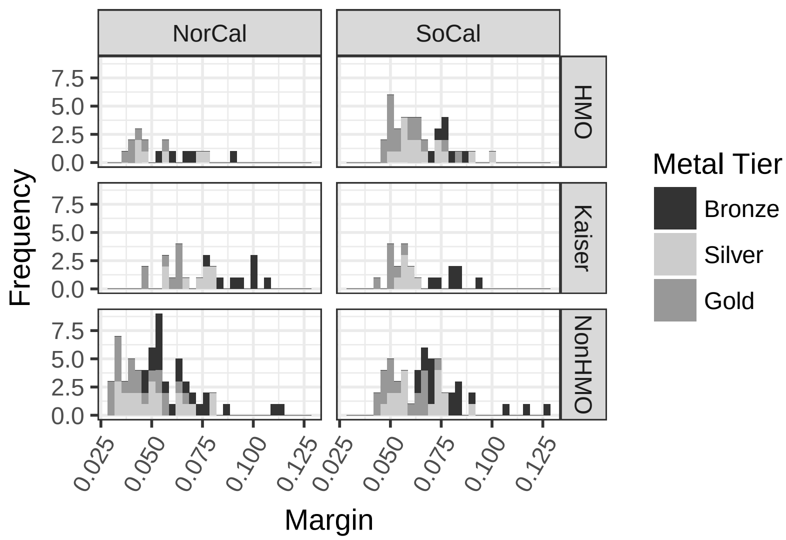

Note: This figure plots the distribution of profit margins for all plans under current ACA regulations. Each "point" in the distribution is a single plan (jt). Plans in regions 1, 9, and 10 omitted.

The evidence presented so far in this section–preferences correlated with costs and slim markups for generous plans–could be suggestive evidence of adverse selection. I now turn to the counterfactual simulations to test this more formally, and find this is not the case.

6 Counterfactual Simulations

6.1 Specifications

Using the above estimated model, I address the questions proposed in this paper. The objectives of this paper are to determine the effect of community rating and risk adjustment in the ACA on market outcomes, and quantify the associated adverse selection. To that end, there are four policy simulations of interest:

1. **No community rating**: This environment assumes prices are set proportional to expected costs such that there is no pooling across genders and premium ratios can exceed 3:1. Since prices reflect costs, there would be no adverse selection and hence no risk adjustment is

needed. This simulation is purely to serve as a benchmark to compare market outcomes for the other policies. In the model, this scenario implies the administrative pricing factor is adjusted to match the variation in expected costs: $\theta_i^p = \theta_i^c$.

2. **Community rating without risk adjustment**: This environment preserves the rating regulations from the ACA (3:1 ratio and age-based pooling) but does not allow for any risk adjustment. In the model, this implies $T_{ij} = 0$ and everything else preserved.

3. **Community rating with ACA risk adjustment ("Baseline")**: This is the current regulatory environment with potential adverse selection from the age rating rules, but offset with the current method of risk adjustment transfers. No policy changes are made to the baseline model.

4. **Community Rating with alternative risk adjustment**: This is the "textbook' method derived in Section 3.4. It benchmarks to individual plan costs and scales relative to risk in the entire population. As derived earlier, this sets $T_{ij}^{PlanPop} = (\theta_i^c/\bar{\theta}^c - \theta_i^p/\bar{\theta}^p)\bar{\theta}^c c_j$.

Simulations are conducted by making changes to the regulatory environment (e.g. changing T_{ij} or θ_i^p), and solving for new equilibrium prices given by the first-order conditions in (9). Market shares and government spending are implied by demand under new equilibrium prices in each counterfactual scenario. Each of these are examined in the next section to evaluate the effect of the policies on market outcomes. In all scenarios, I use the APTC *methodology* as in the current policy. This means that subsidized net premiums for the second lowest silver (SLS) will not be impacted, regardless of how premiums respond to counterfactuals. Most notably, this shields subsidized high-cost consumers from high premiums in the simulation without community rating.[52]

Under each policy, I also provide explicit measures of adverse selection. Recall the earlier discussion in Section 3.4 stating that adverse selection can be determined in this environment by the difference between standardized marginal costs (SMC) and standardized average costs (SAC). Therefore, I calculate adverse selection for any plan j as follows:

$$\begin{cases} SMC_j - SAC_j < 0 & \implies \text{Adverse Selection} \\ SMC_j - SAC_j > 0 & \implies \text{Advantageous Selection} \end{cases}$$

[52]For robustness, I also conducted the simulations holding the APTC *level* fixed. All the qualitative results were unchanged. The main difference is it changes the redistributional effects.

where again I define these objects as:

$$SMC_j(p_j) = \frac{\sum_i s'_{ij}(p_j)\Big(\theta_i^c c_j - T_{ij}\Big)}{\sum_i s'_{ij}(p_j)\theta_i^p}$$
$$SAC_j(p_j) = \frac{\sum_i s_{ij}(p_j)\Big(\theta_i^c c_j - T_{ij}\Big)}{\sum_i s_{ij}(p_j)\theta_i^p}$$

The degree to which $SMC_j - SAC_j$ is negative is the degree of adverse selection for that plan. The benefit of this structural framework is that I can compare the degree of adverse selection under current vs. alternative risk adjustment policies, which allows me to identify the causal effect of risk adjustment on adverse selection.

6.2 Simulation Results

This section reviews the results from the counterfactual simulations outlined above.[53] For each of the 4 different policy environments, I discuss premiums, enrollment, spending, and measures of adverse selection.

Before getting to the results, recall the basic predictions from Section 2. Community rating is predicted to drive down enrollment and increase prices by deterring low-cost consumers. These predictions hold for both the extensive margin (prices/enrollment overall), and the intensive margin (price/enrollment for higher quality plans). Risk adjustment should be able to offset some of this by eliminating the effect of risk composition on prices. ACA risk adjustment is expected to have little or no effect on take up. The alternative risk adjustment might be able to increase take up but the degree is an empirical question based on the amount of residual extensive margin adverse selection.

I begin by examining the effect of each policy on prices in the market. Table 6 presents the (unweighted) average premiums of each policy by plan type. The premium of each plan is measured as the average premium in the population, but at the current baseline price ($\bar{\theta}^p p_{jt}$). First consider the effect of community rating alone (column 2 vs. column 1). The first three rows of the table give premiums by metal tier. Going from no community rating to community rating without risk adjustment, it is clear that prices are effected as one would theorize. The average premium of a Gold/Platinum plan increases from \$466 to \$530, while Bronze plans decline slightly from \$291 to \$288. These reflect a shift of younger consumers away from high-generosity plans to low-generosity plans. For network types, one can see the premiums of all plan types increase under community

[53]For the counterfactual exercises I exclude regions 1, 9 and 10 (about 10% of state enrollment). In these regions, some people only had access to one insurer in 2015. This makes the simulated prices unstable (particularly from the selection at the heart of this paper). Understanding how selection impacts these markets with very little competition is important and a good topic for future work. To keep things relatively simple, this paper removes them from the analysis.

Table 6: Average Premium by Plan Type

	No Comm.	Community Rating		
	Rating	No RA	ACA RA	PlanPop RA
Metal Tier				
1. Bronze	291	288	299	296
2. Silver	362	377	371	366
3. Gold/Plat	466	530	474	467
Network Type				
HMO	370	398	378	368
Kaiser	377	405	385	374
NonHMO	406	444	415	415
Region				
NorCal	450	496	460	467
SoCal	364	391	371	361
Overall	388	420	396	391

Note: This table gives the unweighted mean premium under each simulation for each plan type. Plan-level premiums are based on the average in population (scaled by $\overline{\theta}^p$). Plans in regions 1, 9, and 10 omitted. Units in dollars.

rating. The average HMO premium would increase from \$370 to \$398, while that of Non-HMOs would increase from \$406 to \$444. This suggests that adverse selection impacts both plan types, though disproportionately more for non-HMOs. The last row indicates that the market as a whole is adversely selected. The average premium would increase from \$388 to \$420 with community rating.

Next consider the effect of risk adjustment under the ACA, presented in the third column of the table. Generally speaking, prices in this environment are somewhere between the prior two scenarios. The average premium of a Gold/Platinum plan would decrease with risk adjustment from \$530 to \$474, compared to pure community rating. Interestingly, the average premium for Bronze plans would increase from \$288 to \$299, even higher than under no community rating. This is because risk adjustment requires plans that attract healthier enrollees to make payments, which increases premiums. Examining premiums by network types, all plan types experience lower premiums when adding risk adjustment. Note, however, these are unweighted by enrollment, so this is largely driven by high generosity plans that have relatively small enrollment.

Finally, consider the effect of using the alternative "PlanPop" risk adjustment, presented in the last column. Relative to ACA risk adjustment, the alternative risk adjustment generally reduces premiums for all categories. It significantly reduces premiums for HMOs, while there is virtually no effect for non-HMOs. This is because this method disproportionately lowers risk adjustment payments for lower cost plans, which passes through to premiums. The fact that prices overall don't come down to the levels seen without community rating suggests there is still adverse selection along the extensive margin that puts upward pressure on all prices.

The effects of each policy on enrollment also match theoretical predictions. Since the price effects discussed above largely parallel many of the same patterns, I keep this discussion brief. Tables 8 and 7 present enrollment shares and counts, respectively, for each plan type under each counterfactual. Rows and columns represent the same simulations and plan types discussed in the prior paragraphs.

Beginning with shares in Table 8, all the patterns described on prices are present. Without risk adjustment, community rating decreases market shares of Gold/Platinum plans from 18.9% of the market to just 3.3%. It would also decrease enrollment in non-HMO plans from 48.3% to 41.1%. Note that in many areas HMOs are not available. Hence these effects are much larger within regions where HMOs are present. Adding ACA risk adjustment largely closes this gap. For example, it increases the share of Gold/Platinum plans to 10.5% from 3.3% without risk adjustment–hence roughly half of the difference. ACA risk adjustment also increases the non-HMO share from 41.1% to 45.7%, which is over halfway to the 48.3% without community rating. Finally, examining the last column, one can see that the alternative risk adjustment further restores enrollment towards generous plans as was predicted. Table 7 gives the impact of each policy on the extensive margin by presenting the total enrollment counts. The first column shows that without

Table 7: Market Enrollment by Plan Type

	No Comm. Rating	Community Rating		
		No RA	ACA RA	PlanPop RA
Metal Tier				
1. Bronze	232,979	331,518	230,763	228,795
2. Silver	589,481	509,381	544,626	552,934
3. Gold/Plat	191,934	29,080	90,493	106,947
Network Type				
HMO	254,608	263,370	234,706	235,886
Kaiser	269,620	249,278	235,068	240,049
NonHMO	490,166	357,330	396,107	412,741
Overall	1,014,394	869,979	865,882	888,676

Note: This table gives total enrollment under each simulation for each plan type. Plans in regions 1, 9, and 10 omitted. Units in households.

Table 8: Market Shares by Plan Type

	No Comm. Rating	Community Rating		
		No RA	ACA RA	PlanPop RA
Metal Tier				
1. Bronze	23.0	38.1	26.7	25.7
2. Silver	58.1	58.6	62.9	62.2
3. Gold/Plat	18.9	3.3	10.5	12.0
Network Type				
HMO	25.1	30.3	27.1	26.5
Kaiser	26.6	28.7	27.1	27.0
NonHMO	48.3	41.1	45.7	46.4
Overall	100	100	100	100

Note: This table gives enrollment shares conditional on take-up, under each simulation for each plan type. Plans in regions 1, 9, and 10 omitted. Units in Percentages (%).

community rating, total enrollment would be just over 1 million households in the examined regions. Adding community rating decreases total enrollment to roughly 870 thousand, a 14.2% reduction.[54] Adding ACA risk adjustment has virtually no effect on enrollment, decreasing it a very small amount if anything. The alternative "PlanPop" risk adjustment, conversely, can both restore shares as well as increase total enrollment. This is seen in the last column, where overall enrollment would increase from 866 to 889 thousand households.

Table 9: Risk Adjustment Receipts Per Enrollee

	No Comm. Rating	Community Rating: No RA	Community Rating: ACA RA	Community Rating: PlanPop RA
Metal Tier				
1. Bronze	0	0	-23.44	-9.04
2. Silver	0	0	1.62	16.38
3. Gold/Plat	0	0	33.34	65.14
Network Type				
HMO	0	0	-8.46	5.95
Kaiser	0	0	-3.99	13.22
NonHMO	0	0	3.57	22.73
Region				
NorCal	0	0	-4.32	14.45
SoCal	0	0	-0.52	16.3
Overall	0	0	-1.74	15.7

Note: This table gives the average risk adjustment revenue per household, under each simulation for each plan type. Negative values imply firms make payments into RA system. Plans in regions 1, 9, and 10 omitted. Units in dollars.

Next, I turn to the impact of each policy on transfer amounts and government spending.[55]. Table 9 presents the average risk adjustment amount each plan would receive per enrollee. Notably, Under the ACA method, HMOs *pay* an average risk adjustment amount of \$8.46 per household while PPOs receive \$3.57 on average.[56] Using the alternative risk adjustment as opposed to the method in the ACA would lead to overall risk adjustment payments of \$15.70. This reflects adverse

[54] Recall that I hold the APTC method fixed which protects most high-cost enrollees in the case without community rating. Also, as pointed out, I do not have the framework to measure the welfare consequences of this since I underestimate the true variation in medical costs and ignore dynamics. While total enrollment has declined, community rating enables those with the highest WTPs to get access to insurance.

[55] I only measure risk adjustment and APTCs and ignore how CSRs change–some people might shift to enhanced silver plans from Bronze, affecting the total outlays of government spending, which I do not measure.

[56] The overall average ACA risk adjustment transfer is non-0 despite budget neutrality because regions 1, 9 and 10 are omitted. The fact that it is negative implies that regions 1, 9 and 10, are getting positive transfers on average–i.e. have costlier enrollees. These areas are largely rural.

selection into the market as a whole. In this case, all plan types would receive larger payments (or pay less). For example, HMOs would get \$5.95 per enrollee as opposed to making payments in the current environment.

Table 10: Impact of Risk Adjustment on Government Spending

	APTC	RA	APTC + RA
Baseline			
CR + ACA RA	327,261	438	327,699
Difference (Δ)			
CR Only	-2,422	-438	−2,860
CR + PlanPop RA	-11,526	13,519	1,957

Note: This table gives the impact of each simulation on *monthly* government spending, relative to the current baseline policies. CR is community rating. APTC is the tax credits to subsidize premiums. RA is risk adjustment. Plans in regions 1, 9, and 10 omitted. Units in \$1,000 dollars.

A main result of Table 9 is that the alternative risk adjustment, while it can increase total enrollment and the share in higher quality plans, would require additional funding from the government. Since premium subsidies (APTC) amounts also change between policies, I compare overall government spending in Table 10. The last row of the table presents the impact of switching from the current policy to the alternative risk adjustment. As can be seen, the alternative risk adjustment would require an additional \$13.5 million (monthly) in the exchange. But by reducing risk adjustment payments, it would also lower premiums and hence government spending on APTCs. APTC spending would decline by \$11.5 million, offsetting most of the spending on risk adjustment in the exchange. Hence, the overall budgetary impact of the alternative risk adjustment would be relatively small once factoring in savings from APTC.[57]

To summarize, in absence of any risk adjustment, community rating in the ACA would have a significant impact on the type of enrollment and coverage rates. Risk adjustment successfully offsets most of the effect on market composition–i.e. the chosen network type and metal tier–but no impact on the overall decline in take-up rates. The alternative risk adjustment further restores market shares to the no community rating equilibrium and slightly increase overall enrollment. This risk adjustment would require funding from the government, but most of it would be offset

[57]Recall, off-exchange plans are also combined in the same risk adjustment pool. For those plans, there is no savings in APTCs to offset any positive payments required for risk adjustment. However, the off-exchange population is higher income and likely to have very different risk composition. It could have less extensive market adverse selection, or even "advantageous" selection (Finkelstein and McGarry, 2006; Fang et al., 2008). Moreover, that group by construction has much higher WTPs for insurance. Therefore, the increases to consumer surplus from the alternative risk adjustment are likely much higher for the off-exchange population. More data on this population is needed to do a better cost-benefit analysis.

by savings in the APTCs.

6.2.1 Adverse Selection in the Market

This section describes explicit measures of adverse selection for all plans under the different policy simulations. The discussion in the prior section is suggestive of adverse selection, but the degree is not precisely clear. The metric for adverse selection used in this study is the relative cost of the marginal consumer (standardized for age rating) compared to the average of all inframarginal consumers, i.e. $SMC_j - SAC_j$. Recall a negative value is sufficient evidence of adverse selection for that plan since the marginal consumer is relatively lower risk.

Figure 10: Adverse Selection Under Each Policy

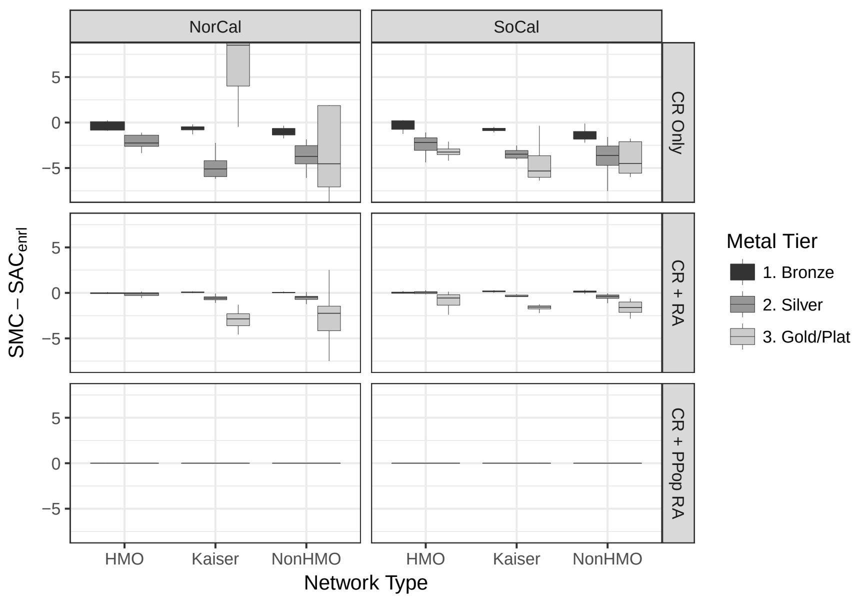

Note: This figure plots the distributions of standardized marginal costs relative to standardized average costs for each plan. A negative value implies adverse selection into the plan, even after any risk adjustment transfers. The unit of analysis is a plan. Outlier plans and plans with fewer than 200 enrollees are excluded. As in other analyses, regions 1, 9, and 10 are also excluded.

Figure 10 plots the distribution of these values for all plans under each policy, omitting the No Community Rating policy which by definition has no adverse selection. The patterns are generally consistent with the prior section and theoretical predictions. First, examine the top panel which represents the market without risk adjustment. Generally, non-HMOs experience more adverse selection, as do higher metal tiers. For example, for a typical Gold/Platinum non-HMO plan, the marginal consumer costs \$5 less than the rest of the plan's enrollment pool. However, the spread of the box plot indicates for some plans this difference is much larger.

The biggest exception to these patterns is Kaiser in Northern CA, where the Gold/Platinum plans appear to positively select from the market–i.e. marginal consumers are actually higher cost. The reason for this is that in Northern CA in particular, there is a large share of low-cost consumers (young males) who have relatively higher incomes and are concentrated in areas that Kaiser plans are offered (non-rural).[58] For these consumers, the income effect dominates adverse selection which leads to generous coverage among low-cost consumers (Finkelstein and McGarry, 2006; Fang et al., 2008).

Moving down to the second panel gives the impact of adding ACA risk adjustment. It is clear that risk adjustment reduces adverse selection, but does not eliminate it entirely. The previously described patterns still hold but are attenuated. For a typical Gold/Platinum Non-HMO plan, the difference between the cost of the marginal and inframarginal consumers is about $-\$2$ after risk adjustment. The pattern on Gold/Platinum Kaiser plans described above appears to be over-corrected by ACA risk adjustment. Finally, note that the bottom panel has no adverse selection by construction of the alternative risk adjustment method.

7 Discussion

In this section, I discuss the implications of the aforementioned findings. The first implication is that the level of competitiveness in the California market is promising for health care reform. Even in markets with a small number of firms, prices are not substantially higher than costs due to the high degree of price-sensitivity. There has been much attention on the lack of insurer participation in many ACA markets.[59] Some have suggested this pattern is a result of adverse selection or other imperfections in the market. However, the simulations suggest risk adjustment in the ACA is quite effective (assuming risk scores in practice are near-perfect predictors of true expected cost as in the model). Hence, these findings support an alternative hypothesis: price-sensitive consumers keep margins low. Beyond a relatively small number of entrants, marginal profits are too low to cover large fixed costs of operating health plans.

Moreover, the high price sensitivity in the market provides an explanation about the types of plans that exist. Given the high degree of plan substitutability for a large share of enrollees, low-cost plans (i.e. narrow networks or tightly managed care) have a strong competitive advantage. A good anecdotal example is UnitedHealthcare in Covered CA. The plan entered a few markets in 2016 with generous networks and hence relatively high costs and prices. The plan garnered very low enrollment and exited the following year. Ho and Lee (2017) measure the welfare implications of narrowing networks more formally. While this level of competitiveness is generally positive for

[58] These might be freelance workers in the technology sector.

[59] See: `https://www.kff.org/health-reform/issue-brief/insurer-participation-on-aca-marketplaces-2014-2017/` or `https://www.kff.org/interactive/counties-at-risk-of-having-no-insurer-on-the-marketplace-exchange-in-2018/`

the ACA, the results might not carry over to other populations (e.g. with high incomes), where elasticities are known to be lower (Ho, 2006; Ericson and Starc, 2015). Additionally, if product variety is extremely low as in some markets nationwide, it is important in and of itself regardless of prices.[60] More work is needed explicitly testing alternative hypotheses on why there are so few entrants and on the types of plans that enter.

Given the results in the simulations, another implication is that risk adjustment is very important in addressing adverse selection between plans. This is true regardless of using the method in the ACA or the "textbook" alternative tested in the simulations. However, risk adjustment that is tied to individual plan costs and that corrects for selection along the extensive margin could have benefits to consumers in multiple ways. First, it would reduce premiums for low-cost plans by reducing risk adjustment payments. Second, it would increase transfers to plans that attract sicker enrollees which could reduce residual adverse selection. It would also reduce adverse selection along the extensive margin to a small degree. It is important to note, however, that there are other efficiency considerations that I cannot measure in this framework. The current method of risk adjustment more strongly incentivizes cost reductions in the long run for high-cost plans. Hence, ACA risk adjustment could lead to more aggressive bargaining or additional medical management strategies, both of which would pass to consumers though lower premiums.

Finally, policies that can increase premiums for any particular group should be expected to reduce rates of take-up. While community rating has important efficiency and distributional effects, it is likely contributing to low rates of coverage among those with low expected health care spending. Hence, policies aimed at universal coverage should pair community rating with subsidies that target low-cost groups (Tebaldi, 2016). The use of a stronger mandate would achieve higher rates of coverage and has other benefits to consumers, such as lower markups (Ericson and Starc, 2015) or reduced adverse selection (Hackmann et al., 2015). However, this study suggests that a stronger mandate also has consequences for those that have low valuations of insurance.[61]

8 Robustness

The findings of this paper are within the context of the particular model employed, which is not without limitations. There are a number of simplifying assumptions needed due to data availability or tractability that impact the results. The most crucial of these is the assumption that expected spending across consumers within a plan varies only by age and gender–hence adverse selection is only caused by gender differences and age differences that exceed the 3:1 pricing ratio. I ignore all

[60] Having a very small number of plans is also problematic in these markets where government subsides are linked to prices (Jaffe and Shepard, 2017).

[61] The overall welfare is more complicated by the fact that consumers without insurance seem to benefit from "uncompensated care" (Finkelstein et al., 2017a), which would be saved under more complete coverage. Finkelstein et al. (2017b) explore other beneficiaries of the ACA reforms, since many consumers seem to have low valuations of insurance.

other health conditions (or preferences) that could determine expected spending. The impact of this is that I am underestimating the true variance in expected spending, which can be large due to consumers with certain chronic high-cost conditions. While this is clearly not accurate, it is less incorrect than in other contexts. The market excludes those that are covered by other government programs such as the elderly, many chronically disabled, and those in poverty–all of which would likely have higher variation in expected health spending. To the extent that I am missing significant variation, I underestimate the degree of adverse selection in the market. This is particularly relevant for the simulation of community rating without risk adjustment, and the degree of extensive margin selection under current ACA risk adjustment. In the first case of community rating without risk adjustment, there would likely be even lower coverage in generous plans. In the case of using "textbook" risk adjustment which corrects for extensive margin selection, the benefits are likely to be even larger (i.e. I currently underestimate the amount of adverse selection along the extensive margin). The government cost of the risk adjustment would also likely increase with the alternative risk adjustment method.[62]

In future iterations of this study, I plan to add unobservable heterogeneity in costs which can be correlated with newly added unobservable preferences. The correlation can be identified from the relationship in premiums across metal tiers (see footnote 62) and variation in premiums under varying market structures (similar to the method employed by Lustig (2010)). Adding unobservable preference heterogeneity will also make the demand side more flexible, and can be identified from adding other years of data. While the current model is imperfect, it is still informative in getting the relative effectiveness of different risk adjustment methods, as well the direct impact of community rating on total enrollment.

Given the limitations discussed above, I do not calculate welfare comparisons with the model. Since I underestimate variation in costs and in preferences, I will be missing important tails in both of those distributions. Moreover, as highlighted by Handel et al. (2015), the benefits from community rating come from intertemporal risk smoothing. I do not model this aspect of insurance and so would not provide a complete measure of the welfare benefits. For these reasons, I focus exclusively on market outcomes. In this study, I use the adverse selection metric as a proxy for welfare. But since adverse selection is only problematic insofar as it leads to welfare reductions, more work is needed examining welfare explicitly.

Another major limitation of this work is that I assume risk scores are perfect measures for expected spending. While they have improved dramatically over time, there could still be health

[62] There is an important empirical finding that suggests the model is imperfectly capturing cost heterogeneity that is correlated with preferences. Once adjusting for actuarial value differences (including moral hazard), the cost of a 21 year-old should be the same across all tiers within a market-carrier. Contrary to this, I find the AV-adjusted costs is increasing in the level of coverage, which is a prediction of Rothschild and Stiglitz (1976) under hidden information. Hence, there is likely unobserved cost heterogeneity which is correlated with the preference for more generous coverage. Fortunately, this precise gap can be used in future iterations of the model to identify the degree of unobserved adverse selection.

(Brown et al., 2014) or preference (Einav et al., 2013, 2016; Shepard, 2016) reasons that this is not the case. A violation in this assumption translates to additional adverse selection in either simulation with risk adjustment. Since this bias enters both cases, the differences should roughly cancel. It would also overestimate the benefits of risk adjustment relative to straight community rating.

Also as a practical issue, the alternative “textbook” risk adjustment described in the analysis is unattainable for a number of reasons. First, it assumes the regulators perfectly know plan costs. Using claims plus some administrative amount could help infer plan costs, but it is likely to be imperfect. In practice, risk adjustment in this style has been based on prices instead of costs, such as in Medicare Advantage. The degree of competitiveness in the market will dictate how much prices and costs differ. If there is sufficient market power, risk adjustment tied to prices could effect firm pricing incentives (Curto et al., 2014; Jaffe and Shepard, 2017). Second, adjusting for risk relative to the population requires knowledge of the risk of the uninsured households, which is not observable in claims. However, regulators could use survey data on health conditions to approximate the relative risk of the insured and uninsured populations.

Finally, it should be pointed out once more that this study is within the context of Covered CA, which differs from other exchanges around the country. The many strategies outlined in Section 2.2 make the California market more competitive, and with higher quality products. If other exchanges differ in these ways, the implications of this study might be different.

9 Conclusion

The ACA establishes a number of regulatory policies aimed at granting access to coverage and minimizing any resulting inefficiencies. This is particularly true in the newly created subsidized exchanges, where private firms compete in a regulated market environment. This paper examines the regulations targeting the supply-side of that market–community rating and risk adjustment. I not only estimate a model that includes these features of the ACA, but I provide a framework to think about adverse selection in this context–a main source of inefficiency that policymakers are concerned about.

Since risk adjustment in the ACA successfully eliminates the majority of adverse selection across plans, how does this reconcile with low rates of coverage and low enrollment in more preferred plans? This study finds that consumers are highly price-sensitive and hence are drawn to low-cost options, regardless of adverse selection. Similarly, community rating, while having large welfare benefits, drives many low-cost consumers out of the market by increasing their premiums. In general, most of the reduction in enrollment from community rating cannot be recovered from risk adjustment, even when adjusting on the extensive margin. While risk adjustment can shape relative market shares to more closely represent those without adverse selection, it does little to overall take-up. This study supports prior findings and suggests that policies aimed at universal coverage need to consider the high price-sensitivity for low-cost consumers (Tebaldi, 2016; Ericson and Starc, 2015) and the relatively low willingness-to-pay for insurance in general (Finkelstein et al., 2017a).

While this study closely examines market outcomes under each policy, there is still much to know about the welfare implications. This is an important area of future work. Both community rating and risk adjustment have clear trade-offs. Community rating redistributes across consumers. The benefits of reducing premium hikes are paired with increasing prices to low-cost consumers, many of whom will no longer buy coverage. For risk adjustment, redistribution happens across firms by "subsidizing" firms that attract high-cost enrollees and "taxing" those that attract low-cost enrollees. The benefit is that generally desirable plans that are exposed to adverse selection can remain in the market and will be affordable to consumers of all risk types. But the consequence is that low-cost products become more expensive, which can also hurt consumers. Finding the optimum for each of these policies needs further investigation and will likely be case-specific. This paper sheds light on these trade offs and sets up a framework to continue this investigation. As the U.S. government increasingly relies on the private market to deliver public benefits, these questions are increasingly important. Whether it be the ACA exchanges, Medicare, or Medicaid, this study shows these supply-side policies can have a significant impact on the health coverage people obtain.

References

George A. Akerlof. The market for "lemons": Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.

Steven Berry, James Levinsohn, and Ariel Pakes. Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890, 1995.

Steven Berry, James Levinsohn, and Ariel Pakes. Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112 (1):68–105, 2004.

Steven T Berry. Estimating discrete-choice models of product differentiation. *The RAND Journal of Economics*, 25(2):242–262, 1994.

Andrew B Bindman, Denis Hulett, Todd P Gilmer, and John Bertko. Sorting out the health risk in california's state-based marketplace. *Health Services Research*, 51(1):129–145, 2016.

Jason Brown, Mark Duggan, Ilyana Kuziemko, and William Woolston. How does risk selection respond to risk adjustment? new evidence from the medicare advantage program. *The American Economic Review*, 104(10):3335–3364, 2014.

Thomas Buchmueller and John DiNardo. Did community rating induce an adverse selection death spiral? evidence from new york, pennsylvania, and connecticut. *The American Economic Review*, 92(1):280–294, 2002.

Gary Claxton, Cynthia Cox, A Damico, L Levitt, and K Pollitz. Pre-existing conditions and medical underwriting in the individual insurance market prior to the aca. *Menlo Park, CA: Kaiser Family Foundation*, 2016.

Vilsa Curto, Liran Einav, Jonathan Levin, and Jay Bhattacharya. Can health insurance competition work? evidence from medicare advantage. Technical Report 20818, National Bureau of Economic Research, 2014.

David M Cutler and Sarah J Reber. Paying for health insurance: The trade-off between competition and adverse selection. *The Quarterly Journal of Economics*, 113(2):433–466, 1998.

Leemore Dafny. Estimation and identification of merger effects: An application to hospital mergers. *The Journal of Law and Economics*, 52(3):523–550, 2009.

Liran Einav, Amy Finkelstein, and Mark R Cullen. Estimating welfare in insurance markets using variation in prices. *The Quarterly Journal of Economics*, 125(3):877–921, 2010.

Liran Einav, Amy Finkelstein, Stephen P. Ryan, Paul Schrimpf, and Mark R. Cullen. Selection on moral hazard in health insurance. *American Economic Review*, 103(1):178–219, February 2013.

Liran Einav, Amy Finkelstein, Raymond Kluender, and Paul Schrimpf. Beyond statistics: The economic content of risk scores. *American Economic Journal: Applied Economics*, 8(2):195–224, 2016.

Keith M Marzilli Ericson. Consumer inertia and firm pricing in the medicare part d prescription drug insurance exchange. *American Economic Journal: Economic Policy*, 6(1):38–64, 2014.

Keith M Marzilli Ericson and Amanda Starc. Pricing regulation and imperfect competition on the massachusetts health insurance exchange. *Review of Economics and Statistics*, 97(3):667–682, 2015.

Hanming Fang, Michael P Keane, and Dan Silverman. Sources of advantageous selection: Evidence from the medigap insurance market. *Journal of Political Economy*, 116(2):303–350, 2008.

Amy Finkelstein and Kathleen McGarry. Multiple dimensions of private information: Evidence from the long-term care insurance market. *American Economic Review*, 96(4):938–958, September 2006.

Amy Finkelstein, Nathaniel Hendren, and Mark Shepard. Subsidizing health insurance for low-income adults: Evidence from massachusetts. Technical Report 23668, National Bureau of Economic Research, 2017a.

Amy Finkelstein, Neale Mahoney, and Matthew J Notowidigdo. What does (formal) health insurance do, and for whom? Technical Report 23718, National Bureau of Economic Research, 2017b.

Michael Geruso and Timothy Layton. Selection in health insurance markets and its policy remedies. Working Paper 23876, National Bureau of Economic Research, September 2017.

Jacob Glazer and Thomas G McGuire. Optimal risk adjustment in markets with adverse selection: An application to managed care. *American Economic Review*, 90(4):1055–1071, 2000.

Austan Goolsbee and Amil Petrin. The consumer gains from direct broadcast satellites and the competition with cable tv. *Econometrica*, 72(2):351–381, 2004.

Gautam Gowrisankaran, Aviv Nevo, and Robert Town. Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 105(1):172–203, 2014.

Martin B Hackmann, Jonathan T Kolstad, and Amanda E Kowalski. Adverse selection and an individual mandate: When theory meets practice. *The American Economic Review*, 105(3): 1030–1066, 2015.

Ben Handel, Igal Hendel, and Michael D Whinston. Equilibria in health exchanges: Adverse selection versus reclassification risk. *Econometrica*, 83(4):1261–1313, 2015.

Berend Hasselman. *nleqslv: Solve Systems of Nonlinear Equations*, 2017. URL `https://CRAN.R-project.org/package=nleqslv`. R package version 3.3.1.

Florian Heiss. Structural choice analysis with nested logit models. *The Stata Journal*, 2(3):227–252, 2002.

David A Hensher. Sequential and full information maximum likelihood estimation of a nested logit model. *The Review of Economics and Statistics*, 68(4):657–667, 1986.

Marek Hlavac. *stargazer: Well-Formatted Regression and Summary Statistics Tables*. Harvard University, Cambridge, USA, 2015. URL `http://CRAN.R-project.org/package=stargazer`. R package version 5.2.

Kate Ho and Robin S Lee. Equilibrium provider networks: Bargaining and exclusion in health care markets. Technical Report 23742, National Bureau of Economic Research, 2017.

Kate Ho, Joseph Hogan, and Fiona Scott Morton. The impact of consumer inattention on insurer pricing in the medicare part d program. Working Paper 21028, National Bureau of Economic Research, 2015.

Katherine Ho. The welfare effects of restricted hospital choice in the us medical care market. *Journal of Applied Econometrics*, 21(7):1039–1079, 2006.

Katherine Ho. Insurer-provider networks in the medical care market. *American Economic Review*, 99(1):393–430, 2009.

John Holahan, Linda J Blumberg, and Erik Wengle. *What Does the Failure of Some Co-ops and the Possible Pullout of United Healthcare Mean for the Affordable Care Act?* Urban Institute, 2016.

Sonia P Jaffe and Mark Shepard. Price-linked subsidies and health insurance markups. Technical Report 23104, National Bureau of Economic Research, 2017.

John Kautter, Gregory C Pope, and Patricia Keenan. Affordable care act risk adjustment: Overview, context, and challenges. *Medicare & Medicaid Research Review*, 4(3), 2014.

Timothy Layton, Ellen J Montz, and Mark Shepard. Health plan payment in us marketplaces: Regulated competition with a weak mandate. Technical Report 23444, National Bureau of Economic Research, 2017.

Timothy J. Layton. Imperfect risk adjustment, risk preferences, and sorting in competitive health insurance markets. *Journal of Health Economics*, Forthcoming.

Josh Lustig. Measuring welfare losses from adverse selection and imperfect competition in privatized medicare. *Manuscript. Boston University Department of Economics*, 2010.

Neale Mahoney and E Glen Weyl. Imperfect competition in selection markets. *The Review of Economics and Statistics*, 99(4):637–651, 2017.

Aviv Nevo. Measuring market power in the ready-to-eat cereal industry. *Econometrica*, 69(2): 307–342, 2001.

Joseph P Newhouse, Willard G Manning, Emmett B Keeler, and Elizabeth M Sloss. Adjusting capitation rates using objective health measures and prior utilization. *Health Care Financing Review*, 10(3):41–54, 1989.

Matthew T Panhans. Adverse selection in aca exchange markets: Evidence from colorado. 2017.

Daniel Polsky and Janet Weiner. The skinny on narrow networks in health insurance marketplace plans. 2015.

Gregory C Pope, Henry Bachofer, Andrew Pearlman, John Kautter, Elizabeth Hunter, Daniel Miller, and Patricia Keenan. Risk transfer formula for individual and small group markets under the affordable care act. *Medicare & Medicaid Research Review*, 4(3), 2014.

R Core Team. *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria, 2017. URL `https://www.R-project.org/`.

Roni Caryn Rabin and Reed Abelson. Health plan cost for new yorkers set to fall 50%. *New York Times*, 2013.

Jo Reynaerts, R Varadha, and John C Nash. Enhencing the convergence properties of the blp (1995) contraction mapping. 2012.

Michael Rothschild and Joseph Stiglitz. Equilibrium in competitive insurance markets: An essay on the economics of imperfect information. *The Quarterly Journal of Economics*, 90(4):629–649, 1976.

Steven Ruggles, Katie Genadek, Ronald Goeken, Josiah Grover, and Matthew Sobek. Ipums-usa, university of minnesota, www.ipums.org, 2017.

Mark Shepard. Hospital network competition and adverse selection: Evidence from the massachusetts health insurance exchange. Working Paper 22600, National Bureau of Economic Research, 2016.

Pietro Tebaldi. Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca. 2016.

Robert Town and Su Liu. The welfare impact of medicare hmos. *The RAND Journal of Economics*, 34(4):719–736, 2003.

Kenneth E Train. *Discrete Choice Methods with Simulation*. Cambridge University Press, 2009.

Ravi Varadhan. *SQUAREM: Squared Extrapolation Methods for Accelerating EM-Like Monotone Algorithms*, 2016. URL `https://CRAN.R-project.org/package=SQUAREM`. R package version 2016.8-2.

Micah Weinberg and Patrick Kallerman. A study of affordable care act competitiveness in california, February 2017.

Hadley Wickham. *ggplot2: Elegant Graphics for Data Analysis*. Springer-Verlag New York, 2009. ISBN 978-0-387-98140-6. URL `http://ggplot2.org`.