Market Structure in the Presence of Adverse Selection

Conor Ryan

Unversity of Minnesota - Economics Department

Preliminary - Work in Progress

JEL: I11, L13 ; Keywords: Insurance, Competition

February 21, 2019

Abstract

Competitive markets with adverse selection lead to inefficient allocations. These failures come, in part, from firms that do not internalize the effect of their own prices on the mix of consumers that purchase—and therefore the costs of—other products in the market. Monopolists do not suffer from this inefficient sorting as they own every product in the market, but create inefficient allocations by charging significant markups. In this paper, I explore the trade-offs between these two distortions and their interaction with policies commonly used to mitigate adverse selection. I use novel choice data from the non-group health insurance market and the HHS-HCC risk prediction model to estimate the joint distribution between preferences and cost and use these estimates to simulate the equilibrium effects of market structure. I show that under certain policy regimes, more concentrated markets can improve allocations for consumers with a high cost to serve and high willingness-to-pay. I simulate equilibrium in current non-group insurance markets and simulate a merger to demonstrate how two policies—the individual mandate penalty and risk adjustment—interact with market structure. In the case of the individual mandate, I find that it leads to higher premiums in the most concentrated markets, the opposite of the intended effect. I find that, while risk adjustment has less effect in concentrated markets, market concentration itself provides an implicit form of risk adjustment.

1 Introduction

The presence of adverse selection, where consumers with the highest willingness-to-pay are also the most costly to serve, can lead to well-known market failures in health insurance markets (Akerlof (1970), Rothschild and Stiglitz (1976), Einav et al. (2010), Hendren (2013), Handel et al. (2015)). These failures come, in part, from firms that do not internalize the effect of their own prices on the consumers that purchase—and therefore the costs of—other products in the market. This externality among firms leads to inefficient sorting of consumers among the menu of offered products (Layton (2017)). Several theoretical papers have suggested that market power can attenuate the welfare losses from adverse selection (Veiga and Weyl (2016), Mahoney and Weyl (2017), Lester et al. (2015)), and indeed a monopolist does not suffer from inefficient sorting, as it internalizes the market-wide distribution of product costs. However, market power comes with its own distortions via markups charged over marginal cost. In this paper, I explore the trade-off between these two sources of inefficiency, and the implications of market structure on the policy solutions for mitigating these distortions. I find that while competition delivers broad welfare gains for consumers through lower markups, these gains are partially offset by adverse sorting. The gains from lower markups are shared generally by all consumers, but the loss from adverse sorting is concentrated among those with the highest willingness-to-pay.

Despite significant heterogeneity in market structure in US health insurance markets, the relationship between market structure, adverse selection, and the policies used to address them has received comparatively little attention (Geruso and Layton (2017)). In the non-group health insurance market, the largest firm had a market share of over 85% in 5 states and less than 33% in another 5 states. In 2014, the Affordable Care Act implemented a suite of reforms designed to address the market failures known to competitive markets but ignored important ways those policies interact with market structure in concentrated markets. For example, I find that a purchasing mandate and penalty can lead to higher prices in concentrated markets, the opposite of the intended effect, and show that risk adjustment is unnecessary in highly concentrated markets.

I model strategic firms that sell differentiated insurance products and set prices to maximize profit according to Nash-Bertrand price competition. Consumer's choose among insurance products and an option to be uninsured, and I allow a consumer's demand to be correlated with the cost of selling her insurance—the key feature of adverse selection. In this environment, I can decompose the total welfare loss in a competitive equilibrium with adverse selection into two sources: (i) non-negative profits and (ii) inefficient sorting. I show that that first distortion is present in any equilibrium, and its magnitude is increasing in market concentration as firms are able to extract more profit. The second source of welfare loss, inefficient sorting, is decreasing in market concentration and entirely absent in a monopoly. I characterize efficient sorting incentives in markets with non-negative profits and discuss how it relates to current risk adjustment policies that intend to mitigate sorting inefficiencies.

To estimate the model, I use new data on household health insurance choices in the non-group health insurance market made through an online insurance broker. The non-group market, where individuals buy insurance directly rather than through an employer or government program, serves about 40 million people. While there are many segments of the US health insurance market, the non-group market demonstrates the most classic manifestations of adverse selection and has substantial geographic variation in market structure. These features make the non-group market a focus for national health policy and an important setting to study the relationship between adverse selection and market structure.

This data is unique in two respects. First, it is a large sample of non-group insurance purchases made through a non-government broker. Estimates from the American Community Survey indicate that 30% of the national non-group market are middle-to-high income—i.e. earn more than 400% of the federal poverty level and are ineligible to receive premium subsidies—while the same group represents only 2% of consumers purchasing insurance through government-run marketplaces (ASPE (2016)). The data are slightly over-representative of this higher income group (50% of my sample). Much of the previous work in estimating insurance demand has found that income and subsidy eligibility are important determinants of elasticity, which suggests that this section of the market may have substantially different preferences and may be important for understanding firm behavior. Second, the data span more than 100 local markets, which allows me to estimate equilibrium outcomes in a diverse cross-section of market structure.

In order to identify the relationship between demand and cost, I use a novel approach that links moments on average costs with demand moments via the Health and Human Services Hierarchical Condition Categories (HHS-HCC) risk prediction model.[1] Data that links non-group health insurance choices with measures of health status are rare, and recent approaches

[1]The HHS-HCC risk prediction model is used to administer the risk adjustment transfer system in the non-group market.

to identifying the relationship rely on simulated variation in demand and average costs (Tebaldi (2017)). I improve on this approach by using HHS-HCC moments in the non-group market to discipline a method of simulated moments approach with moments on the distribution of costs and risk in the Medical Expenditure Panel Survey. This method does not require that I assume anything about the whether firms are behaving optimally.

With credible estimates of the joint distribution of preferences and cost, I simulate an economy with varying levels of market concentration to illustrate the trade-offs between welfare loss as a result of markups and inefficient sorting. In particular, I simulate an economy with a menu of 20 products each offered by a single product firm, with the exception of a single large firm that owns at least two of the products. I vary the level of concentration in the market by increasing the number of products owned by the large firm. I find that consumers are better off in a competitive market, as the gains from lower markups are substantial. However, the welfare gains of competitive markets are smallest for consumers with the highest willingness to pay, which bear the cost of inefficient sorting through high prices for more generous insurance products. When there is a strong purchasing mandate, the consumers with the highest willingness to pay can benefit from moderate levels of concentration.

Since the distortions in competitive markets are different than those in concentrated markets, it follows that government policies intended to correct competitive market distortions may have heterogeneous and unexpected effects across different levels of market concentration. Out of the 109 markets for which I have data, one third have a Herfindahl-Hirschman Index (HHI) of less than 3600, another third have an HHI of between 3600 and 5500, and the final third have an HHI of at least 5500. I find that in the least concentrated markets, the risk adjustment and mandate penalty policies have the intended effect. The risk adjustment policy reduces the premium spread between different levels of generosity through a system of inter-firm transfers, and the individual mandate penalty lowers premiums overall by increasing participation among lower cost and lower willingness-to-pay individuals.

In the most concentrated markets, risk adjustment has very little effect, as large firms recapture much of the transfers. However, market concentration leads to more efficient sorting. Since there is little welfare loss resulting from sorting in concentrated markets, the policy is less necessary. The individual mandate penalty has the opposite of the intended effect, and leads to higher premiums in equilibrium. This is a result of firms with substantial market power that respond to the increase in demand by increasing markups rather than passing through the benefits of lower average costs to consumers.

This paper makes three main contributions. First, I provide a model and intuition for the trade-off between two sources of inefficiency in markets with adverse selection. This builds on a theoretical literature on contract design in markets with adverse selection that documents the ways in which private firms deviate from the socially optimal (e.g., Akerlof (1970), Rothschild and Stiglitz (1976), Veiga and Weyl (2016), Lester et al. (2015)) and an empirical literature measuring the effects of these deviations in health insurance markets (e.g., Einav et al. (2010), Handel et al. (2015), Layton (2017)).

While US health insurance markets are highly concentrated, there has been less focus in the literature on the effects of market power on contract design. Some recent theoretical work uses elasticity estimates from the literature to show that welfare in insurance markets may be U-shaped in the degree of competition and that monopolists have an efficient sorting incentive in the quality of a single product (Mahoney and Weyl (2017), Veiga and Weyl (2016), Lester et al. (2015)). This paper extends these results to a setting with multiple differentiated products that can be easily used for empirical work.

Recent work by Geruso et al. (2018) evaluates the relationship between intensive and extensive margin selection. While their work focuses on a model of perfectly competitive firms that make zero profits, the intuition behind the forces driving consumer sorting is similar to the arguments made here about inefficient sorting and markups (roughly analogous to a tax on the extensive margin). This paper highlights that policy makers should view competition as a lever that can influence the selection properties of the market.

Second, I show that competition may have an ambiguous effect on generous insurance products, even without taking into account effects on cost. Much of the empirical literature on the effects of competition on insurance prices is motivated by the two-sided nature of the market—insurance firms with market power may be able to raise markets but also lower costs through hospital bargaining (Ho and Lee (2017), Dafny et al. (2012)). These papers, as well as recent empirical work on the non-group market (Dafny et al. (2015), Abraham et al. (2017)), show that competition typically leads to lower prices. However, this paper shows that the effects of market power may also be uneven across different product offerings. In particular, the effect of competition on the most generous plan offerings may be small and even positive, before accounting for bargaining effects. This suggests new possible interpretations for previous empirical findings. In the appendix (available by request), I extend Dafny et al. (2015) to show that the price-reducing effects of competition vanish for generous products.

My third contribution is to estimate a structural model to evaluate the heterogeneous effects of the individual mandate penalty and the ACA risk adjustment policy across markets with different levels of concentration. I am contributing to a growing literature on evaluating policies in regulated health insurance markets with a model of imperfect insurance competition (Miller et al. (2018), Jaffe and Shepard (2017), Tebaldi (2017), Ericson and Starc (2015), Starc (2014), Saltzman (2017)), a related literature that studies health insurance firms' specific mechanisms and incentives to engage in risk selection (Aizawa and Kim (2015), Decarolis and Guglielmo (2017)).

I build on this literature by provide estimates of the demand for health insurance using novel data: non-group health insurance purchases through a national, non-government broker. Previous literature on the demand for health insurance in the non-group market focused plans purchased through government-run marketplaces, frequently in California and Massachusetts (Tebaldi (2017), Ericson and Starc (2015), Frean et al. (2017), Shepard (2016), Saltzman (2017), DeLeire et al. (2017)), or addressed only the elasticity of the decision to purchase any insurance (Marquis et al. (2004), Gruber and Poterba (1994)).

In addition to providing demand estimates from new data, I also implement a new approach to identifying the joint distribution of preferences for health insurance and health risk, the key feature of adverse selection. In markets where there the data is available, this relationship can be identified through observing measures of health status (Aizawa and Kim (2015), Shepard (2016), Jaffe and Shepard (2017)). However, this data is uncommon for the non-group market. One approach is to estimate the relationship between a random willingness-to-pay for coverage generosity and firm-level average costs (or optimality conditions) through the simulated distribution of enrollment (Tebaldi (2017)). I improve on this method by applying the HHS-HCC risk prediction model to the Medical Expenditure Panel Survey, which contains information on health status, demographics, and health expenditures in the non-group insurance market. I use these moments, along with risk score moments published by regulators, to robustly estimate the relationship between demand, risk, and cost.

There is a large body of literature on the effects of the individual mandate penalty (Frean et al. (2017), Graves and Gruber (2012), Hackmann et al. (2015), Saltzman (2017), **?**, Geruso et al. (2018)). Much of this work finds that the mandate had an important effect on coverage during the Massachusetts health reform in 2008. However, it may not be generalizable to the national implementation of the penalty in 2014 (Frean et al. (2017)). Hackmann et al. (2015) also find that the Massachusetts health reform led in general to lower markups, though they attribute this change to many of the other market reforms that came with the mandate penalty. To the best of my knowledge, this is the first paper to document the ways in which the effects of a mandate penalty depend on local market structure.

I am also contributing to a literature on how risk adjustment transfer systems relate to firm strategies (Glazer and McGuire (2000), Ellis and McGuire (2007), Geruso and Layton (2015), Brown et al. (2014), Layton (2017), Saltzman (2017), Geruso et al. (2018)). Most of this work focuses on the Medicare Advantage market, where risk adjustment has a much longer history and takes a slightly different form. Layton (2017) shows how the imperfections in the ACA risk prediction can be exploited in competitive markets. Geruso et al. (2018) and Saltzman (2017) explore the welfare implications of the ACA risk adjustment system in conjunction with the individual mandate. In this paper, I outline the welfare maximizing sorting incentive and show that concentrated markets may not require substantial risk adjustment.

In section 2, I present the model and illustrate how market structure relates to two sources of inefficiency. In section 3, I explain the non-group health insurance environment and the data. In section 4, I describe the demand estimation and results, and in section 5, I describe the cost estimation. In section 6, I demonstrate the relationship between market structure and adverse selection through a simulated economy, and in section 7, I estimate the policy counterfactuals.

2 Setting and Data

2.1 The Individual Market Under the ACA

The non-group health insurance market, referred to as the "individual market," is the only available insurance market for any household that does not receive an insurance offer from an employer or from the government, through Medicaid or Medicare. 18 million individuals are covered by insurance in this market, and it is the sole offer of health insurance for nearly 20 million more uninsured individuals. Consumers in this market are under the age of 65, earn above the federal poverty level, and are typically employed at smaller firms, self-employed, or unemployed. Relative to the general population, the market is younger, has lower average incomes, and are predominantly single-person households.

The individual market has been a recent focus of national health care policy precisely because of the issues addressed by this paper. Prior to 2014, the market was characterized by low take-up rates, high prices for generous insurance coverage, and frequent coverage denials. The Affordable Care Act (ACA) attempted to address many of these issues through policies targeting competition and adverse selection. Additionally, many markets continue to be highly concentrated. Markets tend to be most concentrated in rural areas where there are already difficulties with access to health care.

Products offered in the individual market are strictly regulated under the ACA. Every plan must cover a set of mandated benefits, and the maximum allowable out-of-pocket expenditures are limited ($6,600 in 2015). Every plan must fit into one of 5 actuarial value categories: Platinum, Gold, Silver, Bronze, or Catastrophic, which are expected to cover, respectively, 90%, 80%, 70%, 60%, and 57% of expenditures on average. Catastrophic plans typically provide no additional coverage beyond the maximum allowable expenditure limit and do not satisfy the individual mandate requirement for individuals over 30 years old.

Households that earn between 100% and 400% of the federal poverty level (FPL) are eligible for a household-specific premium tax credit that sets the after-subsidy premium of the second-lowest cost Silver plan to be a certain percent of household income (roughly 2% at 100% of FPL, and 9.5% at 400% of FPL). Households that earn between 100% and 250% of the federal poverty level are also eligible for cost-sharing subsidies that increase the generosity of Silver plans.

Given these rules, insurance firms are free to compete in price.[2] Firms are restricted to setting a single base price for each product in a particular "rating area"—geographic divisions within each state that are set by the state insurance regulatory authority. The premiums charged to an individual household are adjusted by the age of each household member according to an "age rating" formula, also set by the state.

Policies to Address Adverse Selection

The goal of health reform in the non-group market was to create markets that would both protect consumers from unaffordable prices on account of their health status and avoid problems related to adverse selection. The ACA helps to create affordable insurance options through subsidies for health insurance premiums, strict regulations on what kinds of products

[2]State insurance regulators must approve year-over-year price increases for continuing products. The allowable premium increase is depends on the discretion of the state, and depends on the experience of the insurance firm over the prior year. The pressure these regulatory bodies put on prices vary across states.

insurance firms could offer, and formulas that specify what the prices of those products could depend on.

To address adverse selection, the primary policy tool is the "Individual Mandate", a requirement to purchase insurance and an associated penalty for being uninsured. By taxing all individuals that do not buy health insurance, the insurance market can supposedly be reassured that a broad sample of consumers will purchase insurance, rather than simply the most costly. From 2016 through 2018, the mandate penalty was the maximum of \$695 or 2.5% of household income, and beginning in 2019, the penalty is \$0. In many economic models, an individual mandate is required for the market to exist at all (Handel et al. (2015), Azevedo and Gottlieb (2017)), but the empirical importance of the mandate is ambiguous (Hackmann et al. (2015), Ryan et al. (2019), Ericson and Starc (2015)). In spite of the mandate, the insured rate among the eligible population is still only about 50%.

The individual mandate is designed to target selection on the extensive margin of purchasing insurance. To address intensive margin selection—the tendency of individuals with high expected costs to choose more generous insurance—the ACA implemented "risk adjustment," a system of risk-based subsidies (taxes) that compensate firms for enrollees with higher (lower) than average expected costs. Risk adjustment is administered between firms on the basis of the average risk of each plan in the market with the intention of allowing firms to be agnostic about their consumers' potential risk, or "eliminate the influence of risk selection on the premiums that plans charge." (Pope et al. (2014), Kautter et al. (2014)). Risk-based subsidies are a common policy instrument to reduce adverse selection in health insurance markets (McGuire et al. (2011), Van de ven and Ellis (2000), Ellis and McGuire (2007)).

The government collects claims data throughout the year from every insurance firm in the market to assess the average risk at the plan level using the HHS-HCC risk prediction methodology. This method attributes to each individual a risk score based on age, sex, and a set of diagnoses codes that are organized into hierarchical condition categories. Plans that have lower than average levels of risk are taxed and plans that have higher than average levels of risk receive subsidies. The formula that determines the taxes and subsidies is constructed to be budget neutral at the state-level: the total taxes across all firms within a state are mechanically equivalent to the total subsidies.

2.2 Choice Data

Consumers in the individual market can purchase insurance by contacting an insurance firm directly, visiting the government-run online marketplace, or shopping for insurance through a third-party marketplace. While only the government can administer subsidies, other vendors can route consumer information through the government portal. Not all plans are offered on all platforms, and insurance firms may elect to list some products on certain platforms and not on others. However, apart from insurers that do not list on the government marketplace at all, the kinds of plans listed by insurers typically have only small differences across platforms.[3]

The data on health insurance purchases come from a third-party online broker. The website sells plans that are offered both on and off the ACA health insurance exchanges. In 2015, the website was authorized to sell subsidized health insurance plans in most states. I observe the choices of subsidized and unsubsidized consumers across 48 states.

The data contain information on the age of the consumer, the first three digits of the consumers' zipcode, the plan purchased by the consumer, and the subsidy received. A single observation in the data represents a household, but I observe only one member's age. I assume that this is the age of the head-of-household, i.e. the purchaser of the plan. However, in order to match the household to its relevant choice set, I have to know the ages of every adult (over the age of 14) in the household. I assume that every household that contains more than one individual contains two adults of the same age, and any additional persons are children under the age of 14.[4]

I observe the income of the consumers with some censoring. I observe the income, or can impute the income from the observed subsidy value, of nearly every individual that receives a subsidy. However, I do not know the income for most individuals that do not receive a subsidy. I make the assumption that these individuals have an income level such that they are not eligible for a subsidy. For the purposes of estimation, this assumption is not terribly restrictive. It requires that every individual eligible for a subsidy receives a subsidy, or at least selects a plan as if they would receive the subsidy for which they are eligible. While

[3] Analysis of the Robert Wood Johnson Foundation HIX 2.0 data on plan offerings shows minimal differences between plan offerings on and off the exchange in premiums or deductibles.

[4] The choice data contains premium information, which is noisy for administrative reasons. While I do not use the information in the analysis, I can see if it suggests that my household composition assumption is close to correct. I find that the correlation between the assigned quoted premium for single households and the listed premium quote (.45) is close to the correlation for families (.49).

there is some evidence that there are a non-trivial amount of consumers that are eligible for subsidies that do not receive them on a monthly basis, all consumers should eventually receive the full value of the subsidy for which they are eligible when they file taxes.

After dropping observations because of missing data or incomplete choice sets, the remaining data includes roughly 75,000 individual and family health insurance choices across 14 states and 109 rating areas.[5]

In Table 1, I summarize the online broker data and compare it to other references on the individual insurance market: the 2015 American Community Survey (ACS) and the Office of the Assistant Secretary for Planning and Evaluation(APSE) at the US Department of Health and Human Services. The ACS survey design offers the broadest depiction of the market across all market segments. ASPE publishes detailed descriptive statistics on purchases made through the federally facilitated HealthCare.gov. Relative to the ACS, enrollment through HealthCare.gov is weighted heavily towards low-income, subsidy eligible consumers. As a result, the plan type market shares reported by ASPE are weighted heavily towards silver plans which have extra cost-sharing benefits at low incomes. The online broker data is disproportionately higher income and younger enrollees. The last panel shows plan type market shares conditioned on earning at least 400 % FPL, and the choices are roughly similar with higher enrollment in Bronze plans through the online broker.

2.3 Cost Data

The Center for Medicaid and Medicare Services (CMS) makes publicly available the state-level financial details of insurance firms in the Individual Market for the purpose of regulating the Medical Loss Ratio.[6] This information includes the number of member-months covered by the insurance firm in the state and total costs. I use this information to calculate average firm-level costs by state. [7]

In addition to average firm-level costs, I use average metal-level costs from the 2016 Premium Rate Filing data. Before each new plan-year, insurance firms must submit requests to state

[5]Choice sets are discussed in section 2.5

[6]Insurance firms in this market are restricted in how much premium revenue they may collect, relative to an adjusted measure of medical costs. This constraint is not typically binding. Excess revenue is returned to consumers via a rebate.

[7]These data also include other market segments, such as "grandfathered" plans that have not been actively offered since 2014. These market segments are relatively small, so they should not have a large affect on the average costs of the firms in my data.

Table 1: Data Description

	Online Broker	ACS	ASPE
		Age Distribution	
Under 18	0.0%	0.6%	8%
18 to 26	17.5%	19.3%	11%
26 to 34	23.7%	16.7%	17%
35 to 44	20.4%	16.5%	17%
45 to 54	20.4%	20.1%	22%
55 to 64	18.0%	26.7%	25%
		Income Distribution	
Under 200% FPL	27.6%	36.1%	68%
200% to 400% FPL	21.8%	33.2%	31%
Over 400% FPL	50.6%	30.6%	2%
		Metal Level Market Shares	
Catastrophic	5.4%		1%
Bronze	35.5%		22%
Silver	46.1%		67%
Gold	9.8%		7%
Platinum	3.2%		3%
	Metal Level Market Shares (Income > 400% FPL)		
Catastrophic	7.8%		7%
Bronze	46.8%		35%
Silver	25.8%		32%
Gold	14.9%		19%
Platinum	4.7%		8%

Notes: The American Community Survey numbers come from heads of household that are insured through the individual insurance market. The ASPE numbers come from the 2015 Open Enrollment Report for enrollments through HealthCare.gov. The age numbers are not adjusted for head of household.

insurance regulators to increase the premiums for products that they will continue to offer. In these filings, insurance firms include information on the cost and revenue experience of the insurance product in the prior year. Thus, I can observe this information for products offered in 2015 that are offered again in 2016 and for which premiums increased.

This provides selected coverage of plans offered in the market. Some firms are missing from the data, and plans with higher than expected costs are more likely to appear in the data. After aggregating to the metal level, I observe average costs for 222 out of 333 possible state-firm-metal combinations, which covers 47 out of 55 firms in my choice data.

2.4 Risk Score Data

The 2015 Medical Conditions File (MCF) of the Medical Expenditure Panel Survey (MEPS) contains self-reported diagnoses codes and can be linked to demographic information in the Population Characteristics file. The publicly available data only list 3-digit diagnoses codes, rather than the full 5-digit codes. I follow McGuire et al. (2014) and assign the smallest 5-digit code for the purpose of constructing the condition categories.[8] I then apply the publicly available coefficients for the 2015 HHS-HCC risk prediction model.

I use MEPS data to construct the distribution of risk scores for demographic subgroups and moments of the relationship between age, risk, and total medical expenditures.

To identify the relationship between risk scores and demand, I use six moments on the risk distribution among market enrollees. CMS publishes annual reports on the results of the risk adjustment transfer program. Since the beginning of the program in 2014, they publish average risk scores by state and total member-months by state. I am using a national distribution of risk scores, so I aggregate this information into a national average risk score of all enrollees.

Beginning in 2017, CMS published average risk scores by metal-level and market segment. I use four moments on the average risk score in Bronze, Silver, Gold, and Platinum plans. In order to make it comparable to my data, I use the average of on and off exchange market segments, and scale the risk scores by the ratio of the 2015 national average risk score to the 2017 national average risk score. I supplement these moments with the national average risk score in 2015 and the average risk score of a group of large firms, imputed from their MLR filings.

2.5 Choice Sets

I observe in the choice data only the ultimate choices made by the consumers, but not the scope of available options. In order to construct choice sets, I use the HIX 2.0 data set compiled by the Robert Wood Johnson Foundation. This data set provides detailed cost-sharing and premium information on plans offered in the individual market between 2014 and 2017. The data is nearly a complete depiction of the market for the entire US, but there

[8] For example, I treat a 3-digit code of '571' as '571.00'. McGuire et al. (2014) find that moving from 5-digit codes to 3-digit codes does not have a large effect on the predictive implications for risk scores.

are some markets in which there is missing cost-sharing information, or insurance firms are missing altogether.

I restrict the analysis to markets in which I observe the entire choice set and can be reasonably confident that the online broker presents nearly the complete choice set of health insurers. Using state-level market shares from the Medical Loss Ratio reporting data, I throw out any markets in which I do not observe any purchases from insurance firms that have more than 5% market share in the state. In this way, I hope to ensure that my sample of choices is not segmented to only a portion of the market.

The choice sets in each market are large. The typical market has about 150 plans to choose from, and these plans do not necessarily overlap with other markets. Since I observe only a sample of choices, there are many plans that I do not observe being chosen. This does not necessarily imply that these plans have a zero market share, but simply that the number of choices is large relative the observed number of choices. The median number of choices per market is 300.

To simplify this problem, I aggregate to the level of firm-metal offerings in a particular market. For example, all Bronze plans offered by a single insurance firm are considered a single product. While firms typically offer more than one plan in a given metal level, the median number of plan offerings per metal level is 3, and the 75^{th} percentile is 5. Thus, this aggregation is not terribly restrictive. Wherever there is more than one plan per category, I aggregate by using the median premium within the category. In my current analysis, the only other product attributes I use in estimation are common to all plans in each category.

2.6 Demographic and Uninsured Information

I use the 2015 American Community Survey (ACS) to construct a rating-area specific "market share" for the outside good, uninsurance, and generate the underlying distribution of demographics for simulating equilibrium. I consider the population of individuals who might consider purchasing individual market health insurance to be any legal US resident that is not eligible for Medicaid, Medicare, and is not enrolled in health insurance through their employer. Technically, any individual can switch from these insurance categories to the individual market at any time, however the insurance plans in the individual market are considerably more expensive and typically require larger amounts of cost sharing, so that kind of switching is likely to be small. I consider an individual that is not enrolled in employer sponsored insurance but has an offer that they chose not to accept to be in the individual market. I treat these consumers as identical to the rest of the population, though by law they are not allowed to receive health insurance subsidies. This population is small (Planalp et al. (2015)).

Medicaid eligibility is determined by state-level eligibility categories which are determined by income, age, and family status. I also assume that any one that is determined as enrolled in Medicaid in the ACS is eligible.[9]

3 Model

3.1 Environment

There are a set of M markets, J insurance contracts, and F firms, indexed by m, j, and f. I will write J^m for the subset of products that are sold in market m; J^f is the subset of products owned by firm f; and J^{mf} are the products in area m that are owned by firm f. Insurance contracts have some characteristics which are local, e.g. network coverage, so I will model products as market specific: $j \in J^m \implies j \notin J^{m'}$ for $m \neq m'$. An insurance product is a fixed tuple of observed and unobserved characteristics, (X_j, ξ_j), and has a base premium p_j.

Consumers

A household, i, located in market m, has a set of characteristics, τ, and preferences θ. The household pays a premium for product j that depends on its characteristics and the base premium, $P(\tau, p_j)$. I will write $P_j(\tau)$ as a shorthand for the household specific premium for product j. There are a continuum of households in each market distributed by $F_m(\tau, \theta)$. Households also have idiosyncratic preferences over products $\{\epsilon_{ij}\}_{j \in J^m}$, which I assume are independently and identically distributed by type I extreme value. The indirect utility that household i receives from purchasing a product j is given by

$$\nu_{ij} = u(P_j(\tau), X_j, \xi_j; \theta_i) + \epsilon_{ij}$$

[9] I follow a methodology outlined by the Government Accountability Office (GAO (2012)) to construct households purchasing units and correct for under-reporting of Medicaid in survey data.

where u is a utility function that depends on θ, and on τ via the premium. For a shorthand, I will write $u_j^{\tau\theta} \equiv u(P_j(\tau), X_j, \xi_j; \theta)$. The share of households with characteristics τ and preferences θ that choose to purchase product j is

$$S_j^{\tau\theta}(\boldsymbol{P}_m) = \frac{e^{u_j^{\tau\theta}}}{1 + \sum_{k \in J^m} e^{u_k^{\tau\theta}}}$$

where $\boldsymbol{P}_m = \{p_j\}_{j \in J^M}$.

Firms

A firm, f, may compete in several markets $M^f \subset M$, and has a profit function defined as

$$\Pi^f = \sum_{m \in M^f} L^m \sum_{j \in J^{mf}} \int_{\tau,\theta} S_j^{\tau\theta}(\boldsymbol{P}_m)\Big(P_j(\tau) - C_m(X_j, \tau, \theta) + T_j(\boldsymbol{P}_m)\Big) dF_m(\tau, \theta),$$

where L^m is the size of market m, and $C_m(X, \tau, \theta)$ represents the market-specific expected marginal cost of enrolling a household with characteristics τ and preferences θ in a product with characteristics X. I will write $c_j^{\tau\theta}$ as a short hand for $C_m(X_j, \tau, \theta)$.

The price functions $P_j(\tau)$ can be written as

$$P_j(\tau) = p_j A(\tau)$$

where $A(\tau)$ is a price modifier that is set by regulation and not firms. The only parameter of the price function that firms can decide is the base premium, p_j.

The transfers, $T_j(\boldsymbol{P}_m)$ represent risk adjustment transfers to firms that depend on the equilibrium outcome of the market via equilibrium prices. [10] The risk adjustment transfers take the form

$$T_j(\boldsymbol{P}_m) = \underbrace{\frac{E[\sum_k S_k^{\tau\theta} c_k^{\tau\theta}]}{E[\sum_k S_k^{\tau\theta}]}}_{\text{Pooled Cost}} - \underbrace{\frac{E[S_j^{\tau\theta} c_j^{\tau\theta}]}{E[S_j^{\tau\theta}]}}_{\text{Average Cost}}$$

[10] Other markets are governed by risk adjustment transfers that more explicitly depend on personal attributes rather than the distribution of risk in the market. However, these transfers could still be written in this "average risk transfer" form.

Equilibrium

Each firm competes in Bertrand-Nash price competition to set base prices $\{p_j\}_{j\in J^f}$ to maximize profit. A competitive equilibrium in a market is a vector of base premiums, $\{p_j\}_{j\in J^m}$ such that the prices of each firm $\{p_j\}_{j\in J^{mf}}$ maximize the profit of firm f, given the prices of every other firm, $\{p_j\}_{j\in J^m/J^{mf}}$. Without any risk adjustment transfers, the equilibrium price for a single product firm is the standard sum of marginal cost and a markup.

$$p_j^* = -\frac{S_j}{S_j'} + MC_j(\boldsymbol{P}^*)$$

where,

$$\begin{aligned} S_j &= \int_{\tau,\theta} S_j^{\tau\theta}(\boldsymbol{P}^*)dF(\tau,\theta) \\ S_j' &= \int_{\tau,\theta} \frac{\partial S_j^{\tau\theta}(\boldsymbol{P}^*)}{\partial p_j} dF(\tau,\theta) \\ MC_j(\boldsymbol{P}^*) &= \frac{1}{S_j}\int_{\tau,\theta} \frac{\partial S_j^{\tau\theta}(\boldsymbol{P}^*)}{\partial p_j} c_j^{\tau,\theta} dF(\tau,\theta) \end{aligned}$$

In the presence of risk adjustment transfers, the equilibrium price can be written as

$$P_j = -\frac{S_j}{S_j'} + \frac{S_j}{\sum_{k\in J^m} S_k}\Psi_j MC_j^{mkt} + \big(1 - \frac{S_j}{\sum_{k\in J^m} S_k}\Psi_j\big)PC$$

where,

$$\begin{aligned} MC_j^{mkt} &= \frac{1}{\sum_k \frac{\partial S_k}{\partial p_j}} \int_{\tau,\theta} \sum_k \frac{\partial S_k^{\tau\theta}(\boldsymbol{P}^*)}{\partial p_j} c_k^{\tau,\theta} dF(\tau,\theta) \\ PC &= \frac{1}{\sum_k S_k} \int_{\tau,\theta} \sum_k S_j^{\tau\theta}(\boldsymbol{P}^*) c_k^{\tau,\theta} dF(\tau,\theta) \\ \Psi_j &= \frac{\sum_k \frac{\partial S_k}{\partial p_j}}{S_j'} \end{aligned}$$

Intuitively, the risk adjustment transfer methodology leads firms to weight the average pooled cost of the market with the market-wide marginal cost with respect to the price of the firm's product, and the weight is related to the market share of the product. The term Ψ_j is the "margin share" of product j: the share of ratio of extensive demand derivative of the entire market with respect to the price of product j to the demand derivative of product j.

To see how optimal prices under risk adjustment vary with market structure, note that as a firm becomes very small ($S_j \to 0$), the firm simply charges a markup over the pooled average cost. If a firm is a monopolist, ($S_j = 1$), then the margin share is also equal to 1, and we get the standard monopolist markup over marginal cost—i.e., risk adjustment has no effect.

3.2 Market Structure and Adverse Selection

In this environment, total utilitarian welfare in a particular market is given by

$$SW_m(\boldsymbol{P}_m) = \int_{\tau,\theta} \underbrace{CS^{\tau\theta}(\boldsymbol{P}_m)}_{\text{Consumer Surplus}} + \underbrace{\sum_{k \in J^m} S_k^{\tau\theta}(P_k(\tau) - C_m(X_k, \tau, \theta))}_{\text{Profit}} \, dF(\tau, \theta)$$

where

$$CS^{\tau\theta}(\boldsymbol{P}_m) = E_{\epsilon_i}\left[\max_{k \in J^m} u(P_j(\tau), X_j, \xi_j; \theta_i) + \epsilon_{ij}\right]$$

In general, the social welfare maximizing price of each product is equal to its marginal cost. When marginal costs are constant in the quantity produced but heterogeneous across consumers served, the social welfare maximizing price is equal to the average cost of the marginal consumer. Unless necessary, I will drop the market subscript, m. Suppose that premiums are constant across individuals, $P_j(\tau) = p_j$. In this case, the vector of social welfare maximizing base premiums, $\boldsymbol{P}^W$, solves

$$\begin{aligned} 0 &= \int_{\tau,\theta} \sum_{k \in J^m} \frac{\partial S_k^{\tau\theta}}{\partial p_j}(p_k^W - C_m(X_k, \tau, \theta)) dF(\tau, \theta) \quad \text{for all } j \\ \boldsymbol{P}^W &= E\left[\frac{\partial \boldsymbol{S^{\tau\theta}}}{\partial \boldsymbol{P}}\right]^{-1} E\left[\left[\frac{\partial \boldsymbol{S^{\tau\theta}}}{\partial \boldsymbol{P}}\right] \boldsymbol{C}^{\tau\theta}\right] \end{aligned}$$

where $\left[\frac{\partial \boldsymbol{S^{\tau\theta}}}{\partial \boldsymbol{P}}\right]$ is the matrix of demand derivatives and $\boldsymbol{C}^{\tau\theta}$ is a vector of the expected costs for each product in the market given τ and θ.[11] In words, the social welfare maximizing

[11] This result does not depend on the specifics of my demand specification. It comes from the results that $\partial CS^{\tau\theta}(\boldsymbol{P}_m)/\partial p_j = S_j^{\tau\theta}(\boldsymbol{P}_m)$, which holds under much less restrictive assumptions on demand (Small and Rosen (1981)).

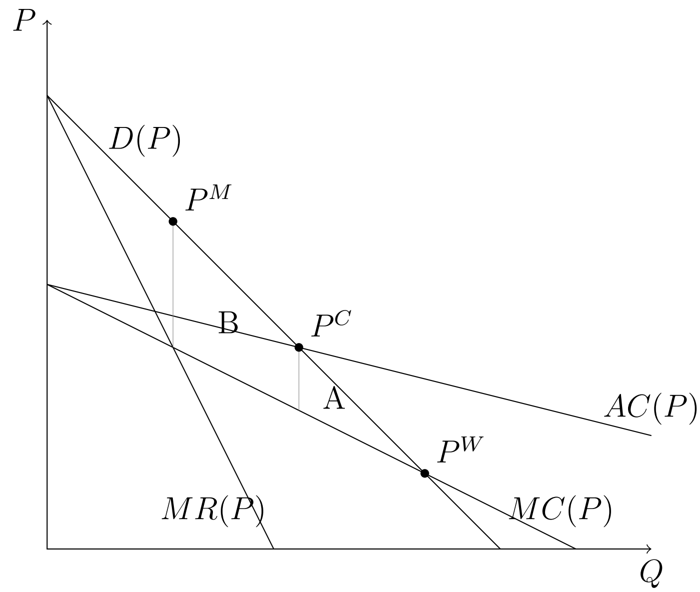

Figure 1: Single Product Adverse Selection

premium for a given product is equal to the average cost of the marginal consumers across all products in the market. For short-hand, I will write,

$$P_j^W = MC_j^{mkt}(\boldsymbol{P}^W)$$

where $MC_j^{mkt}(\boldsymbol{P}^W)$ is the average cost of the market-wide marginal consumer with respect to the price of product j.

Figure 1 displays how the welfare maximizing price relates to equilibrium prices in the case of a single product. The presence of adverse selection implies a downward sloping marginal cost curve ($MC(P)$), and average costs that are greater than marginal cost ($AC(P)$). The social welfare maximizing price is P^W, where marginal cost intersects the demand curve.

The first source of inefficiency in markets with adverse selection come from firms that must earn *non-negative* profits. In Figure 1, it is clear that the welfare maximizing price is below average cost, and therefore cannot be offered by firms that cannot operate at a loss. The zero-profit equilibrium price is P^C, which leads to the a total welfare loss equal to the area of the triangle labeled by A.

Under a monopoly, this source of inefficiency is exacerbated as the monopolist extracts more profits from the market. The monopolist's price, P^M, is determined by where marginal revenue $MR(P)$ intersects marginal cost, and results in additional welfare loss demarcated by B. This intuition for the increasing welfare loss with rising market power resulting from non-negative profits holds more broadly for market structures between perfect competition and monopoly, and in a multi-product setting.

Then second type of inefficiency is *inefficient sorting*. The pricing incentives of profit-seeking firms can lead to inefficient sorting of individuals across the plans available. In general, this is the result of a firms not internalizing the externality of shifting consumers to different plans that may change market-wide costs while maintaining average utility.

For illustration consider the following constrained planning problem, in which a social planner seeks to maximize total consumer surplus, subject to the constraint that market-wide profits exceed some level, $\overline{\Pi}$.

$$\max_{\{p_j\}_{j\in J^m}} \int_{\tau,\theta} CS^{\tau\theta}(\boldsymbol{P}_m) dF(\tau,\theta)$$
$$\text{such that } \int_{\tau,\theta} \sum_{k\in J^m} S_k^{\tau\theta}(P_k(\tau) - C_m(X_k,\tau,\theta)) dF(\tau,\theta) \geq \overline{\Pi}$$

The solution to this problem is a vector of prices, $\boldsymbol{P}^s$, and a Pareto weight (or shadow price) on profits, λ, that solves

$$\underbrace{P_j^s + \frac{\lambda - 1}{\lambda} \boldsymbol{\Psi}^{-1} \frac{S_j}{S_j'}}_{\text{Marginal Social Benefit}} = MC_J^{mkt}$$

If the Pareto weight on profits is equal to 1, this constrained problem is identical to the previous social planner's problem. However, if $\overline{\Pi} \geq 0$, then $\lambda > 1$. In this case, the welfare maximizing price of each product is determined by where the "social marginal benefit" is equal to average cost. This is equivalent to charging an adjusted markup over marginal cost. [12]

In Figure 2, I illustrate efficient sorting for a simple two-product example, a generous product H and less generous product L. This figure an analogous plot to Figure 1, except that the price difference is plotted on the y-axis, $\Delta P = P_H - P_L$, and the quantity purchased of H only

[12] Contrary to typical risk adjustment policy, the welfare-maximizing sorting requires an intervention that targets market-wide marginal costs—rather than average costs—and markups. These prices are not easy to implement for several reasons which include that the value of λ is unknown to regulators. In future work, I will explore a set of transfers and taxes that can implement the optimal prices.

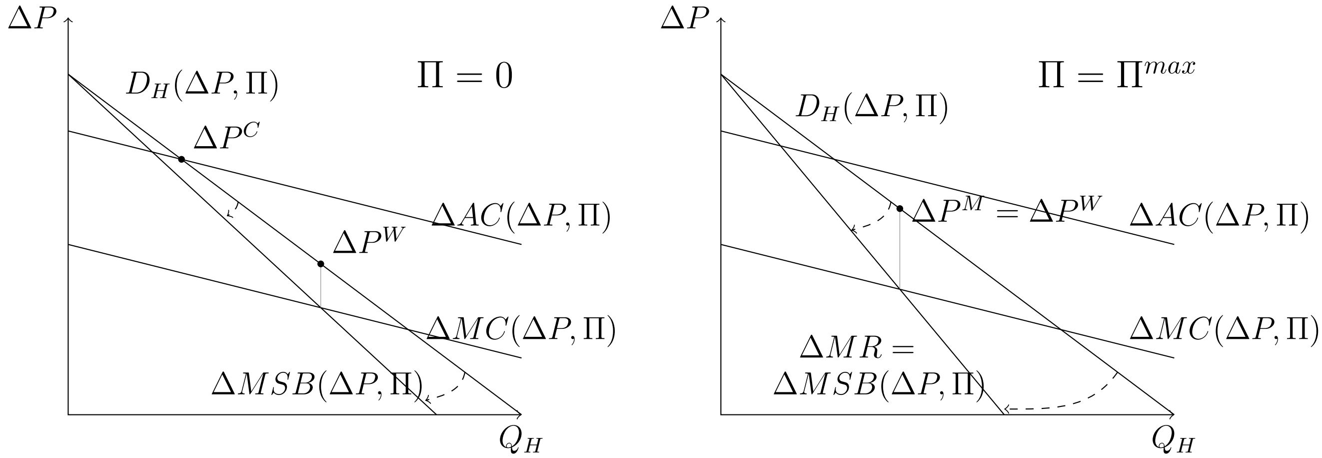

(a) Optimal Sorting with Zero Profits

(b) Optimal Sorting with Monopolist Profits

Figure 2: Two Product Sorting

is on the x-axis. The difference in average costs ΔAC and the difference in marginal costs ΔMC are decreasing as larger enrollment in the H plan narrows the selection differences. In this toy example, the price difference and the profits earned, Π, pin down the equilibrium prices and quantities. In the left panel (2a), I consider the case of perfect competition (free-entry) and zero profits. The competitive ΔP^C is equal to the difference in average costs. However the welfare maximizing ΔP^W (conditional on firms breaking even) is determined by where the social marginal benefit curve, ΔMSB intersects marginal cost. As in the first example, the competitive market sets the price difference to be too large, even conditional on breaking even.

In the right panel (2b), I consider the case of a monopolist earning the maximum level of profit. In this case, ($\lambda \to \infty$), the social marginal benefit curve coincides with the monopolists marginal revenue curve, and the monopolist charges the socially optimal price difference, conditional on earning the maximum level of profits.

This analysis does not suggest that monopolists are benevolent, but rather that the monopolist internalizes all of the competitive externalities. And while a monopolist creates large welfare losses through high markups, a social planner can do no better without making either the firm or the consumers worse off. This is not the case in less concentrated markets.

4 Demand

4.1 Empirical Specification

In my empirical specification, households have characteristics $\tau_i = (a_i, y_i, Z_i, r_i^{HCC})$, where a is an average age-rating of all household members, y is household income, Z is a vector of demographic variables, and an unobserved risk score, r^{HCC}. Households also have preferences $\theta_i = (\gamma_i, \alpha_i, \beta_i)$.

I treat rating areas as geographic markets, and I aggregate all products to the firm-metal-market level. For example, a product is a Bronze plan offered by Aetna in the Georgia's 1^{st} rating area. A product j in market m is a tuple of observed characteristics and an unobserved quality, (X_{jm}, ξ_{jm}), and a base premium p_{jm}. I include 2 observed characteristics: the actuarial value of the plan and whether or not the plan is sold by a "healthy" firm. [13]

I parameterize $u(P_j(\tau_i), X_j, \xi_j; \theta_i)$ as

$$u_{ijm} = \gamma_i + \alpha_i(a_i p_{jm} - s(y_i)) + \beta_i X_{jm} + \xi_{jm}$$

where $s(y)$ is a function that maps income to subsidies. I allow the preference for all insurance products, γ, and the utility-value of money, α_i to be demographic specific, and I allow the preference over observed characteristics, β_i, to depend on a households risk score, r^{HCC}.

$$\begin{aligned}
\gamma_i &= \alpha_0 + \alpha_z' Z_i \\
\alpha_i &= \alpha_0 + \alpha_z' Z_i \\
\beta_i^k &= \beta_0 + \beta_r^k r_i^{HCC}
\end{aligned}$$

Since I do not have data that links risk scores to health insurance choices, I treat risk as an unobserved household characteristic. I assume that risk scores are distributed according to a distribution that can depend on household demographics, Z_i.[14]

$$r_i^{HCC} \sim G(Z_i)$$

[13] I designate a number of firms as "Healthy" that, in 2015, consistently enrolled healthier than average enrollees across all the markets they participate in. These firms include United Healthcare, Cigna, Assurant, and companies associated with Blue Cross Blue Shield.

[14] I use the demographics of the head-of-household as the representative demographics for the household.

4.2 Risk Score Distribution

I use the publicly available Health and Human Services Hierarchical Condition Categories risk adjustment model (HHS-HCC), used in the individual market for the purpose of administering risk adjustment transfers, to link health insurance demand and health status. The HHS-HCC risk adjustment model is designed to predict expected log plan spending on an individual, based on demographics and health condition diagnoses. It is the result of a linear regression of log plan spending on a set of age-sex categories and a set of hierarchical condition categories based on diagnoses codes.

$$\log(\text{Plan Spending}_{it}) = \gamma_0 + \sum_g \gamma_{tg}^{age}\gamma_{tg}^{sex} Age_{ig} Male_{ig} + \sum_{g'} \gamma_{tg'}^{HCC} HCC_{ig'} + \eta_{it}$$

The prediction regressions are performed separately for different types of plans t, where t represents the metal category of the plan. The resulting risk score for an individual is a normalized predicted log spending value. Since all independent variables have a value of either 1 or 0, this is the sum of all coefficients that apply to a particular individual.

$$r_{it} = \underbrace{\sum_g \gamma_{tg}^{age,sex} Age_g Male_g}_{r_{it}^{dem}} + \underbrace{\sum_{g'} \gamma_{tg'}^{HCC} HCC_{g'}}_{r_{it}^{HCC}}$$

I use the HCC component of an individuals Silver plan risk score, $r_{i,Silver}^{HCC}$ as a measure of health status. Unless specifically noted, I will write r_i^{HCC} to refer to the Silver plan HCC risk score component.

Parametric Distribution

I estimate $\hat{G}$ from the 2015 Medical Conditions File (MCF) of the Medical Expenditure Panel Survey. The MCF contains self-reported diagnoses codes and can be linked to demographic information in the Population Characteristics file. The publicly available data only list 3-digit diagnoses codes, rather than the full 5-digit codes. I follow McGuire et al. (2014) and assign the smallest 5-digit code for the purpose of constructing the condition categories and matching the HHS-HCC risk coefficient.[15]

[15] For example, I treat a 3-digit code of '301' as '301.00'. McGuire et al. (2014) find that moving from 5-digit codes to 3-digit codes does not have a large effect on the predictive implications for risk scores.

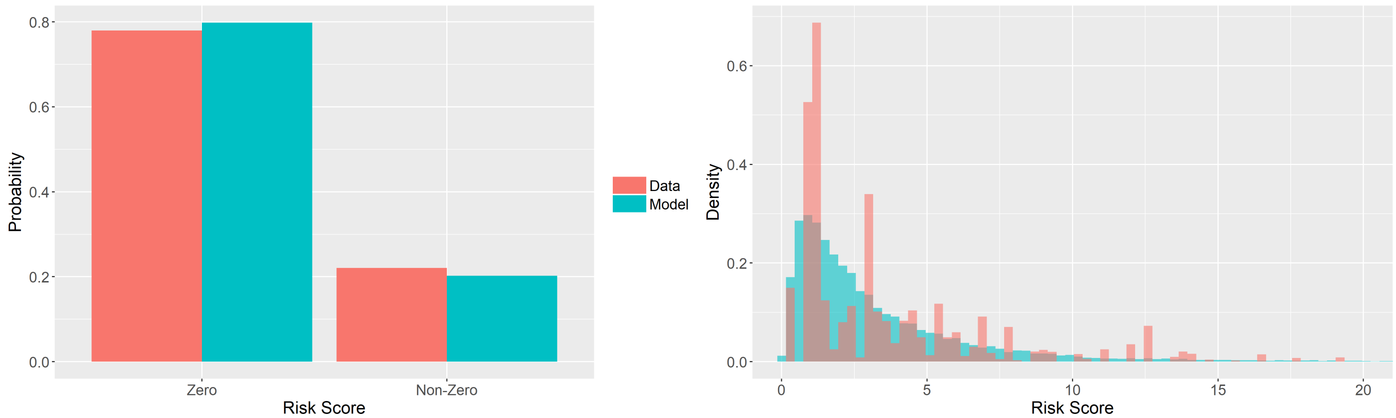

(a) Incidence of Positive Risk Scores (b) Distribution of Positive Risk Scores

Figure 3: Risk Score Distribution Model Fit

In the data, a majority of individuals have no relevant diagnoses, i.e. $r_i^{HCC} = 0$.[16] In order to match this feature, I model the distribution as a mixture of a binomial distribution and log-normal distribution. With some probability $\delta(Z_i)$, the household has a non-zero risk score, and $log(r_i^{HCC}) \sim N(\mu(Z_i), \sigma)$. With probability $1 - \delta(Z_i)$, $r_i^{HCC} = 0$. I allow the probability of having any relevant diagnoses and the mean of the log-normal distribution to vary by two age categories, above and below 45 years old, and two income categories, above and below 400 percent of the federal poverty level.

Table 2 displays the moments of the distributions in the data. To simulate equilibrium, I will use the estimated risk distribution and the demographic distribution from the American Community Survey to predict the population risk distribution, and Figure 3 shows how that prediction compares to the MEPS data.

Table 2: Parametric Distribution of Risk Scores

		$\delta(Z_i)$	$\mu(Z_i)$	σ
Older than 45	Less than 400% FPL	0.318	0.903	0.950
	Greater than 400% FPL	0.258	1.16	0.950
Younger than 45	Less than 400% FPL	0.139	0.568	0.950
	Greater than 400% FPL	0.140	0.556	0.950

Notes: All moments are computed from MEPS data.

[16] I exclude uninsured individuals from the analysis to avoid low diagnoses rates because of infrequent contact with medical providers.

4.3 Estimation

The primary concern in correctly identifying the parameter governing price-elasticity, α_i, is that price may be correlated with the unobserved quality ξ_{jm}. In this environment, the premium regulations provide a source of variation in price which is exogenous to variation in unobserved quality. The age-adjustment on premium, a_i, increases monotonically with age, and strictly increases with every age after 25. Income based subsidies are available to households that earn below 400 percent of the federal poverty level. These subsidies decline continuously within the subsidy eligible range.

I use fixed effects to control for ξ_{jm}, and I allow for progressively greater flexibility in the fixed effects. While this is not a formal test of the exogeneity assumption, it provides a sense of whether the price coefficient estimates are sensitive to the degree that I control for non-price, unobserved quality.

Estimation Moments

I use General Method of Moments (GMM) to combine the likelihood of observed consumer level choices with aggregate moments on risk scores. The first set of moments are the derivatives of the log-likelihood function with respect to the parameters.

$$M_1(\theta) = \frac{\partial \mathcal{L}}{\partial \theta}$$

I include 6 additional moments on the average risk scores of enrollees in Bronze, Silver, Gold, and Platinum plan categories, the average risk score of all enrollees, and the average risk score in "healthy" firms. (The data moments are described in detail in section 2.4). These moments are calculated using total, metal-level dependent risk scores.

$$\bar{R}(\theta) = \begin{bmatrix} E[R_{i,b}|\text{consumer } i \text{ purchases a Bronze plan}] \\ E[R_{i,s}|\text{consumer } i \text{ purchases a Silver plan}] \\ E[R_{i,g}|\text{consumer } i \text{ purchases a Gold plan}] \\ E[R_{i,p}|\text{consumer } i \text{ purchases a Platinum plan}] \\ E[R_{i,t}|\text{consumer } i \text{ purchases any plan}] \\ E[R_{i,t}|\text{consumer } i \text{ purchases a plan offered by a ``healthy'' firm}] \end{bmatrix}$$

$$M_2(\theta) = \bar{R}(\theta) - R^{\text{data}}$$

The estimate $\hat{\theta}$ minimizes the objective function, for a given weighting matrix W,

$$\hat{\theta} = \text{argmin} M(\theta)' W M(\theta)$$
$$\text{where} \quad M(\theta) = \begin{bmatrix} M_1(\theta) \\ M_2(\theta) \end{bmatrix}$$

I perform the two-step GMM estimation, first using the identity matrix for W to recover an estimate of the asymptotic variance of the moments, $\hat{V}$, and then re-estimating with $W = \hat{V}^{-1}$.

Since not all of my moments apply to all observations in the data, I follow Petrin (2001) to compute moments that do apply to each observation—e.g. the product of an indicator function of whether a product is a bronze plan and the expected bronze plan risk score—and a function to translate these universal moments into the relevant estimation moments. If the universal moments are $\tilde{M}$, then the asymptotic variance can be estimated with

$$\hat{V} = E[\nabla H(\tilde{M})\big(\tilde{M}\tilde{M}'\big)\nabla H(\tilde{M})']$$
$$\text{where} \quad H(\tilde{M}(\theta)) = M(\theta)$$

4.4 Results

In Table 3, I present the results from the demand estimation. In specifications 1 through 4, I estimate demand using maximum likelihood, without taking the supplemental risk score moments into account. The specifications use increasingly flexible fixed effects to control for unobserved quality. The price parameters are all statistically significant and relatively stable across all specifications. While, this is not proof that my exogeneity assumption holds, it supports that the identifying variation in price is robust to different methods of controlling for unobserved product quality.

In specifications 5 and 6, I include the risk score moments. Most noticeably, including risk score moments leads to a positive relationship between willingness-to-pay for actuarial value and health risk, as would be expected if there is adverse selection. The price coefficient for the youngest age category, 18 - 30 is lower than in the maximum likelihood estimations, but the other age groups have similar estimates.

The final specification, 6, was not estimated with the efficient weighting matrix. As a result,

Table 3: Demand Estimation Results

	Maximum Likelihood				GMM	
	(1)	(2)	(3)	(4)	(5)	(6)
Annual Premium ($000)	-1.48	-1.36	-1.39	-1.45	-2.80	-1.67
Age 31 - 40	0.209	0.211	0.192	0.204	1.41	-0.387
Age 41 - 50	0.400	0.370	0.351	0.425	1.85	0.497
Age 51 - 64	0.815	0.736	0.745	0.832	2.26	1.10
Family	-0.055	-0.041	-0.049	-0.039	-0.326	-0.055
Subsidized	0.073	0.131	0.117	0.182	-0.032	-0.016
AV	4.27	6.51	3.49	6.59	7.31	12.36
Healthy Firm	(0.020)	(-2.48)	(1.60)	(-5.18)	1.57	(-8.07)
β_r^{AV}	-0.114	-0.084	0.333	(0.023)	0.447	0.521
$\beta_r^{\text{Healthy Firm}}$	0.113	0.138	-0.327	(-0.031)	(-0.004)	-0.023
Fixed Effects						
Age, Family, Income	Y	Y	Y	Y	Y	Y
Firm	**Y**				**Y**	
Firm-Market-Category		**Y**				**Y**
Firm-Market-Age			**Y**			
Firm-Market-Age-Category				**Y**		

Notes: All parameters are statistically significant at 0.1 percent level, unless they are presented in parenthesis. The top row of price coefficients corresponds to the estimate for households that do not fall into any of the listed subgroups (single, high income, 18 to 30 year olds). The price coefficients for other households are obtained by adding the relevant demographic adjustments to the top line. Fixed effects for product category include two categories: above and below 70% actuarial value.

I will use specification 5 as my preferred demand estimation for the remainder of the paper. More progress on GMM estimations with more flexible fixed effects is forthcoming.

I can calculate the implied willingness-to-pay for additional insurance (WTP), i.e. the additional premium a consumer is willing to pay for an increase in the actuarial value of the insurance plan. I find that the median consumer is willing to pay \$48 per month to increase the actuarial value of their insurance by 10%, which is roughly equivalent to switching from a Bronze plan to a Silver plan. The 10th percentile of WTP is \$22 per month, and the 90th percentile is \$112 per month. There is a long right tail, with the 99th percentile of consumers willing to pay \$153 per month.

5 Marginal Cost

5.1 Empirical Model

Instead of imposing that firms are behaving optimally, I use aggregate data on firm's costs and moments on how health care expenditures vary by age and risk to estimate costs through Method of Simulated Moments (MSM). I specify the expected cost function, $C_m(X_j, \tau_i, \theta_i)$, as

$$\log(c_{ijm}) = \phi_0^m + \phi_1 AV_j + \phi_2 Age_i + \phi_3 r_i^{HCC} + \omega_{ijm}$$

This specification assumes that the i.i.d. errors in the cost function, ω_{ijm} are orthogonal to the preference draws in the demand estimation.

$$E[\epsilon_{ijm}\omega_{ijm}] = 0$$

I must assume that the only mechanisms through which cost and preferences can be correlated are through age and risk scores.[17] Unfortunately, there is substantial evidence that there remains significant and predictable variation in health spending that is not captured by age and HCC categories (Brown et al. (2014), Layton (2017)). If the assumption is violated in

[17] An alternative specification could treat expected total medical spending as a household characteristic. Then, I could allow preferences to vary with this characteristic instead of risk scores. This has the advantage of circumventing the exogeneity assumption through the technicality of putting individual level spending directly in the utility function. However, the principle concern that residual costs unobservable to the econometrician are correlated with demand errors would remain.

a manner consistent with adverse selection, my estimates for expected cost conditional on observables will be biased upward.[18] The parameter most affected by this bias is likely to be the parameter on actuarial value, which can be interpreted as an outcome variable possibly correlated with the error in the cost regression through correlation with preference draws. I will use moments on age and risk to discipline regression and show that my estimates are consistent with cost differences implied by differences in coverage generosity.

5.2 Estimation

I use moments on average cost by firm and state from the 2015 Medical Loss Ratio reporting data, and average cost by firm, state and metal-level from 2016 Rate Filing data. (For more detailed description of the data, see section 2.3). I supplement these average cost moments with 10 moments on the relative expenditures of 5-year age groups from 20 to 65 year olds, and the relative expenditure of individuals with a non-zero HHS-HCC risk score and those with a risk score of 0.

I use the estimated demand parameters to simulate product enrollment and use MSM to match simulated moments to those from the data. I use a two-step estimation and universal moments to estimate an efficient weighting matrix, as described in section 4.3.

5.3 Results

Table 4 displays the results of the cost estimation. Specifications 1 through 3 are an OLS regression of average product costs on the simulated average age and risk of enrollees. While the estimations are not directly comparable, the comparison between the OLS and MSM estimates demonstrate how adding additional moments from MEPS helps to discipline estimates of individual level costs, and in particular the parameter on actuarial value.

In Table 5, I provide some validation for the actuarial value parameter estimate which may

[18] For illustration, suppose I estimate $\hat{\phi}$ to solve for a single product and single observable type,

$$\frac{E[S_{ij}c_{ij}]}{S_j} - AC^{data} = 0$$
$$E[S_{ij}c_{ij}] = S_j AC^{data}.$$

This is equivalent to

$$S_j E[c_{ij}] - \text{cov}(S_i, c_{ij}) = S_j AC^{data}.$$

Table 4: Cost Estimation Results

	OLS (Product Moments Only)			MSM
	(1)	(2)	(3)	(4)
Age	0.031***	0.025**	0.019	0.304
Risk	-0.265**	0.097	-0.076	0.063
Actuarial Value	6.684***	3.78***	5.17	1.95
State	N	Y	Y	Y
Firm	N	N	Y	N

Note: Standard errors are not yet available for the MSM estimation. For the OLS regressions, statistical significance is reflected as below.

'*' - 5% Confidence Level; '**' - 1% Confidence Level; '***' - 0.1% Confidence Level

Table 5: Relative Cost Comparison

	Model Prediction	ACA Guidelines
Bronze	-	-
Silver	1.22	1.20
Gold	1.47	1.44
Platinum	1.80	1.725

Note: Calculations compare the expected cost ratio of each metal level to Bronze in the predicted model, and according to the according to the ratio of actuarial values and IDF, as specified by ACA regulations.

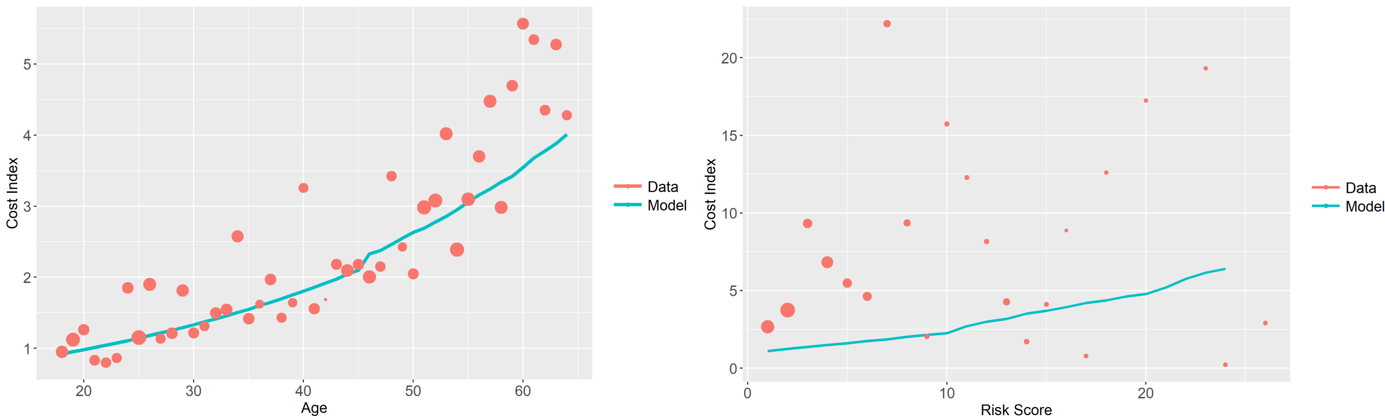

(a) Age-Cost Curve Fit

(b) Risk-Cost Curve Fit

Figure 4: Cost Model Fit to Age and Risk

be susceptible to endogeneity concerns. I compare the implied relative average costs between each metal level and a bronze plan, holding fixed the population of enrollment to the difference in cost implied by Affordable Care Act guidelines.[19] The guidelines suggests two plans that differ in actuarial value, but not in enrollment, should differ in cost only through the actuarial value and an "Induced Demand Factor" (IDF) that accounts for increased demand for medical services. The IDF ranges from 1.0 for a Bronze plan to 1.15 for a Platinum plan. As seen in Table 5, the relative costs predicted by the model are close to those predicted by the actuarial value and the IDF. I am able to closely match the age-cost curve in the data, but my estimation understates the relationship between risk and cost (Figure 4).

The results of the cost estimation also reflect a significant degree of adverse selection, and the relationship is strongest in the right tail of the willingness-to-pay and cost distributions. Consumers in the top decile of willingness-to-pay for an additional 10 percentage points in actuarial value are in the top 34% of cost to provide that additional coverage. Those in the bottom decile of willingness-to-pay are in the bottom 51% of cost to cover.

5.4 Optimal Strategies

In estimating the parameters of demand and marginal cost, I have made no assumptions that firms are setting prices to optimally maximize profit. However, I can check whether or not

By assuming that the covariance term is 0, I will overestimate the expected cost conditional on observables. I mitigate this issue by implicitly estimating the degree of covariance that is captured through age and risk.

[19] For the purposes of risk adjustment, the government specifies guidelines for "allowable costs," or cost differences without accounting for selection.

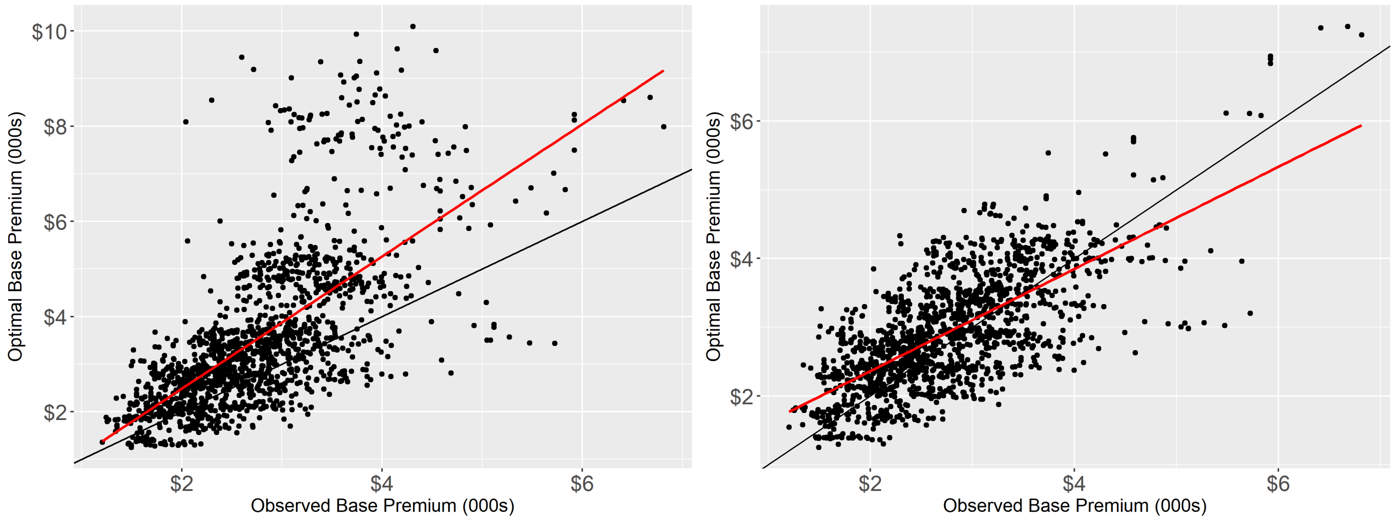

(a) No Risk Adjustment

(b) Perfect Adjustment

Figure 5: Observed Premiums Compared to Optimal Premiums

the parameters I have estimated are consistent with profit-maximizing firms. In Figure 5, I plot the optimal base premium implied by estimated parameters and the firms' first order conditions under two scenarios, compared to the actual base premiums set by the firms.

The left panel, (5a), compares optimal price if there were no risk adjustment policy, and the right panel (5b) compares the optimal price if risk adjustment perfectly predicted firms expected costs and firms had unbiased expectations. The no-risk-adjustment scenario over-predicts the premiums of high-priced contracts, and the risk-adjustment scenario under predicts the premiums of high-priced contracts. This is consistent with a risk adjustment system that imperfectly adjusts for firm costs. Firms also may have had incorrect expectations on the variance of risk adjustment payments, which could have led to higher price spreads.

6 Equilibrium Effects of Market Structure

In this section, I simulate an economy using the estimated joint distribution of preferences and cost, and the demographic distribution of non-group insurance market consumers from the American Community Survey (ACS).

I use the national sample of consumers eligible to purchase non-group insurance from ACS with the corresponding weights.[20] I draw a single risk score for each household from the distribution of the head-of-household, and assign the household the corresponding preferences based on demographics and risk. Each household then has an implied expected cost of being enrolled in plan with a particular actuarial value, generating a joint distribution of preferences, risk, and cost.

In this economy, there are 20 available products, consisting of 10 products with 60% actuarial value and 10 products with 80% actuarial value. I will refer to these products as "low" coverage and "high" coverage, respectively. Each product is owned by a single-product firm with the exception of the Large Firm, which may own more than one of the products. Each firm sets the price of its product(s) in Nash-Bertrand competition. I assume that there is full community-rating, i.e. each product has only a single price which every consumer must pay to enroll in the insurance. Firms cannot turn down any willing purchaser.

I assign every firm, including the Large Firm, an identical firm specific preference, which is calibrated to generate an insured rate of 65% when every product is sold by a single-product firm.[21] I scale the size of the Large Firm by increasing the number of products that it owns from 1 out of 20 to 18 out of 20. If the Large Firm owns more than one product, it is always a symmetric portfolio of plans, consisting of an equal number of high and low coverage products. When I display "market share," I am referring to the share of insurance products sold, not counting the outside good of uninsured.

6.1 Market Structure and Prices

The primary relationship between market structure and the equilibrium outcomes of price and marginal costs are displayed in Figure 6. In the left panel, (6a), I show the effect of the size of the large firm on the equilibrium prices and marginal costs of the large firm's high coverage and low coverage insurance products. When all firms are small, the price of the high coverage product is nearly twice as large as the low product, and the price difference is driven by differences in the marginal cost of each product. When the Large Firm has substantial market share, the marginal costs of offering each product narrow considerably. The marginal cost of the high coverage product falls by more than 10%, primarily due to the large firm internalizing that high cost enrollees will switch to the firms' other products.

[20] Section 2.6 has more details on how I determine eligibility.

[21] The results do not change qualitatively given different numbers of products offered or different calibrated rates of insured.

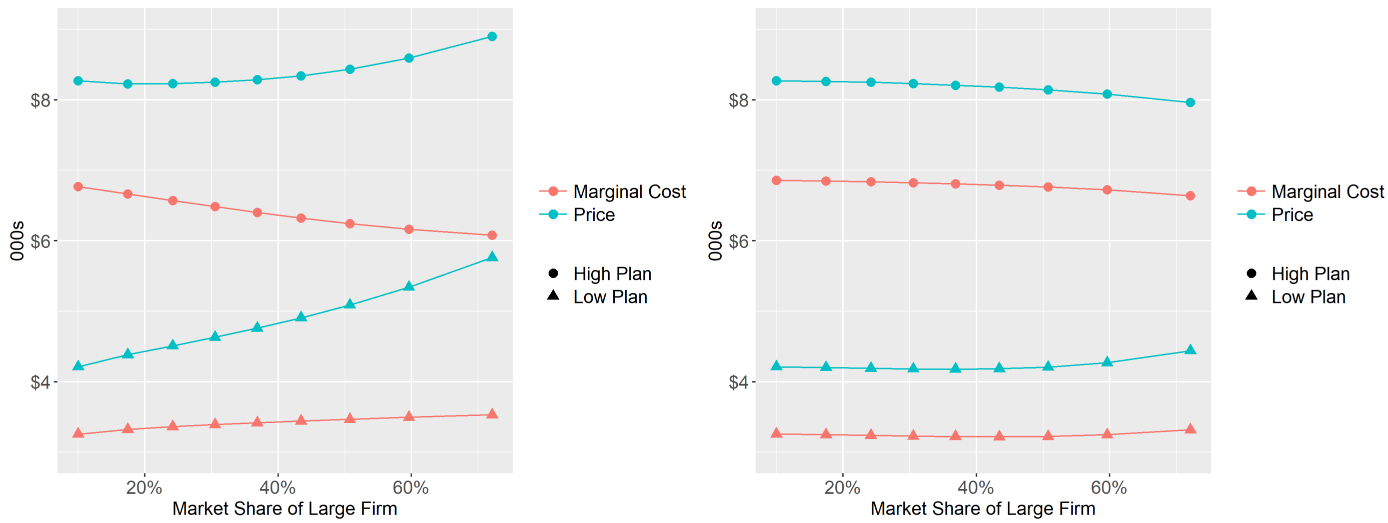

(a) Large Firm

(b) Small Firm

Figure 6: Effect of Market Structure on Equilibrium Prices

However, the prices of both products increase with the size of the Large Firm. The increases in markups that come with market power outweigh the reductions in marginal cost for the high coverage product. The right panel, (6b), shows how the best response prices of the small firms change in response to the increasing market power of the Large Firm. For the most part, the margins of the small firms do not change. However, because prices are strategic complements, the premiums set by the small firms do slightly narrow.

6.2 Market Structure and the Mandate Penalty

In the presence of the individual mandate, consumers are more likely to purchase insurance, and more likely to switch to another insurance plan rather than become uninsured in the event of a price increase. This affects both the markups and marginal costs of strategic firms.

In Figure 7, I display the same comparison between market structure, prices, and marginal costs under a scenario where consumers must pay a $1,500 penalty for deciding not to purchase insurance.[22] Since consumers are more likely to switch within the market, the size of the Large Firm has a larger effect on marginal costs, leading both the marginal costs for the high coverage and low coverage plans to decline. While markups increase, the declines

[22]This is close to the optimal value of the mandate penalty value ($1,461) suggested by Hackmann et al. (2015).

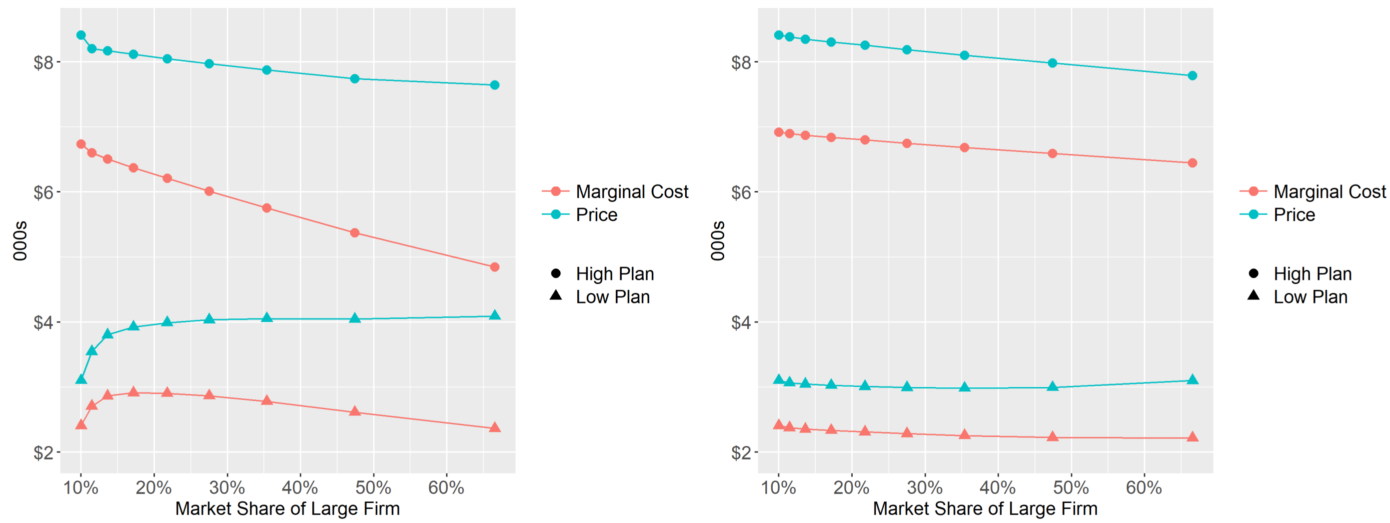

(a) Large Firm

(b) Small Firm

Figure 7: Effect of Market Structure on Equilibrium Prices
with a $1,500 Mandate Penalty

in marginal cost are enough to lead to lower prices for the high coverage product when the Large Firm owns a greater share of products.

From the perspective of policy-makers deciding on an optimal mandate penalty, the ownership structure of the insurance market is important, since the mandate effects both markups and costs. In Figure 8, I show how the effects of a mandate on consumer welfare depend on market structure. In competitive markets, when the Large Firm owns a small number of products, consumer welfare is monotonically increasing for modest values of the individual mandate penalty. The reductions in premiums offset the burden of the tax. However, if the market is very concentrated, small values of the individual mandate do not have a large enough effect on premiums to offset the welfare loss from the tax. The Large Firm doesn't fully pass-through market-wide average cost reductions to the consumer, and markups increase in response to greater demand for insurance. Greater mandate penalties provide benefits for consumers but have declining benefits at very large values.

6.3 Market Structure and Risk Adjustment

Next, I simulate how the effects of risk adjustment interact with market structure. In Figure 9, I illustrate how the effects of risk adjustment fade gradually with increasing market concentration. The policy still has an impact on the small firms in the market, but the large

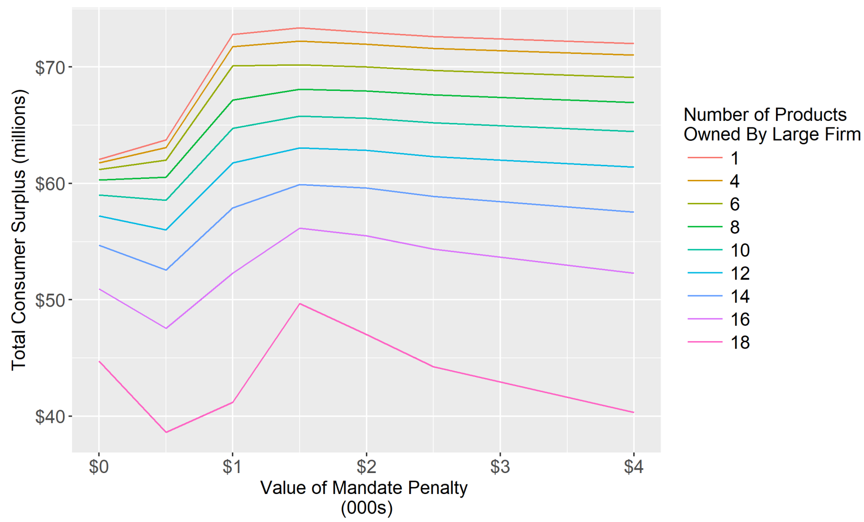

Figure 8: Effect of Mandate and Market Structure on Consumer Surplus

firms internalize the effect of their prices on the pooled cost of the market cost.

As a result, risk adjustment provides little benefit to consumers when the market is concentrated. In fact, the increased price in the low coverage option can outweigh the small benefits of the high coverage good even for the riskiest consumers. In Figure 10, I show the average change in consumer welfare for the low risks in the market and the top 1% of the riskiest consumers as a result of risk adjustment. While risk adjustment delivers benefits for high risk consumers in competitive markets, it can lead to welfare losses for even the highest risks in concentrated markets.

7 Policy Analysis

In this section, I simulate equilibrium in the non-group market in the 109 geographic markets for which I have data. I use the ACS to generate the demographic distribution in each local market, or rating area. Rather than draw a single risk score for each household, I duplicate each demographic observation into 1000 separate households, each with $1/1000^{th}$ of the ACS weight, and draw 1000 separate risk scores. I then assign preferences and plan-specific costs according to the demographic and risk characteristic of each household to generate the join distribution of preferences and cost.

I simulate equilibrium under four separate policy regimes. In the "baseline" regime, I assume that both risk adjustment and the individual mandate are in effect, as well as the other

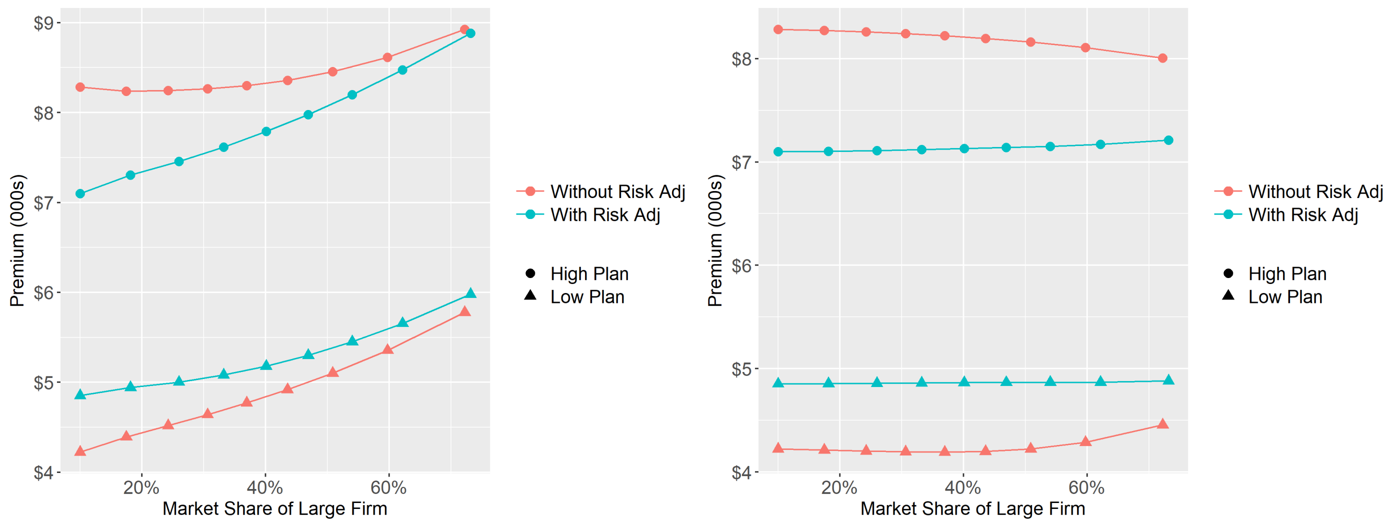

(a) Large Firm

(b) Small Firm

Figure 9: Effect of Risk Adjustment on Equilibrium Prices

rules and regulations of the ACA. This includes price-linked subsidies, age-rating, and the medical-loss ratio requirements.[23]

I then re-solve for equilibrium under three additional policy regimes: no individual mandate, no risk adjustment, and neither risk adjustment nor the individual mandate. All of these scenarios are reasonable future policy regimes: the individual mandate penalty is set to $0 beginning in 2019, and the risk adjustment methodology is being challenged in court. However, this analysis is intended to isolate the effects of changing policies, holding constant all other aspects of the way that firms set prices. Notably, I model a "perfect" risk adjustment system, where the policy makers and firms both have un-biased expectations over the relationship between risk and cost. And I do not account for any other pressures that state regulators may put on insurance prices in the event that these policies are repealed.

The purpose of this analysis is two fold. The first is to demonstrate the heterogeneity of policy effects across different levels of market concentration. The second is to show the complementarity between certain policy combinations and market concentration.

[23]In 2015, the ACA also included two other risk protection programs: reinsurance and risk-corridors. I assume that the reinsurance program only affects the household-specific cost liability to the insurance firms, does not directly enter the firms' problem. I don't model the risk-corridor program, which provides profit risk-sharing across firms. Sacks et al. (2017) find that the risk corridor program leads to lower prices among some firms.

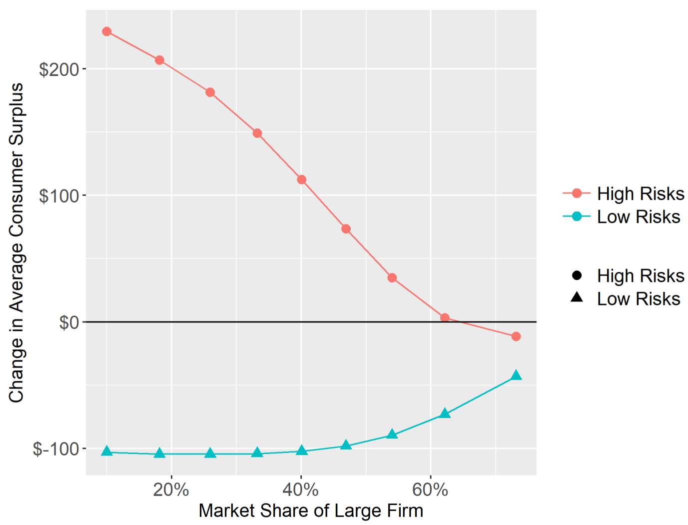

Figure 10: Effect of Mandate and Market Structure on Consumer Surplus

Cross-sectional Analysis

In Table 6, I show the results of each policy regime by market concentration. Market concentration is determined by baseline market shares. In the baseline scenario, 36 markets have a Herfindahl-Hirschman Index (HHI) of less than 3600, 36 markets have an HHI of between 3600 and 5500, and 37 markets have an HHI of at least 5500.

Panel A demonstrates the effects of risk adjustment and the individual mandate in the least competitive markets. In this setting, we see the intended and the complementarity of the individual mandate and risk adjustment. Risk adjustment, both with and without the individual mandate, has an economically significant effect on reducing premiums across all categories. However, in the presence of risk adjustment, the individual mandate has a relatively small effect on prices. In comparing column (1) to the baseline, repealing the individual mandate would primarily only raise premiums for Gold plans. However, if there were no risk adjustment, the individual mandate has a much larger effect. Comparing column (3) to column (2), premiums across all product categories are 21 - 26% greater without the individual mandate in place. This suggests that risk adjustment is a crucially important feature of competitive markets.

In Panels B exhibits many of the same trends of the most competitive markets, though the magnitudes are smaller. In Panel C, the least competitive markets, the effects of both risk adjustment and the individual mandate are more ambiguous. Risk adjustment leads to lower prices among the more generous insurance categories, but higher prices among the cheaper

Table 6: Cross-Section: Effects of Risk Adjustment and the Mandate Penalty

Panel A: Less than 3000 HHI

	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y	N	N
Mandate	Y	N	Y	N
Bronze Premium	3.58	3.57	3.67	4.64
Silver Premium	4.12	4.13	4.52	5.64
Gold Premium	4.89	4.96	6.52	7.94

Panel B: Between 3000 and 6000 HHI

	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y	N	N
Mandate	Y	N	Y	N
Bronze Premium	3.94	3.89	4.00	3.97
Silver Premium	4.47	4.45	4.88	4.92
Gold Premium	5.42	5.46	6.98	7.35

Panel C: Greater than 6000 HHI

	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y	N	N
Mandate	Y	N	Y	N
Bronze Premium	4.15	4.08	4.03	3.92
Silver Premium	4.75	4.65	4.93	4.82
Gold Premium	5.45	5.41	6.27	6.30

Notes: All values are levels in thousands of USD per year, with the exception of the insurance rate. All averages are computed using the baseline enrollment distribution. Platinum plans are excluded since they are not offered in every market. I do not control for market compositional effects across the HHI categories.

Bronze plans. Moreover, the effect of the importance individual mandate is less dependent on the presence of risk adjustment. Since market concentration provides better implicit risk sorting, the presence of a policy adjusting prices between plans is both less effective and less necessary.

In the least competitive markets, the individual mandate also has the opposite of the intended effect, leading to higher premiums rather than lowering them as in the competitive markets. When markets are very concentrated, firms respond to increased demand for insurance with higher markups than outweigh the pressure to lower prices that come from broadening the pool of insured. This result is intuitive in light of other research that show modest reductions in average cost associated with modest changes in the number of insured.

Merger Analysis

To isolate the effect of market concentration, I simulate the consummation of two mergers proposed in 2015: Aetna proposed to acquire Humana, and Anthem Blue Cross proposed to acquire Cigna. The Department of Justice sued successfully to block both mergers in court due to competitive concerns both in the non-group market, as well as in the employer-sponsored and Medicare Advantage markets. In my counter-factual, I observe 21 geographic markets across 4 states that would have been affected by one or both of these mergers. The markets in my data were not the primary markets of concern for the Department of Justice. This mitigates concerns that these markets that market characteristics or firm-strategies would have been substantially impacted, and is consistent with modest price effects of the merger. The median pre-merger HHI in affected markets is 4788, and the median change in HHI is 200.

The results in Table 7 highlight the two first-order predictions of increased concentration in markets with adverse selection: Prices will increase on average, but price increases will be concentrated among products with lower insurance coverage. Panel A shows the price effects on the products of the merging parties—Aetna, Human, Anthem Blue Cross, and Cigna—and Panel B shows the effects on all products in the markets affected by the merger.

The intuition of the results of the cross-sectional results can also be applied here. These results also suggest a substitutability between market concentration and risk adjustment. In markets without risk adjustment, columns (2) and (3), the mergers result in either a modest reduction or slight increase in consumer surplus, depending on whether the individual

Table 7: Merger Analysis: Effect of Market Concentration

Panel A: Products of the Merging Parties

	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y	N	N
Mandate	Y	N	Y	N
Bronze Premium	4.1%	3.0%	2.0%	1.1%
Silver Premium	2.1	1.4	1.0	0.7
Gold Premium	0.8	0.7	-0.6	1.7

Panel B: All Products

	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y	N	N
Mandate	Y	N	Y	N
Bronze Premium	0.5%	0.3%	0.1%	0.8%
Silver Premium	0.3	0.1	0.1	0.3
Gold Premium	0.1	0.0	-0.1	-0.1

Avg Consumer Surplus	-3.4%	-1.6%	-0.5%	0.7%
Insurance Rate	-0.6	-1.2	-0.1	0.5

Notes: All values are percent changes in price, consumer surplus, or the insurance rate that result from the simulated merger. Panel A shows changes for only the merging parties' products, and panel B takes all products into account. The pre-merger averages are computed using pre-merger enrollment and post-merger averages are computed using post-merger enrollment. Thus, the percent changes include changes in consumer choice.

mandate is in effect. Increased market concentration leads to better sorting among the plans for consumers, a benefit that helps to mitigate the harms resulting from higher markups. In the case of policy regime (3), this reduction in the variance of prices is enough to more than compensate consumers for price increases on average.

The negative interaction between the individual mandate and market concentration also shows in these results. In policy regimes where the individual mandate is in effect—the baseline scenario and column (2)—the merger results in higher prices and lower consumer welfare than the counter-factual regimes where the individual mandate is absent.

8 Conclusion

I combine household level choices in the non-group health insurance market and the HHS-HCC risk prediction model with aggregate moments on cost and risk to identify the joint distribution between the willingness-to-pay for health insurance and the expected cost to an insurance firm. I find that there is significant heterogeneity in willingness-to-pay, and a strong relationship between a household's willingness-to-pay and the cost to cover that household, the key feature of adverse selection.

I use these estimates in a framework of imperfect competition to demonstrate how market structure relates to the two forms of distortion in markets with adverse selection: non-negative profits and adverse sorting. I show that under certain policy regimes, concentrated markets can improve allocations for high willingness-to-pay consumers over competitive markets.

I also demonstrate how market structure interacts with two common policies intended to correct the distortions of adverse selection—a penalty for being uninsured and risk adjustment transfers—by simulating equilibrium in four policy regimes with both, either, or neither policies in effect. I show the relationship between market structure and these policies both across the cross-section of markets and by simulating a proposed merger.

Taken together, these results suggest first and foremost that policy makers must be aware of market concentration when designing policies to help consumers. For instance, allowing states to set the value of the individual mandate would provide greater ability to optimally set high penalties where markets are competitive and low penalties where they are not.

Another take-away is that market concentration can act as a form of risk adjustment. In fact, a monopoly will implement the optimal “risk adjusted” pricing without the prompting of any policy. In many areas, this degree of concentration is likely more costly than the benefits of perfect risk adjustment would merit. However, it does imply that policy makers need not be concerned with a risk adjustment policy that can influence the behavior of a monopolist.

Limitations and Extensions

My intention is use a sufficient amount of data to estimate a relatively simple framework in which to narrowly examine the relationship between market structure and adverse selection. In doing so, my analysis certainly suffers from measurement error common in empirical work, but especially pronounced in the art of predicting medical expenses. I provide a methodology for estimating the relationship between preferences and cost without using data on individual level costs, but without access to such data, I am certainly only able to capture a portion of the relationship.

I am considering only one dimension of a many-faceted market. I am modeling the interaction between households and health insurance companies, while taking the relationships between households and health care providers, and the bargaining interactions between insurance firms and providers as exogenous and invariant to changes in policy/market structure. There is a innovative body of research that investigate these other portions of the market (Ho and Pakes (2014), **?**). In particular, Gowrisankaran et al. (2015) show how market structure influences the bargaining positions between insurers and hospitals.

I am also abstracting from a number of eccentricities of health care demand, where consumers are deciding between large numbers of complex products and have only a tenuous understanding of their own risks (Handel (2013), Abaluck and Gruber (2016)).

In future work, I hope to extend this research by using a model of imperfect competition in which firms may endogenously choose the generosity of their contracts. This margin can be important for determining plan characteristics such as network breadth and drug formularies, and is likely important when considering the welfare effects of adverse selection and market structure.

References

Abaluck, Jason and Jonathan Gruber, “Evolving Choice Inconsistencies in Choice of Prescription Drug Insurance,” *American Economic Review*, aug 2016, *106* (8), 2145–2184.

Abraham, Jean, Coleman Drake, Jeffrey S Mccullough, · Kosali Simon, B Jean, and Marie Abraham, “What drives insurer participation and premiums in the Federally-Facilitated Marketplace?,” *International Journal of Health Economics and Management*, 2017, *17*, 395–412.

Aizawa, Naoki and You Suk Kim, “Advertising and Risk Selection in Health Insurance Markets,” *Finance and Economics Discussion Series*, mar 2015, *2015* (101), 1–50.

Akerlof, George A., “The Market for ”Lemons”: Quality Uncertainty and the Market Mechanism,” *The Quarterly Journal of Economics*, 1970, *84* (3), 488.

ASPE, “Health Insurance Marketplace 2016 Open Enrollment Period: Final Report,” *Department of Health and Human Services, Office of the Assistant Secretary for Planning and Evaluation*, 2016, pp. 1–48.

Azevedo, Eduardo M. and Daniel Gottlieb, “Perfect Competition in Markets With Adverse Selection,” *Econometrica*, jan 2017, *85* (1), 67–105.

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston, “How does risk selection respond to risk adjustment? New evidence from the Medicare Advantage Program,” *American Economic Review*, 2014, *104* (10), 3335–3364.

Dafny, Leemore, Jonathan Gruber, and Christopher Ody, “MORE INSURERS LOWER PREMIUMS Evidence from Initial Pricing in the Health Insurance Marketplaces,” *American Journal of Health Economics*, 2015, *1* (1), 53–81.

_ **, Mark Duggan, and Subramaniam Ramanarayanan**, “Paying a Premium on Your Premium? Consolidation in the US Health Insurance Industry,” *American Economic Review*, 2012, *102* (2), 1161–1185.

de ven, Wynand P.M.M. Van and Randall P. Ellis, *Risk adjustment in competitive health plan markets*, Vol. 1, Elsevier Science B.V., 2000.

Decarolis, Francesco and Andrea Guglielmo, “Insurers’ response to selection risk: Evidence from Medicare enrollment reforms,” *Journal of Health Economics*, dec 2017, *56*, 383–396.

DeLeire, Thomas, Andre Chappel, Kenneth Finegold, and Emily Gee, "Do individuals respond to cost-sharing subsidies in their selections of marketplace health insurance plans?," *Journal of Health Economics*, dec 2017, *56*, 71–86.

Einav, Liran, Amy Finkelstein, and Mark R. Cullen, "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics*, aug 2010, *125* (3), 877–921.

Ellis, Randall P. and Thomas G. McGuire, "Predictability and predictiveness in health care spending," *Journal of Health Economics*, jan 2007, *26* (1), 25–48.

Ericson, Keith M Marzilli and Amanda Starc, "Pricing Regulation and Imperfect Competition on the Massachusetts Health Insurance Exchange," *Review of Economics and Statistics*, 2015, *97* (3), 667–682.

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

GAO, "CHILDREN'S HEALTH INSURANCE: Opportunities Exist for Improved Access to Affordable," *Government Accountability Office*, 2012, pp. 1–45.

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

_ **and Timothy Layton**, "Upcoding: Evidence from Medicare on Squishy Risk Adjustment," 2015.

_ **, Timothy J Layton, Grace Mccormack, and Mark Shepard**, "Trade-offs between Extensive and Intensive Margin Selection in Competitive Insurance Markets," 2018.

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

Gowrisankaran, Gautam, Aviv Nevo, and Robert Town, "Mergers When Prices Are Negotiated: Evidence from the Hospital Industry," *American Economic Review*, jan 2015, *105* (1), 172–203.

Graves, John A and Jonathan Gruber, "How Did Health Care Reform in Massachusetts Impact Insurance Premiums?," *American Economic Review: Papers & Proceedings*, 2012, (3), 508–513.

Gruber, Jonathan and James Poterba, "Tax Incentives and the Decision to Purchase Health Insurance: Evidence from the Self-Employed," *The Quarterly Journal of Economics*, aug 1994, *109* (3), 701–733.

Hackmann, Martin B, Jonathan T Kolstad, and Amanda E Kowalski, "Adverse Selection and an Individual Mandate: When Theory Meets Practice," *American Economic Review*, 2015, *105* (3), 1030–1066.

Handel, Ben, Igal Hendel, and Michael D. Whinston, "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk," *Econometrica*, 2015, *83* (4), 1261–1313.

Handel, Benjamin R, "Adverse Selection and Inertia in Health Insurance Markets :," *American Economic Review*, 2013, *103* (7), 2643–2682.

Hendren, Nathaniel, "Private Information and Insurance Rejections," *Econometrica*, sep 2013, *81* (5), 1713–1762.

Ho, Kate and Ariel Pakes, "Hospital choices, hospital prices, and financial incentives to physicians?," *American Economic Review*, 2014, *104* (12), 3814–3840.

_ **and Robin S Lee**, "INSURER COMPETITION IN HEALTH CARE MARKETS," *Econometrica*, 2017, *85* (2), 379–417.

Jaffe, Sonia and Mark Shepard, "Price-Linked Subsidies and Health Insurance Markups," *National Bureau of Economic Research Working Paper Series*, 2017, *No. 23104.*

Kautter, John, Gregory C. Pope, and Patricia Keenan, "Affordable Care Act Risk Adjustment: Overview, Context, and Challenges," *Medicare & Medicaid Research Review*, 2014, *4* (3), 1–11.

Layton, Timothy J., "Imperfect risk adjustment, risk preferences, and sorting in competitive health insurance markets," *Journal of Health Economics*, dec 2017, *56*, 259–280.

Lester, Benjamin, Ali Shourideh, Venky Venkateswaran, and Ariel Zetlin-Jones, "Screening and Adverse Selection in Frictional Markets," *Journal of Political Economy (Forthcoming)*, 2015.

Mahoney, Neale and E Glen Weyl, "Imperfect Competition in Selection Markets," Technical Report 4, National Bureau of Economic Research 2017.

Marquis, M. Susan, Melinda Beeuwkes Buntin, José J. Escarce, Kanika Kapur, and Jill M. Yegian, "Subsidies and the demand for individual health insurance in California," *Health Services Research*, oct 2004, *39* (5), 1547–1570.

McGuire, Thomas G., JOSEPH P. Newhouse, and ANNA D. SINAIKO, "An Economic History of Medicare Part C," *Milbank Quarterly*, jun 2011, *89* (2), 289–332.

_ , Joseph P. Newhouse, Sharon-Lise Normand, Julie Shi, and Samuel Zuvekas, "Assessing incentives for service-level selection in private health insurance exchanges," *Journal of Health Economics*, may 2014, *35*, 47–63.

Miller, Keaton, Amil Petrin, Robert Town, and Michael Chernew, "Optimal Managed Competition Subsidies," 2018.

Petrin, Amil, "Quantifying the Benefits of New Products: The Case of the Minivan," apr 2001.

Planalp, Colin, Julie Sonier, and Brett Fried, "State-Level Trends in Employer-Sponsored Health Insurance: A State-by-State Analysis," *Robert Wood Johnson Foundation*, 2015, pp. 1–83.

Pope, Gregory C, 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*, 2014, *4* (3), 1–23.

Rothschild, Michael and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *The Quarterly Journal of Economics*, 1976, *90* (4), 629.

Ryan, Conor, Roger Feldman, and Stephen Parente, "Estimating the Demand for Individual Health Insurance," 2019.

Sacks, Daniel W, Khoa Vu, Tsan-Yao Huang, and Pinar Karaca-Mandic, "The Effect of the Risk Corridors Program on Marketplace Premiums and Participation," 2017.

Saltzman, Evan, "Demand for Health Insurance: Evidence from the California and Washington ACA Marketplaces," *Working Paper*, 2017, pp. 1–46.

Shepard, Mark, "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," 2016.

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

Starc, Amanda, "Insurer Pricing and Consumer Welfare: Evidence from Medigap," 2014.

Tebaldi, Pietro, "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," 2017.

Veiga, André and E. Glen Weyl, "Product design in selection markets," *Quarterly Journal of Economics*, may 2016, *131* (2), 1007–1056.