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

December 9, 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 a 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 more concentrated markets can improve allocations for consumers with a high cost to serve and high willingness-to-pay by reducing the prices of the most generous insurance products. 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. I find that the individual mandate penalty increases the ability of large firms to charge high markups, increasing the average price in concentrated markets and increasing the price effect of potential mergers. I also find that market concentration provides an implicit form of risk adjustment, providing a possible benefit from mergers and reducing the necessity of a risk adjustment policy in concentrated markets.

1 Introduction

In the presence of adverse selection, where consumers with the highest willingness-to-pay are also the most costly to serve, health insurance markets can suffer from well-known failures (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 little attention (See Geruso and Layton (2017) for a broad review). The largest firm in the non-group health insurance market has a market share of over 85% in 5 states and less than 33% in another 5 states. In 2014, the Affordable Care Act (ACA) 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.

I model strategic firms that sell differentiated insurance products and set prices to maximize profit according to Nash-Bertrand price competition. Consumers choose among insurance products and an option to be uninsured, and I allow each 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 increase markups. 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, provides the only health insurance offers for about 40 million people. While not the largest segment of the US health insurance system, the non-group market demonstrates the most classic manifestations of adverse selection and has substantial geographic variation in market structure. Even after the market reforms implemented by the ACA, this market still features uninsurance rates around 50%, high prices, and lower levels of insurance coverage than other markets. 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-

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

group health insurance product choices with measures of health status are rare, and recent approaches to identifying the relationship rely on simulated variation in demand and average costs or assuming that prices are set optimally as specified by the model (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 the estimated supply and demand for health insurance, I simulate two counterfactual exercises under four policy regimes: the baseline policy that has both risk adjustment and the mandate penalty in effect, two regimes in which each policy is no longer in effect, and one regime where neither policy is in effect. In the first exercise, I explore the effect of these policies across the cross section of geographic markets. 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.

In a second exercise, I simulate two proposed mergers that were blocked by the US Department of Justice in 2015—Aetna proposed to acquire Humana, and Anthem Blue Cross proposed to acquire Cigna. In this exercise, I can more directly show how market concentration interacts with certain policy regimes. I find that, as a result of the merger, the prices of the less generous plan offerings increase by more than more generous insurance plans. In all policy regimes other than the baseline, the merger leads to higher average premiums for Bronze plans but lower average premiums for Silver and Gold plans. This demonstrates two key relationships between current policies and market structure: i) the individual mandate increases the ability for concentrated firms to raise price and (ii) market concentration, similar to risk adjustment policies, decreases the spread of prices between high actuarial and low actuarial value insurance plans.

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), Aizawa and Kim (2015), 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.

Products offered in the individual market are strictly regulated under the Affordable Care Act (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, frequently through reduced deductibles and out-of-pocket spending limits.

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

[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.

The primary policy tool to address adverse selection on the extensive margin of purchasing insurance 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%.

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. (2014a)). 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 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

[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).

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 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 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

[5] Choice sets are discussed in section 2.3

Table 1: Data Description

	Online Broker	ACS	ASPE
		Age Distribution	
Under 18	0.0%	0.0%	8%
18 to 26	17.5%	11.3%	11%
26 to 34	23.7%	16.9%	17%
35 to 44	20.4%	20.0%	17%
45 to 54	20.4%	24.2%	22%
55 to 64	18.0%	27.1%	25%
		Income Distribution	
Under 200% FPL	27.6%	29.5%	68%
200% to 400% FPL	21.8%	34.5%	31%
Over 400% FPL	50.6%	36.1%	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.

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. The only other product attributes I use in estimation are common to all plans in each category.

2.4 Market-level Demographic and Uninsured Information

The 2015 American Community Survey (ACS) provides the demographic distribution used in cost estimation and counterfactual simulations, and also to supplement the choice data used for estimation in two dimensions. First, the choice data does not contain any information on households that decide not to purchase insurance. Second, the choice data is a sample of all non-group insurance purchases in a market and is not informative for the relative size of the insured and uninsured populations. The ACS contains information on a representative sample of households in each geographic market, whether or not those households are eligible participants in the non-group market, and whether or not the households are insured. For more detailed information on constructing households and determining market eligibility, see appendix section A.1.

The estimation data is created by treating the choice data as a random sample from the population of insured individuals, conditional on subsidy eligibility and geographic market. [6] Each observation from the choice data within a particular subsidy eligibility category and market is given a equal weight such that the weights sum to the size of the population as determined by the ACS. The ACS also provides a weighted random sample of the uninsured population, where the weights are given by the ACS personal survey weights.

2.5 Consumer Risk Data

The 2015 Medical Expenditure Panel Survey (MEPS) is used to construct the distribution of risk among consumers and compute moments relating covered medical expenses to consumer risk and age.

The MEPS Medical Conditions File (MCF) contains self-reported diagnoses codes which can be linked to information on household demographics, insurance coverage, and medical expenses in the Full Year Consolidated File. I apply the HHS-HCC risk prediction model coefficients, published by CMS, to the self-reported diagnoses to compute risk scores. For details on the processing of the MEPS data, see appendix section A.2.

[6] The specifics of the weighting mechanics do not significantly alter the price elasticity and risk preference estimates. They are important for how well the model predicts untargeted moments like aggregate insurance rates and the firm first order conditions.

To identify the relationship between risk scores and demand, I use aggregate 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. Since MEPS contains a nationally representative distribution of risk scores, I target the national average risk score in the non-group market in 2015.

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.

In order to allow for firm-level quality to vary with consumer risk, I target the average risk score among groups of firms using data Medical Loss Ratio data. Firms are divided into four groups based on whether they are in the top 25^{th} percentile of average risk, and whether they are large (high average quality) or small (low average quality). I divide large and small firms based on whether the firm enrolls at least 5% of the insured population in each state. See the appendix for more detail on processing the Medical Loss Ratio data (appendix section A.4) and computing firm-level average risk (appendix section A.5)

2.6 Firm Cost Data

In addition to person level medical expenditure moments, this paper also uses firm level cost moments to estimate firm-specific marginal costs from two data sources.

First, the Center for Medicare and Medicaid Services (CMS) tabulates the Premium Rate Filings that insurance firms must submit to state insurance regulators to increase the premiums for products 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 and projections for the following year. For detail on the processing of this data, see appendix section A.3.

I use projected firm-level average cost and the average ratio of experienced costs across metal levels for all firms. [7] Unfortunately, the rate filing data do not fully cover every firm. As a

[7]Using projected average costs and experienced ratios lead to the best fit for untargeted firm first order conditions. This could possibly be because the product-level projections are distorted by the firms incentives

result, firm-level average costs are supplemented by Medical Loss Ratio (MLR) filings.

CMS makes publicly available the state-level financial details of insurance firms in the Individual Market for the purpose of regulating the MLR.[8] This information includes the number of member-months covered by the insurance firm in the state and total costs. For the firms that do not appear in the rate filing data, I adjust the average cost by the average ratio of projected costs to the average experienced costs imputed from the MLR data. In 2015, this ratio is 71.5% for firms in my sample. For more detail on the processing of this data, see appendix section A.4.

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. The outside good also has a price, in the form of a purchase penalty. However, the characteristics of uninsurance, (X_0, ξ_0) are normalized to have no value to consumers.

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, $P(\tau, p_j)$, that depends on its characteristics through age-rating regulation, income-based subsidies, and the size of the household. Importantly, the premium is not conditional on any direct measure of health status. I will write $P_j(\tau)$ as a shorthand for the household specific premium for product j. There are a continuum of

to meet the rate review requirements. While the decision to use projected or experienced costs does affect the marginal cost estimation, it does not qualitatively impact the results.

[8] 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.

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

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

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

$$S_j^{\tau\theta}(\boldsymbol{P}_m) = \frac{e^{u_j^{\tau\theta}}}{e^{u_0^{\tau\theta}} + \sum_{k\in J^m} e^{u_k^{\tau\theta}}}$$

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

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

where L^m is the size of market m, and $C_m(X,\tau,\theta)$ represents the firm-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 $C_f(X_j,\tau,\theta)$.

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. [9] The risk adjustment transfers take the form

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

[9] 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.

These transfers are positive and represent a tax on a particular insurance plan when the enrolled population of the plan has a lower expected cost than the overall population. The transfers are negative and represent a subsidy to a plan if the enrolled population is more costly on average than the population.

Price Regulation

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

$$P_j(\tau) = p_j A(\tau) - B(p_{2LS}, \tau)$$

where $A(\cdot)$ and $B(\cdot)$ are price modifier that are set by regulation and not firms. The multiplicative term, A, alters the base premium based on the size and age composition of the household. The additive term, B, represents a household specific subsidy that depends on household income, and the price of the second lowest-priced Silver plan, p_{2LS}.

For the main results of this paper, I will assume that household subsidies are vouchers that do not depend on the price of any firm's products, i.e. $B(p_{2LS}, \tau) = B(\tau)$. This isolates the mechanisms of interest from firms' strategic interactions with the household subsidy and government spending. I present results with price-linked subsidies in Appendix section (NEED TO ADD).

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, $\boldsymbol{P}^* = \{p_j^*\}_{j \in J^m}$ 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,

$$S_j = \int_{\tau,\theta} S_j^{\tau\theta}(\boldsymbol{P}^*) dF(\tau,\theta)$$

$$S_j' = \int_{\tau,\theta} \frac{\partial S_j^{\tau\theta}(\boldsymbol{P}^*)}{\partial p_j} dF(\tau,\theta)$$

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

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,

$$MC_j^{mkt} = \frac{1}{\sum_k \frac{\partial S_k}{\partial p_j}} \int_{\tau,\theta} \sum_k \frac{\partial S_k^{\tau\theta}(\boldsymbol{P}^*)}{\partial p_j} c_k^{\tau,\theta} dF(\tau,\theta)$$

$$PC = \frac{1}{\sum_k S_k} \int_{\tau,\theta} \sum_k S_j^{\tau\theta}(\boldsymbol{P}^*) c_k^{\tau,\theta} dF(\tau,\theta)$$

$$\Psi_j = \frac{\sum_k \frac{\partial S_k}{\partial p_j}}{S_j'}$$

Intuitively, the risk adjustment transfer 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, consider the case for a firm that is very smal, $S_j \to 0$. In this case, $\frac{S_j}{\sum_{k\in J^m} S_k} \to 0$, and the firm charges a markup over the pooled average cost. If a firm is a monopolist, $S_j = 1$ and $J^m = \{j\}$, then $\Psi_j = 1$ and $\frac{S_j}{\sum_{k\in J^m} S_k} = 1$. In this case, 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. In order to avoid addressing the cost and benefits of government spending on health insurance subsidies, I will assume that government subsidies are fixed. For simplicity, I will treat prices as constant across all consumers, $P_j(\tau) = p_j$. In this case, the vector of social welfare maximizing base premiums, $\boldsymbol{P}^W$, solves

$$0 = \int_{\tau,\theta} \sum_{k\in J^m} \frac{\partial S_k^{\tau\theta}}{\partial p_j}(p_k^W - C_m(X_k,\tau,\theta))dF(\tau,\theta) \quad \text{for all } j$$

$$\boldsymbol{P}^W = E\left[\frac{\partial \boldsymbol{S}^{\boldsymbol{\tau\theta}}}{\partial \boldsymbol{P}}\right]^{-1} E\left[\left[\frac{\partial \boldsymbol{S}^{\boldsymbol{\tau\theta}}}{\partial \boldsymbol{P}}\right]\boldsymbol{C}^{\tau\theta}\right]$$

where $\left[\frac{\partial \boldsymbol{S}^{\boldsymbol{\tau\theta}}}{\partial \boldsymbol{P}}\right]$ is the matrix of demand derivatives, $\boldsymbol{C}^{\tau\theta}$ is a vector of the expected costs for each product in the market given τ and θ, and the E operator stands in for the average over τ and θ.[10] In words, the social welfare maximizing 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 $MC_j^{mkt}(\boldsymbol{P}^W)$ as the average cost of the market-wide marginal consumer with respect to the price of product j.

[10]This result does not depend on the specifics of a 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)).

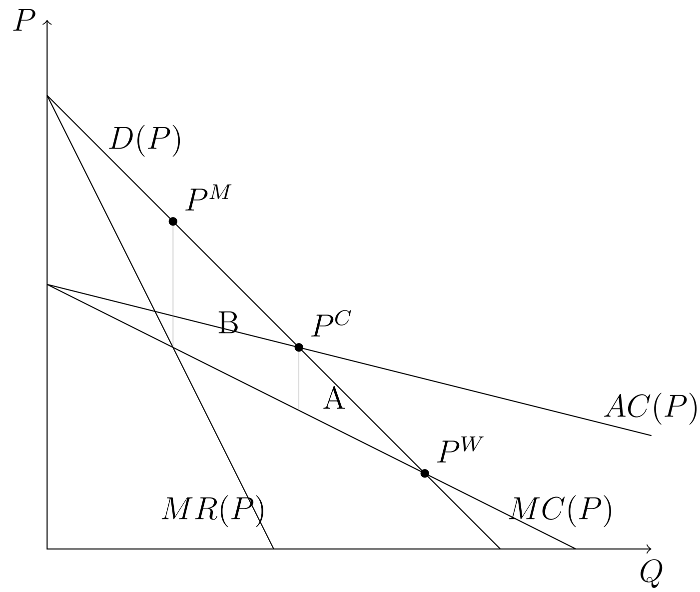

Figure 1: Single Product Adverse Selection

The first source of inefficiency in markets with adverse selection come from firms that must earn *non-negative* profits. Figure 1 displays a single product market where adverse selection leads to a downward sloping marginal cost curve $MC(P)$. In this case, the welfare maximizing price, P^W, is below average cost. 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 and the welfare loss grows to include the area demarcated by B.

Then second type of inefficiency is *inefficient sorting*. Firms setting prices for their own products alter the costs of their competitors by affecting the distribution consumers that purchase their competitors products. This between firm externality leads inefficient sorting of individuals across the plans available.[11]

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, $\overline{\Pi}$.

[11] The externality that one product imposes on another can be either positive or negative depending on the relative risk composition and prices of the products.

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

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

$$
\begin{aligned}
P_j+\lambda\frac{1-S_0}{-\partial S_0/\partial P_j}&=MC_j^{mkt}\\
\boldsymbol{P}^s+\lambda E\left[\frac{\partial \boldsymbol{S^{\tau\theta}}}{\partial \boldsymbol{P}}\right]^{-1}\boldsymbol{S}&=E\left[\frac{\partial \boldsymbol{S^{\tau\theta}}}{\partial \boldsymbol{P}}\right]^{-1}E\left[\left[\frac{\partial \boldsymbol{S^{\tau\theta}}}{\partial \boldsymbol{P}}\right]\boldsymbol{C}^{\tau\theta}\right]\\
\underbrace{P_j^s+\frac{\lambda-1}{\lambda}\frac{1}{\boldsymbol{\Psi_j}}\frac{S_j}{S_j'}}_{\text{Marginal Social Benefit}}&=MC_J^{mkt}
\end{aligned}
$$

where Ψ_j is the margin share defined in section 3.1. 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 "marginal social benefit" is equal to average cost. This is equivalent to charging an adjusted markup over market-wide 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 is 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 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.

The left panel (2a) displays 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, ΔMSB is equal to marginal cost. The competitive market sets the price difference to be too

[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.

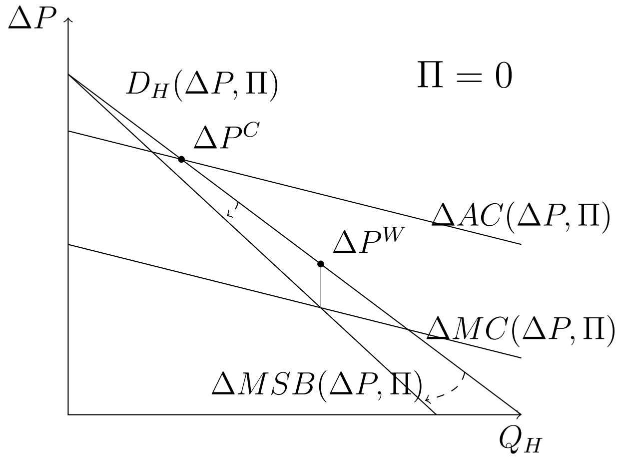

(a) Optimal Sorting with Zero Profits

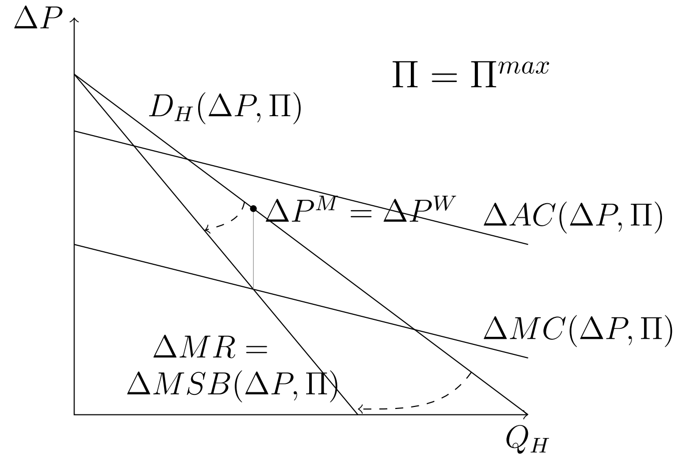

(b) Optimal Sorting with Monopolist Profits

Figure 2: Two Product Sorting

large.

In the case of a monopolist earning the maximum level of profit, (2b), the social marginal benefit curve coincides with the monopolists marginal revenue curve. As a result, the monopolist charges the socially optimal price difference, conditional on earning the maximum level of profits.

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 r^{HCC} is an unobserved risk score. 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, one 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}. The observed characteristics include the actuarial value of the insurance plan and three categorical variables: whether or not the firm enrolled less than 5% of the insured population of the market in 2015, whether or not the firm's enrolled population was among the top quartile of average risk, and the interaction of these two categories. These variables allow for risk-related preferences over firms.

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

$$\begin{aligned} u_{ijm} &= \gamma_i + \alpha_i(a_i p_{jm} - s(y_i)) + \beta_i X_{jm} + \xi_{jm} \\ u_{i0m} &= \alpha_i m(y_i) \end{aligned}$$

where $s(y)$ is a function that maps income to subsidies and $m(y)$ maps income to the mandate penalty. I allow the preference for insurance, γ_i, 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}.

$$\begin{aligned} \gamma_i =& \gamma_0 + \gamma_z' Z_i + \gamma_r r_i^{HCC} \\ \alpha_i =& \alpha_0 + \alpha_z' Z_i \\ \beta_i^k =& \beta_0 + \beta_r^k r_i^{HCC} \end{aligned}$$

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.[13]

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

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 plan spending on an individual, based on demographics and health condition diagnoses. It is the result of a linear regression of relative plan spending on a set of age-sex categories and a set of hierarchical

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

condition categories based on diagnoses codes.

$$\frac{\text{Plan Spending}_{it}}{\text{Avg. Plan Spending}_t} = \gamma_0 + \sum_g \gamma_{tg}^{age,sex} Age_{ig} Male_{ig} + \sum_{g'} \gamma_{tg'}^{HCC} HCC_{ig'} + \eta_{it}$$

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 relative 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 $\hat{G}$, the distribution of risk scores, 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.[14]

In the data, a majority of individuals have no relevant diagnoses, i.e. $r_i^{HCC} = 0$.[15] In order to match this feature of the data, I model a discrete probability that an individual has non-zero risk score and, conditional on having a non-zero risk score, a continuous distribution of risk scores. With some probability $\delta(Z_i)$, the household has a non-zero risk score drawn from a log-normal distribution, i.e. $log(r_i^{HCC}) \sim N(\mu(Z_i), \sigma)$. 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

[14]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.

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

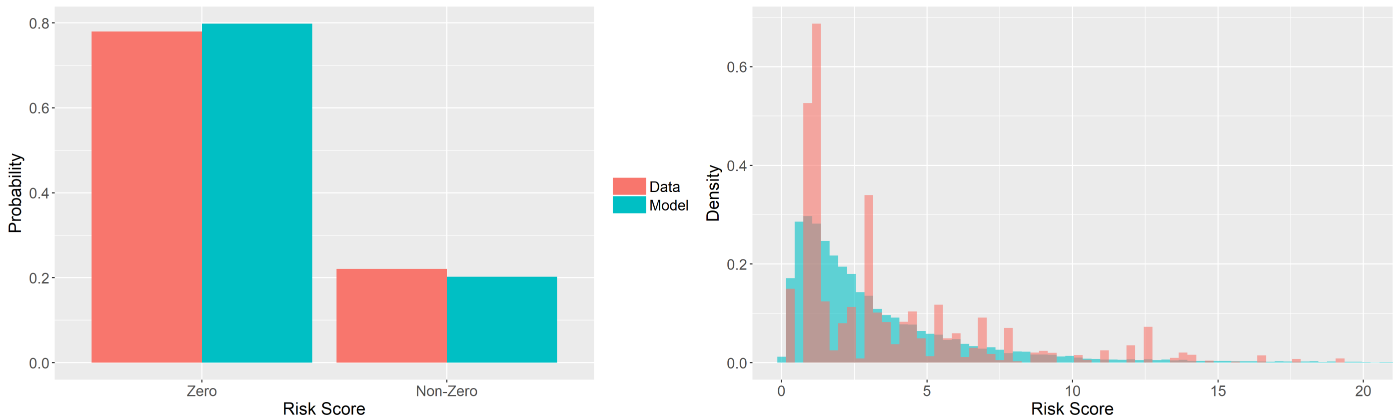

(a) Incidence of Positive Risk Scores

(b) Distribution of Positive Risk Scores

Figure 3: Risk Score Distribution Model Fit

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.

4.3 Estimation

This model has two primary identification concerns. First, plan premiums price may be correlated with the unobserved quality ξ_{jm}, leading to biased estimates of α_i. In this environment, the premium regulations provide a source of variation in price which is exogenous to variation in unobserved quality (Tebaldi (2017)). The age-adjustment on premium, a_i, increases monotonically and non-linearly 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 am able to allow price sensitivity to also depend on age and income, but only in broad buckets. Intuitively, the variation in price within each demographic bucket defined by Z_i identifies α for that particular demographic.

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.

The second concern is the identification of the risk coefficients, $(\gamma_r, \{\beta_r^k\})$. These parameters are incorporated into the estimation equations in the same manner as variance parameters for distributions of unobserved consumer preferences (e.g. Berry et al. (1995)). However, since I have data on the distribution of risk in the market and moments on the average risk of individuals that choose certain products, I am able to incorporate these "macro" moments to ensure that the model captures the appropriate risk-related substitution patterns and improve identification (Petrin (2001)).

The demand model targets nine moments on the distribution of consumer risk: the average risk score of all insured consumers, the average risk score of enrollees in the Bronze, Silver, Gold, and Platinum plan categories, the average risk score among firms in the top quartile of risk for both large and small firms, and the risk score for the remainder of large and small firms.[16]

Estimation Procedure

To estimate the demand model, this paper follows Imbens and Lancaster (1994) to combine maximum likelihood with macro moments. The model includes two sets of moments: the gradient of the log-likelihood function with respect to the parameters and the difference between the simulated risk moment targets and data values.

Due to the large number of fixed effects that control for unobserved product characteristics—each specification of the model contains at least one thousand parameters—it is computationally difficult to robustly locate the minimum of the objective function. Instead, I use a two-step estimation procedure that divides the parameter space into risk and non-risk related

[16]I denote a large firm as a firm that had at least a 5% share of the insured population in the state where it operated in 2015.

parameters: $\theta_r = (\gamma_r, \beta_r)$ and $\theta_{-r} = (\gamma_0, \gamma_z, \alpha_0, \alpha_z, \beta_0, \xi_{jm})$.

The log-likelihood gradient moments are also divided by risk and non-risk parameters. Let $m_1(\theta)$ represent the gradient of the log-likelihood function with respect to θ_{-r}, $m_2(\theta)$ represent the gradient with respect to θ_r, and $m_3(\theta)$ be the difference between simulated aggregate risk moments and those in the data.

$$
\begin{aligned}
m_1(\theta) &= \frac{\partial \ln \mathcal{L}(\theta)}{\partial \theta_{-r}} \\
m_2(\theta) &= \frac{\partial \ln \mathcal{L}(\theta)}{\partial \theta_r} \\
m_3(\theta) &= \{R_q^{\text{data}} - E[r_{ij}|\text{consumer } i \text{ purchases a plan } j \in J^q]\}
\end{aligned}
$$

where the q^{th} risk moment applies to a group J^q of products. These groups consist of all products, products in each metal level, products of firms that are small and non-risky, firms that are large and non-risky, firms that are small and risky, and firms that are large and risky. (For the data definitions of large and risky, see section 2.5). When computing the risk scores that related to data moments, I also must use product specific risk moments. The risk score is increasing the generosity of the metal-level. I approximate this by separately estimating the risk distribution for each metal level and assigning risk scores for an individual, $\{r_{ij}\}$, such that the risk scores of each product occupies the same point in the CDF of the metal-level specific risk score distribution.

The intuition of the two-stage procedure is to separate the parameters that require extra moments for identification from those that can be solved for using the well-behaved likelihood function. The estimate for risk parameters, $\hat{\theta}_r$, solves

$$\hat{\theta}_r = \text{argmin } M((\tilde{\theta}_{-r}(\theta_r), \theta_r))' W M((\tilde{\theta}_{-r}(\theta_r), \theta_r))$$

where $\tilde{\theta}_{-r}(\theta_r)$ represent the non-risk parameters that maximize the likelihood of the data given the risk parameters θ_r.

This two-step procedure is in the same spirit of the δ inversion implemented in Berry et al. (1995) to separate the mean utility of each product from the parameters governing preference heterogeneity. Instead of matching aggregate market shares, I find the best fit parameters to the likelihood of observing the micro data, given the guess of parameters governing risk preferences. For more detail on the estimation procedure and an argument for the consistency of the estimate, see appendix section B.

4.4 Results

In Table 3, I present the results from the demand estimation. 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 the maximum likelihood specification, the identifying variation that would allow risk preferences on the constant, actuarial value, and the groups of firms to be separately identified is very weak.[17] In these specifications, I restrict risk preferences to actuarial value.

In the next specifications, I include the risk score moments in the GMM estimation. The main difference in the GMM estimation results is the separate identification of risk preferences across insurance, actuarial value, and the firm categories. The parameter estimates for price sensitivity and mean actuarial value preference are relatively consistent across the maximum likelihood and GMM estimations. The final specification, GMM-2, has the most flexible fixed effects to control for mean differences in product quality and will be used as the preferred specification.

5 Marginal Cost

5.1 Empirical Model

I use aggregate data on firm costs and moments on how health care expenditures vary by age and risk to estimate costs through Method of Simulated Moments (MSM). This method does not require that firms are playing optimal strategies according to the specification of

[17]Identification only comes from substitution patterns. However, since the preferences are perfectly correlated, the parameters affect the likelihood function nearly co-linearly.

Table 3: Demand Estimation Results

	Maximum Likelihood		GMM	
	(LL-1)	(LL-2)	(GMM-1)	(GMM-2)
Annual Premium ($000)	-1.46	-1.26	-1.42	-1.25
Age 31 - 40	0.24	0.24	0.24	0.24
Age 41 - 50	0.34	0.29	0.35	0.30
Age 51 - 64	0.69	0.55	0.67	0.55
Family	-0.17	-1.13	-0.14	-0.13
Subsidized	0.09	0.21	0.12	0.22
AV	4.40	9.36	4.57	9.49
Risk Preferences				
Constant			(-0.00)	0.07
AV	1.19	0.90	0.61	0.61
High Risk Firm			0.01	0.01
Small Firm			-0.04	-0.06
Small and High Risk			0.06	0.06
Fixed Effects				
Age, Family, Income	Y	Y	Y	
Firm	**Y**		**Y**	
Firm-Market-Category		**Y**		**Y**

Notes: All parameters are statistically significant at 0.1 percent level, except for those displayed in parenthesis. Currently, the efficient standard errors are currently used for GMM, instead of a standard error method that accounts for the two stage estimation strategy. 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.

the model. I specify the expected cost function, $C_f(X_j, \tau_i, \theta_i)$, as

$$\log(c_{ijm}) = \psi_f + \phi_1 AV_{jm} + \phi_2 Age_i + \phi_3 r_i^{HCC} + \omega_{ijm}$$

where ϕ_f is a firm-state specific fixed effect, AV_{jm} is the actuarial value of the product, Age_i is the average age of the household, and r^{HCC} is the risk score of household. 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.[18] 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 endogeneity is consistent with adverse selection, and the remaining variation in health status is orthogonal to age and HCC categories, then the coefficient on actuarial value will be biased upward.[19] The result of this bias is to attribute some portion of the selection differences of cost to product differences of cost. In the context of this study, this leads to conservative conclusions about the implications of adverse selection.

[18] 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 this particular exogeneity assumption, but the principle concern that residual costs unobservable to the econometrician are correlated with demand errors would remain.

[19] 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

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

I assume that, conditional on age and risk score, this covariance term is 0. If there is an endogeneity problem consistent with adverse selection, this covariance term would be positive and increasing in plan generosity, leading to an upward bias in the estimated coefficient on adverse selection.

Reinsurance

I include the ACA reinsurance in predicting costs. The reinsurance program operates on the individual level. In 2015, the federal government covered 45% of an insurance firms annual liabilities for a particular individual that exceeded an attachment point, $\underline{c} = \$45,000$, and up to a cap, $\bar{c} = \$250,000$.

$$
\begin{aligned}
c_{ijm}^{cov} &= \min\big(\max(c_{ijm} - \underline{c}, 0), \bar{c} - \underline{c}\big) \\
c_{ijm}^{exc} &= \max(c_{ijm} - \bar{c}, 0) \\
c_{ijm}^{rein} &= \min(c_{ijm}, \underline{c}) + 0.45 c_{ijm}^{cov} + c_{ijm}^{exc}
\end{aligned}
$$

Estimation

The MSM estimation procedure targets four sets of moments which each identify four sets of parameters. The age and risk parameters are identified using moments from the Medical Expenditure Panel Survey (section A.2). For clear identification, the estimation targets age moments among adults that have a risk score of zero. The moments are computed as the ratio of average covered expenditures within 5 year age brackets for adults between 25 and 64 years old to the average covered expenditures of 20 to 24 year olds. The cost parameter on risk is identified using the ratio of average covered expenditures among adults with a positive risk score to those with a risk score of zero.

The parameter on actuarial value is identified using the ratio of experienced cost of each metal level to Bronze plans from the 2016 rate filing data. And conditional on these three cost parameters, ϕ, the firm-specific cost parameter, ψ, is set to precisely match the projected average cost in the 2015 rate filing data. See section A.3 for more detail on the data.

When simulating moments that match data from the insurance firm rate filings, I use the reinsurance adjusted cost, c_{ijm}^{rein}. The moments from the Medical Expenditure Panel Survey are computed using total covered expenses across all insured individuals. Thus, I use the predicted cost c_{ijm} to compute these moments.

Cost is estimated using two-stage Method of Simulated Moments to obtain the efficient weighting matrix. The estimated demand parameters are used to simulated the distribution of consumer age and risk throughout products in each market, using ACS data as the

Table 4: Cost Estimation Results

	(GMM-1)	(GMM-2)
Age	0.47	0.47
Risk	0.10	0.10
Actuarial Value	4.20	3.62
State-Firm	Y	Y

Note: All parameter estimates are significant at the 1% level. The specifications correspond to the different demand estimation specifications that are used to simulate the moments.

population of possible consumers (see section 2.4). For a detailed description of the cost estimation procedure, see appendix section C.

5.2 Results

Table 4 displays the results of the cost estimation. The table presents results for each GMM demand specification used to simulate the moments targeted by the cost estimation.

In Table 5, I present the targeted and estimated moments used in the cost estimation. The age and risk moments are matched more closely than the metal level ratio moments. In particular, the cost specification leads to over estimates of the cost of covering individuals with platinum coverage. The combination of ordered risk preferences, age preferences, and log-linear costs in actuarial value lead to the implication that the difference in average costs among expensive and generous plans (Gold and Platinum) is very large. The model also predicts silver plan average costs that are very similar to gold plan average costs, which is a result of high cost selection into silver plans among individuals who receive cost-sharing reductions. In extensions of this work, more flexible specifications in actuarial value.

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 the estimated parameters are consistent with profit-maximizing firms in Nash-Bertrand competition. In Figure 4, I plot the marginal revenue and marginal cost implied by estimated parameters under the baseline policy regime, which includes a mandate penalty, risk adjustment, and reinsurance.

Table 5: Cost Estimation Fit of Cost-Ratio Moments

	Data	Model Fit (GMM-1)	Model Fit (GMM-2)
Age ($r^{HCC} = 0$)			
18 - 24	1.0	1.0	1.0
25 - 29	1.33	1.29	1.29
30 - 34	1.44	1.62	1.61
35 - 39	2.08	2.48	2.41
40 - 44	2.98	2.31	2.28
45 - 49	1.74	2.91	2.88
50 - 54	3.49	3.18	3.23
55 - 59	2.98	3.90	4.02
60 - 64	3.57	3.96	4.12
Risk			
$r^{HCC} = 0$	1.0	1.0	1.0
$r^{HCC} > 0$	3.56	3.35	3.35
Metal Level			
Bronze	1.0	1.0	1.0
Silver	2.28	3.11	2.47
Gold	3.80	3.02	2.89
Platinum	4.28	8.13	7.39

Note: This table displays the targeted and estimated cost ratios that are used to identify the marginal cost parameters in Table 4. In each category—age, risk, and metal level—the ratios are defined relative to the first row. The first row of each category is equal to one by construction. The two columns of estimated moments represent the two demand estimation specifications used to simulate the moments.

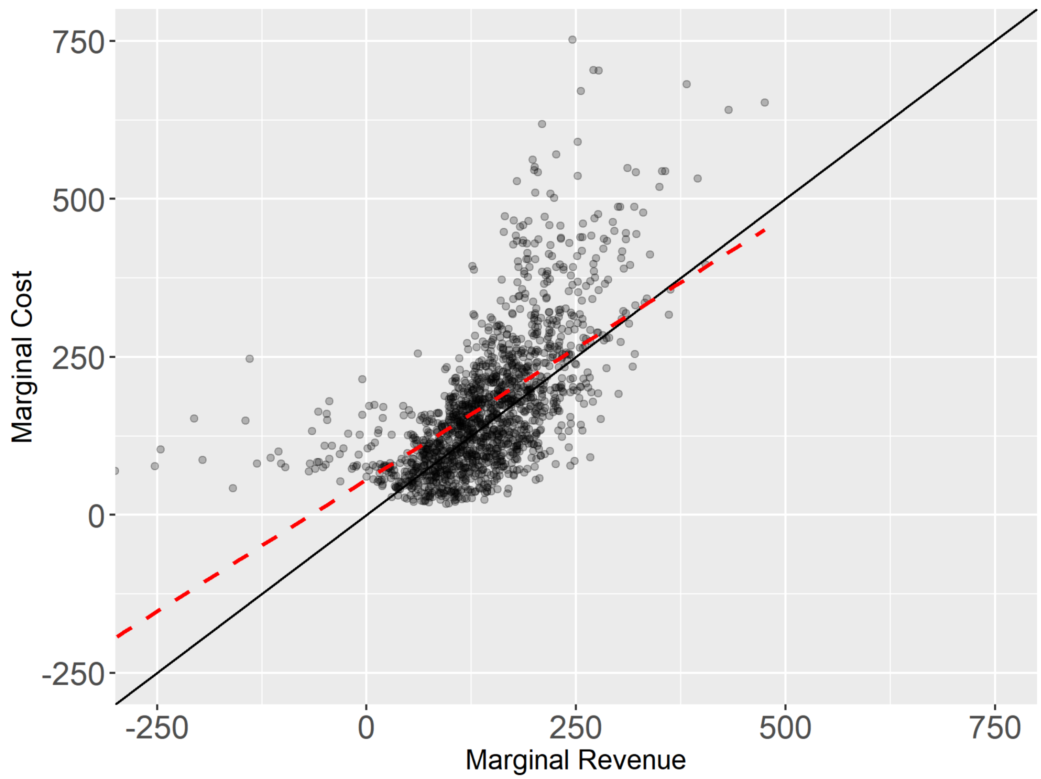

Figure 4: Marginal Revenue vs Marginal Cost in Baseline Model

On average, firms are setting marginal revenue close to marginal cost. However, the largest deviations come from firms in very concentrated markets. The median of estimated marginal cost less marginal revenue in the most competitive two-thirds of markets (markets with an HHI of less than 5200) is \$4.99 per month, and the mean is \$10.2. In the most concentrated third of markets, the median difference is \$34.2 per month and the mean is \$54.0. A possible explanation for marginal costs that exceed the implied marginal revenue in very concentrated markets is that state insurance agencies are successful in negotiating lower markups on behalf of consumers. This mechanism will not be modeled in this paper, but influences how the results should be interpreted for near monopoly markets.

6 Policy Analysis

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 rules and regulations of the ACA. This includes price-linked subsidies, age-rating, reinsurance, and the medical-loss ratio requirements.[20] I assume that firms do not know whether or not the benchmark plan to which subsidies are linked is a plan that they offer. Jaffe and Shepard (2017) show that this leads to lower effective elasticities and higher markups.

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.

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 3350, 36 markets have an HHI of between 3350 and 5200, and 37 markets have an HHI of at least 5200.

Panel A demonstrates the effects of risk adjustment and the individual mandate in the most 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 substantially for Silver and Gold plans, and leads to only small increases in the premiums of Bronze plans. However, in the presence of risk adjustment, the individual mandate has a relatively small effect on prices. In comparing column (2) to the baseline, repealing the individual mandate would raise average premiums by between 4 and 7 percent. However, if there were no risk

[20] In 2015, the ACA also included another other risk protection programs: 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.

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

Panel A: Less than 3350 HHI					
	Observed	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y	N	Y	N
Mandate	Y	Y	Y	N	N
Bronze Premium	2.08	1.90	1.84	1.97	1.88
Silver Premium	2.57	2.71	3.68	2.91	5.03
Gold Premium	2.85	2.90	6.82	3.06	9.02

Panel B: Between 3350 and 5200 HHI					
	Observed	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y	N	Y	N
Mandate	Y	Y	Y	N	N
Bronze Premium	2.32	2.10	1.19	2.15	2.01
Silver Premium	2.86	3.03	3.34	3.25	4.23
Gold Premium	3.29	3.04	4.12	3.20	6.71

Panel C: Greater than 5200 HHI					
	Observed	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y	N	Y	N
Mandate	Y	Y	Y	N	N
Bronze Premium	2.01	2.64	2.57	2.51	2.43
Silver Premium	2.76	3.70	3.75	3.63	3.70
Gold Premium	3.14	3.83	4.60	3.76	4.63

Notes: All values are levels in thousands of USD per year, with the exception of the insurance rate. All model averages are computed using the baseline enrollment distribution. The observed averages are computed using the observed 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.

adjustment, the individual mandate has a much larger effect. Comparing column (3) to column (1), average premiums for Silver and Gold plans are 32 - 37% greater without the individual mandate in place. This suggests that risk adjustment is a crucially important feature of providing firms an incentive to provide generous plans at affordable prices.

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 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 Panel C, 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.

An important consideration is how well this model of profit-maximizing firms in Nash-Bertrand competition matches the observed prices. Figure 4 shows that marginal costs are slightly higher than marginal revenue, and Table 6 shows that the model substantially over-estimates equilibrium prices in the most concentrated markets. One potential reason is that state insurance regulators negotiate more fiercely when facing near monopolists. If states are effective in negotiating lower markups, it is possible that the benefits of more concentrated markets may outweigh the harms.

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-

Table 7: Merger Analysis: Effect of Market Concentration

Panel A: Products of the Merging Parties

	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	N	Y	N
Mandate	Y	Y	N	N
Bronze Premium	6.3%	7.4%	1.3%	4.0%
Silver Premium	2.7	-1.8	-0.4	-3.5
Gold Premium	2.0	-2.5	-1.1	-3.5

Panel B: All Products

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

Notes: All values are percent changes in price 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.

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. While the validity of this exercise does not require that these 21 markets be randomly selected, the non-group market was not the primary market of concern for the Justice Department. This provides some confidence that these mergers were not driven by unobserved features of these 21 markets, and is consistent with modest price effects of the merger. The median pre-merger HHI in affected markets is 3733, and the median change in HHI is 197.

The results in Table 7 highlight the two first-order predictions of increased concentration in markets with adverse selection: Prices increase on average, and the price spread between high and low actuarial value insurance will decrease. These effects are present across all policy regimes. 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 also hold here. First, the individual mandate increases the ability of large firms to raise price. In comparing the baseline column to column (2) and column (1) to column (3), the presence of the individual mandate increases the price effects of the merger, by several percentage points in each product category.

Second, market concentration mitigates the externalities between firms that drive large spreads in prices between high and low actuarial value plans, in a similar manner to risk adjustment. In all policy regimes, the price increases among Bronze plans are smaller than the price increases in Gold plans, and in three out of the four regimes, the average premiums for Silver and Gold plans decline. This suggests that in markets without risk adjustment of an individual mandate penalty in effect, market concentration can provide a benefit to some consumers, and in particular the consumers that are willing to pay for high actuarial value insurance.

7 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.

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.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, jul 1995, *63* (4), 841.

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.

Hayashi, Fumio, *Econometrics*, Princeton, New Jersey: Princeton University Press, 2000.

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.

Imbens, Guido and Tony Lancaster, "Combining Micro and Macro Data in Microeconometric Models," *Review of Economic Studies*, 1994, *61* (4), 655–680.

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.

_ **, Gregory Pope, Sara Freeman, Lindsey Patterson, Michael Cohen, and Patricia Keenan**, "The HHS-HCC Risk Adjustment Model for Individual and Small Group Markets under the Affordable Care Act," *Medicare & Medicaid Research Review*, 2014, *4* (3), E1–E46.

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.

A Data Processing

A.1 American Community Survey

This paper uses the 2015 American Community Survey (ACS) to match the demographic distribution of the uninsured population and the income distribution of the insured population in each market. 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)).

In order to address under-reporting of Medicaid enrollment, I define any parent that receives public assistance, any child of a parent that receives public assistance or is enrolled in Medicaid, any spouse of an adult that receives public assistance or is enrolled in Medicaid or any childless or unemployed adult that receives Supplemental Security Income payments as being enrolled in Medicaid. Besides Medicaid and CHIP enrollment, an individual is considered eligible for either program if his or her household income falls within state-specific eligibility levels. If an individual is determined to be eligible for Medicaid through these means but reports to be enrolled in private coverage, either non-group coverage or through an employer, I assume that these individuals are enrolled in Medicaid. This accounts for those that confuse Medicaid managed care programs with private coverage, and Medicaid employer insurance assistance.

I follow a methodology outlined by the Government Accountability Office (GAO (2012)) to construct health insurance purchasing units. This method first divides households as identified in the survey data into tax filers and tax dependents, linking tax dependents to particular tax filers. A tax filing household, characterized by the single filer or joint filers and their dependents, is generally considered to be a health insurance purchasing unit. In some cases, certain members of a tax household will have insurance coverage through another source, e.g. an employer or federal program. In this case, the health insurance purchasing unit is the subset of the household that must purchase insurance on the non-group market.

A.2 Medical Expenditure Panel Survey

The 2015 Medical Conditions File (MCF) of the Medical Expenditure Panel Survey (MEPS) contains self-reported diagnoses codes. 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. For example, I treat a 3-digit code of '571' as '571.00'. This implies that many conditions in the hierarchical risk prediction framework are censored. However 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.

I link the MCF to the Full Year Consolidated File to identify the age and sex of the individual, and then apply the 2015 HHS-HCC risk prediction methodology (Kautter et al. (2014b)). The risk coefficients are published by CMS and publicly available.

A.3 Rate Filing Data

The rate filing data are divided into two files—a firm-level worksheet and a plan-level worksheet—and contain information on the prior year experience of the plan and the projected experience of the plan in the coming year.

To construct moments on the ratio of average cost across metal level categories, I use the prior year experience submitted in the 2016 rate filings data. To recover the average cost after reinsurance, I subtract the experienced total allowable claims that are not the issuer's obligation and the experienced risk adjustment payments from the total allowable claims.

The ratio of average cost across each metal level category is computed as the weighted average of every within firm ratio. I compute the average cost across all plans within each metal level category in each firm, and then compute the weighted average of the ratios across each firm. Each step is weighted using the reported experienced member months. The model moments are constructed in the same manner.

To estimate firm average costs, this paper takes advantage of the firm's projected costs for the 2015 plan year. During the first several years of the market, insurance firms experienced higher than projected costs, which led many firms to exit the market in the first three years. In order to capture this expectation in the strategies of the firms, I use the projected firm level average cost from the 2015 plan year firm-level rate filing data. I compute post-reinsurance projected costs by subtracting projected reinsurance payments from "projected incurred claims, before ACA Reinsurance and Risk Adjustment."

Some firms do not appear in the risk filing data. I compute the projected average cost for those firms by adjusting the experienced average cost reported in the Medical Loss Ratio filings by the average ratio of projected to experienced claims. In 2015, the average ratio of project to experienced claims for firms in my sample is 71.5%.

A.4 Medical Loss Ratio Data

This paper uses two pieces of information from the Medical Loss Ratio filings: average cost and average risk adjustment transfers.

Firms are defined by operating groups at the state level. Some firms submit several medical loss ratio filings under for different subsidiaries in a given state. I group these filings together.

Average cost is defined as total non-group insurance claims divided by total non-group member months, current as of the first quarter of 2016. This computation includes claims and member months that may not be a part of the non-group market as it is characterized in this analysis. For instance, grandfathered insurance plans that are no longer sold to new consumers are included. These are likely to be a small portion of the overall market.

To compute the average risk adjustment payment, some adjustment to the qualifying member months is required. Unlike medical claims, grandfathered plans (and other similar non-ACA compliant plans) are not included in the risk adjustment system. Dividing the total risk adjustment transfer by the total member months will bias the average transfer towards zero.

The interim risk adjustment report published by CMS includes the total member months for every state. And the MLR filings separately list the risk-corridor eligible member months, which are a subset of the risk adjustment eligible member months. I define "potentially non-compliant" member months as the difference between risk-corridor eligible member months and total member months. I reduce the potentially non-compliant member months of all firms in each state proportionally so that total member months is equal to the value published by CMS, with two exceptions. First, firms that opted not to participate in the ACA exchange in that state have zero risk-corridor eligible member months. I do not reduce the member months of these firms, as I cannot isolate the potentially non compliant months. Second, if the risk-corridor eligible member months exceed the total member months published by CMS, I assume that the risk-corridor eligible member months are exactly equal to the risk adjustment eligible member months.

A.5 Computing Firm-level Risk

In this paper, I use firm-level risk transfers to infer the equilibrium distribution of risk across firms. With a bit of simplification, the ACA risk transfer formula at the firm level can be written as

$$T_f = \left[\frac{\bar{R}_f}{\sum_{f'} S_{f'} \bar{R}_{f'}} - \frac{\bar{A}_f}{\sum_{f'} S_{f'} \bar{A}_{f'}} \right] \bar{P}_s$$

where $\bar{R}_f$ is the firm level of average risk and $\bar{A}_f$ is the firm level average age rating, where the average is computed across all the firms products and weighted by members, a geographic adjustment, and a metal-level adjustment. S_f is the firm's state-level inside market share, and $\bar{P}_s$ is the average total premium charged in the state.

Every element of this formula is data available in the Interim Risk Adjustment Report on the 2015 plan year, except for the plan-level market shares, the plan-level average age rating, and the plan-level average risk. As a simplification, I assume that the average age rating is constant across all firms, and that the weighting parameters in the risk component are negligible. I compute the implied firm-level average risk as

$$\bar{R}_f = \left(\frac{T_f}{\bar{P}_s} + 1 \right) \bar{R}$$

where the risk transfer T_f is the average firm-level risk adjustment transfer from MLR data, $\bar{P}_s$ is the average state level premium reported in the interim risk adjustment report, and $\bar{R}$ is the national average risk score reported in the interim risk adjustment report.[21]

[21] The formula implies that the state average risk score should go in place of the national average. However,

Another potential method to capture the relative risk of firms is simply to target the risk adjustment transfer itself, T_f, while everything else depends on the parameters of the demand model. In smaller samples of the data, I have found that this does not substantially alter the results of the estimation but introduces nonlinearities in the moment calculations that make the task of finding a minimum to the GMM objective function considerably more difficult.

B Demand Estimation Procedure

B.1 The GMM Objective Function

The GMM objective function is a composite of aggregate data moments and the first derivatives of the likelihood function. The likelihood function is defined as

$$\mathcal{L} = \Pi_i \Pi_{j \in J^{M(i)}} \left(\int_r S_{ijr} \; g(r; Z_i) dr \right)^{Y_{ij}}$$

where $J^{M(i)}$ is the set of products that are available to consumer i, S_{ijr} is the predicted probability that consumer i will choose product j, conditional on having a risk score of r, $g(\cdot; Z)$ is the probability density of the risk score, conditional on the demographic vector of consumer i. The risk score integral can be broken into two parts.

$$\int_r S_{ijr} \; g(r; Z_i) dr = (1 - \delta(Z_i)) S_{ij0} + \int_{r>0} S_{ijr} \; g(r; Z_i) dr$$

where $\delta(Z) = \int_{r>0} g(r; Z_i) dr$ is the probability that a risk score for an individual is greater than 0. The distribution of positive risk scores is log-normal, and I compute the integral by simulation using 1,000 halton draws.

The parameter set θ is divided into two sets, $\theta_r = (\gamma_r, \beta_r)$ and $\theta_{-r} = (\gamma_0, \gamma_z, \alpha_0, \alpha_z, \beta_0, \xi_{jm})$. The moments are divided into three sets. $m_1(\theta)$ represents the gradient of the log-likelihood function with respect to θ_{-r}, $m_2(\theta)$ is the gradient with respect to θ_r and $m_3(\theta)$ is difference between simulated aggregate risk moments and those in the data.

I do not allow the risk distribution among consumers to vary by geography (other than through composition). I use the national risk score to abstract from these geographical differences.

$$
\begin{aligned}
m_1(\theta) &= \frac{\partial \ln \mathcal{L}(\theta)}{\partial \theta_{-r}} \\
m_2(\theta) &= \frac{\partial \ln \mathcal{L}(\theta)}{\partial \theta_{r}} \\
m_3(\theta) &= \{R_q^{\text{data}} - E[r_{ij}|\text{consumer } i \text{ purchases a plan } j \in J^q]\}
\end{aligned}
$$

For a positive definite weighting matrix, W, the GMM objective function is given by

$$
Q(\theta) = m(\theta)' \; W \; m(\theta)
$$

$$
\text{where} \quad m(\theta) = \begin{bmatrix} m_1(\theta) \\ m_2(\theta) \\ m_3(\theta) \end{bmatrix}
$$

B.2 Two-Stage Estimation Procedure

The size of the parameter set θ_{-r} is very large due to the large number of fixed effects used to control for ξ_{jm}. Because of the large parameter space, it is computationally difficult to robustly locate the minimum of the GMM objective function, $Q(\theta)$, for even simple weighting matrices. However, the likelihood function is smooth and concave in the non-risk parameters, θ_{-r}, for large portions of the parameter space. Locating the maximum of the log-likelihood function with respect to θ_{-r}, given any choice for θ_r is computationally expedient. Notice also that, at the maximum of the likelihood function, $m_1(\theta) = 0$.

Let $\tilde{\theta}_{-r}(\theta_r)$ be the parameters θ_{-r} that maximize the likelihood of the data conditional on the parameters θ_r.

$$
\tilde{\theta}_{-r}(\theta_r) = \text{argmax}_{\theta'_{-r}} \ln \mathcal{L}(\theta'_{-r}, \theta_r)
$$

Then, instead of minimizing the objective function, $Q(\theta)$, I estimate θ by minimizing

$$
\tilde{Q}(\theta_r) = \tilde{m}(\theta_r)' \; W \; \tilde{m}(\theta_r)
$$

$$
\text{where} \quad \tilde{m}(\theta_r) = \begin{bmatrix} m_2(\tilde{\theta}_{-r}(\theta_r), \theta_r) \\ m_3(\tilde{\theta}_{-r}(\theta_r), \theta_r) \end{bmatrix}
$$

The function $\tilde{Q}$ is minimized using an iterated two-stage procedure. For an initial starting parameter, θ_r^0, I find $\theta_{-r}^0 = \tilde{\theta}_{-r}(\theta_r^0)$. Holding fixed θ_{-r}^0, I find the next iteration of θ_r^1 by minimizing $\tilde{Q}(\theta_r)$. Then I find $\theta_{-r}^1 = \tilde{\theta}_{-r}(\theta_r^1)$ and the procedure repeats. This two-stage iteration repeats until $|\theta_{-r}^{n-1} - \tilde{\theta}_{-r}(\theta_r^n)| < \epsilon$, where ϵ is some small number such that θ_{-r}^{n-1} is within the likelihood maximization tolerance levels of the true maximum $\tilde{\theta}_{-r}(\theta_r^n)$.

The minimum of $\tilde{Q}(\cdot)$ and the maximum of $\mathcal{L}(\cdot, \theta_r)$ are each found using a newton-raphson methodology. Since computing the hessian matrix is computationally expensive, the hessian is initially computed and subsequently updated using the BFGS algorithm. The optimization algorithm is greedy in its search for higher (or lower) function values. If a better point cannot be found with the initial newton step, it conducts a brief line search by reducing the magnitude of the step uniformly until a better point is found. If a better point cannot be found, as is sometimes the case with the approximated hessian matrix, the true hessian is recomputed and the step is repeated.

I perform this two-stage estimation procedure twice. 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 $h(\theta)$, then the asymptotic variance can be estimated with

$$\hat{V} = E[\nabla_h M(h(\theta))\big(MM'\big)\nabla_h M(h(\theta))']$$
$$\text{where} \quad M(h(\theta)) = m(\theta)$$

B.3 Consistency of Demand Estimation

Imbens and Lancaster (1994) show criteria under which a GMM estimation that combines moments on aggregate data moments with the first order conditions of the log-likelihood function is consistent and efficient. The two-stage estimation detailed in the previous section imposes the constraint that $m_1(\theta) = 0$. Were the system just-identified, this would not be a constraint and indeed a property of any solution. But since the number of moments exceed the number of parameters via the addition of macro moments, the solution to the unconstrained problem may have a $m_1(\theta)$ not equal to 0.

However, the two-stage estimate is still consistent. The argument stems from two observations. First, note that the true parameter, θ_0, meets the constraint. Under the specification and identification assumptions of GMM, the true parameter θ_0 has $m(\theta_0) = 0$. Let $\bar{\Theta}$ be the compact parameter space of the original GMM problem, and let $\Theta \subset \bar{\Theta}$ be the subset of parameters, also compact, that meet the constraint imposed by the two stage constraint, $\Theta = \{\theta \mid \theta_{-r} = \tilde{\theta}_{-r}(\theta_r)\}$. Since θ_0 has $m_1(\theta_0) = 0$, it is clear that the non-risk parameters solve the first order conditions of the maximum likelihood function conditional on the risk parameters. Thus, $\theta_0 \in \Theta$.

Second, for any parameter vector $\theta \in \Theta$, the moments in the standard GMM problem are equivalent to those in the adjusted problem, by construction.

$$\begin{aligned}
m_1(\tilde{\theta}_{-r}(\theta_r), \theta_r) &= m_1(\theta) \\
m_2(\tilde{\theta}_{-r}(\theta_r), \theta_r) &= m_2(\theta) \\
m_3(\tilde{\theta}_{-r}(\theta_r), \theta_r) &= m_3(\theta)
\end{aligned}$$

Thus, the two-stage procedure used in this paper can be viewed as a restriction of the parameter space that does not exclude the true parameter vector, θ_0. Proposition 7.7 of Hayashi (2000) shows the conditions under which non-linear GMM is consistent, and all of those conditions can be applied in this setting.

Another argument for consistency can be made using the weighting matrix of the standard GMM problem. Let $\hat{\theta} \in \Theta$ be the solution to the two-stage minimization problem. For any other parameter vector in the original parameter space, $\theta \in \bar{\Theta}$, there is a diagonal weighting matrix that places a large enough weight on the moments, m_1, such that θ does not lead to a lower value of the objective function that $\hat{\theta}$. If L is the supremum of all such "large enough weights," then a diagonal weighting matrix that places weight L on moments in m_1 and weight 1 on moments in m_2 and m_3 will give $\hat{\theta}$ as the solution to the original GMM problem. Since this weighting matrix is positive definite, the estimate is consistent.

C Cost Estimation Procedure

The cost parameters are estimated by matching a number of moments on firm-level costs and individual-level costs. I constrain the estimation to match precisely the projected-firm level average costs. The remaining cost parameters are estimated to fit moments three sets of moments: the ratio of the average cost of each metal level to the average cost of a bronze plan, the ratio of the average cost of each age group to the average cost of a 21-year old conditional on having a risk score of zero, and the ratio of the average cost of individuals with a positive risk score to those with a risk score of 0.[22] See appendix section A.2 through A.4 on constructing these moments from the data.

Matching Firm Moments

Let $\bar{C}_f^{obs}$ be the observed projected firm-level average cost. I set $\tilde{\psi}(\phi)$ such that these moments are matched exactly. Without incorporating reinsurance, $\tilde{\psi}(\phi)$ can be computed analytically.

$$\bar{C}_f^{obs} = e^{\psi_f} \frac{1}{\sum_{j \in J^f} S_j} \sum_{j \in J^f} \int_i S_{ij} e^{\phi_1 AV_{jm} + \phi_2 Age_i + \phi_3 r_i^{HCC}} dF(i)$$

$$\tilde{\psi}_f(\phi) = \log \left(\frac{1}{\sum_{j \in J^f} S_j} \sum_{j \in J^f} \int_i S_{ij} e^{\phi_1 AV_{jm} + \phi_2 Age_i + \phi_3 r_i^{HCC}} dF(i) \right) - \log(\bar{C}_f^{obs})$$

When incorporating reinsurance, the parameters ψ can no longer be separated from ϕ because they interact in determining how much reinsurance an individual receives. Instead, $\tilde{\psi}$ can be found by iteration.

$$\tilde{\psi}_f^{n+1} = \tilde{\psi}_f^n + \left[\log \left(\frac{1}{\sum_{j \in J^f} S_j} \sum_{j \in J^f} \int_i S_{ij} c_{ijm}^{rein}(\psi_f, \phi) dF(i) \right) - \log(\bar{C}_f^{obs}) \right]$$

Without any reinsurance, this iteration method gives the analytic result at $n = 1$ given any feasible starting point, ψ^0. The reinsurance payments are not particularly sensitive to ψ which affects average payments and have less affect affect on the tails targeted by reinsurance. As a result, $\tilde{\psi}$ can typically be precisely computed with only a handful of iterations.

[22] I have also experimented with including moments on risk adjustment transfers for groups of firms, which does not substantially affect the results.

Method of Simulated Moments

I will write the moments as $d(\phi)$ to represent the remaining moments on the cost ratios by metal level, age, and risk, incorporating the predicted parameters of $\tilde{\psi}(\phi)$. $\hat{\phi}$ is estimated by minimizing, for a weighting matrix W,

$$\hat{\phi} = \text{argmin}_{\phi} \; d(\phi)' W d(\phi)$$

The minimum of the function is found using the non-gradient Neldermead methodology. I estimate $\hat{\phi}$ in two stages. In the first stage, I use the identity weighting matrix and obtain estimates of the variance of the moments, V. In the second stage, I use $W = V^{-1}$. Similar to the demand estimation, the moments do not necessarily apply to every observation of the data. I use the same procedure from Petrin (2001) to compute the variance of the moments (see section B.2).