

Market Power in the Presence of Adverse Selection

Conor Ryan *

University of Minnesota, Department of Economics

Keywords: Insurance, Competition, Adverse Selection

Abstract

In the presence of adverse selection, firms have an incentive to offer products that appeal to low-cost consumers and lead high-cost consumers to purchase insurance elsewhere. This incentive is worst in highly competitive markets and absent in a monopoly. Through this interaction, regulations to address adverse selection and competition policy should be considered complements. This paper estimates a model of imperfect competition in the non-group health insurance market that directly characterizes the relationship between adverse selection and market power using novel choice data from a private online marketplace and a risk prediction model. The largest welfare cost comes from high markups. Current policies targeting adverse selection are successful in reducing the welfare cost of extensive margin selection and inefficient sorting. However, these policies also increase the potential harm of mergers, and in the most concentrated markets, the firms recapture much of the added surplus through higher markups.

1 Introduction

The welfare costs of adverse selection and market power are well studied, but past research and policy making has given less attention to their interaction. To see the relationship, consider the incentive for an insurance firm to offer a product that appeals to low-risk consumers and leads high-risk consumers to purchase insurance elsewhere. This incentive problem is worst in highly competitive markets and absent in a monopoly. Through this interaction, regulations to address adverse selection and competition policy should be considered complements. This paper develops a model of welfare in the health insurance market that directly characterizes the relationship between adverse selection and market power. Market power has welfare benefits by reducing inefficient sorting of consumers among available plan options, a symptom of adverse selection. However, because market

*Email: ryan0463@umn.edu; I would like to express appreciation for the guidance and helpful comments from Tom Holmes, Amil Petrin, Naoki Aizawa, Stephen Parente, Roger Feldman, Vitoria Rabello de Casto, Sam Bailey, and Jordan Pandolfo, as well as all the participants of the Minnesota Applied Microeconomics Workshop.

power also carries the welfare cost of higher markups, the magnitude and net direction of the effects are an empirical question. I estimate the model for the non-group market using novel choice data from a private online marketplace and a risk prediction model to relate preferences to marginal cost. The largest welfare cost in the non-group market comes from high markups. Current policies targeting adverse selection are successful in reducing the welfare cost of extensive margin selection and inefficient sorting. However, these policies also increase the potential harm of lower competition via mergers. In the most concentrated markets, insurance firms recapture much of the added surplus through higher markups.

Adverse selection is a first-order concern in the non-group health insurance market, frequently called the “individual market.” The Affordable Care Act (ACA) was passed in 2010 in response to high rates of uninsurance, limited insurance coverage, and frequent coverage denials in the market (Obama (2009)). Many of the regulations implemented by the law are in direct response to the symptoms of adverse selection identified in the literature (e.g., Cutler and Zeckhauser (2000), Van de Ven and Ellis (2000), Gruber (2008)). For instance, the individual mandate taxes individuals who do not purchase health insurance to mitigate extensive margin selection. The ACA also implemented a risk adjustment policy that taxes firms with healthier-than-average consumers and subsidizes those with above-average risks in order to mitigate inefficient sorting.

The ACA also included efforts to increase the degree of competition in the market.¹ However, the local insurance markets regulated by the ACA still vary widely in their market concentration. The largest firm in the non-group health insurance market has a market share of over 85% in five states and less than 33% in another five states. Since the enactment of the ACA, market concentration has increased, and some local markets are served by a single insurance firm.

I study market power in the presence of adverse selection using a model of the non-group health insurance market with strategic firms that compete in price with a fixed set of insurance products. Consumers make a discrete choice over the available options, which depends on the characteristics of the offers and an individual-specific health risk. The health risk of the consumer affects the costs of the firms, but firms cannot condition prices on this health risk, which creates an asymmetric information problem and adverse selection. The products are differentiated, which provides each firm with market power, i.e., an ability to charge a markup over marginal cost.

¹For example, the ACA provided start-up grants to new entrants and funded an online platform on which consumers could browse all available plans.

The total welfare loss in an equilibrium can be decomposed into three sources: (i) extensive selection, (ii) markups, and (iii) inefficient sorting. The first distortion is the classic welfare cost of adverse selection: the average expected cost of insured consumers is greater than the expected cost of the marginal consumer (Einav et al. (2010)). The second distortion, markups, further drives a wedge between price and marginal cost. This distortion is increasing in the degree of market concentration and exacerbates the extensive selection problem. The final distortion is the inefficient sorting of consumers among the available insurance options, a symptom of adverse selection on the intensive margin of which insurance product to purchase. This distortion is the result of an externality between firms—a firm can affect its competitors costs through the prices of its own products—and is decreasing in the degree of market concentration.

The welfare effect of market power can be represented through this decomposition. An increase in market power will decrease welfare through the markup channel and increase welfare through the sorting channel—a reduction in the cream-skimming incentive between firms. Because of these two channels, the net effect on both the individual and market level is ambiguous.

To estimate the model, I use new data on household health insurance choices in the non-group health insurance market made through a private marketplace (Ryan et al. (2021)). These data are unique in two respects. First, the data contain a substantial fraction of both low- and high-income consumers, in contrast to recent work using government-run marketplace data which tend to be predominantly low-income (ASPE (2016)). Second, the data span more than 100 local markets as defined by regulated rating areas, which allow me to estimate equilibrium outcomes in a diverse cross-section of market structure.

To identify the relationship between demand and cost, I use a novel approach that combines standard discrete choice demand techniques with moments that link demand to average costs via the Health and Human Services Hierarchical Condition Categories (HHS-HCC) risk prediction model.² Data that link non-group health insurance product choices with measures of health status are rare, and recent approaches rely on simulated variation in demand and average costs or the assumption that prices are set optimally as specified by the model (Tebaldi (2020)). 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, which does

²The HHS-HCC risk prediction model is used to administer the risk adjustment transfer system in the non-group market.

not require any optimality conditions.³

With the estimated supply and demand of health insurance, I measure the welfare costs of each distortion in equilibrium with and without the budget-neutral risk adjustment policy of the ACA. In the absence of a risk adjustment policy, the welfare costs of markups and inefficient sorting are of similar magnitudes—\$21.4 and \$20.9 per person per month, respectively. While the risk adjustment transfers do not optimally price the externality between firms, the policy succeeds in reducing the welfare cost of inefficient sorting to \$4.4 per person per month. The least concentrated markets see the largest benefits to total welfare where the potential cost of the distortion is also largest.

To demonstrate the complementarity between selection regulation and competition policy, I simulate a merger between Aetna and Humana, which was proposed in 2015 but blocked by the Justice Department. In the absence of a risk adjustment policy, the merger would have increased both consumer and producer surplus in the least concentrated markets by reducing the welfare cost of inefficient sorting.⁴ In the most concentrated markets, where this welfare cost is already small, the merger is harmful for consumers. In the presence of a risk adjustment policy, the merger has only a small effect on the sorting incentives among firms and all consumers are harmed by the increase in markups. The average harm to consumers from the proposed merger is larger in all markets under the risk adjustment policy than without it.

I evaluate equilibrium outcomes under four policy regimes: the baseline policy that has both risk adjustment and the mandate penalty in effect, two regimes in which one of the policies is no longer in effect, and one regime in which neither policy is in effect. 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 the most and least comprehensive insurance products, and the individual mandate penalty increases the total number of insured and lowers premiums. In the most concentrated markets, risk adjustment is ineffective and the individual mandate is harmful. Firms with substantial market power that respond to the increase in demand caused by the fee on being uninsured by increasing prices with larger markups rather than passing along the benefits of lower average costs to consumers.

³A similar risk adjustment system exists for Medicare (CMS-HCC), which can be more easily observed and has been used in a similar demand specification (Aizawa and Kim (2018), So (2019)).

⁴This conclusion supports the theoretical results of Veiga and Weyl (2016) that suggest that welfare could be U-shaped in market concentration.

Relation to the Literature

This paper makes three main contributions. First, I provide a model and intuition for the trade-off between two sources of inefficiency—markups and inefficient sorting—in markets with adverse selection. I build 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. (2019)) 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 U.S. health insurance markets are highly concentrated, there has been less focus in the literature on the effects of market power on adverse selection and policy design. Veiga and Weyl (2016) show in a theoretical model that a monopolist has an optimal sorting incentive when choosing the quality of a single product offering. This paper shows an analogous result in multi-product markets with fixed qualities. 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 (Mahoney and Weyl (2017), Veiga and Weyl (2016), Lester et al. (2019)). This paper takes these ideas to an empirical setting with multiple differentiated products in the non-group health insurance market, a focus of the policy debate around adverse selection and competition. This paper also builds on Geruso et al. (2018) and Saltzman (2021)—which evaluate the relationship between intensive and extensive margin selection—by introducing the relationship between these welfare costs and market power.

Second, I build on a literature that uses structural models of differentiated products to analyze the welfare impacts of policies addressing adverse selection and market concentration in health insurance markets. This draws from a large literature on estimating the demand for insurance (Gruber and Poterba (1994), Town and Liu (2003), Marquis et al. (2004), Handel and Kolstad (2015), Handel et al. (2019), Geruso (2017), DeLeire et al. (2017), Frean et al. (2017), Drake (2019), Ryan et al. (2021)). There is a growing literature on evaluating policies in regulated health insurance markets with a model of imperfect insurance competition (Miller et al. (2019), Jaffe and Shepard (2020), Shepard (2016), Tebaldi (2020), Ericson and Starc (2015), Starc (2014), Saltzman (2021)), and a related literature that studies health insurance firms' specific mechanisms and incentives to engage in risk selection (Cao and McGuire (2003), Brown et al. (2014), Newhouse et al. (2015), Newhouse et al. (2013), Aizawa and Kim (2018), Decarolis and Guglielmo (2017), Geruso et al. (2019)).

In addition to providing new evidence on the demand for health insurance, I 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 in which the data are available, this relationship can be identified through observing measures of health status (Aizawa and Kim (2018), So (2019), Shepard (2016), Jaffe and Shepard (2020)). However, these data are 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 (2020)). This paper does not assume optimality and instead combines demand data with cost and risk moments by applying the HHS-HCC risk prediction model to the Medical Expenditure Panel Survey (MEPS).

There is a substantial empirical literature on the effects of competition (Cutler and Reber (1998), Town (2001), Dafny et al. (2012)). Much of the recent work in this area 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 (Capps et al. (2003), Gowrisankaran et al. (2015), Ho and Lee (2017)). 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. 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 comprehensive plan offerings may be small and even positive, before accounting for bargaining effects.

Finally, this paper contributes to a large body of literature that studies the effects of policies designed to address adverse selection. 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 (2019), Ericson and Starc (2015), 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. Hackmann et al. (2015) find that the Massachusetts health reform led, in general, to lower markups, though they do not separate the effect of the mandate penalty from the other concurrent market reforms. This paper extends this insight by demonstrating how the effect of the individual mandate depends on the 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 (2020), Brown et al. (2014), Aizawa and Kim (2018), Layton (2017), Saltzman (2021),

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. This paper explicitly characterizes the incentive among strategic firms that leads to inefficient sorting and assesses the degree to which the risk adjustment policy implemented by the ACA mitigates this incentive.

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 decompose the sources of welfare loss in the non-group insurance market, and in section 7, I estimate the policy counterfactuals.

2 Model

2.1 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 products are assumed to be 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, uninsurance, also has a price in the form of a penalty for being uninsured. 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 , $P(\tau, p_j)$, the amount of which depends on the household's age, income, and size due to age-rating regulation and income-based subsidies. Importantly, the premium is not conditional on any direct measure of health status. I will write $P_j(\tau)$ for the household-specific premium for product j . There are a continuum of households in each market distributed by $F_m(\tau, \theta)$. Households also have idiosyncratic

preferences over products $\{\varepsilon_{ij}\}_{j \in J^m}$, 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

$$v_{ij} = u(P_j(\tau), X_j, \xi_j; \theta_i) + \varepsilon_{ij} \quad (1)$$

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

$$S_j^{\tau\theta}(p_m) = \frac{e^{u_j^{\tau\theta}}}{e^{u_0^{\tau\theta}} + \sum_{k \in J^m} e^{u_k^{\tau\theta}}} \quad (2)$$

where $p_m = \{p_j\}_{j \in J^m}$.

Firms

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

$$\Pi^f = \sum_{m \in M^f} L^m \sum_{j \in J^{mf}} \int_{\tau, \theta} S_j^{\tau\theta}(p_m) \left(P_j(\tau) - C_f(X_j, \tau, \theta) - T_j(p_m) \right) dF_m(\tau, \theta), \quad (3)$$

where L^m is the size of market m , and $C_f(X, \tau, \theta)$ 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, $T_j(p_m)$ represent risk adjustment transfers to firms that depend on the equilibrium outcome of the market via equilibrium prices.⁵ The risk adjustment transfers take the form

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

These transfers are positive and represent a tax on a particular insurance plan when the

⁵Other 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.

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

$$P_j(\tau) = p_j A(\tau) - B(p^{2lps}, \tau)$$

where $A(\cdot)$ and $B(\cdot)$ are price modifiers 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^{2lps} .

For the main results of this paper, I will assume that the firm does not consider the effect of its own prices on the subsidies in equilibrium, i.e., $\frac{\partial B(p^{2lps}, \tau)}{\partial p^{2lps}}$, for the firm which offers the second lowest-priced Silver plan. This isolates the mechanisms of interest from firms' strategic interactions with government spending. In Appendix Section D.1, I relax this assumption in an extension.

2.2 Market Structure and Adverse Selection

Consider the problem of a constrained social planner that chooses prices subject to a constraint on promising a total profit of $\bar{\Pi}$ to the insurance industry.⁶

$$\begin{aligned} & \max_{\{p_j\}_{j \in J^m}} \int_{\tau, \theta} CS^{\tau\theta}(p_m) dF(\tau, \theta) & (4) \\ \text{such that } & \int_{\tau, \theta} \sum_{k \in J^m} S_k^{\tau\theta} (P_k(\tau) - c_k^{\tau\theta}) dF(\tau, \theta) \geq \bar{\Pi} \end{aligned}$$

where

$$CS^{\tau\theta}(p_m) = E_{\varepsilon_i} \left[\max_{k \in J^m} u(P_k(\tau), X_k, \xi_k; \theta_i) + \varepsilon_{ik} \right]$$

⁶The results of this section do not depend on the specifics of a demand or consumer surplus specification, only that $\partial CS^{\tau\theta}(p_m) / \partial p_j = -S_j^{\tau\theta}(p_m)$, which holds under much less restrictive assumptions on demand (Small and Rosen (1981)).

For exposition of the welfare mechanisms, I make several simplifications. Suppose that prices are constant across all consumers, $P_j(\tau) = p_j$. Consumers may receive subsidies, but those subsidies are fixed. I will drop the market subscript m , unless necessary. Consider a single product in this market, j . The consumer welfare maximizing price, given some guaranteed industry profit, solves

$$p_j + \frac{\lambda - 1}{\lambda} \frac{S_j}{S'_j} = \frac{E \left[\frac{\partial S_j^{\tau\theta}}{\partial p_j} C_j^{\tau\theta} \right]}{S'_j} \quad (5)$$

where λ is the Pareto welfare weight of profit. This is similar to the standard first order condition of a firm, where the difference between price and marginal cost is a markup term given by $\frac{S_j}{S'_j}$. In this case, the markup term is modified by the Pareto weight, λ . Under utilitarian welfare with equal weight between consumers and producers ($\lambda^W = 1$), the markup term vanishes and the welfare maximizing price is equal to the average cost of the marginal consumer.

In markets with adverse selection, the marginal cost curve is downward sloping in quantity and always below average cost, which leads to the first source of inefficiency: *extensive selection*. Firms are not covering costs at λ^W , and the break-even price, p_j^{Ext} , solves equation 5 with $\lambda^{Ext} > \lambda^W = 1$. The resulting welfare cost is the classic welfare cost of adverse selection in which some consumers with a willingness to pay that exceeds their marginal cost do not purchase insurance (Einav et al. (2010)).

The second source of inefficiency is standard to any market in which firms have market power: *markups*. Market power allows firms to extract more surplus in the form of profit, and the welfare maximizing price conditional on granting this additional profit can be expressed as a solution to 5 with $\lambda^{Mkup} > \lambda^{Ext}$. A competitive equilibrium in which a monopolist selects its prices to maximize profit is represented as $\lambda \rightarrow \infty$.

If there is only a single product in the market, these two mechanisms are a full description of difference between the competitive equilibrium and the welfare maximizing allocation. In the case of multiple differentiated products, there is a third welfare cost: *inefficient sorting*. In setting the prices of its own products, a firm alters the costs of its competitors by affecting the distribution consumers (and expected costs). This externality between firms leads to inefficient sorting of individuals across the plans available and inefficiently high total costs in the market.

In a multi-product setting, the welfare maximizing price of product j conditional on a given level of industry profit is given by

$$p_j + \frac{\lambda - 1}{\lambda} \frac{S_j}{S'_j} = \underbrace{\frac{E \left[\frac{\partial S_j^{\tau\theta}}{\partial p_j} c_j^{\tau\theta} \right]}{S'_j}}_{\text{Private Marginal Cost}} - \underbrace{\sum_{k \neq j} \frac{E \left[\frac{\partial S_k^{\tau\theta}}{\partial p_j} (p_k - c_k^{\tau\theta}) \right]}{S'_j}}_{\text{Sorting Externality}} \quad (6)$$

The sorting externality term represents the effect of a change in the price of product j on the earnings of other products in the market. If there is no selection concern, a price increase for one firm's product has a positive externality on other products through the diversion of profitable consumers. In an environment with adverse selection, this term has an ambiguous sign because some consumers have $c_k^{\tau\theta} > p_k$.

This term is closely related to why change a change in market structure has an ambiguous effect on equilibrium prices. In merger analysis, the sorting externality is referred to as the upward-pricing-pressure that would result from a merger-to-monopoly (Farrell and Shapiro (2010)). It characterizes the additional incentive for firms to raise prices through the reduced opportunity cost that a consumer will purchase the product of a competitor. In the presence of adverse selection, it is possible that some portion of diverted consumers are not profitable. The sign of the externality term between the merging products determines whether a merger between those products is likely to increase or reduce the price of each of the products.

2.3 Policy Effects

In this paper, I focus on two policies to address adverse selection. The first policy is the individual mandate penalty, a tax on being uninsured. The intention is to increase the demand for insurance, thereby reducing the marginal cost of providing insurance and lowering the price. However, in the presence of market power, an increase in the demand for insurance has two effects: a decrease in the marginal cost and an increase in the markup. In a single product market where everyone has a marginal benefit of insurance greater than the marginal cost, the total welfare will unambiguously increase as the portion of the population that purchases insurance goes up. However, the increasing markup will limit gains for consumers and potentially reduce consumer welfare through higher equilibrium prices. The effect of mandate penalty in the context of possibly increasing markups was the

focus of Hackmann et al. (2015), which studied the implementation of a similar mandate penalty in Massachusetts prior to the ACA.⁷

The second policy is a risk adjustment policy. This policy is a transfer between firms that is equal to the difference between the firm's own average cost and the implied average cost of the firm if it were to insure the same risk balance as the market as a whole (Pope et al. (2014)).⁸

$$T_j(P) = \frac{E[\sum_k S_k^{\tau\theta} c_k^{\tau\theta}]}{E[\sum_k S_k^{\tau\theta}]} - \frac{E[S_j^{\tau\theta} c_j^{\tau\theta}]}{E[S_j^{\tau\theta}]}$$

Pooled Cost
Average Cost

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

$$p_j^* + \frac{S_j}{S_j'} = \Psi_j \frac{E\left[\left(\sum_k \frac{\partial S_k^{\tau\theta}}{\partial p_j}\right) c_j^{\tau\theta}\right]}{\sum_k \frac{\partial S_k}{\partial p_j}} + (1 - \Psi_j) \frac{E[\sum_k S_k^{\tau\theta} c_k^{\tau\theta}]}{\sum_k S_k} \quad (7)$$

where,

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

There are two important features of equilibrium under risk adjustment. First, the transfers adjust the private incentive of the firm according to how the marginal cost of its products deviates from the market-wide average cost. The policy-induced incentive is not the optimal sorting incentive in equation 6 that penalizes or reward firms based on the profitability of their marginal consumers.

Second, this particular policy converges to the firm's own private incentive as the market share of a particular product increases or if one firm merges with others in the market. The policy follows the importance of intensive selection by fading out with market concentration.

⁷Geruso et al. (2018) also show that an increase in market-wide demand can increase the degree of inefficient sorting.

⁸The implemented policy has to approximate this transfer using a risk-scoring system, but I will assume for theoretical simplicity that the regulator has full information about consumer risk.

3 Non-group Market Data

The non-group insurance market is the only source of health insurance for any individuals or households that do not receive an offer for insurance through their employer or a government program. Consumers can purchase insurance by contacting an insurance firm directly, visiting the government-run marketplace, or shopping for insurance through a third-party marketplace. 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.⁹

Since the implementation of the ACA, all insurance products in this market must fit within one of five categories known as “metal” levels: Catastrophic, Bronze, Silver, Gold, and Platinum, listed in increasing level of generosity. Households (or individuals) may purchase products that are offered in their local rating area for a price that depends on the size and age composition of the household, the household income, and whether or not the members are smokers. Households that earn 100% of the federal poverty level (FPL) receive a subsidy that is sufficient for the household to buy the second-lowest price Silver plan in their rating area for roughly 2% of their household income. This subsidy declines non-linearly to 9.5% for households that earn 400% of FPL, and subsidies are zero for households that earn greater.¹⁰ Households that earn less than 250% of FPL receive additional subsidies to cover reduced cost-sharing. Insurance prices are adjusted by an age-rating factor for each member of the household which, in 2015, increases from 0.635 for children under the age of 21 to 3 for a 64 year old. Some states add additional premium increases of up to 50% for household members that smoke.

3.1 Choice Data

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

⁹Analysis 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.

¹⁰In recent years, California has extended subsidies to higher income households as well.

The data contain information on the age of the consumer, the first three digits of the consumers' zip code, the household's income, 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 oldest member of 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 21.¹¹ Appendix Section A.1 contains more detail on sample selection, missing data, and constructing the relevant choice set for each household.

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

The data from the private marketplace are a selected sample of all the consumers facing a particular firm. Using the same data set, Ryan et al. (2021) find that income is a primary determinant of driving selection into the private online market place. In order to create a sample of consumers that is representative of the consumer population facing firms in this market, I treat the choice data as a random sample conditional on subsidy eligibility and geographic market. Each observation from the choice data within a particular subsidy eligibility category and market is given an equal weight such that the weights sum to the size of the population as determined by the 2015 American Community Survey (ACS). The ACS also provides a sample of the uninsured population. For more detailed information on processing the ACS, see Appendix Section A.2.¹²

In Table 1, I summarize the data and compare it to other data on the non-group insurance market: the ACS and data reported by the Office of the Assistant Secretary for Planning and Evaluation (ASPE) at the U.S. 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

¹¹The choice data contains information on the premium paid for a subset of the observations. In combination with the base premium of the purchased product, the premium paid can be used to impute household composition. Using the median base premium in the selected firm and metal-level, I construct an imputed household age-rating measure. The correlation between this imputation and the more simple age-rating rule applied to the rest of the sample is 0.90. The results are robust to alternative assumptions about age rating.

¹²The weights do not significantly alter the price elasticity and risk preference estimates from demand estimation. They are important for how well the model predicts untargeted moments like aggregate insurance rates and the firm first-order conditions.

	ACS	ASPE	Private Marketplace	
			Un-weighted	Weighted
			<u>Age Distribution</u>	
Under 18	0.0%	9.0%	0.0%	0.0%
18 to 25	7.6%	11.3%	11.1%	11.4%
26 to 34	17.2%	17.5%	30.8%	29.1%
35 to 44	22.2%	16.8%	21.4%	20.1%
45 to 54	25.3%	20.9%	19.9%	20.5%
55 to 64	27.7%	23.3%	16.8%	19.0%
			<u>Income Distribution</u>	
Under 250% FPL	32.1%	76.1%	30.8%	43.0%
250% to 400% FPL	24.5%	15.4%	9.1%	13.4%
Over 400% FPL	43.4%	8.5%	60.1%	43.6%
			<u>Metal Level Market Shares</u>	
Catastrophic		1.1%	5.0%	3.6%
Bronze		24.2%	39.2%	36.0%
Silver		66.4%	41.8%	48.8%
Gold		6.6%	11.1%	9.4%
Platinum		1.7%	2.9%	2.2%

Notes: The table compares the weighted and unweighted distribution of consumers in the estimation data sample relative to other data sources on the non-group market. The age distributions reported are for the head-of-household with the exception of ASPE, which is the individual-level distribution.

Table 1: Data Description

type market shares reported by ASPE are weighted heavily towards Silver plans that have extra cost-sharing benefits at low incomes. While the private marketplace is tilted towards higher-income and younger households, the ACS weighting moves the demographic distributions and market shares closer to those in the other data sources. Ryan et al. (2021) investigate these relationships in more detail and show that the market shares, conditional on income and geography, are quite close to those reported by ASPE.

3.2 Cost Data

To identify the relationship between marginal cost and demand, the key feature of adverse selection, I use moments on consumer medical risk in both the demand and cost estimations. The 2015 Medical Expenditure Panel Survey (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 Center for Medicare and Medicaid Services (CMS), to the self-reported diagnoses to compute risk scores. For details on the processing of the MEPS data, see Appendix Section A.3.

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 consumers of different risk to value firms differently, I target the average risk score among groups of firms using data Medical Loss Ratio data.¹³ Firms are divided into four groups based on whether they are in the top 25th percentile of average risk, and whether they are large (high average value to consumers) or small (low average value to consumers). 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.5) and computing firm-level average risk (Appendix Section A.6).

In the cost estimation, I estimate marginal costs from the simulated distribution of the age and risk of consumers in each insurance product and a combination of moments on the relative costs of individuals by age and risk, the average cost of insurance product categories, and the average costs of each firm. The individual level moments come from MEPS (Appendix Section A.3), the product category level data come from rate filings to state insurance regulators (Appendix Section A.4), and the average firm-level costs come from the Medical Loss Ratio data (Appendix Section A.5).

¹³These additional moments may capture consumer value of broad networks, for example.

4 Demand

4.1 Empirical Specification

In the 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 $\theta_i = (\gamma_i, \alpha_i, \beta_i)$.

Geographic markets are defined as the rating areas, and all products are aggregated to the firm-metal-market level. For example, one product is a Bronze plan offered by Aetna in the Georgia's 1st 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

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

where $B(y)$ 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. The preference over observed characteristics, β_i , depends on a households risk score, r^{HCC} .

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

Risk is treated as an unobserved household characteristic. Risk scores are distributed ac-

cording to a distribution that can depend on household demographics, Z_i .¹⁴

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

4.2 Risk Score Distribution

The risk scores in the demand model correspond to the output of the 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. 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 condition categories based on diagnoses codes.

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

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. Because all independent variables in the regression take a value of either 1 or 0, the risk score is equal to the sum of all coefficients that apply to a particular individual.

$$r_{it} = \underbrace{\sum_g \gamma_{tg}^{age,sex} \text{Age}_g \text{Male}_g}_{r_{it}^{dem}} + \underbrace{\sum_{g'} \gamma_{tg'}^{HCC} \text{HCC}_{g'}}_{r_{it}^{HCC}}$$

Unless specifically noted, r_i^{HCC} will refer to the Silver plan HCC risk-score component and represent standard a measure of health status across all product types.

Parametric Distribution

The distribution of risk scores, \hat{G} , is estimated 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

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

file. The publicly available data only list three-digit diagnoses codes, rather than the full five-digit codes. I follow McGuire et al. (2014) and assign the smallest five-digit code for the purpose of constructing the condition categories and matching the HHS-HCC risk coefficient.¹⁵

In the data, a majority of individuals have no relevant diagnoses, i.e., $r_i^{HCC} = 0$.¹⁶ In order to match this feature of the data, the distribution combines a discrete probability that an individual has a non-zero risk score and a continuous distribution of positive 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 income categories, above and below 400 percent of the federal poverty level.

Table 2 displays the moments of the risk score distributions for each metal level in the data. Figure 1 compares the risk distribution in the MCF with the simulated risk distribution in the estimation sample.

Age	Income (% of FPL)	$\delta(Z_i)$	Bronze		Silver		Gold		Platinum	
			$\mu(Z_i)$	σ	$\mu(Z_i)$	σ	$\mu(Z_i)$	σ	$\mu(Z_i)$	σ
≤ 45	$\leq 400\%$	0.16	0.31	1.32	0.47	1.07	0.57	0.96	0.66	0.91
	$> 400\%$	0.13	0.40	1.32	0.53	1.07	0.61	0.96	0.71	0.91
> 45	$\leq 400\%$	0.33	0.71	1.32	0.80	1.07	0.86	0.96	0.95	0.91
	$> 400\%$	0.24	0.65	1.32	0.74	1.07	0.81	0.96	0.89	0.91

Notes: This table displays three aspects of the distribution of HHS-HCC risk scores in the 2015 Medical Conditions File of the MEPS. The first column displays the portion of risk scores that are positive for four categories divided by age and income. The next columns display the mean and variance of the log of the risk score for each metal-level specific risk score. The mean depends on these same demographic groups, and the variance is calculated across the whole population.

Table 2: Parametric Distribution of Risk Scores

¹⁵For example, I treat a three-digit code of '301' as '301.00'. McGuire et al. (2014) find that moving from five-digit codes to three-digit codes does not have a large effect on the predictive implications for risk score estimation. In this case, there is measurement error as the model used was originally estimated on 5-digit codes.

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

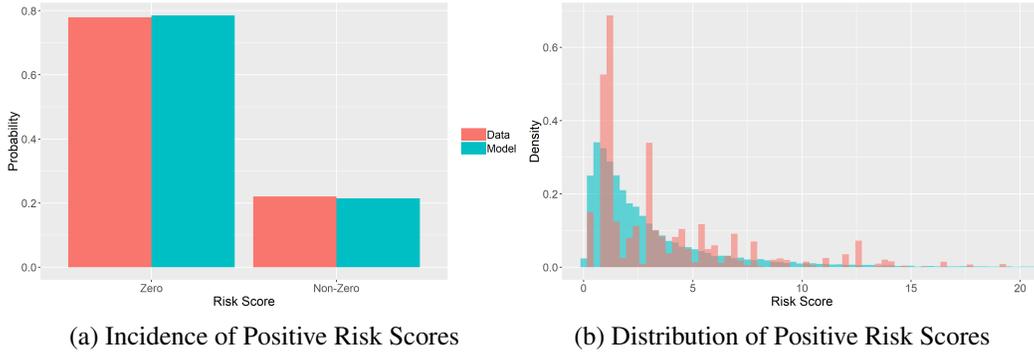


Figure 1: Risk Score Distribution Model Fit

Note: The data distribution comes from applying the HHS-HCC risk prediction methodology to the distribution of self-reported diagnoses in the 2015 Medical Conditions File of the MEPS. The model distribution comes from predicting the distribution risk scores in the ACS sample. In both cases, the risk score of the Silver metal-level is displayed.

4.3 Estimation

This model has two primary identification concerns. First, a plan premium's 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 (2020)). 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 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, because 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 (2002)).

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.¹⁷

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 1,000 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 parameters: $\theta_r = (\gamma_r, \{\beta_r^k\})$ 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,

¹⁷I 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.

products of firms that are large and non-risky, products of firms that are small and risky, and products of firms that are large and risky. (For the data definitions of large and risky, see Section 3.2). When computing the risk scores that are comparable to data moments, the simulated moments must also account for metal-level differences in product-specific risk scores. To approximate the relationship between metal level and risk score, the risk distributions of each metal level are assumed to be perfectly correlated. An individual is assigned a set of product-specific risk scores, $\{r_{ij}\}$, such that the risk scores of each product occupies the same point in the CDF of the metal-level-specific risk-score distribution. In the data, the correlation between the risk scores of any two metal levels is at least 0.998.

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

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

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, the algorithm finds the best fit parameters that maximize 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

Table 3 presents the results from the demand estimation. The GMM specifications are supplemented with maximum-likelihood specifications that do not target risk-score moments. The maximum-likelihood specifications arrive at similar results as the GMM specifications, with the exception of a larger estimate of the price sensitivity of families. The maximum-likelihood estimation cannot identify different preference parameters that relate to the unobserved risk score without additional moments. As a result, it includes only an unobserved preference for actuarial value that depends on the risk score distribution and finds a stronger relationship between risk and willingness to pay for coverage. The discrepancy appears for

	Maximum Likelihood		GMM		
	(LL-1)	(LL-2)	(GMM-1)	(GMM-2)	(GMM-3)
<i>Premium</i>	-1.46	-1.26	-1.31	-1.25	-1.23
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Age 31 - 40	0.24	0.24	0.24	0.26	0.26
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Age 41 - 50	0.34	0.29	0.31	0.31	0.30
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Age 51 - 64	0.69	0.55	0.60	0.55	0.54
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Family	-0.17	-1.13	-0.09	-0.07	-0.07
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Subsidized	0.09	0.21	0.12	0.19	0.20
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
AV	4.40	9.36	4.23	9.35	9.48
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
<i>Risk Preference</i>					
Constant			0.06	0.09	0.08
			(0.01)	(0.00)	(0.00)
AV	1.19	0.90	0.62	0.60	0.61
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
High Risk Firm			0.01	0.01	0.00
			(0.01)	(0.00)	(0.00)
Small Firm			-0.06	-0.05	-0.06
			(0.00)	(0.00)	(0.00)
Small & High Risk			0.05	0.06	0.07
			(0.01)	(0.01)	(0.02)
Fixed Effects					
Age, Fam., Inc.	Y	Y	Y		Y
Firm	Y				
Firm-Market			Y	Y	
Firm-Category				Y	
Firm-Mkt-Cat.		Y			Y

Notes: 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. Premiums are in thousands of dollars per year.

Table 3: Demand Estimation Results

two reasons. First, identification comes only from substitution patterns, which could suggest that there is larger preference variation that is not related purely to health risk. Second, the restriction of a single dimension of heterogeneity puts more emphasis on the actuarial value parameter. Together, these results suggest that substitution patterns in the data are consistent with health risk being an important, unobserved aspect of demand. The additional moments on risk score provide additional identification, allow for more detailed heterogeneity in demand, and allow for better targeting of important aspects of the market that are relevant for counterfactual simulations, such as the average risk level of a firm.

The specifications use increasingly-flexible fixed effects to control for unobserved quality. When the fixed effects account for product category, the estimated willingness to pay for additional actuarial value increases substantially. This increase indicates that consumers are more willing to switch from Bronze to Silver or Gold to Platinum metal levels, but may be less willing to jump from Silver to Gold plans. In the analysis that follows this section, the GMM-2 is used as the preferred specification. For more fine levels of fixed effects, the simulation sample, based on the ACS, has consumers with options that do not appear in the choice data. Due to sample size, the estimation data may not have low-income households that are offered a Silver plan with a cost-sharing reduction in a particular rating area, for instance. In that case, the fixed effect for that product category would not have a value. The GMM-2 specification is used to avoid dropping these observations. The specification GMM-3, which estimates the Cartesian product of the fixed effects in GMM-2, has nearly identical results.

In the preferred specification, the mean consumer willingness to pay for a 10% increase in the actuarial value of an insurance plan is \$98.6 per month. This actuarial increase is roughly equivalent to switching from a Bronze plan to a Silver plan. The average price difference to consumers between Bronze and Silver plans is about \$62 per month. There is substantial variation in willingness to pay. The 10th percentile of willingness to pay is \$73.6 per month, and the 90th percentile is \$141 per month. A \$10 increase in the premium of a product decreases enrollment by an average of 9.1%, and a \$10 increase in the monthly premium of all insurance products decreases the probability that the average consumer will purchase insurance by 0.4%. These semi-elasticities are similar to other estimates in the literature (Tebaldi (2020), Saltzman (2019)).

5 Marginal Cost

5.1 Empirical Model

The average cost of covering a particular household with a particular insurance product is estimated through Method of Simulated Moments (MSM) using moments on average firm costs and health care expenditures by age and risk. This method does not require the assumption that firms are playing optimal strategies according to the specification of 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_i^{HCC} is the risk score of household. This specification assumes that the identically and independently distributed errors in the cost function, ω_{ijm} , are orthogonal to the preference draws in the demand estimation.

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

This assumption implies that the only mechanisms through which cost and preferences are correlated are through age and risk scores.¹⁸ If this assumption is violated and the remaining endogeneity is consistent with adverse selection, 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

¹⁸An 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. Doing so 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.

¹⁹For illustration, suppose I estimate $\hat{\phi}$ to solve for a single product and single observable type,

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

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.

attribution leads to conservative conclusions about the implications of adverse selection.

Reinsurance

In 2015, the ACA implemented a transitional reinsurance program, which mitigates a portion of the liability to insurance firms of very-high-cost enrollees. This policy was important in limiting the amount of realized adverse selection facing insurance firms and is included in cost estimation in order to match the post-reinsurance average firm costs. The federal government covered 45% of an insurance firm's annual liabilities for a particular individual that exceeded an attachment point, $\underline{c} = \$45,000$, and up to a cap, $\bar{c} = \$250,000$. For an individual with a cost c_{ijm} , the insurance firm is liable for the cost c_{ijm}^{rein} under the reinsurance policy.

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

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 (Appendix Section A.3). For clear identification of costs by age separate from risk score, 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 five-year age brackets for adults between 25 and 64 years old to the average covered expenditures of adults between 20 and 24 years old. 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 Appendix Section A.4 for more detail on the data.

When simulating moments that match data from the insurance firm rate filings, I use

	(GMM-1)	(GMM-2)
Age	0.47 (0.02)	0.47 (0.01)
Risk	0.10 (0.00)	0.10 (0.00)
Actuarial Value	4.20 (0.04)	3.84 (0.05)
State-Firm	Y	Y

Note: This table displays the estimates of the marginal cost function. The specifications correspond to the different demand estimation specifications that are used to simulate the moments. Marginal costs are not estimation for the final specification, GMM-3, since this specification cannot be used in counterfactual analyses.

Table 4: Cost Estimation Results

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 MSM to obtain the efficient weighting matrix. The estimated demand parameters are used to simulate the distribution of consumer age and risk throughout products in each market, using ACS data as the population of possible consumers (see Appendix Section A.2). 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. The estimation suggests a substantial amount of adverse selection. One measure of adverse selection is the relationship between consumer price elasticity and consumer cost. The problems associated with adverse selection will be more severe if the covariance is strong. These estimates imply that the mean cost in the lowest decile of semi-elasticity is \$84 per month and the mean cost in the highest decile of semi-elasticity is \$850 per month. The means of the top and bottom decile of costs, without considering elasticities, are \$35 and \$1200 per month, respectively. These results suggest that consumer semi-elasticity explains a substantial amount of the variation in expected consumer costs and that adverse selection is an important feature of this market.

	Data	Model Fit	
		(GMM-1)	(GMM-2)
Age ($r^{HCC} = 0$)			
18 - 24	1.0	1.0	1.0
25 - 29	1.34	1.29	1.32
30 - 34	1.44	1.62	1.65
35 - 39	2.08	2.48	2.46
40 - 44	2.98	2.31	2.32
45 - 49	1.74	2.91	2.90
50 - 54	3.49	3.18	3.26
55 - 59	2.98	3.90	3.99
60 - 64	3.57	3.96	4.11
Risk			
$r^{HCC} = 0$	1.0	1.0	1.0
$r^{HCC} > 0$	3.57	3.35	3.30
Metal Level			
Bronze	1.0	1.0	1.0
Silver	2.28	3.11	2.56
Gold	3.80	3.02	2.93
Platinum	4.28	8.13	7.62

Note: This table displays the targeted and estimated cost ratios that are used to identify the marginal cost estimation. 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. Marginal costs are not estimation for the final specification, GMM-3, since this specification cannot be used in counterfactual analyses.

Table 5: Cost Estimation Fit of Cost-Ratio Moments

Table 5 presents 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 overestimates 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 much greater than the difference in average cost among the less comprehensive options (Silver and Bronze). The model also predicts Silver plan average costs that are very similar to Gold plan average costs, which are a result of high-cost selection into Silver plans among individuals who

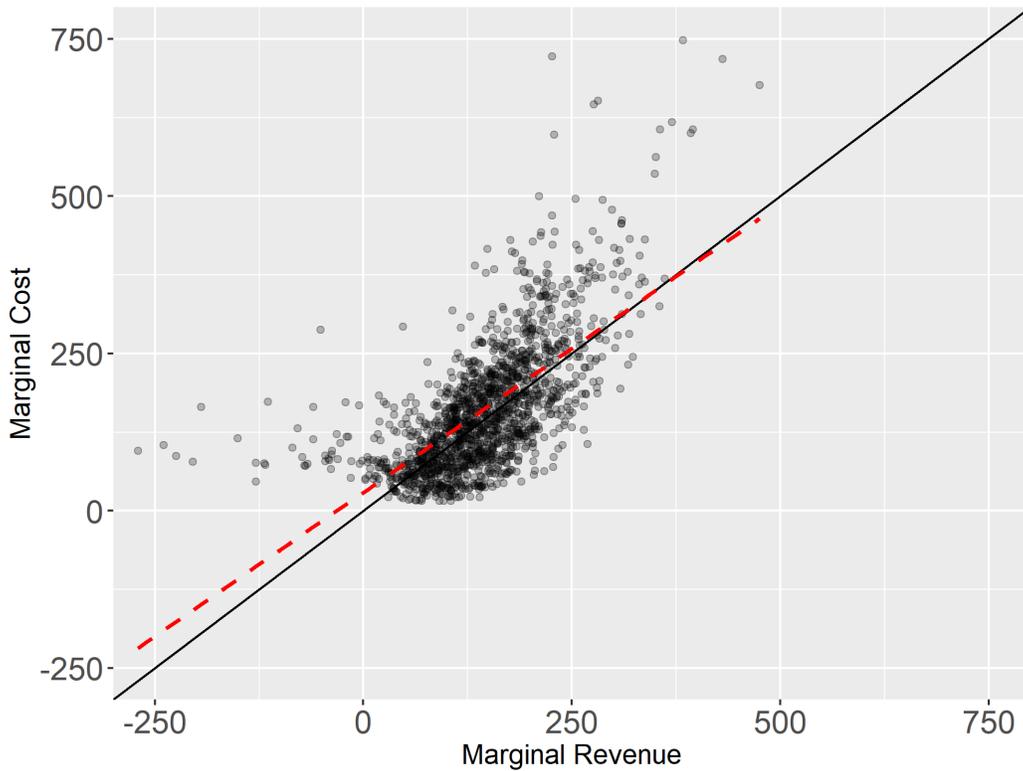


Figure 2: Marginal Revenue vs Marginal Cost in Baseline Model

Note: The product-level marginal cost and marginal revenue predicted by the estimated model are roughly equal on average. Each dot represents a product. The model does relatively well with products that are close to the mean marginal revenue and costs but struggles to fit the outliers.

receive cost-sharing reductions.

In estimating the parameters of demand and marginal cost, firms are not assumed to be setting prices to optimally maximize profit. Figure 2 plots the marginal revenue and marginal cost implied by estimated parameters under the baseline policy regime, which includes a mandate penalty, risk adjustment, and reinsurance. This plot can be used to assess whether the estimated model implies that firms are behaving as profit maximizers by setting prices to equate marginal revenue and marginal cost.

On average, the baseline model suggests that firms are setting marginal revenue close to marginal cost. 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 Welfare Analysis

In this section, I decompose the total welfare cost in the non-group insurance market. The decomposition is computed by comparing the total welfare for each market at four price vectors: (i) the total welfare maximizing price, p^A ; (ii) the consumer welfare maximizing price conditional that average revenue is at least average cost, p^B ; (iii) the consumer welfare maximizing price conditional that aggregate profit is at least equilibrium aggregate profit, p^C ; (iv) the competitive equilibrium price, p^D . All four prices are computed with and without the risk adjustment policy of the ACA.²⁰ Total welfare is measured as the sum of consumer surplus and aggregate profit, and each welfare cost is computed as follows:

$$\begin{aligned} \text{Extensive Selection Cost} &= SW(p^A) - SW(p^B) \\ \text{Markup Cost} &= SW(p^B) - SW(p^C) \\ \text{Intensive Selection Cost} &= SW(p^C) - SW(p^D) \end{aligned}$$

To cleanly demonstrate the relationship between welfare and market structure in the presence of adverse selection, consider the problem of a social planner that seeks to maximize consumer and producer surplus—possibly with a constraint on producer surplus—without consideration of the government’s budget constraint. The level of premium subsidies is held fixed at the pre-merger equilibrium level when computing the welfare maximizing prices and post-merger prices in order to cleanly characterize the relationship between welfare and market structure without changing the generosity of the subsidy. After holding the generosity level fixed, net government spending still depends on consumer choices, which in turn depends on prices. To avoid contending with the many possible objectives a government may have in mind when setting this policy, such as the social cost of uninsur-

²⁰The welfare maximizing price, p^A , and consumer welfare maximizing break-even price, p^B , do not depend on the risk adjustment policy. In both cases, the social planner optimally adjusts for all risks. The consumer welfare maximizing price that is constrained by equilibrium profit depends only on the risk adjustment policy through the policy’s effect on profit.

ance or aspects of consumer welfare that do not affect demand, the changes in government spending on subsidies and government revenue through the mandate penalty are ignored in this section. Appendix Section D.2 shows how this analysis changes if the planner considers a dollar of government revenue identical to a dollar of consumer surplus.

The only policy analyzed in this section is the budget-neutral risk adjustment policy. Section 7 contains a more general welfare analysis with both risk adjustment and the individual mandate.

Cross-sectional Analysis

Table 6 shows the decomposition of welfare costs across markets at different levels of concentration. The first welfare cost, extensive selection, is a symptom of adverse selection and does not depend on market structure.²¹ The welfare cost of markups increases as firms gain market power and increase the aggregate profit earned. This source of welfare loss is absent if firms are operating at break-even prices. The final source of welfare loss is attributable to inefficient sorting that results from selection on the intensive margin, as described in Section 2.2. This source of welfare loss is absent if the market is fully monopolized.

While the qualitative relationship of each source of welfare loss is evident from theory, the relative magnitudes and the implication for the welfare costs of market concentration are ambiguous. Using the estimated demand and supply of insurance in the non-group market, I measure the average per-person welfare cost of each source in each market. The data contain 109 markets, which span HHI values of 1,890 to 10,000—the value that indicates a monopoly.

The welfare cost of extensive margin selection is small, with an average of \$3.28 per person. The low magnitude is likely a result of the subsidy structure and individual mandate, which mitigate the number of individuals who go uninsured due to a high price of insurance. While the rate of uninsurance is still quite high, around 50%, these estimates suggest that this reflects consumer willingness to pay that is lower than the marginal cost of coverage.²²

In the absence of a risk adjustment policy, the average welfare costs of markups and inefficient sorting are large and of similar magnitude, \$21.4 and \$20.9 per person, respec-

²¹The degree of adverse selection in a market affects this first category of loss and could also affect the equilibrium market structure, though this is not the focus of this paper.

²²Finkelstein et al. (2019) find that consumer willingness to pay is substantially below the marginal cost of coverage in Massachusetts.

HHI Level	Ext. Sel.	Without Risk Adj.		With Risk Adj.		# of	
		Markup	Sorting	Markup	Sorting	Mkts	Mkt Size
< 3,000	2.49	16.8	28.0	21.7	2.10	23	2,099
3,000 - 3,500	3.22	17.4	31.4	23.4	4.13	20	1,662
3,500 - 4,000	3.69	19.7	15.2	25.4	4.58	8	1,295
4,000 - 4,500	3.07	17.2	17.6	23.2	3.92	12	1,956
4,500 - 5,500	3.04	17.4	19.9	23.7	3.59	13	1,723
5,500 - 6,500	4.24	23.8	19.5	27.6	6.58	14	2,222
6,500 - 9,000	3.13	32.7	12.8	32.5	10.5	14	1,011
9,000 - 10,000	3.33	68.6	5.14	69.5	5.20	5	351
All Markets	3.28	21.4	20.9	26.2	4.72	109	12,320

Notes: In the absence of a risk adjustment policy, the welfare cost of inefficient sorting and markups are of similar magnitudes and inversely correlated. This table displays the average welfare cost from extensive selection, markups, and inefficient sorting within each category of market concentration, measured by the HHI. It also shows the welfare costs under a the revenue neutral risk adjustment policy. All values in the left panel are in dollars per person per month. The averages are weighted by market population. The right panel displays the total number of markets and total population in thousands of the markets in each HHI category. Welfare is measured as the sum of consumer and producer surplus.

Table 6: Welfare Costs in the Cross-Section

tively. As shown in Section 2.2, the two welfare costs have opposite associations with market concentration. In markets with an HHI of less than 3,500, the average welfare cost of markups is less than the welfare cost of inefficient sorting. In the more concentrated markets, the welfare cost of markups dominates. However, inefficient sorting is an important source of welfare loss in even highly concentrated markets.

The risk adjustment policy primarily targets inefficient sorting and is successful at reducing the welfare cost of inefficient sorting to low levels, \$4.72 per person. While the risk adjustment policy does not optimally tax the sorting externality, it is still successful in reducing the welfare cost of inefficient sorting in the least concentrated markets to negligible levels.²³ The welfare cost of markups increases as a result of the risk adjustment policy, which reflects that firms are able to capture some of the gain in total surplus. The effectiveness of the risk adjustment policy declines with market concentration, where the welfare cost of inefficient sorting is already small.

²³The risk adjustment policy considered in this paper abstracts from the ways in which the administered taxes and subsidies differ from the conceptual transfers targeted by the policy. For instance, I do not consider the possibility that some risks may be unobservable to the government, or that the particular parameters of the policy may differ from the theoretical optimal parameters.

Merger Analysis

One channel to prevent firms from accumulating market power is through strict scrutiny of mergers. In this section, I simulate the effect of a proposed—and subsequently blocked—merger between Aetna and Humana on the non-group insurance market in Georgia, where both firms had substantial market share. This merger affects 10 local markets in Georgia. Prior to the merger, five of the markets have an HHI of less than 3800 and five of them have an HHI of greater than 3800. The resulting change in HHI ranges from 211 to 2430, with a median change of 839.

The welfare effect of the merger is presented in Table 7. The effects are split into the most and least concentrated markets pre-merger and the results are presented both with and without a risk adjustment policy. The intuition of these results follows the previous section. In the absence of a risk adjustment policy, less concentrated markets have a high welfare cost of inefficient sorting. And in these markets, consumers can be made better off as the result of an increase in market power that lowers the welfare cost of inefficient sorting. This can be seen in the first row of the Table 7, where both consumers and producers benefit from a merger.

When the welfare cost of inefficient sorting is low, either because the market is already highly concentrated or because of a risk adjustment policy, there is little to gain from a merger and consumers suffer from higher markups. One implication of these results is that the presence of a risk adjustment should be paired with an increased scrutiny of mergers in the health insurance industry.

7 Policy Analysis

This section studies equilibrium under four separate policy regimes. The baseline regime assumes 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.²⁴ While the level of subsidies does depend on price level in the market, as determined by the price-linking formula specified by the ACA, firms treat the subsidies as given rather than possibly linked to a plan that they

²⁴In 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.

	Change in Welfare		Source of Change		# of	
	Δ CW	Δ PS	Δ Markup	Δ Sorting	Mkts	Mkt Size
W/out Risk Adj.						
HHI<3800	0.20	1.4	-2.4	4.0	5	1,279
HHI>3800	-2.7	2.0	-2.6	1.9	5	365
With Risk Adj.						
HHI<3800	-3.1	0.7	-2.3	-0.16	5	1,279
HHI>3800	-6.0	2.0	-3.6	-0.42	5	365

Notes: In the absence of a risk adjustment policy, a merger in an un-concentrated market can increase both consumer and producer welfare. This table displays the average change in producer and consumer surplus as a result of a merger between Aetna and Humana in the state of Georgia. It also decomposes the total welfare change into a portion attributable to higher markups and a portion attributable to less (or more) inefficient sorting. All values in the left two panels are in dollars per person per month. The averages are weighted by market population. The right panel displays the total number of markets and total size of all markets in each HHI category.

Table 7: Effect of a Merger on Producer and Consumer Welfare

offer. Appendix Section D.1 shows how these results change if firms behave strategically with respect to the subsidy formula.²⁵ The risk adjustment policy included in this model is a “perfect” risk adjustment system, where the policy makers and firms both have full information about the relationship between risk and cost and an accurate measure of risk. It does not assume that risk is perfectly predictable but that all predictable risks are accounted for in the adjustment.

Next, I consider three policy regimes: i) no individual mandate, ii) no risk adjustment policy, and iii) neither the risk adjustment policy nor the individual mandate. The purpose of this analysis is twofold. The first is to demonstrate the heterogeneity of policy effects across different levels of market concentration. The second is to show how the effects of market concentration depend on the policy regime.

Cross-sectional Analysis

Table 8 shows 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. For brevity, only the results

²⁵Jaffe and Shepard (2020) show that this leads to lower effective elasticities and higher markups.

Panel A: Less than 3350 HHI					
	Observed	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y		Y	
Mandate	Y	Y	Y		
Prices					
Bronze	172	158	152	163	155
Silver	212	226	272	240	392
Gold	237	243	480	256	623
Welfare					
Consumers		48.2	41.6	71.5	67.5
Producers		47.9	50.8	35.6	39.1
Government		-15.7	-22.8	-36.6	-46.7
Insurance Rate	53.9	47.8	45.7	37.2	36.0
Panel B: Greater than 5200 HHI					
	Observed	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y		Y	
Mandate	Y	Y	Y		
Prices					
Bronze	166	214	211	204	199
Silver	230	311	321	306	316
Gold	265	331	401	326	396
Welfare					
Consumers		38.7	39.4	71.3	73.5
Producers		69.0	70.9	48.4	50.1
Government		-58.4	-66.8	-76.5	-85.4
Insurance Rate	46.7	43.0	43.6	35.2	36.3

Notes: In less concentrated markets, the Individual Mandate reduces average prices and the risk adjustment policy narrows the price difference between high and low coverage levels. In concentrated markets, the Individual Mandate leads to higher prices and the Risk Adjustment policy has little effect. All price levels represent the base premium per month, before age rating and subsidies are applied. Welfare calculations are in dollars per month, and the insurance rate is the percent of eligible consumers that purchase insurance. Model averages are computed using the baseline enrollment distribution. The observed averages are computed using the observed enrollment distribution. There are no controls for market compositional effects across the HHI categories.

Table 8: Cross-Section: Effects of Risk Adjustment and the Mandate Penalty

for the least concentrated third and most concentrated third of markets are displayed.

Panel A demonstrates the effects of risk adjustment and the individual mandate in the

most competitive markets. The intended complementarity of the individual mandate and risk adjustment is evident in this setting. 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 3 and 6 percent. However, if there were no risk adjustment, the individual mandate has a much larger effect. Comparing column (3) to column (1), average premiums for Silver and Gold plans are 44% and 30% 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.

The estimates of welfare in each policy regime tell a similar story. The risk adjustment policy provides a substantial increase in consumer welfare, with a small decline in profit earned by the firms. The individual mandate decreases consumer welfare overall, because the benefits through lower prices are not enough to offset the incidence of the tax. However, the insurance rate does increase substantially in competitive markets, the primary objective of the tax.

In Panel B, 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 comprehensive insurance categories, but higher prices among the cheaper Bronze plans. Moreover, the effect of the individual mandate is less dependent on the presence of risk adjustment. Similarly, consumer and producer welfare are both slightly lower under the risk adjustment policy, both with and without the individual mandate. Since market concentration provides better implicit risk sorting, the presence of a policy adjusting prices between plans is both less effective and less necessary. The individual mandate also has the opposite of the intended effect on prices than in competitive markets, leading prices to increase slightly. 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. While the insurance still increases significantly under the individual mandate, it is to a lesser degree than in the competitive markets. 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 (Panhans (2019)).

An important consideration is how well this model of profit-maximizing firms in

Nash-Bertrand competition matches the observed prices. Figure 2 shows that marginal costs are slightly higher than marginal revenue, and Table 8 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 outcome of a proposed merger between Aetna and Humana. The merger was initially proposed in 2015 and the Department of Justice subsequently sued successfully to block both mergers in court due to competitive concerns both in the non-group market, as well as in the employer-sponsored and Medicare Advantage markets. In the counter-factual, 10 local markets in Georgia would have been affected by the merger. The median pre-merger HHI in affected markets is 3739, and the median change in HHI is 839.

The results in Table 9 highlight the two first-order predictions of increased concentration in markets with adverse selection: Prices increase on average across all coverage levels, and the price spread between high and low actuarial value insurance decreases. These effects are present across all policy regimes. Bronze plans have the highest price increases, and increase in price across all policy regimes. Silver and Gold plans increase in price by a smaller amount, and decrease in price in the absence of either a risk adjustment policy. The increased market power that results from the merger decreases the welfare cost of inefficient sorting, which drives large spreads in prices between high and low actuarial value plans. Market power has a similar effect to a risk adjustment policy in this manner. The details of the welfare decomposition in the baseline model is presented in Section 6. These results suggest that the benefits to consumers are larger in the absence of a risk adjustment policy or the individual mandate.

The intuition behind in the individual mandate results hold in the merger analysis as well: the mandate penalty increases the market power of the participating firms. 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. In the absence of the individual mandate, consumers are better off on average after the merger and the number of insured remains roughly constant. This highlights the substitutability of market concentration and selection regulations like

	Baseline	(1)	(2)	(3)
Risk Adjustment	Y		Y	
Mandate	Y	Y		
Merging Party Prices				
Bronze	8.0%	8.7%	2.9%	5.0%
Silver	3.5	-2.8	0.3	-4.0
Gold	2.6	-2.3	-0.6	-2.9
All Other Prices				
Bronze	-0.2	-0.5	-0.3	-0.7
Silver	-0.2	-1.7	-0.3	-1.4
Gold	-0.7	-2.0	-0.5	-1.7
Welfare				
Consumers	-1.0	-3.1	1.0	-0.2
Producers	6.4	1.7	2.3	-1.2
Government	-11	3.8	-2.9	2.5
Insurance Rate	-1.6	-1.1	0.1	-0.1

Notes: The risk adjustment and individual mandate policies increase the harm from mergers. All values are percent changes that result from the simulated merger. The top panel shows changes for only the merging parties' products, and the middle panel takes all products into account. The last panel shows changes in consumer, producer, and government surplus. Positive values of change in government surplus correspond to a net reduction in federal spending. 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.

Table 9: Merger Analysis: Effect of Market Concentration

risk adjustment and the individual mandate.

8 Conclusion

This paper combines 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. The estimates are employed in a framework of imperfect competition over a fixed set of products to demonstrate how market power interacts with three sources of welfare loss in markets with adverse selection: i) extensive margin selection, ii) markups, and iii) inefficient sorting. Market power reduces total welfare through large markups but also reduces the welfare cost of inefficient sorting. This result is similar to the-

oretical results on the efficient sorting incentive of monopolists that choose product quality in a market with adverse selection (Veiga and Weyl (2016)). I show that under certain market conditions, an increase in market power can improve average consumer welfare through a reduction in the welfare cost of inefficient sorting that is not wholly offset by higher markups.

This paper shows how market structure interacts with two common policies intended to correct the distortions of adverse selection—the individual mandate, a penalty for being uninsured, and risk adjustment transfers—both across the cross-section of markets and by simulating a proposed merger. The policy analysis results in two main takeaways. First, the individual mandate increases the market power of the participating insurance firms. Markets that are very concentrated have higher prices under an individual mandate than without it. The increase in markups that results from the mandate outweighs the reduction in average cost caused by an increase in the percentage of consumers that purchase insurance. Similarly, mergers in markets with the individual mandate will have larger price effects than if the mandate were not in place.

Second, the risk adjustment transfer policy is not effective at altering the equilibrium price in highly concentrated markets. However, highly concentrated markets do not suffer from a large welfare cost of inefficient sorting, the primary target of a risk adjustment policy. Therefore, while the policy is less effective, it is also less necessary. Moreover, in the absence of a risk adjustment policy, a merger accomplish a similar goal of narrowing the price difference of high and low quality insurance products through increasing market power.

Taken together, these results suggest that competition policy and regulations targeting adverse selection should be considered complements. Markets that are more competitive require more attention to the problems caused by adverse selection, and markets with strong selection regulations should be kept as competitive as possible.

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 power.

References

- 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,” *American Economic Review*, mar 2018, *108* (3), 828–867.
- 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.
- 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.
- Cao, Zhun and Thomas G. McGuire**, “Service-level selection by HMOs in Medicare,” *Journal of Health Economics*, nov 2003, *22* (6), 915–931.
- Capps, Cory, David Dranove, and Mark Satterthwaite**, “Competition and market power in option demand markets,” *RAND Journal of Economics*, 2003, *34* (4), 737–763.
- Cutler, D. M. and S. J. Reber**, “Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection,” *The Quarterly Journal of Economics*, may 1998, *113* (2), 433–466.
- Cutler, David M. and Richard J. Zeckhauser**, “Chapter 11 The anatomy of health insurance,” in “Handbook of Health Economics,” Vol. 1, Elsevier, 2000, pp. 563–643.
- 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.
- Drake, Coleman**, “What Are Consumers Willing to Pay for a Broad Network Health Plan?: Evidence from Covered California,” *Journal of Health Economics*, 2019, 65, 63–77.
- Einav, Liran, Amy Finkelstein, and Mark R Cullen**, “Estimating welfare in insurance markets using variation in prices,” *The quarterly journal of economics*, 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.
- Farrell, Joseph and Carl Shapiro**, “Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition,” *The B.E. Journal of Theoretical Economics*, 2010, 10.
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard**, “Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts,” *American Economic Review*, apr 2019, 109 (4), 1530–1567.
- 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, “Demand heterogeneity in insurance markets: Implications for equity and efficiency,” *Quantitative Economics*, 2017, 8, 929–975.

– **and Timothy Layton**, “Upcoding: Evidence from Medicare on Squishy Risk Adjustment,” <https://doi.org/10.1086/704756>, jan 2020, 128 (3), 984–1026.

– , **Timothy J Layton, Grace McCormack, and Mark Shepard**, “Trade-offs between Extensive and Intensive Margin Selection in Competitive Insurance Markets,” 2018.

– , **Timothy Layton, and Daniel Prinz**, “Screening in contract design: Evidence from the ACA health insurance exchanges,” *American Economic Journal: Microeconomics*, may 2019, 11 (2), 64–107.

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, “Covering the uninsured in the United States,” sep 2008.

– **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 and Jonathan T Kolstad**, “Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare,” *American Economic Review*, 2015, 105 (8), 2449–2500.
- Handel, Benjamin R., Jonathan T. Kolstad, and Johannes Spinnewijn**, “Information frictions and adverse selection: Policy interventions in health insurance markets,” *Review of Economics and Statistics*, may 2019, 101 (2), 326–340.
- Hayashi, Fumio**, *Econometrics*, Princeton, New Jersey: Princeton University Press, 2000.
- Ho, Kate 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 Microeconomic Models,” *Review of Economic Studies*, 1994, 61 (4), 655–680.
- Jaffe, Sonia and Mark Shepard**, “Price-Linked Subsidies and Imperfect Competition in Health Insurance,” *American Economic Journal: Economic Policy*, aug 2020, 12 (3), 279–311.
- Kautter, John, 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,” <https://doi.org/10.1086/700730>, jan 2019, 127 (1), 338–377.
- Mahoney, Neale and E. Glen Weyl**, “Imperfect Competition in Selection Markets,” *The Review of Economics and Statistics*, oct 2017, 99 (4), 637–651.

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, 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,” 2019.

Newhouse, Joseph P, J Michael McWilliams, Mary Price, Jie Huang, Bruce Fireman, and John Hsu, “Do Medicare Advantage Plans Select Enrollees in Higher Margin Clinical Categories?,” *Journal of Health Economics*, 2013, 32 (6).

Newhouse, Joseph P., Mary Price, J. Michael McWilliams, John Hsu, and Thomas G. McGuire, “How much favorable selection is left in medicare advantage?,” *American Journal of Health Economics*, nov 2015, 1 (1).

Obama, Barak, “Obama’s Health Care Speech to Congress,” sep 2009.

Panhans, Matthew, “Adverse Selection in ACA Exchange Markets: Evidence from Colorado,” *American Economic Journal: Applied Economics*, 2019, 11 (2), 1–36.

Petrin, Amil, “Quantifying the Benefits of New Products: The Case of the Minivan,” *Journal of Political Economy*, jul 2002, 110 (4), 705–729.

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: Evidence from a Private Online Marketplace,” 2021.
- 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,” *Journal of Health Economics*, 2019, 63, 197–222.
- , “Managing adverse selection: underinsurance versus underenrollment,” *The RAND Journal of Economics*, jun 2021, 52 (2), 359–381.
- 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.
- So, Jaemin**, “Adverse Selection, Product Variety, and Welfare,” *Working Paper*, 2019.
- 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,” 2020.
- Town, Robert**, “The welfare impact of HMO mergers,” *Journal of Health Economics*, nov 2001, 20 (6), 967–990.
- and **Su Liu**, “The Welfare Impact of Medicare HMOs,” *The RAND Journal of Economics*, 2003, 34 (4), 719.
- Veiga, André and E. Glen Weyl**, “Product design in selection markets,” *Quarterly Journal of Economics*, may 2016, 131 (2), 1007–1056.

Online Appendix

A Data Processing

A.1 Choice Sets

The choice data contain only the ultimate choices made by the consumers, 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 in 2015. The data set is nearly a complete depiction of the market for the entire United States, but there are some markets in which some cost-sharing information is missing, or insurance firms are absent altogether.

A household's choice set depends on the age composition of its members and the household income. Since I observe only one age of the household, I use a simple rule to impute the age composition: any household with more than one individual contains two adults of the same age and additional persons are under the age of 21. For a subsample where I can infer the age composition based on their charged premium, this simple rule has a correlation with the inferred age composition of 0.9. The income information also contains some missing values. For subsidized consumers, income can be imputed from the observed subsidy value and the household size. I use this imputed income for subsidized consumers with missing income information. However, doing so is not possible for the consumers that do not receive a subsidy. I assume that those in the data without a reported subsidy amount have an income greater than the subsidy qualification threshold.

I restrict the analysis to markets in which I observe characteristics of the entire choice set and can be reasonably confident that the private marketplace 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 set in each market is large. The typical market has about 150 plans to choose from, and these plans do not necessarily overlap with other markets. Because I observe only a sample of choices, there are many plans that I do not observe being chosen. The lack of observed choices does not necessarily imply that these plans have a zero market

share and may be due to the fact that the number of options is large relative to 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 three, and the 75th percentile is five. 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.

A.2 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. The population of individuals who might consider purchasing individual market health insurance is 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. An individual that is not enrolled in employer sponsored insurance but has an offer that they chose not to accept is assumed to be in the individual market. These consumers are treated 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, 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 are assumed to be 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, they are assumed to be enrolled in Medicaid. This accounts for those that confuse Medicaid managed care programs with private coverage, and Medicaid

employer insurance assistance.

This paper follows the Government Accountability Office methods (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.3 Medical Expenditure Panel Survey

The Medical Expenditure Panel Survey (MEPS) is a nationally representative household survey on demographics, insurance status, and health care utilization and expenditures. In this paper, MEPS provides moments on the distribution of risk scores in the insured population and the relative costs of households by the age and risk score of the head of household and the risk. All moments are constructed using all surveyed households with health insurance in order to avoid the effect of access barriers on the reported expenditures, utilization, and diagnoses.

The 2015 Medical Conditions File (MCF) of 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. (2014)). The risk coefficients are published by CMS and publicly available.

A.4 Rate Filing Data

The Center for Medicare and Medicaid Services (CMS) tabulates the Premium Rate Filings that insurance firms must submit to state insurance regulators if they intend 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.

The data contain information on the firms' projected costs and experienced average costs. I use projected firm-level average cost and the average ratio of experienced costs across metal levels for all firms.²⁶ Unfortunately, the rate filing data do not fully cover every firm. As a result, firm-level average costs are supplemented by Medical Loss Ratio 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."

²⁶Using 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 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.

For the firms that 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.5 Medical Loss Ratio Data

CMS makes publicly available the state-level financial details of insurance firms in the Individual Market for the purpose of regulating the MLR.²⁷ This information includes the number of member-months covered by the insurance firm in the state and total costs.

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 scale 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

²⁷Insurance 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.

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.6 Computing Firm-level Risk

This paper 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.²⁸

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

²⁸The formula implies that the state average risk score should go in place of the national average. However, 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.

demand model. In smaller samples of the data, I have found that this does not substantially alter the results of the estimation but introduces non-linearities 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

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

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, $\theta_r = (\gamma_r, \{\beta_r^k\})$ 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.

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

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

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

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

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 .

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

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

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

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

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 $|\theta_{-r}^{n-1} - \tilde{\theta}_{-r}(\theta_r^n)| < \varepsilon$, 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 (2002) 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

$$\hat{V} = E[\nabla_h M(h(\theta))(MM')\nabla_h M(h(\theta))']$$

where $M(h(\theta)) = m(\theta)$

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 than $\hat{\theta}$. 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. The estimation is constrained to precisely match the projected-firm level average costs. The remaining cost parameters are estimated to fit 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 0.²⁹ See Appendix Section A.3 through A.5 on constructing these moments from the data.

²⁹I have also experimented with including moments on risk adjustment transfers for groups of firms, which does not substantially affect the results.

Matching Firm Moments

Let \bar{C}_f^{obs} be the observed projected firm-level average cost. The firm-specific cost parameters, $\tilde{\psi}(\phi)$, can be set 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 effect on the tails targeted by reinsurance. As a result, $\tilde{\psi}$ can be precisely computed with only a handful of iterations.

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 ,

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

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 (2002) to compute the variance of the moments (see Section B.2).

D Robustness

D.1 Price-Linked Subsidies

The main analyses of this paper maintain the assumption that, while the equilibrium subsidies are determined by a price-linked formula, the participating insurance firms do not incorporate this formula in their price setting strategy. This is equivalent to firms behaving as though the formula is using the prices of some other firm as the benchmark. This assumption is more realistic in markets in which there are many firms than those with only one or two competitors.

To relax this assumption, this section follows Jaffe and Shepard (2020) and assumes that the price elasticity internalized by the firm that owns the benchmark silver plan accounts for the fact that an increase in the price of the benchmark silver plan will also increase the subsidy for all subsidy eligible households in the market. This has two important implications. First, the price elasticity for the benchmark silver plan is zero for all subsidized households: the net price paid by subsidized households does not change with the base premium. Second, an increase in the price of the benchmark plan can increase total market-wide enrollment by causing an effective price decrease for all other insurance products. The quantity of consumers that enroll all other products may exceed the quantity of consumers that dis-enroll from the benchmark plan. These two features lead to a substantial inflation of silver plan premiums.

In order to smooth the computation of equilibrium, I assume that firms have an expectation over the probability that each silver plan they offer is the benchmark premium. Let p^{2lps} represent the second lowest-price silver plan, i.e. the benchmark premium specified by regulation. All silver plans in the market are assigned a probability that each plan is the benchmark plan, π_j , given by

$$\pi_j = \frac{e^{-\chi|p_j - p^{2lps}|}}{\sum_k e^{\chi|p_k - p^{2lps}|}}. \quad (8)$$

The parameter χ governs the certainty with which firms' know if they offer the benchmark premium. In the limiting case of a very large χ , this probability distribution collapses to certainty. In the results in this section, I set $\chi = 0.1$, which corresponds roughly to a firm knowing with 99% probability that its plan is the benchmark silver plan if the absolute price difference of the next closest silver plan is more than \$40. In the baseline equilibrium, the

Panel A: Less than 3350 HHI					
	Observed	Baseline	(1)	(2)	(3)
Risk Adjustment	Y	Y		Y	
Mandate	Y	Y	Y		
Prices					
Bronze	172	161	152	164	153
Silver	212	263	300	277	322
Gold	237	246	320	254	357
Panel B: Greater than 5200 HHI					
	Observed	Baseline	(1)	(2)	(3)
Prices					
Bronze	166	224	227	204	216
Silver	230	360	378	306	371
Gold	265	341	384	326	376
Panel C: Merger Analysis					
Merging Party Prices					
Bronze		10.7%	12.4%	6.7%	9.3%
Silver		8.6	4.5	6.1	3.2
Gold		5.3	4.7	2.5	2.3
Welfare					
Consumers		1.4	3.2	2.7	3.3
Producers		11	10	11	9
Government		-16	-14	-11	-10
Insurance Rate		-1.7	0.32	1.1	2.7

Notes: Panels A and B of this table display the same results as shown in Table 8, with the assumption that firms incorporate the price-linked subsidy formula in their strategies. All price levels represent the base premium per month, before age rating and subsidies are applied. Welfare calculations are excluded here. Panel C shows merger analysis information that mirrors Table 9. Values display the percent change as a result of a merger in each policy scenario. Negative values in government surplus change represent an increase in net federal spending.

Table 10: Effects of Selection Policies with Full Information on Price-Linked Subsidies

benchmark plans in 59 out of 109 markets are assigned probabilities greater than 70%, and in 90 markets the probabilities exceed 50%. Increasing the certainty parameter does not substantially alter the results of this section.

A summary of the policy results when firms have full information on the price-linked subsidy formula is displayed in Table 10. There are two main differences between the results shown here and those in Tables 8 and 9. First, the overall price level is higher. Second, the average price of Silver plans is much greater. These are intuitive outcomes of a reduction in the price elasticity of the benchmark Silver plan. In particular, the average price of Silver plans is greater than the average price of Gold plans in the baseline scenario, despite having lower costs. In the presence of risk adjustment, the profitability of shifting a consumer from a Silver plan to a Gold plan benefits from metal-level specific risk adjustment which increases payments to a plan for the same consumer, if that consumer is enrolled in a higher metal tier. This, in conjunction with the subsidy effect of the benchmark Silver plan, leads many firms to offer expensive Silver plans and cheaper Gold plans.

While the landscape of average premiums is different, the key results on the relationship between market power and adverse selection regulations still holds. The individual mandate leads to lower prices in the most competitive markets but higher prices in the least competitive markets. And the risk adjustment policy narrows the gap between the most expensive and least expensive plans (with the additional complication of the Silver-Gold plan inversion), but the policy does not have as large of an effect in the most competitive markets.

In the merger analysis, it still holds that price increases are concentrated in the least generous plans, which lead to a narrowing the premium gap between Bronze and Gold plans. However, the merger results in premium increases across all plans in every policy scenario. When firms have full information about the subsidy formula, the benefit from a merger not only includes the normal sales recapture, but also an increased probability in offering the benchmark silver plan. This larger increase in market power leads to larger price effects. However, because the majority of consumers do not experience large net premium increases and also benefit from smaller relative price differences between Bronze and Gold plans, average consumer welfare increases in every policy scenario. This increase coincides with a large increase in producer surplus and a still larger increase in government spending.

D.2 Welfare and Government Spending

Section 6 decomposes the sources of welfare loss in a competitive equilibrium with respect to a social planner that is seeking to maximize consumer and producer surplus, but does not consider government spending. Since this paper does not seek to characterize an optimal policy, this provides a useful illustration of the relationship between market power and welfare. This section relaxes this assumption to show how the welfare decomposition would change if the social planner considered government spending on subsidies net of the tax revenue gained from the individual mandate.

Suppose that the social planner considers a reduction in a dollar of government spending equivalent to an additional dollar of consumer or producer surplus. The planner treats government subsidies determined in equilibrium as fixed vouchers—i.e. cannot affect the amount of the subsidy—but does consider the subsidy as a part of the net cost of providing insurance to a particular customer. Recall that $c_j^{\tau\theta}$ represents the household-product specific cost of coverage and let b^τ be the household-specific premium subsidy. Then, the solution to the constrained planner problem given in equation 5 becomes

$$p_j + \frac{\lambda - 1}{\lambda} \frac{S_j}{S'_j} = \frac{E \left[\frac{\partial S_j^{\tau\theta}}{\partial p_j} (C_j^{\tau\theta} + b^\tau) \right]}{S'_j} \quad (9)$$

Including the subsidies in the welfare calculation results in an increase in the marginal cost of insurance for households that are eligible for subsidies. The main implication of this change is that the extensive margin distortions of the competitive equilibria disappears. The total welfare maximizing price vector in every local market also has positive total profits. Therefore, there is no longer any extensive margin welfare distortion that results from firms setting prices that cover average costs. Intuitively, this is a result of consumer willingness to pay that does not exceed the true cost of coverage.

The results of the welfare decomposition in the cross-section of geographic markets is displayed in table 11. Beyond eliminating the extensive margin selection distortion, incorporating government spending into the social planner problem also significantly reduces the distortion from equilibrium markups. This occurs for the same reason as the extensive selection distortion is zero: the efficient level of profits is positive and, in some cases, close to the equilibrium level of profits.

The results concerning the inefficient sorting distortion are similar to those when government spending is not taken into account. The welfare cost of inefficient sorting is still

HHI Level	Without Risk Adj.		With Risk Adj.		# of Mkts	Mkt Size
	Markup	Int. Sel.	Markup	Int. Sel.		
< 3,000	1.79	16.12	2.97	1.26	23	2,099
3,000 - 3,500	1.10	14.73	2.23	1.75	20	1,662
3,500 - 4,000	0.33	5.56	0.33	1.23	8	1,295
4,000 - 4,500	0.94	7.23	1.76	1.48	12	1,956
4,500 - 5,500	0.90	7.09	1.13	1.42	13	1,723
5,500 - 6,500	2.61	9.59	3.15	4.80	14	2,222
6,500 - 9,000	4.38	7.06	4.02	5.86	14	1,011
9,000 - 10,000	21.2	2.56	21.3	2.06	5	351
All Markets	2.20	9.84	2.78	2.42	109	12,320

Notes: This table displays the average welfare cost from extensive selection, markups, and inefficient sorting within each category of market concentration, measured by the Herfindahl-Hirschman Index (HHI). It also shows the welfare benefit (displayed as a negative cost) from the revenue neutral risk adjustment policy. All values in the left panel are in dollars per person per month. The averages are weighted by market population. The right panel displays the total number of markets and total population in thousands of the markets in each HHI category. Welfare is measured as the sum of consumer and producer surplus. This table does not display the extensive selection distortion, which is zero in all markets.

Table 11: Welfare Costs in the Cross-Section, Including Government Spending

substantial in the absence of a risk adjustment policy and declines with market concentration. When the risk adjustment policy is present, the inefficient sorting distortion is small, and does not have a strong relationship with market concentration.

The results displayed in this section should be interpreted with caution. The generosity of the premium subsidies leads many households to have a "marginal cost" of insurance that exceeds their willingness to pay. This is consistent with the literature on the value to consumer of government provided insurance Finkelstein et al. (2019). However, consumers may face obstacles or biases that decrease the demand for insurance, which may be internalized by the government in the creation of the subsidy policy. The intention of the government policy is clearly to increase the probability that these households will decide to purchase insurance. This is at odds with the social planner in this section, which finds a net welfare loss in insuring these households. A welfare function that discounts federal spending relative to consumer surplus would likely be better able to rationalize the current levels of spending.³⁰ While estimating the correct discount rate is outside of the scope of

³⁰Even in this setting, where one dollar of federal revenue is equivalent to one dollar of consumer surplus, the total welfare when the current law policy is in place is greater than total welfare absent any subsidy or mandate policy.

this paper, this exercise informs the direction and possible magnitude of changes to the main results of the paper that result from taking the federal budget into account.