Mergers in the Presence of Adverse Selection

Conor Ryan [*]

Abstract

In the presence of adverse selection, market power has a welfare benefit. Firms have an incentive to offer products that appeal to low-cost consumers and lead high-cost consumers to purchase insurance elsewhere. This inefficient sorting distortion is worst in highly competitive markets and absent in a monopoly. Market power also carries the welfare cost of higher markups. Due to these two opposing forces, the both the magnitude and the direction of the effect of a merger on total welfare is an empirical question. In this paper, I show how this trade-off can be captured in an empirically tractable discrete choice model and apply the model to a proposed merger in the individual insurance market regulated by the ACA. I find that, absent any taxes or transfers to address adverse selection, the merger leads to improvements in both consumer and producer surplus. When accounting for the taxes and transfers under the ACA that intend to address adverse selection, the inefficient sorting distortion is small in even the most competitive markets, and the merger reduces consumer and total welfare. This highlights that antitrust enforcement and other policies that encourage competition are complements to regulations targeting adverse selection.

1 Introduction

In the presence of adverse selection, market power has a welfare benefit: it reduces inefficient sorting. To see this, consider that a firm has an incentive to offer a product that appeals to low-risk consumers and encourages high-risk consumers to purchase from a competitor. This distortion declines with fewer competitors and is absent in a monopoly. However, market power also carries the welfare cost of higher markups. Due to these two opposing forces, the effect on total welfare of a particular merger, entrant, or other policy that affects competition in a market with adverse selection is an empirical question.

In this paper, I show how this trade-off can be captured in an empirically tractable discrete choice model. I apply the model to the health insurance market for individuals regulated by Affordable Care Act (ACA) to ask: what is the effect of mergers in this market on total welfare, and how do these merger effects depend on the policies in place to address adverse selection?

[*]Pennsylvania State University, Department of Economics, conor.ryan@psu.edu

First, in the absence of any taxes or transfers to address adverse selection, a merger leads to improvements in both consumer and producer surplus. Because the welfare benefits of further concentration is diminishing, the benefits are greatest for mergers with less pre-merger market concentration. Second, the current system of taxes and transfers under the ACA substantially reduces inefficient sorting in less concentrated markets. Because the potential benefit from a merger is already being addressed by these policies, mergers under the current policy environment are harmful across all levels of concentration, reducing both consumer surplus and total welfare. This highlights that antitrust enforcement and other policies that encourage competition are complements to regulations targeting adverse selection.

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)). In this paper, I study two such policies: the individual mandate taxes which individuals who do not purchase health insurance to mitigate extensive margin selection, and 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.

I study mergers in the presence of adverse selection using a model of strategic firms that compete in price with a fixed set of differentiated insurance products. A merger between any two products creates new incentives in setting the prices of those products due to recapturing the sales that would otherwise be diverted from one set to the other in the event of a price increase. Absent any selection, this incentive puts positive pressure on prices because the recaptured sales are profitable. The merged firm incorporates this new positive incentive, and the merger results in greater prices.

In the presence of adverse selection, the consumers that are potentially diverted from

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

one set of products to the other may not be profitable. If this set of consumers have an expected cost that greater than pre-merger price of a product, then the incentive for the merging product is negative: the recaptured sales represent loss not profit. Intuitively, the resulting downward pressure on prices comes from removing the incentive to keep a high price that deters these expensive, marginal consumer. A merger may therefore result in lower prices for some products.

Ambiguous merger price effects creates ambiguity in the effect of a merger on total welfare. Mergers in this model will always result in greater producer surplus, but may also lead to greater consumer surplus.[2] The direction and magnitude of the merger incentives depend on the character of intra-product selection described above, and the ultimate effects of the merger depend on the degree to which these incentives pass through to prices. While each potential merger may be unique, the potential that a merger in a market with adverse selection might increase total welfare (and lower prices) depends directly on the welfare cost of inefficient sorting. I present a framework to decompose the welfare loss into its two sources: markups and inefficient sorting.

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 contract 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 cross-section of markets with diverse levels of concentration.

To identify the key selection parameters, 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.[3] I use moments on the average HHS-HCC risk score for product categories and the relative risk score of insurance firm beneficiaries to identify how product preferences vary with risk. I combine these estimates with data on average firm costs and moments on the distribution of costs and risk in the Medical Expenditure Panel Survey to

[2]Mergers in a market characterized by Nash-Bertrand equilibrium among differentiated products always result in greater profits.

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

capture how medical risk is related to costs for each firm.

With the estimated supply and demand of health insurance, I measure the welfare costs of each distortion under each of four policy scenarios: baseline current law which includes the risk adjustment policy and the individual mandate, and three additional scenarios in which each of these two policies is removed individually and together. In the absence of a risk adjustment policy, the welfare costs of markups and inefficient sorting are of similar magnitudes—\$22.1 and \$23.2 per person per month, respectively. The least concentrated markets have the greatest welfare cost due to sorting and the least welfare cost due to markups. 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 \$4.9 per person per month.

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. In the most concentrated markets, where the pre-merger welfare cost of sorting is smaller, the merger is harmful for consumers but still improves total welfare. In the absence of both risk adjustment and the individual mandate, the merger is beneficial for consumers across all levels of competition. In the baseline policy scenario, 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, and larger under the individual mandate.

These results do not suggest that selection regulations like the individual mandate and risk adjustment are harmful for consumers. However, it is important for policy makers to consider that these policies *increase* the harm from reductions in competition. Alternatively, in markets where no such regulations exist, policy makers should be cognizant that additional concentration can be beneficial not only for total surplus but even for consumer surplus. In this manner, selection regulation and competition policy are complements.

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. Some recent theoretical work has shown that welfare in markets with adverse selection may be U-shaped in the degree of competition (Mahoney and Weyl (2017), Veiga and Weyl (2016), Lester et al. (2019)). The potential for welfare benefits from increased concentration highlights the importance of an empirically tractable model that can capture this trade-off. This paper presents such a model and allows for flexibility in between-firm selection, the key determinant of whether a particular merger will improve welfare.

In addition to empirical tractability, this paper extends the results of the literature to a setting where the product characteristics are fixed, but firms compete by setting the prices of a menu of products. 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. 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, and in particular, 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.

2 Model

2.1 Environment

There are a set of differentiated insurance contracts J, which are owned by F firms, indexed by j and f. I will write J^f for the subset of products owned by firm f. A product is characterized by a price p, an observed characteristic governing the generosity of the insurance x, and an unobserved characteristic ξ.

Consumers

There are a continuum of households, indexed by i and distributed by $F(i)$. Households make a discrete choice among the set of insurance products and are heterogeneous in their preferences for the price, α_i, and preferences for insurance generosity, β_i. Households have an additive idiosyncratic preferences over products $\{\varepsilon_{ij}\}_j$, 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

$$v_{ij} = \alpha_i p_j + \beta_i x_j + \xi_j + \varepsilon_{ij} \tag{1}$$

I will write the probability that household i chooses product j, given p, x, and ξ as $S_{ij}(p_j, p_{-j})$ where $p_{-j} = \{p_k\}_{j \neq k}$.

For each product j, a consumer i costs c_{ij} to insure. They key measures of selection between products is given by the covariance between the cross-derivatives of demand and the cost to insure. Given the logit specification of demand, this covariance is given by

$$\mathrm{Cov}\left(\frac{\partial S_{ij}}{\partial p_k}, c_{ij}\right) = \mathrm{Cov}\left(-\alpha_i s_{ij} s_{ik}, c_{ij}\right) \tag{2}$$

Equilibrium

The equilibrium vector of prices $p^* = \{p_f^*\}_{\forall f}$ solves the Nash-Bertrand competitive equilibrium between the firms such that for every f,

$$p_f^* \in \arg\max_{p_f} \int_i \sum_{j \in J^f} S_{ij}\left(p^*\right)\left(p_j - c_{ij}\right) di$$

2.2 Decomposing Channels of Welfare Loss

In this section, I define concretely the two channels through which a change in the market structure of in a market with adverse selection might affect welfare. First, mergers carry the traditional welfare cost of greater markups, which I will refer to as the *markup* channel. Second, a merger may improve total welfare by incorporating better selection externalities into product prices, which I will refer to as the *sorting* channel.

The total welfare loss is a market is characterized by the difference between a benchmark optimal social welfare and the welfare attained in competitive equilibrium. I will then characterize the contribution of the *markup* channel as the difference between the benchmark optimal social welfare and the maximum attainable welfare given that the industry earns the equilibrium level of profit. The remaining difference between this constrained optimum and the equilibrium welfare is then attributed to the *sorting* channel.

The benchmark optimal social welfare is the maximum possible utilitarian welfare that can be decentralized with a vector of product-level prices and consumers choosing optimally among those products.In a setting with multiple products, this already represents an important restriction from the first-best allocation and a potentially large welfare cost of adverse selection. However, the magnitude of this cost is unrelated to the market structure and this restriction is maintained throughout the paper.

The social welfare function, $SW(\cdot)$, is given by the sum of consumer surplus and producer profits.[4]

$$SW(p) = \int_i CS_i(p)dF(i) + \int_i \sum_{k \in J} S_{ik}(p_k - c_{ik})dF(i) \tag{3}$$

where

$$CS_i(p) = E_{\varepsilon_i}\left[\max_{k \in J} v_{ik}\right]$$

The social welfare maximizing price of a particular product is equal to the average cost of the marginal consumer of that product plus the total marginal benefit (or cost) of an increase in the price of j to other firms in the market, labeled as the total sorting externality.

[4] The results of this section do not depend on the specifics of a demand or consumer surplus specification, only that $\partial CS_i(p)/\partial p_j = -S_{ij}(p)$, which holds under much less restrictive assumptions on demand (Small and Rosen (1981)).

Let the utilitarian welfare maximizing vector of prices be p^W.

$$p_j^W = \underbrace{\frac{E\left[\frac{\partial S_{ij}}{\partial p_j} c_{ij}\right]}{S_j'}}_{\text{Private Marginal Cost}} \underbrace{- \sum_{k \neq j} \frac{E\left[\frac{\partial S_{ik}}{\partial p_j}(p_k - c_{ik})\right]}{S_j'}}_{\text{Total Sorting Externality}} \tag{4}$$

Next, consider the problem of a constrained social planner that chooses product-level prices subject to a constraint on promising a total profit of $\bar{\Pi}$ to the insurance industry.

$$\max_{\{p_j\}_{j \in J}} \int_i CS_i(p) dF(i) \tag{5}$$
$$\text{such that } \int_i \sum_{k \in J} S_{ik}(p_k - c_{ik}) dF(i) \geq \overline{\Pi}$$

The consumer welfare maximizing price of product j conditional on a given level of industry profit is given by

$$\underbrace{p_j + \frac{\lambda - 1}{\lambda} \frac{S_j}{S_j'}}_{\text{Marginal Social Benefit}} = \underbrace{\frac{E\left[\frac{\partial S_{ij}}{\partial p_j} c_{ij}\right]}{S_j'}}_{\text{Private Marginal Cost}} - \underbrace{\sum_{k \neq j} \frac{E\left[\frac{\partial S_{ik}}{\partial p_j}(p_k - c_{ik})\right]}{S_j'}}_{\text{Total Sorting Externality}} \tag{6}$$

where λ is equivalent to a Pareto welfare weight on profit.

Under utilitarian welfare ($\lambda = 1$), the markup term vanishes. In the presence of adverse selection, any market with non-negative profit will have $\lambda > 1$. The marginal cost curve is downward sloping in quantity and always below average cost, and the industry profit at welfare maximizing prices is negative. Even in perfectly competitive markets, firms must charge a markup over marginal cost in order to break even. When firms have market power, this welfare loss is further exacerbated through greater markups. Let the constrained efficient vector of prices be P^{CE}.

The welfare cost of markups is the difference between the unconstrained maximum welfare and this constrained efficient welfare, $SW(P^W) - SW(P^{CE})$. The welfare cost of inefficient sorting is the reduction in welfare of moving from the constrained efficient problem to a competitive equilibrium, $SW(P^{CE}) - SW(P^*)$.

The welfare cost of inefficient sorting is due to the combination of competition and differentiated products in the presence of adverse selection, which can be illustrated through the two cases where it is absent. First, if the market is monopolized by a single firm, the monopolist fully internalizes the sorting externalities. In this case, equation 6 converges to the monopolist's first order condition with $\lambda \to \infty$. Second, the sorting externality will also be zero if there is a single, homogeneous product. Even in the perfectly competitive case, there can be no inefficient sorting because there is no between-product selection.

With some oversimplification, a merger leads to an increase in industry wide profits. This increase in profits leads to an increase in the welfare cost of markups and a decrease in the welfare cost of sorting. This is informative for when mergers might be welfare improving. If concentration is already quite high and industry profits are large, the welfare cost of inefficient sorting is small, and additional concentration is unlikely to improve total welfare. If concentration is low and industry profits are small, the welfare cost of sorting may be large, and additional concentration may improve welfare.

More rigorously, whether a particular merger may have a net-positive effect on total welfare depends on the between-product selection of the merging products. In the next section, I explore this in detail.

2.3 Incentives Created by a Merger

A merger between any two sets of products creates new incentives in setting the prices of those products, which we can characterize using *Generalized Pricing Pressure* (GePP). The definition of GePP is the difference between the pre-merger and post-merger first order conditions for a particular product's price, both normalized to be quasi-linear in marginal cost (Jaffe and Weyl (2013)). Consider a merger between two single product firms which own the products j and k. The post-merger first order condition for product j is as follows:

$$0 = \underbrace{p_j + \frac{S_j}{S_j'} - \frac{E\left[\frac{\partial S_{ij}}{\partial p_j} c_{ij}\right]}{S_j'}}_{\text{Pre-Merger First Order Condition}} + \underbrace{\frac{\frac{\partial S_k}{\partial p_j}}{S_j'}\left(p_k - E\left[c_{ik}\right]\right) + \frac{\text{Cov}\left(\frac{\partial S_{ik}}{\partial p_j}, p_k - c_{ik}\right)}{S_j'}}_{\text{GePP}_{jk}} \tag{7}$$

where S_j is the market-wide market share of product j and S_j' is the own-price derivative: $\frac{\partial S_j}{\partial p_j}$.

The term, $GePP_{jk}$, captures the incentive for j created by the merger with k. Absent any selection, an increase in the price of j has a positive externality on k by diverting profitable consumers. The covariance term is 0, and GePP reduces to the product of the diversion ratio between j and k and the average markup for k. This is typically called upward pricing pressure, as it always captures a positive externality between firms. As the result of a merger, the merged firm recaptures this positive externality (diverted profitable consumers) and has an incentive to increase prices above the pre-merger level.

In the presence of adverse selection, this term may not be positive. As increase in the price of j may still divert consumers to k, but those consumers may not be profitable to k. Crucially, it doesn't depend on any aggregate measure of adverse selection but rather the specific selection from j into k that results from a price increase in j. This is captured by the covariance term. When the total GePP is negative, the merger creates an incentive to reduce the price.

For intuition, consider a case in which every consumer of product j has a higher expected cost than every consumer of product k, and at the current prices, every consumer indifferent between j and k is unprofitable to k. In this case, a *decrease* in the price of j is beneficial for k because it attracts unprofitable consumers away from its product. After the merger, j internalizes this benefit and has an incentive to reduce its price.

The potential for negative price effects opens the possibility that a merger increases total welfare. To a first order approximation, the effect of a merger on consumer surplus is given by

$$\Delta CS = -\sum_k \Delta p_k S_k$$

where Δp_k is the effect of the merger on the price of product k. The ambiguity in price effects creates ambiguity in the effect on consumer surplus. Because producer surplus is increasing as a result of a merger, this creates the potential that a merger is welfare increasing.[5]

2.4 Risk Adjustment in the Affordable Care Act

The ACA includes a risk adjustment transfer policy specifically intended to mitigate between-firm adverse selection. The government administers a transfer between firms that is equal to the difference between the firm's own average cost and the implied average cost of the

[5] Even without adverse selection, the standard UPP also characterizes a sorting welfare externality due to potentially asymmetric costs, which raises the possibility that re-allocation can increase total welfare at the cost of consumer surplus. However, in models of Nash-Bertrand price competition, asymmetric costs are already adequately reflected in prices, limiting the potential gains from a merger. Kao and Menezes (2007)

firm if it were to insure the same risk balance as the market as a whole (Pope et al. (2014)).[6]

$$T_j(p) = \underbrace{\frac{E[\sum_k S_{ik}c_{ik}]}{E[\sum_k S_{ik}]}}_{\text{Pooled Cost}} - \underbrace{\frac{E[S_{ik}c_{ij}]}{E[S_{ij}]}}_{\text{Average Cost}}$$

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

$$p_j^* + \frac{S_j}{S_j'} = \Psi_j \frac{E\left[(\sum_k \frac{\partial S_{ik}}{\partial p_j})c_{ij}\right]}{\sum_k \frac{\partial S_j}{\partial p_j}} + \left(1 - \Psi_j\right)\frac{E[\sum_k S_{ik}c_{ik}]}{\sum_k S_k} \tag{8}$$

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.

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

[6]The 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.

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.[7]

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.[8] 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.

The ACA includes two key policies to address adverse selection. To address adverse selection on the extensive margin of purchasing insurance, the ACA implemented the "Individual Mandate", a requirement to purchase insurance and an associated penalty for being uninsured. By taxing all individuals that do not buy health insurance, the insurance market can supposedly be reassured that a broad sample of consumers will purchase insurance, rather than simply the most costly. From 2016 through 2018, the mandate penalty was the maximum of $695 or 2.5% of household income. In 2015, the year of the data for this paper, the mandate was at half this level, and beginning in 2019, the penalty became $0.

To address intensive margin selection—the tendency of individuals with high expected costs to choose more generous insurance—the ACA implemented "risk adjustment," a system of risk-based subsidies (taxes) that compensate firms for enrollees with higher (lower) than average expected costs. The government collects claims data throughout the year from every insurance firm in the market to assess the average risk at the plan level using the HHS-HCC risk prediction methodology. This method attributes to each individual

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

[8] In recent years, California has extended subsidies to higher income households as well.

a risk score based on age, sex, and a set of diagnoses codes that are organized into hierarchical condition categories. Plans that have lower than average levels of risk are taxed and plans that have higher than average levels of risk receive subsidies. The formula that determines the taxes and subsidies is constructed to be budget neutral at the state-level: the total taxes across all firms within a state are mechanically equivalent to the total subsidies.

Risk-based subsidies are a common policy instrument to reduce adverse selection in health insurance markets (McGuire et al. (2011), Van de ven and Ellis (2000), Ellis and McGuire (2007)). The intention is to "eliminate the influence of risk selection on the premiums that plans charge," and see Section 2.4 for more detail on how risk adjustment works in a model of imperfect competition (Pope et al. (2014), Kautter et al. (2014a)).

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.

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.[9] 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

[9]The 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.

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.[10]

			Private Marketplace	
	ACS	ASPE	Un-weighted	Weighted
			Age Distribution	
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%
			Income Distribution	
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%
			Metal Level Market Shares	
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

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 Plan-

[10]The 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.

ning 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 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 risk adjustment transfers between firms, which are measures of the relative risk of each firms beneficiaries within a state. [11] 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).

4 Demand

4.1 Empirical Specification

In the empirical specification, households in nmarket m 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 indicator variables that include three age buckets, whether or not the household includes only one person, and whether or not the household is subsidy eligible. Households have an unobserved risk score, r^{HCC}. Households also have preferences $\theta_i = (\alpha_i, \beta_i)$.

The general model in Section 2 is specified as follows.

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

where $B(y)$ is a function that maps income to subsidies and $M(y)$ maps income to the penalty for choosing not to buy health insurance. I allow the preference for the utility-value of money, α_i, to be demographic specific. Observed characteristics X_{jm} include the actuarial rating of the plan and a firm fixed effect. The preference over observed character-

[11]These additional moments may capture consumer value of broad networks, for example.

istics, β_i, depends on a households risk score, r^{HCC}.

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

Risk is treated as an unobserved household characteristic. Risk scores are distributed according to a distribution that can depend on household demographics, Z_i.[12]

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

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.

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

The prediction regressions are performed separately for different types of plans t, where t represents the metal category of the plan. The resulting risk score for an individual is a normalized predicted relative-spending value. 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.

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

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.

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

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 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.[13]

In the data, a majority of individuals have no relevant diagnoses, i.e., $r_i^{HCC} = 0$.[14] 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.

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

[13] For 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.

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

Age	Income (% of FPL)	$\delta(Z_i)$	Bronze $\mu(Z_i)$	Bronze σ	Silver $\mu(Z_i)$	Silver σ	Gold $\mu(Z_i)$	Gold σ	Platinum $\mu(Z_i)$	Platinum σ
$\leq$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

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 "micro" moments to ensure that the model captures the appropriate risk-related substitution patterns and improve identification (Petrin (2002)).

The demand model targets eighty 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 each of the firm-state combinations in the data.

To estimate the demand model, I follow Grieco et al. (2021) to combine a mirco-data log-likelihood function with product-level GMM moments. The current results are presented using the identity weighting matrix, but estimation with the efficient weighting matrix is in progress.

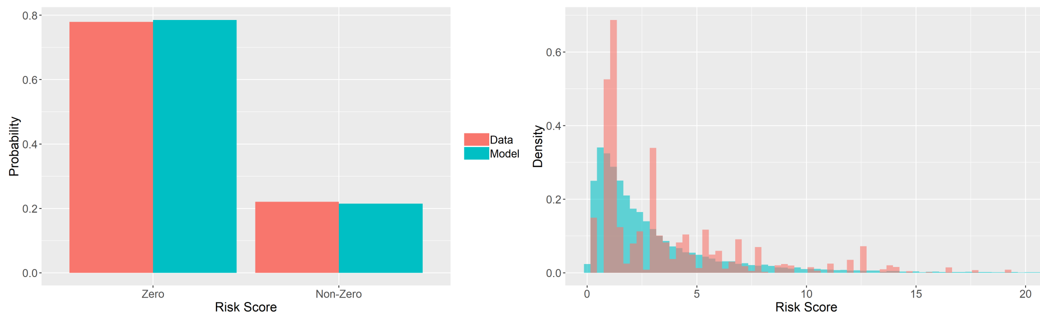

(a) Incidence of Positive Risk Scores (b) Distribution of Positive Risk Scores

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.

4.4 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 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.

In the preferred GMM specification, the mean consumer willingness to pay for a 10% increase in the actuarial value of an insurance plan is \$143 per month. This actuarial

	Maximum Likelihood		
	(LL-1)	(LL-2)	(GMM)
Premium	-1.46	-1.26	-1.23
	(0.00)	(0.00)	
Age 31 - 40	0.24	0 .24	0.26
	(0.00)	(0.00)	
Age 41 - 50	0.34	0.29	0.37
	(0.00)	(0.00)	
Age 51 - 64	0.69	0.55	0.59
	(0.00)	(0.00)	
Family	-0.17	-1.13	0.04
	(0.00)	(0.00)	
Subsidized	0.09	0.21	0.24
	(0.00)	(0.00)	
AV	4.40	9.36	10.97
	(0.00)	(0.00)	
Risk Preference			
AV	1.19	0.90	0.60
	(0.00)	(0.00)	
Firm - Risk Interaction			Y
Fixed Effects			
Age, Fam., Inc.	Y	Y	Y
Firm	Y		
Firm-Market			Y
Firm-Category			Y
Firm-Mkt-Cat.		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. Standard errors for the GMM estimation are not yet completed.

Table 3: Demand Estimation Results

increase is roughly equivalent to switching from a Bronze plan to a Silver plan (or Silver to Gold). 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 \$92.4 per month, and the 90th percentile is \$228 per month. The average own-price elasticity of consumers is -4.1, and the semi-elasticity of purchasing any insurance at all is -0.03, i.e. a \$10 increase in monthly price of every insurance product will decrease insurance enrollment by 3%. These 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, $C_f(X_j, \tau_i)$, as

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

where ϕ_f is a firm-state specific fixed effect, AV_{jm} is the actuarial value of the product, Age_i is the average age of the household, and r^{HCC} is the risk score of household. This specification assumes that the 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.[15] If this assumption is violated and the remaining endogeneity is consistent with adverse selection, then the coefficient on actuarial

[15] An alternative specification could treat expected total medical spending as a household characteristic. Then, I could allow preferences to vary with this characteristic instead of risk scores. 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.

value will be biased upward.[16] 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 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.

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

Estimation

The MSM estimation procedure targets four sets of moments which each identify four sets of parameters. The age and risk parameters are identified using moments from the Medical Expenditure Panel Survey (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

[16] For 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.

	(GMM)
Age	0.42
Risk	0.10
Actuarial Value	3.40
State-Firm	Y

Note: This table displays the estimates of the marginal cost function. Standard errors are not yet completed.

Table 4: Cost Estimation Results

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 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 B.

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. These estimates imply that the mean cost in the lowest decile of the own-price elasticity magnitude is $337 per month and the mean cost in the highest decile is $266 per month. The means of the top and bottom decile of costs, without considering elasticities, are $54 and $448 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
Age ($r^{HCC} = 0$)		
18 - 24	1.0	-
25 - 29	1.34	1.29
30 - 34	1.44	1.59
35 - 39	2.08	2.37
40 - 44	2.98	2.17
45 - 49	1.74	2.69
50 - 54	3.49	2.96
55 - 59	2.98	3.67
60 - 64	3.57	3.73
Risk		
$r^{HCC} = 0$	1.0	-
$r^{HCC} > 0$	3.57	3.52
Metal Level		
Bronze	1.0	-
Silver	2.28	2.37
Gold	3.80	3.09
Platinum	4.28	6.49

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

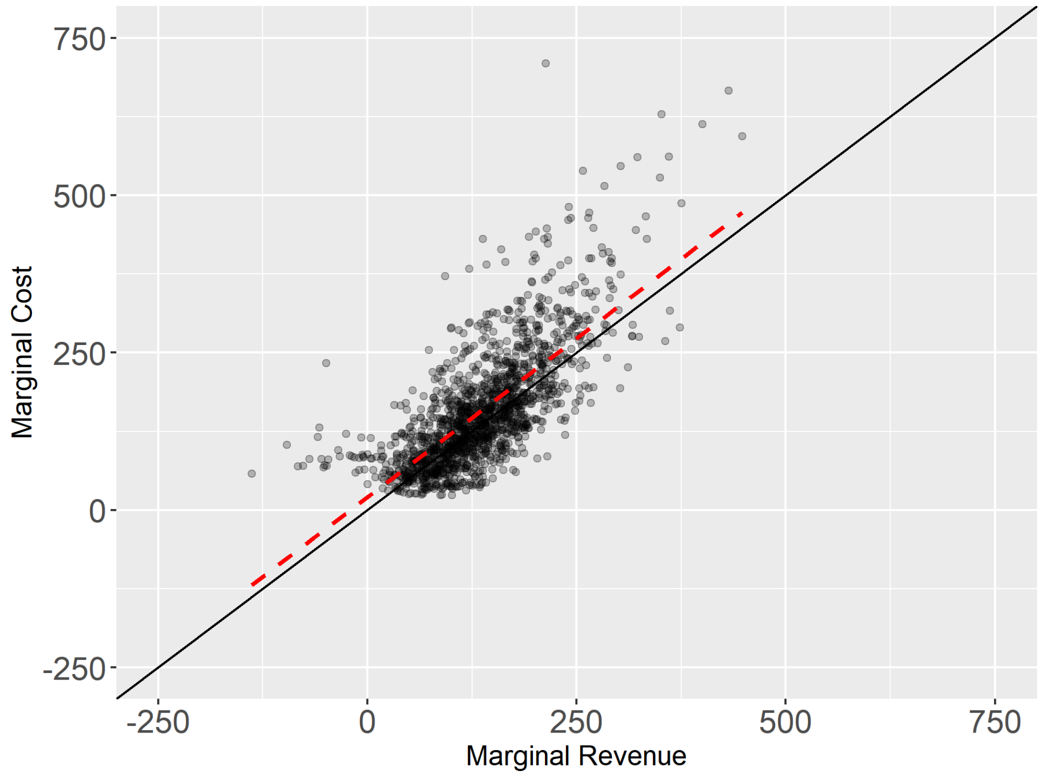

Figure 2: Marginal Revenue vs Marginal Cost in Baseline Model

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.

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 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, Market Concentration, and Mergers

In this section, I perform three exercises motivated by Section 2. First, I consider the 109 local markets in the cross section to quantify each welfare cost—markups and sorting—and study the relationship between these welfare costs, market concentration, and selection regulations. Second, I study a proposed merger in Georgia to study the welfare effects of a merger, and how those effects might depend on initial levels of concentration and selection regulations.

For these first two exercises, I hold fixed the baseline level of premium subsidies and ignore any changes in government spending for all welfare questions. In many markets in my data, consumers benefit less than a dollar for each additional dollar of government spending, a result consistent with other work on government sponsored health insurance (Finkelstein et al. (2019)). To avoid comparing to a benchmark where the optimal outcome is zero government spending and zero insurance enrollment, I treat the government's subsidy policy as fixed and outside the consideration of the planner. In the final exercise, I simulate merger effects when the subsidies are price-linked, as in current law.

6.1 Cross-section of Welfare Costs

I apply the welfare decomposition from Section 2.2 to each of the 109 local markets in the data. I do so under four policy scenarios using combinations of two policies designed to address adverse selection. The first policy is risk adjustment, described in detail in Section 2.4. The second policy is the "individual mandate," which is a tax on being uninsured. In the baseline scenario, both policies are in effect. The other three scenarios consist of removing each policy individually and then together. Note that this section is focused on decomposing welfare cost and does not address differences in total welfare, which is the subject of other work in this area (Saltzman (2021), Geruso et al. (2018)).

Table 6 shows the decomposition of welfare costs across markets at different levels of concentration (defined using Herfindahl-Hirschman Index (HHI) ranging from 0 to 1) in each policy scenario. Across each policy scenario the average welfare cost of markups is about $20-$30 per person per month. Generally, the welfare cost of markups is greater in more concentrated markets where firms have greater market power, and this effect is most pronounced at the highest levels of concentration.

In policy scenarios (1) and (3), which do not include a risk adjustment policy, the average welfare cost of sorting is greater than that of markups and concentrated among the most competitive markets. In comparing scenarios (1) and (3), the individual mandate policy reduces the magnitude of the welfare cost of sorting relative to markups, but the risk adjustment policy is more effective. While the risk adjustment policy does not optimally address the sorting inefficiency, it succeeds in greatly reducing the welfare cost of sorting. In the policy scenarios that include a risk adjustment policy, including the baseline, the welfare cost of sorting is small with no strong correlation with market concentration.

6.2 Welfare Decomposition of the Effects of a Merger

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 0.37 and five of them have an HHI of greater than 0.37. The change in HHI that results from the merger ranges from 0.02 to 0.26, with a median change of 0.08.

Table 7 contains the welfare and price effects of the merger in each of the four policy regimes described in the previous section. The effects are split into the most and least con-

	Baseline		(1)		(2)		(3)	
Risk Adj.	Y				Y			
Mandate	Y		Y					
	Markup	Sorting	Markup	Sorting	Markup	Sorting	Markup	Sorting
HHI								
< 0.3	27.1	3.3	17.4	31.8	25.9	9.9	13.6	43.8
0.3 - 0.4	29.2	2.8	21.4	23.7	30.0	6.7	17.5	36.1
0.4 - 0.5	27.8	6.4	20.9	25.6	26.9	11.8	21.4	27.1
0.5 - 0.6	31.1	3.8	20.3	25.9	26.8	10.0	16.4	31.7
0.6 - 0.7	27.1	7.8	20.7	19.7	22.3	12.4	15.6	25.0
0.7 - 0.8	23.3	7.8	26.2	7.9	19.5	9.5	19.9	12.0
0.8 - 0.9	37.4	7.0	33.4	10.7	32.9	9.0	29.9	10.9
0.9 - 1.0	41.9	6.0	39.6	7.5	37.8	5.3	35.9	6.2
Mean	29.7	4.9	22.1	23.2	27.0	9.7	18.5	29.6

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

centrated markets pre-merger. 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 relationship appears in the bottom panel of the table, where positive numbers indicate an increase in welfare (reduction in cost) via each channel.

In the baseline policy scenario, mergers are beneficial for producers and harmful for consumers, and total welfare falls in both categories of initial market concentration. The additional concentration does reduce the welfare cost of sorting by \$0.2 per consumer per month in the less concentrated markets and \$0.8 in the more concentrated markets. This can also be seen in the price effects of the merger. Bronze plans, which the least comprehensive and lowest priced plans, face the largest price increases. In other words, the price difference that consumers face in order to buy more comprehensive insurance is declining. This effect is offset by increase in the markup. All product categories face price increases, and total welfare falls by \$1.6 and \$1.3 per consumer per month due to greater markups.

In all policy regimes other than the baseline, mergers improve total welfare. The most drastic improvements come in the least concentrated markets in the policy scenarios where there is no risk adjustment policy. In these cases, total welfare improves via the sorting channel by \$13.5 and \$16.1 per person. These gains are partially offset by greater costs of markups as firms seek to capture the additional surplus. But consumer surplus improves, and in the scenario without any selection regulation, consumers benefit more than firms as a result of the merger.

These results do not suggest that selection regulations like the individual mandate and risk adjustment are harmful for consumers. However, it is important for policy makers to consider that these policies *increase* the harm from reductions in competition. Alternatively, in markets where no such regulations exist, policy makers should be cognizant that additional concentration can be beneficial not only for total surplus but even for consumer surplus. In this manner, selection regulation and competition policy are complements.

6.3 Effects of a Merger under Price-Linked Subsidies

While the previous two exercises held fixed the level of premium subsidies provided by the government, the true policy links subsidies to the equilibrium prices in the market. Specifically, the subsidies available to a household depends on the second-lowest priced silver plan in a household's choice set.

The first order effect of this policy is to increase the equilibrium prices in the market (Jaffe and Shepard (2020)). An additional implication is to increase the price effects of mergers. A newly merged firm not-only has the benefit of recapturing the sales diverted to its competitors, but is also now more likely to offer the second-lowest priced silver plan, creating additional upward pricing pressure.

I model this feature of the market allowing firms to internalize the link between prices and subsidies through a reduced form expectation that a silver plan they offer is the second lowest-price silver plan. For a given equilibrium price vector, let p^{2lps} represent price of the actual second-lowest price silver plan. Silver plans in the market are assigned a probability with which the firm knows it is the second-lowest priced plan given by

$$\text{Prob}(p_j = p^{2lps}) = \frac{e^{-\chi|p_j - p^{2slps}|}}{\sum_k e^{\chi|p_k - p^{2slps}|}}. \tag{9}$$

	Baseline		(1)		(2)		(3)	
Risk Adj.	Y				Y			
Mandate	Y		Y					
HHI	<0.37	>0.37	<0.37	>0.37	<0.37	>0.37	<0.37	>0.37
Price Effect								
Bronze	7.4	7.6	3.5	5.2	-2.3	3.6	-6.3	1.6
Silver	2.6	1.1	-11.8	1.1	-9.7	-6.0	-35.5	-3.1
Gold	3.4	1.3	-32.4	-13.2	-8.2	-6.2	-31.7	-20.5
Welfare Effect	-1.4	-0.6	9.4	1.5	2.1	0.9	11.9	1.8
CS	-1.8	-1.1	4.7	-0.1	1.6	0.5	7.6	0.7
Profit	0.4	0.5	4.7	1.6	0.5	0.4	4.2	1.1
Decomposition								
Markups	-1.6	-1.3	-4.2	-1.7	-1.6	-1.3	-4.3	-1.3
Sorting	0.2	0.8	13.5	3.2	3.7	2.2	16.1	3.1

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

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 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.

Table 8 contains the welfare and price effects of the merger in the two extreme policy regimes: baseline and when neither risk adjustment nor the individual mandate is in place. With the exception of the least concentrated markets in the baseline scenario, both firms and consumers receive greater surplus as a result of the merger.[17] When subsidies are linked

[17]This is also true of the two other policy scenarios that are not shown here in the interest of brevity.

	Baseline		(3)	
Risk Adj.	Y			
Mandate	Y			
HHI	<0.37	>0.37	<0.37	>0.37
Price Effect				
Bronze	10.2	11.1	7.7	5.6
Silver	15.9	28.9	1.8	6.7
Gold	7.4	7.2	-7.5	-3.5
Welfare Effect				
CS	-1.1	2.7	0.9	1.4
Profit	3.3	7.7	1.7	3.6
Government	-4.2	-15.9	-3.2	-6.8

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 8: Effect of a Merger with Price-linked Subsidies

to prices, there is a new channel through which both consumers and producers can benefit from a merger: consumers are compensated for increased prices with greater subsidies. This incentive can be seen in the baseline scenario price effects that are concentrated in Silver plans. Thus, mergers make all parties better off at the expense of the government.

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," *Quarterly Journal of Economics*, aug 2010, *125* (3), 877–921.

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

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

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.

Grieco, Paul L E, Charles Murry, Joris Pinkse, and Stephan Sagl, "Efficient Estimation of Random Coefficients Demand Models using Product and Consumer Datasets," 2021.

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.

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.

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

Jaffe, Sonia and E. Glen Weyl, "The First-Order approach to merger analysis," *American Economic Journal: Microeconomics*, nov 2013, *5* (4), 188–218.

_ and Mark Shepard, "Price-Linked Subsidies and Imperfect Competition in Health Insurance," *American Economic Journal: Economic Policy*, aug 2020, *12* (3), 279–311.

Kao, Tina and Flavio Menezes, "Welfare Enhancing Mergers under Product Differentiation," 2007.

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

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

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

Lester, Benjamin, Ali Shourideh, Venky Venkateswaran, and Ariel Zetlin-Jones, "Screening and Adverse Selection in Frictional Markets," *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, and Anna D. Sinaiko, "An Economic History of Medicare Part C," *Milbank Quarterly*, jun 2011, *89* (2), 289–332.

__ , __ , **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.

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.

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,” *Workiing 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.

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 75^{th} 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. (2014b)). 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. [18] 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."

[18] Using projected average costs and experienced ratios lead to the best fit for untargeted firm first order conditions. This could possibly be because the product-level projections are distorted by the firms incentives 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.[19] 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

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

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. Forthcoming work will relax the assumption that the age distribution is constant across firms. 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.[20]

Another potential method to capture the relative risk of firms is simply to target the

[20] The 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.

risk adjustment transfer itself, T_f, while everything else depends on the parameters of the demand model. In smaller samples of the data, I have found that this does not substantially alter the results of the estimation but introduces non-linearities in the moment calculations that make the task of finding a minimum to the GMM objective function considerably more difficult.

B 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.[21] See Appendix Section A.3 through A.5 on constructing these moments from the data.

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,

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

$\tilde{\psi}$ 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,

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

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 **??**).