

# Mergers in the Presence of Adverse Selection

Conor Ryan \*

## Abstract

In the presence of adverse selection, there are two channels through which a merger affects welfare: a reduction in inefficient sorting and an increase in markups. Whether the benefit from improved sorting is sufficient to offset the cost of greater markups is an empirical question that depends on the merger. I capture this trade-off in a tractable discrete choice model and apply the model to potential mergers in the non-group insurance market regulated by the ACA. Even in the presence of transfers to address adverse selection, 13% of mergers lead to greater consumer surplus. In markets where the sorting distortion is greater than \$7.5 per person, more than 1 out of 3 mergers improve consumer surplus. This highlights that antitrust enforcement should take the degree of adverse selection into account when assessing the potential harm from a merger.

## 1 Introduction

A merger between two firms in a market can improve social welfare. A common argument is that cost-synergies between the firms—generated by economics of scale, network effects, or other kinds of production complementarities—allow the firms to earn both a greater profit and increase total output. This possibility also exists in markets that suffer from adverse selection.

In the presence of adverse selection, there are two channels through which a merger affects welfare: a reduction in *inefficient sorting* (a positive welfare effect) and an increase in markups (a negative welfare effect). To understand the inefficient sorting channel, 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. Whether or not this welfare benefit is sufficient to offset the welfare cost of greater markups, both in terms of consumer surplus and social welfare, is an empirical question that depends on the merger (or other potential changes in the level of competition).

In this paper, I show how this trade-off can be captured in an empirically tractable discrete choice model that flexibly incorporates the key feature that determines the welfare

---

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

effect of a merger: between firm selection. I apply the model to the non-group (individual) health insurance market in the United States. There are three main findings. First, even with a policy in place to mitigate selection between firms, there are potential mergers that lead to greater consumer surplus and social welfare. Second, most mergers that would generate substantial welfare benefits are in markets which suffer from large welfare costs due to inefficient sorting. In the absence of a policy to address selection, the welfare cost of inefficient sorting is greater, and thus many more mergers are beneficial. This comparison demonstrates a complementarity between selection regulations and antitrust enforcement. Finally, I show that a generalized pricing pressure measure that accounts for selection forms an effective screen for harmful mergers.

The setting, non-group health insurance, is an important and policy-relevant market to study questions of selection and competition. Adverse selection and its consequences are a first order concern that motivated many elements of the ACA. The health reform law targets symptoms of selection that have been identified in the literature (Obama (2009), Cutler and Zeckhauser (2000), Van de Ven and Ellis (2000), Gruber (2008)). Competition is also a focus of policy-makers. Local insurance markets vary widely in their market concentration. The largest firm has a market share of over 85% in five states and less than 33% in another five states. More importantly, managed competition in insurance markets is a common tool to provide health insurance in many market segments in the U.S. and around the world

I model the insurance market as 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 diverted from 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 pricing pressure comes from a reduced incentive to keep a high price that might deter 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 so-

cial welfare and consumer surplus. The direction and magnitude of the merger incentives depend on the degree 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. The incentives due to a merger can be characterized with a straight-forward extension of Generalized Pricing Pressure (GePP), a common tool to make predictions of the direction and rough approximations of the magnitude of the price effects of a merger (Jaffe and Weyl (2013), Miller et al. (2017)).

While each potential merger may be unique, the possibility that a merger in a market with adverse selection might increase social welfare (and lower prices) depends directly on the welfare cost of inefficient sorting. I present a decomposition of welfare loss in selection markets with imperfectly competitive firms and differentiated products into two sources: inefficient sorting (conditional on markups) and markups (conditional on efficient sorting).<sup>1</sup> Markets in which the welfare cost of inefficient sorting is large are candidate markets where mergers may likely be beneficial to consumers.

Holding fixed the information friction that creates selection, the welfare loss can be decomposed

To estimate the model, I use data on household health insurance choices in the non-group health insurance market made through a private marketplace in 2015 (Ryan et al. (2021)). These data are unique in two respects. First, the data contain a substantial fraction of both low- and high-income consumers, in contrast to other work using government-run marketplace data which tend to be predominantly low-income (ASPE (2016)). Second, the data span more than 100 local markets, which allow me to estimate equilibrium outcomes in a cross-section of markets with diverse levels of concentration and investigate a broad set of potential mergers.

To identify the key selection parameters, I use an 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.<sup>2</sup> 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

---

<sup>1</sup>Throughout the paper, I hold fixed the information friction that creates adverse selection and the associated welfare loss.

<sup>2</sup>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 has been used in a similar demand specification (Aizawa and Kim (2018), So (2019)).

the distribution of costs and risk in the Medical Expenditure Panel Survey to 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 two policy scenarios: a baseline policy which includes risk adjustment transfers that target between-firm selection and a policy environment without these transfers. Because the transfers do not optimally target the inefficient sorting externality, the welfare cost of inefficient sorting is still positive. In the absence of the risk adjustment policy, the welfare costs of inefficient sorting are nearly twice as large, but in both scenarios the cost of sorting is much less than the welfare loss due to markups.

To evaluate when mergers may be beneficial to consumers and social welfare, I simulate every potential horizontal merger across all 107 local markets in the data. In the baseline policy scenario, 15% of 1186 merger-market combinations lead to greater social welfare and 13% lead to greater consumer surplus. In the absence of the risk adjustment policy, 22% improve social welfare and 17% improve consumer surplus. Even among the largest mergers (in terms of pre-merger market share), about 1-in-20 lead to greater consumer surplus in the baseline policy scenario.

The markets where mergers are most likely to be beneficial are those where the welfare cost of inefficient sorting is greatest. For markets where the cost is between \$7.5 and \$10 per person per month, 27% of mergers lead to greater consumer surplus in the baseline policy scenario. And in markets where the cost exceeds \$10 per person per month, 72% of mergers lead to greater consumer surplus.

From the perspective of antitrust enforcement, it is useful to have a measure that can reliably predict when a merger will reduce consumer surplus. In standard markets without selection, a common index for consumer harm is an upward pricing pressure (UPP) measure net of cost efficiencies (Farrell and Shapiro (2010)). In the presence of adverse selection, the derived GePP measure is the analogous prediction of consumer harm. However, GePP is challenging to compute without a full model of intra-product selection. I show that antitrust agencies can still use the standard (and easier to compute) UPP with a higher threshold for harm. In the absence of a risk-adjustment policy (or any case where selection is greater), this threshold should be still greater. This makes it clear that the degree of adverse selection appears as a channel for improved "efficiency" through a merger.

Taken together, these results show a complementarity between regulations that target adverse selection (in this case, a risk adjustment policy) and merger enforcement or policies to encourage competition. In the absence of selection regulations, merger enforcement

need not be as aggressive. Many mergers are beneficial to both the firms and consumers, and the optimal level of market concentration might be a somewhat concentrated oligopoly. However, since mergers also lead to greater markups, they are a costly way to address adverse selection. An alternative is to directly regulate selection with policy. These policies should be paired with aggressive merger enforcement or other policies that encourage competition in order to deliver the benefits of a competitive market.

## **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 possibility 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. 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.

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. Chade et al. (2022) demonstrate that a monopolist can approximately solve a multi-dimensional screening problem with endogenous product qualities simply by setting the prices of a small number of contracts with fixed characteristics. An implication is that the economic mechanisms of endogenous contracts also present in settings with fixed contracts, and this paper provides one example. 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. Intuitively, the sorting incentive appears in the incremental cost to purchase the next most generous insurance contract in the menu rather than an optimal level of insurance.

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)).

Recent work by Kong et al. (2023) investigates the relationship between the endogenous market structure and adverse selection in the non-group health insurance market in Massachusetts, and the authors find that adequately regulating adverse selection endogenously leads to more competitive markets. These results provides additional rationale for why antitrust enforcement and selection regulations are complements. The “natural” level of competition is greater when adverse selection is addressed, so mergers should draw more scrutiny. It also highlights the importance of local firm heterogeneity and the corresponding heterogeneity in potential merger effects. Even in the scenario most favorable to competition—when risk is assumed to be perfectly measured and adjusted under the ACA risk adjustment policy—there remains an inefficient sorting distortion and room for some mergers to benefit consumers and total welfare.

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 in most of the US. One approach is to estimate the relationship between an unobserved 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). I follow the method proposed by Grieco et al. (2021) to combine a micro-data log-likelihood function with product-level GMM moments.

There is a substantial empirical literature on the effects of competition in insurance markets (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. I show that risk adjustment policies are both less effective and less necessary in more concentrated markets, highlighting a kind of complementarity between policies that address adverse selection and policies that promote competition.

## 2 Model

### 2.1 Environment

Firms, index by  $f \in F$ , each sell a set of insurance contracts, indexed by  $j \in J^f \subset J$ , at a price  $p_j$ . For the full vector of prices, I will write  $\mathbf{p} = \{p_j\}_{j \in J}$ .

There are a continuum of households, indexed by  $i$ . The probability that a particular household  $i$  purchases a product  $j$  is characterized by  $S_{ij}$ , a product-specific demand curve. For ease of notation, I will drop the  $i$  subscript to denote aggregate functions, e.g.  $S_j = \int_i S_{ij} di$ . Each consumer  $i$  costs  $c_{ij}$  to insure with a product  $j$ .

The key measure of selection between products is given by the change in the average cost of a particular product ( $AC_j$ ) due to the change in the price of another ( $p_k$ ).

$$AC_j(\mathbf{p}) = \frac{1}{S_j(\mathbf{p})} \int_i S_{ij}(\mathbf{p}) c_{ij} di$$

$$\frac{\partial AC_j}{\partial p_k} = \frac{1}{S_j(\mathbf{p})} \int_i \frac{\partial S_{ij}}{\partial p_k} c_{ij} di - \frac{\frac{\partial S_j}{\partial p_k}}{S_j} AC_j$$

Selection arises from the relationship between demand,  $S_{ij}$ , and cost,  $c_{ij}$ . The change in average cost ( $\frac{\partial AC_j}{\partial p_k}$ ) is non-zero if and only if the covariance between consumer elasticities ( $\frac{\partial S_{ij}}{\partial p_k}$ ) and consumer costs ( $c_{ij}$ ) is non-zero.

### Firms and Equilibrium

A firm,  $f$ , competing in a particular market as a profit function defined as,

$$\Pi^f(\mathbf{p}) = \sum_{j \in J^f} (\mathbf{p}) S_j(\mathbf{p}) (\bar{A}_j p_j - AC_j(\mathbf{p})) \quad (1)$$

where  $\bar{A}_j$  is the average rating factor for consumers in product  $j$ . The rating factor transforms the base price,  $p_j$ , into the total price charged to a consumer and hence revenue earned from a sale. In practice, this factor depends only on age and is determined by state regulation. For ease of exposition in this section, I will assume  $\bar{A}_j \equiv 1$  for all products and prices.

The equilibrium vector of prices  $\mathbf{p}^*$  solves the Nash-Bertrand competitive equilibrium



between the firms such that for every  $j$ ,

$$p_j^* \in \arg \max_{p_j} \Pi^{f(j)}(\{p_j, p_{-j}\}).$$

## 2.2 Incentives Created by a Merger

A merger between two firms will be modeled as a new entity which jointly maximizes the profit over the existing products offered by the two firms pre-merger.<sup>3</sup> Generalized Pricing Pressure (GePP) can characterize the incentive that a merger between any two sets of products creates in setting the prices of those products. 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:

$$0 = \underbrace{\left[ p_j + \frac{S_j}{\frac{\partial S_j}{\partial p_j}} \left( 1 - \frac{\partial AC_j}{\partial p_j} \right) - AC_j \right]}_{\text{Pre-Merger First Order Condition}} + \text{GePP}_{jk} \quad (2)$$

$$\text{GePP}_{jk} = \frac{-\frac{\partial S_k}{\partial p_j}}{\frac{\partial S_j}{\partial p_j}} (p_k - AC_k) + \frac{S_k}{\frac{\partial S_j}{\partial p_j}} \frac{\partial AC_k}{\partial p_j} \quad (3)$$

The term,  $\text{GePP}_{jk}$ , captures the price incentive for  $j$  created by the merger with  $k$  and consists of two components. The first term in Equation (3) is the standard upward pricing pressure that generates unilateral incentives to raise prices following a merger (Farrell and Shapiro (2010)). The newly merged firm now internalizes that some fraction of consumers lost due to a price increase ( $\frac{\partial S_k}{\partial p_j} / \frac{\partial S_j}{\partial p_j}$ ) will switch to its newly acquired product  $k$  and generate profit,  $p_k - AC_k$ .

In the presence of selection, the incentive from a merger also contains a second component: the effect of the diverted consumers on the acquired product's cost. An increase in the price of  $j$  diverts consumers to  $k$ , but those consumers may also increase the average cost of  $k$ . If the effect on average cost (second term) is negative and large enough to out-

<sup>3</sup>This model does not preclude free exit for some products, as firms can set arbitrarily high prices. It does preclude the costly introduction of new products. I argue in Section 6.1 that this leads the positive consumer welfare results to be conservative.

weigh the recapture of diverted consumers (second term), the merger creates an incentive to reduce price.

In a similar manner that standard merger effects depend on the specific substitution between the merging products ( $\frac{\partial S_k}{\partial p_j} / \frac{\partial S_j}{\partial p_j}$ ), merger effects in the presence of adverse selection also depends on the specifics of selection between the merging products, ( $\frac{\partial A_k}{\partial p_j} / \frac{\partial S_j}{\partial p_j}$ ). The key to whether a particular merger will lead to an incentive to raise prices for a particular product is whether or not the consumers on the margin between that product and the acquired products are profitable to the acquirer.

For another perspective on the intuition, consider the pricing incentives in a competitive market with adverse selection. One reason to charge a high price is to encourage unprofitable consumers to select a competitors product instead. This is particularly important for products that tend to attract the highest cost consumers, e.g. generous insurance products. This can create a feedback loop with very high prices set by all competitors offering such products in equilibrium, and very low quantity sold for these types of products. In a merger-to-monopoly, this incentive will disappear, leading to downward pressure on prices for these products.

The potential for negative price effects opens the possibility that a merger will increase total social 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  $\Delta p_k$  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 possibility that a merger is welfare increasing.<sup>4</sup>

---

<sup>4</sup>Even without adverse selection, standard upward pricing pressure also characterizes a sorting welfare externality due to asymmetric costs, which raises the possibility that internalizing this incentive can increase total social welfare. However, in models of Nash-Bertrand price competition, asymmetric costs are already adequately reflected in prices, limiting the welfare gains from a merger (Kao and Menezes (2007)).

## 2.3 Two Welfare Costs

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

The social welfare loss in a market is the difference between a benchmark optimal social welfare and the welfare attained in competitive equilibrium. 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 and variation in consumer costs, 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.<sup>5</sup>

$$SW(\mathbf{p}) = \int_i CS_i(\mathbf{p}) di + \sum_{k \in J} S_k(p_k - AC_k) \quad (4)$$

where

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

The social welfare maximizing price of a particular product is equal to the average cost of the product plus three terms related to sorting. The first term is the private effect on that particular product. The second term represents the sorting externality (in the single-product firm case) of the price of product  $j$  on the costs of all other products in the market.

---

<sup>5</sup>The results of this section do not depend on the specifics of a demand or consumer surplus specification, only that  $\partial CS_i(\mathbf{p}) / \partial p_j = -S_{ij}(\mathbf{p})$ , which holds under much less restrictive assumptions on demand (Small and Rosen (1981)).

And the final term represents the standard business stealing externality.

$$p_j^W = AC_j + \underbrace{\frac{S_j}{\frac{\partial S_j}{\partial p_j}} \frac{\partial AC_j}{\partial p_j}}_{\text{Private Sorting}} + \underbrace{\sum_{k \neq j} \frac{S_k}{\frac{\partial S_j}{\partial p_j}} \frac{\partial AC_k}{\partial p_j}}_{\text{Sorting Externality}} - \underbrace{\sum_{k \neq j} \frac{\frac{\partial S_k}{\partial p_j} (p_k - AC_k)}{\frac{\partial S_j}{\partial p_j}}}_{\text{Business Stealing}} \quad (5)$$

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

$$\begin{aligned} & \max_{\{p_j\}_{j \in J}} \int CS_i(\mathbf{p}) di \\ & \text{such that } \sum_{k \in J} S_k(p_k - AC_k) \geq \bar{\Pi} \end{aligned} \quad (6)$$

The constrained efficient price of product  $j$  conditional on a given level of industry profit is given by

$$p_j^{CE} + \frac{\lambda - 1}{\lambda} \frac{S_j}{\frac{\partial S_j}{\partial p_j}} = \underbrace{AC_j + \frac{S_j}{\frac{\partial S_j}{\partial p_j}} \frac{\partial AC_j}{\partial p_j}}_{\text{Private Marginal Cost}} + \underbrace{\sum_{k \neq j} \frac{S_k}{\frac{\partial S_j}{\partial p_j}} \frac{\partial AC_k}{\partial p_j} - \sum_{k \neq j} \frac{\frac{\partial S_k}{\partial p_j} (p_k - AC_k)}{\frac{\partial S_j}{\partial p_j}}}_{\text{Total Externality}} \quad (7)$$

where  $\lambda$  is equivalent to a Pareto welfare weight on profit.

Under utilitarian welfare ( $\lambda = 1$ ), the markup term vanishes. In the standard model of adverse selection, where average cost is everywhere greater than marginal cost, any competitive equilibrium with non-negative profit will have  $\lambda > 1$ . When firms have market power, this welfare loss is further exacerbated through greater markups.

The welfare cost of markups is the difference between the unconstrained maximum welfare and this constrained efficient welfare,  $SW(\mathbf{p}^W) - SW(\mathbf{p}^{CE})$ . It represents the smallest decline in welfare necessary for firms in the market to earn the equilibrium level of profit. This welfare cost is intuitively related to the classic output restriction due to a markup in the single product case (Einav et al. (2010)). Because consumers are efficiently sorted among the available products, the welfare cost of markups is directly related to the quantity of insurance provided.

The welfare cost of inefficient sorting is the difference in welfare between the constrained efficient optimum and the competitive equilibrium,  $SW(\mathbf{p}^{CE}) - SW(\mathbf{p}^*)$ . At both the constrained efficient price vectors,  $\mathbf{p}^{CE}$ , and the equilibrium price vector,  $\mathbf{p}^*$ , industry

profits are equivalent. The welfare cost comes from a reduction in consumer surplus that results from the pricing externalities between firms.

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 7 converges to the monopolist's first order condition with  $\lambda \rightarrow \infty$ . 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 and a decrease in the sorting externality. The increase in profits leads to an increase in the welfare cost of markups and the decrease in the sorting externality leads to a decrease in the welfare cost of sorting. This highlights the relevant trade off and suggests easily observable conditions for when mergers might be welfare improving. If concentration is already quite high, the welfare cost of inefficient sorting is small and additional concentration is unlikely to improve social welfare. If concentration is low, the welfare cost of sorting may be large and additional concentration may improve welfare. This is related to the U-shaped relationship between competition and welfare that has been identified in theoretical models of selection (Veiga and Weyl (2016), Lester et al. (2019)).

However, in a setting with differentiated products, it is difficult to characterize a clean result such as the theoretically optimal level of competition. Whether a particular merger may have a net-positive effect on social welfare depends on the the extent of adverse selection between the particular merging products, as outlined in Section 2.2. In Sections 4 and 5, I pursue an empirical strategy that will uncover this property of consumer demand.

## **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 firm if it were to insure the same risk balance as the market as a whole (Pope et al. (2014)).<sup>6</sup> (For more details on the policy specifics, see Section 3.)

---

<sup>6</sup>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.

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

In the presence of risk adjustment transfers, the firm then faces a new average cost,  $AC_j^T(\mathbf{p}) = AC_j(\mathbf{p}) - T_j(\mathbf{p})$ . The equilibrium price can be written as

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

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 7 that penalizes or reward firms based on the profitability of their marginal consumers. Therefore, it is not guaranteed to eliminate the welfare cost of inefficient sorting.

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 the sorting distortion 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

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.<sup>7</sup>

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. 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.

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.<sup>8</sup> Households that earn less than 250% of FPL also receive additional subsidies to cover reduced cost-sharing.

To address selection between products, 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 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

---

<sup>7</sup>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.

<sup>8</sup>In recent years, California has extended subsidies to higher income households as well.

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. 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 107 rating areas.

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.<sup>9</sup>

The data from the private marketplace are a selected sample of all the consumers facing a particular firm. Using the same data set, Ryan et al. (2021) find that income is a primary determinant of driving selection into the private online market place. In order to create a sample of consumers that is representative of the consumer population facing firms in this market, I treat the choice data as a random sample conditional on subsidy eligibility and geographic market. Each observation from the choice data within a particular subsidy eligibility category and market is given an equal weight such that the weights sum to the size of the population as determined by the 2015 American Community Survey (ACS)<sup>10</sup>. The ACS also provides a sample of the uninsured population.

---

<sup>9</sup>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.

<sup>10</sup>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.



Appendix Section A.1 contains descriptive statistics of the data sample, as well as more detail on sample selection, missing data, and constructing the choice sets. For more detailed information on processing the ACS, see Appendix Section A.2.

## **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. I target 5 moments that CMS publishes in annual reports on the results of the risk adjustment transfer program: the national average risk score for enrolled beneficiaries and the average risk score of consumers in Bronze, Silver, Gold, and Platinum plans. CMS only began to publish risk scores by metal-level in 2017. In order to make it comparable to my data, I use the average of on- and off-exchange market segments, and scale the risk scores by the ratio of the 2015 national average risk score to the 2017 national average risk score.

I also target the risk distribution across firms using risk adjustment transfers. These transfers are related to the relative risk of the enrolled beneficiaries for each firm within a state (See Section 2.4). In Appendix Section A.4, I detail how these moments are constructed from Medical Loss Ratio data submitted to the government.

I combine complement these moments on the risk distribution with similar moments on costs. I match moments on the relative costs of individuals by age and risk, which come from MEPS (Appendix Section A.3). Additionally, I match moments on the average cost of insurance product categories and the average costs of each firm. Product category level data come from rate filings to state insurance regulators (Appendix Section A.5), and the average firm-level costs come from the Medical Loss Ratio data (Appendix Section A.4).

## 4 Demand

### 4.1 Empirical Specification

Households have characteristics  $(a_i, y_i, Z_i, r_i^{HCC})$ , 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.<sup>11</sup> Households have an unobserved risk score,  $r^{HCC}$ .

Households make a discrete choice among the set of insurance products that are available in their market,  $m$ . Households are heterogeneous in their preferences for the price,  $\alpha_i$ , and insurance characteristics,  $\beta_i$ . Products have an unobservable quality,  $\xi$ , and households have additive idiosyncratic preferences over products  $\varepsilon$ , 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

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

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. Observed characteristics  $X_{jm}$  include the actuarial rating of the plan and a firm fixed effect.

I allow the preference for the utility-value of money,  $\alpha_i$ , to be demographic specific. The preference over observed characteristics,  $\beta_i$ , depends on a household risk score,  $r^{HCC}$ .

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

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

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

Because of the distributional assumption on  $\varepsilon$ , the probability that an individual will

---

<sup>11</sup>I use the demographics of the head-of-household as the representative demographics for the household.

purchase a particular product,  $S_{ij}$ , is given by the standard multinomial logit formula.

## 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 non-group 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 derived from diagnoses codes.

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

The prediction regressions are performed separately for different types of plans  $t$ , where  $t$  represents the metal category of the plan. The resulting risk score for an individual is a normalized predicted relative-spending value. Because all regressors take a value of either 1 or 0, the risk score is equal to the sum of all coefficients that apply to a particular individual.

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

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

### Parametric Distribution

The distribution of risk scores,  $\hat{G}$ , is estimated from the 2015 Medical Conditions File (MCF) of the Medical Expenditure Panel Survey. The MCF contains self-reported diagnoses codes and can be linked to demographic information in the Population Characteristics 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.<sup>12</sup> See Appendix Section A.3 for detail on processing the data.

In the data, a majority of individuals have no relevant diagnoses, i.e.,  $r_i^{HCC} = 0$ . 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.,  $r_i^{HCC} \sim \text{Lognormal}(\mu(Z_i), \sigma)$ . 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 for the head of household and two income categories—above and below 45 years old, and above and below 400 percent of the federal poverty level.

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

Age	Income (% of FPL)	$\delta(Z_i)$	<u>Bronze</u>		<u>Silver</u>		<u>Gold</u>		<u>Platinum</u>	
			$\mu(Z_i)$	$\sigma$	$\mu(Z_i)$	$\sigma$	$\mu(Z_i)$	$\sigma$	$\mu(Z_i)$	$\sigma$
$\leq 45$	$\leq 400\%$	0.15	2.86	19.7	2.99	19.5	3.11	19.8	3.31	20.8
	$> 400\%$	0.13	3.02	19.7	3.22	19.5	3.22	19.8	3.40	20.8
$> 45$	$\leq 400\%$	0.31	3.49	19.7	3.73	19.5	3.73	19.8	3.97	20.8
	$> 400\%$	0.24	3.25	19.7	3.46	19.5	3.46	19.8	4.67	20.8

Note: 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 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 1: Parametric Distribution of Risk Scores

### 4.3 Estimation

This model has two primary identification concerns. First, a plan premium's price may be correlated with the unobserved quality  $\xi_{jm}$ , leading to biased estimates of  $\alpha_i$ . In this environment, the premium regulations provide a source of variation in price, which is exoge-

<sup>12</sup>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.




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 of risk scores in the same MEPS sample. In both cases, the distribution of positive Silver metal-level risk scores are displayed.

nous to variation in unobserved quality (Tebaldi (2020)). The age-adjustment on premium,  $a_i$ , increases monotonically and non-linearly with age, and strictly increases with every age after 25. Income-based subsidies are available to households that earn below 400 percent of the federal poverty level. These subsidies decline continuously within the subsidy-eligible range. I am able to allow price sensitivity to also depend on age and income, but only in broad buckets. Intuitively, the variation in price within each demographic bucket defined by  $Z_i$  identifies  $\alpha$  for that particular demographic.

I use fixed effects to control for  $\xi_{jm}$ , 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 product-level 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; and the average risk score of each firm relative to the average risk score in the state for each firm-state combination in the data. Let  $l$  index the moments, let  $n$  index the  $N = 500$  draws from the unobserved distribution of risk, and let  $I(j)$  be the set of consumers that have product  $j$  in their choice set. For each group of products,  $J_l$ , I compute the moments as

$$M_l = \frac{\sum_{j \in J_l} \sum_{i \in I(j)} \sum_{n=1}^N w_i S_{inj} r_{inj}}{\sum_{j \in J_l} \sum_{i \in I(j)} \sum_{n=1}^{500} w_i S_{inj}} - R_m^{data}$$

where  $r_{inj}$  is a product-specific risk draw to match the definition of the moments in the data and  $w_i$  is the weight consumer  $i$  (see Section 3.1 for more details on weighting).

To estimate the demand model, I follow Grieco et al. (2021) to combine a micro-data log-likelihood function with product-level GMM moments. The parameters maximize the sum of the log-likelihood of observed choices less the weighted moment objective value.

$$\hat{\theta} \in \operatorname{argmax}_{\theta} \sum_i \sum_j Y_{ij} \log\left(\frac{1}{N} \sum_n S_{inj}\right) - M'WM \quad (9)$$

The estimation proceeds in two steps. First, I use the identity matrix as the weighting matrix,  $W$ . Second, I set the diagonal of the weighting matrix equal to the inverse of the moment variances evaluated at the parameters estimated in the first stage. Because the moments do not apply to all consumers in the data, I cannot directly compute the moment variances. Instead, I follow Petrin (2002) by computing the variance of a separate set of moments that can be used to construct the intended moments for estimation. In this case, the predicted choice probabilities,  $S_{ij}$ , and the average predicted risk for each product,  $\frac{1}{N} \sum_n S_{inj} r_{inj}$  are sufficient. The variance of the targeted moments can then be computed

using the delta method.

This estimation procedure is analogous to a GMM estimation that uses the first order conditions of the likelihood function as moments (Grieco et al. (2021)). Using the likelihood function in place of an additional set of moments allows the estimation to maintain the desirable convergence and identification properties of maximum likelihood estimation. However, to compute standard errors, I exploit the analogous GMM framework and compute the typical GMM standard errors where the weighting matrix is a block diagonal matrix with the Hessian of the likelihood function in one block and the moment weighting matrix  $W$  in the other.

## 4.4 Results

Table 2 presents the results from the demand estimation. The GMM specifications are supplemented with maximum likelihood estimations (MLE) 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. The restriction of a single dimension of heterogeneity puts more emphasis on the actuarial value parameter rather than risk-related product differentiation among firms. 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 between firm selection.

The specification used throughout the rest of the paper is GMM-2. I estimate three GMM specifications to demonstrate the sensitivity of parameter estimates to the level of fixed effect that controls for cross-product heterogeneity. The most detailed specification (GMM-3) includes fixed effects for every firm-market-category, where category indicates whether the insurance plan is a high coverage plan that covers more than 80% actuarial value. It is a challenge to use this specification in counterfactual simulations, because not all firm-market-category combinations are chosen. Instead, I use specification GMM-2

	MLE		GMM		
	(MLE-1)	(MLE-2)	(GMM-1)	(GMM-2)	(GMM-3)
<i>Premium</i>	-1.46	-1.26	-2.07	-1.35	-1.32
	(0.00)	(0.00)	(0.02)	(0.02)	(0.01)
Age 31 - 40	0.24	0.24	0.30	0.28	0.29
	(0.00)	(0.00)	(0.02)	(0.02)	(0.02)
Age 41 - 50	0.34	0.29	0.62	0.44	0.43
	(0.00)	(0.00)	(0.02)	(0.02)	(0.02)
Age 51 - 64	0.69	0.55	1.20	0.71	0.70
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Family	-0.17	-1.13	0.01	0.06	0.04
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
Subsidized	0.09	0.21	0.34	0.29	0.29
	(0.00)	(0.00)	(0.01)	(0.01)	(0.01)
<i>Actuarial Value (AV)</i>	4.40	9.36	7.03	11.98	11.70
	(0.00)	(0.00)	(0.05)	(0.08)	(0.08)
<i>Risk Preference</i>					
AV	1.19	0.90	0.59	0.55	0.54
	(0.00)	(0.00)	(0.00)	(0.00)	(0.00)
Firm - Risk Interaction			Y	Y	Y
Fixed Effects					
Age, Fam., Inc.	Y	Y	Y	Y	Y
Firm	Y		Y		
Firm-Market				Y	
Firm-Category				Y	
Firm-Mkt-Cat.		Y			Y

Note: The top row of price coefficients corresponds to the estimate for households that do not fall into any of the listed subgroups (single, high income, 18 to 30 year olds). The price coefficients for other households are obtained by adding the relevant demographic adjustments to the top line. Premiums are in thousands of dollars per year.

Table 2: Demand Estimation Results



which has very similar parameter estimates.

The median consumer willingness to pay for a 10% increase in the actuarial value of an insurance plan is \$139 per month. This actuarial increase is roughly equivalent to switching from a Bronze plan to a Silver plan (or Silver to Gold). The median price difference to consumers between Bronze and Silver plans is about \$52 per month. There is substantial variation in willingness to pay. The 10th percentile of willingness to pay is \$94.9 per month, and the 90th percentile is \$293 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 Cost

### 5.1 Empirical Model

The expected 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 costs as

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

where  $\psi_f$  is a firm-state specific fixed effect,  $AV_{jm}$  is the actuarial value of the product,  $Age_i$  is the average age of the household, and  $r_i^{HCC}$  is the risk score of household. This specification assumes that the identically and independently distributed errors in the cost function,  $\omega_{ijm}$ , are orthogonal to the preference draws in the demand estimation.

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

The only mechanisms through which cost and preferences are correlated are through age and risk scores.<sup>13</sup> If this assumption is violated and the remaining endogeneity is consistent

---

<sup>13</sup>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

with adverse selection, then the coefficient on actuarial value will be biased upward.<sup>14</sup> 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.

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

## Estimation

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

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.

<sup>14</sup>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.

Expenditure Panel Survey (Appendix Section A.3). For clear identification of costs by age separate from risk score, the estimation targets age moments among adults that have a risk score of zero. The moments are computed as the ratio of average covered expenditures within five-year age brackets for adults between 25 and 64 years old to the average covered expenditures of adults between 20 and 24 years old. The cost parameter on risk is identified using the ratio of average covered expenditures among adults with a positive risk score to those with a risk score of zero. This helps to separate sorting-related costs from firm-specific or product-specific costs.

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,  $\phi$ , the firm-specific cost parameter,  $\psi$ , is set to precisely match the projected average cost in the 2015 rate-filing data. See Appendix Section A.5 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 3 displays the results of the cost estimation. The table presents results for two GMM demand specifications used to simulate the moments targeted by the cost estimation. The main specification used in the counterfactual results is GMM-2.<sup>15</sup> The estimation implies a substantial amount of variation in consumer costs. The mean cost among the lowest decile (least costly) consumers is \$46.8 per month and the mean cost among the greatest decile (most costly) consumers is \$1,120 per month.

The standard mechanism of adverse selection is present. The 50 percent most elastic

---

<sup>15</sup>The specification GMM-3 is not included. The average firm-level costs cannot be simulated in the same way because not all firm-market-category fixed effects present in the choice sets are chosen in the estimation data.

	(GMM-1)	(GMM-2)
Age	0.44 (0.01)	0.44 (0.01)
Risk	0.11 (0.00)	0.11 (0.00)
Actuarial Value	4.22 (0.02)	4.26 (0.02)
State-Firm	Y	Y

Note: This table displays the estimates of the marginal cost function. Standard errors are computed using the GMM formula, accounting for demand estimation error in the simulation.

Table 3: Cost Estimation Results

consumers with respect to purchasing any insurance are 13% less costly than the least elastic consumers. The consumers that are infra-marginal in the insurance purchase decision are more expensive than the consumers that are more marginal to leaving the insurance market. This is the standard, extensive margin adverse selection Einav et al. (2010).

However, the important features of selection that drive merger effects are not the extensive margin decisions, but substitution patterns among products. Using the own-price elasticity, the 50 percent most elastic consumers are 20% *more* costly than the least elastic consumers. While the consumers most likely to leave the insurance market are less costly on average, the consumers most likely to *switch products* are more costly. This highlights the importance of selection dynamics internal to the insurance market, and the potential for mergers to alleviate resulting distortions.

Table 4 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).

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

	Data	Model Fit	
		GMM-1	GMM-2
Age ( $r^{HCC} = 0$ )			
18 - 24	1.0	-	-
25 - 29	1.34	1.31	1.32
30 - 34	1.44	1.53	1.58
35 - 39	2.08	2.38	2.39
40 - 44	2.98	2.12	2.10
45 - 49	1.74	2.61	2.62
50 - 54	3.49	2.87	2.89
55 - 59	2.98	3.92	3.86
60 - 64	3.57	3.89	3.88
Risk			
$r^{HCC} = 0$	1.0	-	-
$r^{HCC} > 0$	3.57	3.27	3.26
Metal Level			
Bronze	1.0	-	-
Silver	2.28	1.71	1.78
Gold	3.80	3.42	3.33
Platinum	4.28	7.37	7.05

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 4: Cost Estimation Fit of Cost-Ratio Moments

includes risk adjustment and reinsurance. It demonstrates how well the estimated model can fit the data, i.e. the difference between model-predicted marginal revenue and model-predicted 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

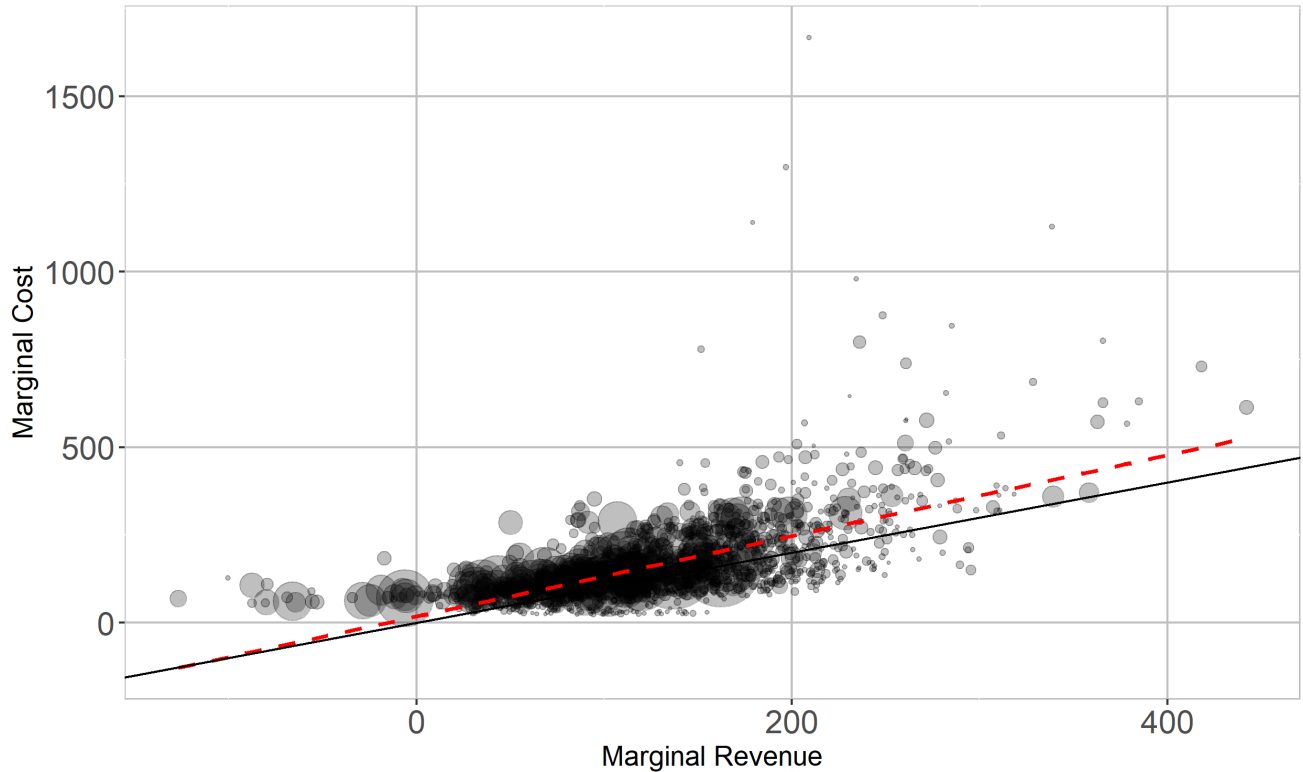


Figure 2: Marginal Revenue vs Marginal Cost in Baseline Model

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

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 in the case that regulators can effectively discipline markups, the results that follow will underestimate the positive effects of consolidation.

## 6 The Welfare Effects of Mergers

In this section, I simulate every potential horizontal merger between firms that compete in at least one local market and compute the effects on welfare. Because competitive equilibrium is not assumed in the estimation of demand and supply, I first re-solve the baseline equilibrium. Next, I solve the post-merger equilibrium for each potential merger. In the data, there are 243 potential bilateral, horizontal mergers between competing firms, each

of which affect an average of 4.8 local markets.<sup>16</sup>

The key interest of this paper is how the effect of mergers depends on the presence and degree of adverse selection. In order to demonstrate this, I repeat this merger-analysis exercise twice: once with the baseline policies in place, and once without the ACA risk-adjustment policy.

A merger between two firms is characterized as jointly maximizing the profit over a set of products that is fixed in both the pre-merger and post-merger equilibrium. Firms can only change their product offerings in a very limited way. The “exit” of a product is possible by setting a price prohibitively high, and “entry” is similarly possible through a price reduction if the pre-merger equilibrium for a product has such a prohibitively price. Because one possible benefit of a merger is more reasonable prices of generous insurance products, this leads to conservative estimates of the benefits from a merger (see Figure 3).

In the following analysis, I make two important assumptions about subsidies. First, I ignore any changes in government spending in the welfare computation. At the estimated parameters, the average consumer surplus generated from a dollar of additional government spending is less than a dollar, 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 very little insurance enrollment, I treat the government’s subsidy policy as fixed and taken as given by the planner who cares only about consumer and producer surplus.

Second, I assume that price subsidies are fixed and treated as vouchers by both the consumers and firms. In reality, subsidies are tied to an order statistic of the equilibrium prices in each market: the second-lowest price silver plan. As has been previously studied, this leads to greater upward pressure on prices (Jaffe and Shepard (2017)). When the analysis is repeated allowing firms to internalize this policy, consumers benefit from more than half of all mergers due to the rising subsidy. In a social welfare analysis that also accounts for the increases in government spending, the fraction of mergers that improve total surplus is similar in magnitude to those presented below. However, these results are sensitive to what weight is placed on government dollars relative to consumer and producer surplus. For more details, see Appendix Section C

---

<sup>16</sup>In markets with adverse selection, there may be multiple equilibria. To verify that the simulated effects of a merger are due to the changing market structure rather than equilibrium selection, I resolve for the pre-merger equilibrium from the post-merger prices. In every case, the solution returns to the original pre-merger equilibrium.

## 6.1 Mergers Can Lead to Greater Welfare

Local markets for individual insurance are quite concentrated. Table 5 shows the distribution of firms and the Herfindahl–Hirschman Index (HHI) in the pre-merger equilibrium of the model, both in the baseline policy scenario and in an environment without risk adjustment. The 2010 Horizontal Merger Guidelines (written in collaboration between the US Department of Justice and Federal Trade Commission) consider markets with an HHI greater than 2,500 to be “highly concentrated” and merit extra scrutiny for merger review. In the pre-merger equilibrium of the baseline policy environment, 91% of the markets exceed this threshold. Using the shares computed from the estimation data, all but 2 markets exceed this threshold.<sup>17</sup>

Because the markets for non-group insurance are concentrated, the sorting problems due to adverse selection are not a significant problem for social welfare. The welfare cost of markups is much greater than the welfare cost of sorting, even in the case where there no risk adjustment is in place.<sup>18</sup>

	25 <sup>th</sup> Percentile	Median	75 <sup>th</sup> Percentile
<b>Baseline</b>			
Pre-merger Firms	3	4	6
Pre-merger HHI	2919	4336	5373
Welfare Cost of Markups	\$24.7	\$28.1	\$32.8
Welfare Cost of Sorting	\$1.6	\$3.2	\$5.0
<b>No Risk Adjustment</b>			
Pre-merger Firms	3	4	6
Pre-merger HHI	2943	4252	5403
Welfare Cost of Markups	\$23.3	\$27.0	\$32.5
Welfare Cost of Sorting	\$3.8	\$5.9	\$8.6

Note: HHI values are computed from the re-solved pre-merger equilibria. Welfare costs are measured in dollars per-person per-month. Quartiles are computed across the 107 local markets.

Table 5: Markets are typically concentrated

<sup>17</sup>I do not consider the possibility of new firm entry. Since insurance markets are tightly regulated and entry is somewhat difficult, this is a realistic assumption in the short-run.

<sup>18</sup>Because average costs are typically greater than marginal costs, the welfare cost of markups as defined in this paper is relative to a social optimal benchmark that typically has negative profits. I could alternatively define the welfare cost of markups relative to a zero-profit benchmark. In this case, the welfare cost of markups is \$3-\$5 lower, but still much greater than the welfare cost of sorting.



	Baseline			No Risk Adjustment		
	Number of Mergers	Fraction $\Delta SW > 0$	Fraction $\Delta CS > 0$	Number of Mergers	Fraction $\Delta SW > 0$	Fraction $\Delta CS > 0$
Total	1186	0.15	0.13	1186	0.22	0.17
<hr/>						
<u><math>\Delta</math> HHI</u>						
<100	702	0.20	0.18	692	0.29	0.21
100 - 200	164	0.08	0.06	173	0.17	0.14
200 - 1000	141	0.07	0.06	143	0.09	0.08
>1000	179	0.05	0.05	178	0.12	0.09
<hr/>						
<u>Sorting Cost</u>						
<\$5	930	0.10	0.08	260	0.05	0.03
\$5 - \$7.5	102	0.21	0.19	383	0.19	0.14
\$7.5 - \$10	129	0.30	0.27	291	0.25	0.20
>\$10	25	0.76	0.72	252	0.42	0.33

Note: In both policy environments, many mergers lead to improvements in consumer surplus and social welfare. This table displays the fraction of mergers with positive welfare effects in the baseline and no-risk-adjustment policy environments. The top line displays the average across all mergers, and the following two panels breakout the results by the size of the merger and the pre-merger welfare cost of sorting. The change in HHI is computed using pre-merger market shares to reflect pre-merger size of merging firms, and the sorting cost is measured in dollars per consumer per month.

Table 6: Many Mergers are Predicted to Improve Consumer Surplus and Social Welfare

Despite the high levels of initial concentration and low welfare cost of sorting, many mergers in both policy environments are predicted to improve *both* consumer and social welfare. The fraction of mergers that improve welfare is displayed in Table 6, broken down by the size of the merger (measured by the change in HHI predicted by pre-merger market shares) and the welfare cost of sorting in the pre-merger equilibrium. I follow the Merger Guidelines and classify the size of mergers into three categories: those unlikely to be of concern (change in HHI of less than 100), those that are potentially concerning (change in HHI of between 100 and 200), and those that are presumed to be harmful to consumers (change in HHI of greater than 200).<sup>19</sup>

There are two important facts to learn from Table 6. First, across all dimensions,

<sup>19</sup>These thresholds are applied to markets that are already concentrated and are guides for scrutiny rather than hard rules.

*some* mergers are beneficial to consumers. Even in the current policy environment where policies are in place to address adverse selection and among large mergers likely to draw intense antitrust scrutiny, consumers are better off in 1 out of 20 markets. While this is a small number of markets, it demonstrates the heterogeneity in merger effects. Just as heterogeneity in consumer substitution patterns can generate heterogeneous merger effects in apparently similar mergers, so too can heterogeneity in consumer selection patterns.

Second, the mergers that lead to greater consumer and social welfare are typically in markets with larger pre-merger welfare costs of sorting, as shown in the third panel of Table 6. Only in markets with a welfare cost of sorting of at least \$5 do an economically significant fraction of mergers benefit consumers. Intuitively, in markets where welfare costs of sorting are low, there is limited room for additional concentration to improve welfare.

The mechanism through which a merger might improve consumer surplus is by reducing the spread between generous insurance products (Gold and Platinum plans) and the less generous options (Bronze and Silver plans). Figure 3 demonstrates how this narrowing occurs in the price spread in each policy scenario: typically through significant price increases in Bronze and Silver plans and lower or negative price changes in Gold and Silver plans.

Intuitively, this is analogous to a firm increasing a fixed price for insurance while decreasing the marginal price for increasing the generosity of the insurance plan. Improving the efficiency of consumer sorting is about setting the efficient marginal price of additional insurance on the extensive margin, which is often less than the equilibrium outcome in markets with adverse selection. Figure 3 demonstrates that this is exactly the prediction from the model, and other empirical work that investigates the effect of competition on prices in the industry find similar trends (Abraham et al. (2017)). This result is related to intuition from a two-part tariff environment, where total surplus can increase if the intensive margin tariff falls closer to marginal cost despite an increase in the extensive margin markup.

Figure 4 shows the distribution of merger effects on consumer surplus in each of the three categories of merger size designated in the Merger Guidelines, plotted against the welfare cost of sorting pre-merger. Each dot represents a merger-market in a particular policy environment.

Among the smallest mergers with a change in HHI of less than 100, it is rare for a merger to lead to significant consumer harm, and occasionally a merger leads to substantial consumer benefits. It is unsurprising that these merger-markets do not generate much consumer harm. However, the presence of large benefits to consumers means that markets




Figure 3: Mergers Decrease Price Spread Between Most and Least Generous Plans

Note: A key mechanism through which mergers reduce inefficient sorting is by reducing the spread plan of different generosity. This figure shows the distribution of price effects across all merger-market-products in the simulation, divided by policy environment. The plan categories are ranked from least to most generous: catastrophic (catas), bronze, silver, gold, and platinum (plat). The dark black line represents the median effect, the box contains the inter-quartile range, and the lines extend to the 10<sup>th</sup> and 90<sup>th</sup> percentiles.

with less overlap in market share should still be considered as a source of possible benefits from a merger.

Even among larger mergers that would be considered potentially harmful under current guidelines, there exist mergers that generate substantial benefits to consumers.<sup>20</sup> And regardless of the change in HHI, mergers in markets with a pre-merger sorting cost greater than \$10 per person per month are generally beneficial to consumers with economically

<sup>20</sup>When using social welfare instead of consumer welfare, medium sized mergers are also unlikely to generate much harm with only 4% of merger-markets generating more than a \$0.50 reduction in welfare per person per month.

significant magnitudes.




Figure 4: Mergers Improve Welfare in Markets with Large Welfare Costs of Sorting

Note: Markets where the welfare cost of sorting is larger are more likely to have mergers that improve social welfare, and greater welfare costs of sorting lead to greater improvements to welfare. This figure shows the effect of each merger on social welfare in both the baseline and no-risk-adjustment policy scenarios by the welfare cost of sorting. Each dot represents a single merger-market. Sorting cost is displayed on a log scale. Both the welfare cost of sorting and the effect on social welfare are measured in dollars per person per month.

## 6.2 Screening For Mergers in the Presence of Adverse Selection

The harm from a merger comes not from the pre-merger market shares specifically but rather the substitution patterns between the merging firms. Farrell and Shapiro (2010) argue

that, while not a perfect predictor of actual price effects, their measure of upward pricing pressure (UPP) (net of cost-efficiencies) can accurately predict the direction of the effect of a merger on prices, and this logic is relatively easily extended to effects on consumer surplus (Jaffe and Weyl (2013)).<sup>21</sup> Miller et al. (2017) demonstrate that UPP can be an effective screen for harmful merger effects under many demand functions, including logit, discrete-choice demand.

In this section, I examine merger screens in the spirit of this literature and consider two measures of pricing pressure. The first is the GePP measure re-defined below in Equation (10). As derived in Section 2.2, this measure captures the full incentive of a merger in the presence of adverse selection. The second measure, is the more typical UPP measure proposed by Farrell and Shapiro (2010) and relatively easily computed by antitrust agencies. This measure is equal to the product of the average diversion ratio and the profit margin, as shown underlined in Equation (10). I follow Farrell and Shapiro (2010) and Miller et al. (2017) in considering the ratio of the pricing pressure measure to pre-merger prices as a screening index.

$$\text{GePP}_{jk} = \underbrace{\frac{-\frac{\partial S_k}{\partial p_j}}{\frac{\partial S_j}{\partial p_j}} (p_k - AC_k)}_{\text{UPP}_{jk}} + \frac{S_k}{\frac{\partial S_j}{\partial p_j}} \frac{\partial AC_k}{\partial p_j} \quad (10)$$

To highlight the importance of the screen, I consider only mergers that are *presumed* to be harmful due to a change in HHI of greater than 200.<sup>22</sup> While most of these mergers lead to significant harm to consumers, 6% of mergers in the baseline policy scenario and 10% of mergers in the no-risk-adjustment policy scenario lead to greater consumer welfare than pre-merger. The goal is how to screen for the harmful mergers in this group without investigating or blocking mergers that benefit consumers.

Consistent with the original discussion of UPP, the full GePP measure is an imperfect prediction of the magnitude of the effects of a merger on consumer welfare, but it is an accurate of the direction of the effect. In Figure 5, I plot the change in consumer surplus

<sup>21</sup>The Merger Guidelines also adopt this view: “[t]he Agencies rely more on the value of diverted sales than on the level of the HHI for diagnosing unilateral price effects in markets with differentiated products.”

<sup>22</sup>The 2010 Horizontal Merger Guidelines state: “Mergers resulting in highly concentrated markets that involve an increase in the HHI of more than 200 points will be presumed to be likely to enhance market power.”

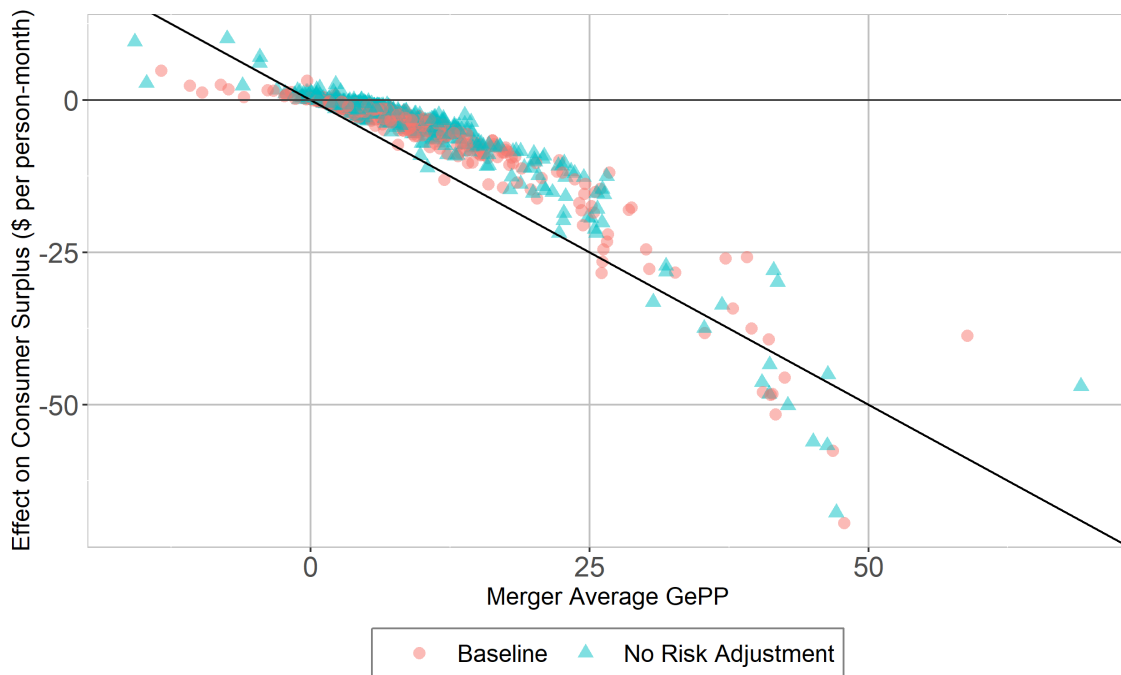


Figure 5: GePP Predicts Direction of Consumer Surplus Effect

Note: Average GePP forms a good prediction for the direction of the effect on consumer surplus. This figure compares the effect of a merger in a particular market relative to the average GePP across all products of the merging parties. Each dot represents a single merger-market. Sorting cost is displayed on a log scale. Both the welfare cost of sorting and the effect on social welfare are measured in dollars per person per month. The dark line represents the 45-degree line.

relative to the average GePP created by the merger. GePP is a conservative screen in the sense that antitrust authorities can safely allow mergers with a negative average GePP, as none of those mergers are harmful. As GePP grows larger, the merger deserves more scrutiny.

It is clear that GePP is the best measure at hand to predict which mergers are likely to be harmful and which beneficial. However, antitrust practitioners may not have the data or time to estimate inter-firm selection patterns in the process of merger review. In Figure 6, I consider the effectiveness both the full GePP measure and the standard UPP measure (averaged across merging products) as a screen for the direction of the effect of a merger on consumer surplus.<sup>23</sup>

Figure 6 shows the fraction of mergers that benefit consumers that would be “inves-

<sup>23</sup>UPP can be computed from information on diversion ratios and profit margins, which can be roughly approximated from relatively high-level information.

tigated” under a particular measure and screening threshold. On the x-axis are screening thresholds which select for investigation only mergers with an average pricing pressure measure that exceeds that level. The data points then represent that percent of mergers that benefit consumers among those mergers that exceed that screening threshold.

If the full GePP measure is used to screen mergers, it is unlikely that beneficial mergers will be investigated with any screening threshold that is greater than 0. UPP can also still be an effective screen. Using a threshold of 0.05, there are no beneficial mergers in the baseline policy scenario and very few without risk adjustment. However, it is important to note that the condition that *any* consumer harm results from the merger may be a lower than typical bar for investigation.

This perspective highlights the relationship between adverse selection and efficiencies in merger analysis. In the case of efficiencies that result from a merger, the bar is greater for upward pricing pressure to lead to significantly higher prices after the merger. Figure 6 shows that this is also the case in the presence of adverse selection, and when adverse selection is worse (as in the no-risk-adjustment policy scenario), the degree to which a UPP makes a wrong prediction about the *direction* of the effect of the merger is greater. When using UPP in the presence of adverse selection, antitrust agencies could raise the bar when screening for harmful mergers.

## 7 Conclusion

This paper demonstrates that additional concentration due to a merger in an insurance market, or any industry suffering from adverse selection, leads to a welfare trade-off. A reduction in inefficient sorting creates a welfare benefit, but the greater markups that come with more market power leads to a welfare reduction. The net effect of these two forces is an empirical question.

In the non-group insurance market, this trade-off is economically relevant. Even in the presence of transfers that are intended to address adverse selection, more than 1 out of 10 mergers lead to an improvement in consumer surplus. In markets where the welfare distortion due to sorting is greater than \$7.5 per person, more than 1 out of 3 mergers improve consumer surplus.

These results provide important insight for policy makers. From the perspective of antitrust enforcement, the degree of adverse selection in a market should be considered when evaluating a merger. The reduction in inefficient sorting due to additional concentration is a

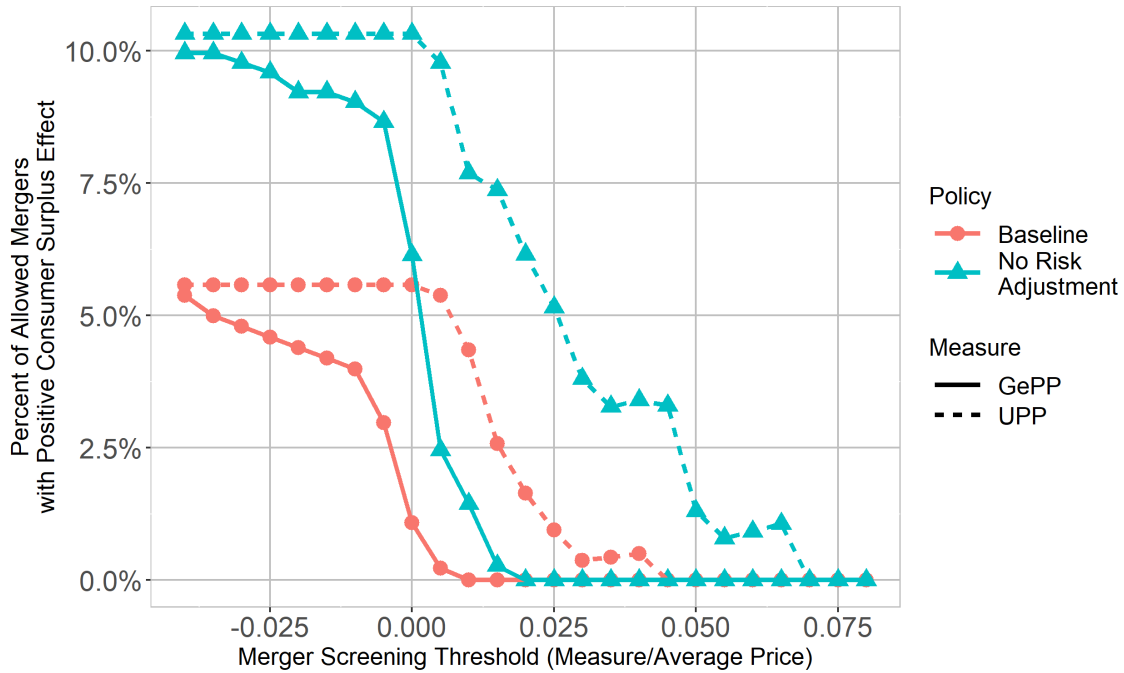


Figure 6: Traditional Screening Methods for Merger Harm Can Still Apply

Note: In the presence of adverse selection, a merger screen for investigation based on UPP risks miscalculating the direction of the effect of a merger. Each dot represents the percent of mergers which exceed a threshold of each pricing measure (GePP and UPP) and lead to greater consumer surplus, displayed for the baseline and no-risk-adjustment policy environment.

kind of merger efficiency that might lead some mergers to provide benefits for consumers. And the extent to which policies are in place to correct for the harmful incentives created by selection is important for evaluating the degree of this potential benefit due to a merger.

This paper builds on a large theoretical and empirical literature and is itself only a small additional step towards understanding managed competition in an environment with adverse selection. I show that complex selection mechanisms can be studied rigorously even in settings where contract characteristics are fixed (Chade et al. (2022)). More work remains to fully understand the effect of mergers in this market. For instance, firm entry or exit in the context of mergers (Caradonna et al. (2023)), adverse selection selection, and their interaction is an important area for future research.



## 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.
- Caradonna, Peter, Nathan Miller, and Gloria Sheu**, “Mergers, Entry, and Consumer Welfare,” *SSRN Electronic Journal*, 2023.
- Chade, Hector, Victoria Marone, Amanda Starc, and Jeroen Swinkels**, “Multidimensional Screening and Menu Design in Health Insurance Markets,” *SSRN Electronic Journal*, 2022, (April).
- 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.

— **and Sarah J. Reber**, “Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection,” 1998.

**Dafny, Leemore, Jonathan Gruber, and Christopher Ody**, “More Insurers Lower Premiums: Evidence from Initial Pricing in the Health Insurance Marketplaces,” *American Journal of Health Economics*, 2015, 1 (1), 53–81.

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

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

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

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

**Drake, Coleman**, “What Are Consumers Willing to Pay for a Broad Network Health Plan?: Evidence from Covered California,” *Journal of Health Economics*, 2019, 65, 63–77.

**Einav, Liran, Amy Finkelstein, and Mark R Cullen**, “Estimating welfare in insurance markets using variation in prices,” *The quarterly journal of economics*, 2010, 125 (3), 877–921.

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

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

**Farrell, Joseph and Carl Shapiro**, “Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition,” *The B.E. Journal of Theoretical Economics*, 2010, 10.

**Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard**, “Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts,” *American Economic Review*, apr 2019, *109* (4), 1530–1567.

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

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

**Geruso, Michael**, “Demand heterogeneity in insurance markets: Implications for equity and efficiency,” *Quantitative Economics*, 2017, *8*, 929–975.

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

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

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

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

**Gowrisankaran, Gautam, Aviv Nevo, and Robert Town**, “Mergers when prices are negotiated: Evidence from the hospital industry,” *American Economic Review*, jan 2015, *105* (1), 172–203.

**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 Health Insurance Markups,” *National Bureau of Economic Research Working Paper Series*, 2017, No. 23104.
- **and —**, “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.
- Kong, Edward, Timothy J Layton, and Mark Shepard**, “Adverse Selection and ( un ) Natural Monopoly in Insurance Markets,” 2023.
- 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.
- Miller, Nathan H., Marc Remer, Conor Ryan, and Gloria Sheu**, “Upward pricing pressure as a predictor of merger price effects,” *International Journal of Industrial Organization*, 2017, *52*, 216–247.
- 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,” *Working Paper*, 2019.
- Starc, Amanda**, “Insurer Pricing and Consumer Welfare: Evidence from Medigap,” 2014.
- Tebaldi, Pietro**, “Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA,” 2020.
- Town, Robert**, “The welfare impact of HMO mergers,” *Journal of Health Economics*, nov 2001, 20 (6), 967–990.
- and **Su Liu**, “The Welfare Impact of Medicare HMOs,” *The RAND Journal of Economics*, 2003, 34 (4), 719.

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

# Appendix

## A Data Processing

### A.1 Processing the Choice Data

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 non-group 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.

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.

In Table 7, I summarize the data sample used in estimation and compare it to other data on the non-group insurance market: the ACS and data reported by the Office of the Assistant Secretary for Planning and Evaluation (ASPE) at the U.S. Department of Health and Human Services. The ACS survey design offers the broadest depiction of the market across all market segments. ASPE publishes detailed descriptive statistics on purchases made through the federally-facilitated HealthCare.gov. Relative to the ACS, enrollment through HealthCare.gov is weighted heavily towards low-income, subsidy-eligible consumers. As a result, the plan 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.



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

Note: 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 7: Data Description

### Choice Sets

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.

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<sup>th</sup> 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**

Data on the size and demographic distributions of both the uninsured and insured populations in each market come from 2015 American Community Survey (ACS). The population of individuals who might consider purchasing non-group health insurance is any legal US resident that is not eligible for Medicaid, Medicare, and is not enrolled in health insurance through their employer. 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 non-group market. These consumers may be ineligible for subsidies but can often obtain waivers to get the same treatment as those without an employer offer. This population is small (Planalp et al. (2015)), and I treat them identically to the rest of the non-group market.

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 households. 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 household is the subset of the tax 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. 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 a head of household under the age of 65.

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 Medical Loss Ratio Data

CMS makes publicly available the state-level financial details of insurance firms in the non-group market for the purpose of regulating the MLR.<sup>24</sup> 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 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.

---

<sup>24</sup>Insurance firms in this market are restricted in how much premium revenue they may collect, relative to an adjusted measure of medical costs. In 2015, this constraint is not often binding. Excess revenue is returned to consumers via a rebate.

## 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. In reality, variation in the average age rating is not very large, and incorporating this variation in the moment matching dramatically increases the computational burden.

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.<sup>25</sup> In the estimation, I target the difference between  $\bar{R}_f$  and an adjusted average risk score for the state that accounts for grandfathered insurance products not sold through the marketplace.

---

<sup>25</sup>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.

## A.5 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 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. I use projected firm-level average cost and the average ratio of experienced costs across metal levels for all firms. Using projected average costs for the firms leads to the best fit for the first order conditions, which are not imposed in estimation. This may be because it more accurately represents firms' expectations when setting their costs. While the decision to use projected or experienced costs does affect the marginal cost estimation, it does not qualitatively impact the results.

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. I use the projected firm level average cost from the 2015 plan year firm-level rate filing data. I compute post-reinsurance projected costs by subtracting projected reinsurance payments from “projected incurred claims, before ACA Reinsurance and Risk Adjustment.”

Some firms do not appear in the risk filing data. For these firms, 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%.

## 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.

### 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  $\psi$  can no longer be separated from  $\phi$  because they interact in determining how much reinsurance an individual receives. Instead,  $\tilde{\psi}$  can be found by iteration.

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

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

### Method of Simulated Moments

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

by minimizing, for a weighting matrix  $W$ ,

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

The minimum of the function is found using the Neldermead method. 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 ??).

## C Effects of a Merger under Price-Linked Subsidies

Price-linked subsidies have two important effects in the context of mergers and market power. First, it allows firms a greater ability to exploit their market power. Greater prices are partially covered by the government rather than consumers, which reduces consumers' effective elasticity and leads to greater markups (Jaffe and Shepard (2020)).

Second, when two firms merge, the price effect is greater not only due to the reduced elasticity of consumers, but also due to the increased probability that the merged firm will control the linked product that governs the subsidy. The merged firm now internalizes more of the subsidy policy, and as a result, it is as if the merged firm faces *less elastic* consumers than pre-merger, even among the same consumers and policy environment.

Without considering selection, the first-order effects of a merger when firms internalize the price-linked nature of subsidies are that government spending increases substantially, firms capture some of this increased spending as an increase in profits, and subsidized consumers are protected against—and in some cases can benefit from—higher prices. The key group that is harmed due to a merger are higher-income, un-subsidized consumers, which make up a smaller portion of the market.

These effects are important in the context of this paper, because they greatly reduce the harm to consumers from increased markups. Consumers may benefit from mergers through less inefficient sorting without bearing the full burden of greater markups.

In this section, I repeat the main results of the paper, allowing for the subsidies to adjust with prices and for firms to internalize their probability of controlling the price-linked product, i.e. the silver plan with the second lowest price. I follow the methodology



of Jaffe and Shepard (2020) and require the equilibrium to be an ex-post best-response. The firm that controls the second-lowest-price silver plan sets the optimal price conditional on the knowledge that the plan is linked to the market-wide subsidy level.

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

$$\pi_j = \frac{e^{-\chi|p_j - p^{2lps}|}}{\sum_k e^{\chi|p_k - p^{2lps}|}}. \quad (11)$$

The parameter  $\chi$  governs the certainty with which firms' know if they offer the benchmark premium. In the limiting case of a very large  $\chi$ , 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 at least \$40. At the observed prices, the benchmark plans in 53 out of 107 markets are assigned probabilities greater than 70%, and in 88 markets the probabilities exceed 50%. With a greater certainty parameter, the equilibrium is more difficult to solve but it does not substantially alter the results of this section.

The price-linked model performs similarly in rationalizing the observed equilibrium. As shown in Figure 2, average marginal revenue and average marginal costs are similar in the baseline model used in the body of the paper. The price-linked model has similar findings, but the marginal revenue variation matches slightly less of the marginal cost variation—33% versus 37%.

Table 8 displays the main results in the price-linked model. In both policy environments, more than half of all mergers are beneficial to consumer, and more than 1 out of 5 of the largest mergers are still beneficial to consumers.

I present two total welfare measures: one that excludes government spending (as in the body of the paper) and another that includes spending. Because the primary costs of a merger in this environment are borne by the government, the vast majority of mergers increase the combined surplus accrued by consumers and firms. In the case that a dollar of government spending is associated with extra resource costs due to distortionary taxes, it may be the case that no mergers produce any welfare benefit. Similarly, if a dollar of government spending in this market is valued at less than a dollar due to preferences for

	Number of Mergers	Fraction $\Delta CS > 0$	Fraction $\Delta SW > 0$ (Govt Spending Excl.)	Fraction $\Delta SW > 0$ (Govt Spending Incl.)
Baseline				
Total	1186	0.53	0.82	0.08
$\Delta$ HHI				
<100	536	0.76	0.89	0.09
100 - 200	140	0.58	0.86	0.07
200 - 1000	318	0.32	0.80	0.08
>1000	185	0.21	0.62	0.05
No Risk Adjustment				
Total	1186	0.51	0.82	0.16
$\Delta$ HHI				
<100	542	0.72	0.89	0.17
100 - 200	136	0.51	0.82	0.10
200 - 1000	318	0.33	0.80	0.14
>1000	190	0.21	0.64	0.19

Note: When accounting for price-linked subsidies, mergers are generally beneficial to consumers. This table displays the fraction of mergers with positive welfare effects in the baseline (top panel) and no-risk-adjustment policy (bottom panel) environments. The top line in each panel displays the average across all mergers, and the following rows breakout the results by the size of the merger. The change in HHI is computed using pre-merger market shares to reflect pre-merger size of merging firms, and the sorting cost is measured in dollars per consumer per month.

Table 8: Many Mergers are Predicted to Improve Consumer Surplus and Social Welfare

redistribution, the fraction of mergers that produce a welfare benefit may be somewhere in between these two estimates.<sup>26</sup>

While more mergers are beneficial in the absence of risk adjustment, suggesting a similar role for inefficient sorting, it is hard to disentangle the mechanisms created by adverse selection from those caused by the price-linked subsidies. Even in the absence of any adverse selection, mergers that lead to greater subsidies can be beneficial for both

<sup>26</sup>As mentioned in Section 6, the total surplus generated by an additional dollar of government spending is less than a dollar.

consumers—many prices fall in absolute terms due to greater subsidies—and for firms—due to greater profit at costs borne primarily by the government. These dynamics are important considerations for this market but outside of the primary scope of this paper.