Imperfect Competition in Selection Markets*

Neale Mahoney† E. Glen Weyl‡

June 8, 2016

Abstract

Policies to correct market power and selection can be misguided when these forces co-exist. We build a model of symmetric imperfect competition in selection markets that parameterizes the degree of market power and selection. We use graphical price-theoretic reasoning to characterize the interaction between these forces. Using a calibrated model of health insurance, we show that the risk adjustment commonly used to offset adverse selection can reduce the amount of coverage and social surplus. Conversely, in a calibrated model of subprime auto lending, realistic levels of competition can generate an oversupply of credit, implying greater market power is desirable.

Keywords: selection, imperfect competition, mergers, risk-adjustment, risk-based pricing
JEL classifications: D42, D43, D82, I13, L10, L41

---

*Weyl acknowledges the financial support of the Ewing Marion Kauffman foundation which funded the research assistance of Kevin Qian. Mahoney acknowledges financial support from the Neubauer Family Foundation. We are grateful to Miguel Espinosa, Mark Sands, André Veiga; seminar participants at the 2014 AEA Meetings, Chicago Booth, the 2014 IIOC, the University of Tokyo, the Searle Antitrust Conference and the NBER Industrial Organization and Insurance meetings; the editor Amitabh Chandra and an anonymous referee for their feedback and to Tim Bresnahan Joshua Gans, Henry Mak, and Michael Whinston for excellent discussions. All errors are our own.

†University of Chicago Booth School of Business and NBER

‡Microsoft Research and University of Chicago

1 Introduction

In health insurance markets, risk adjustment is increasingly used to offset the adverse selection that occurs when consumers with higher medical costs select more generous health plans (e.g., Brown et al., 2014). Reducing adverse selection, however, may be misguided when insurance plans have market power. The reason is that firms facing adverse selection have an incentive to lower their prices to encourage lower cost "young invincibles" to buy their product. Risk adjustment, precisely because it offsets adverse selection, undermines this incentive and can lead to higher prices and lower social surplus.

Conversely, in consumer lending markets, some degree of market power can be helpful. In a perfectly competitive market, lenders have an incentive to reduce down-payment requirements to attract profitable infra-marginal customers from their rivals. These lower down-payments draw in high-risk marginal borrowers, to whom loans are socially wasteful. A monopolist lender would internalize these "cream-skimming" externalities, potentially increasing social surplus.

Thus, selection and imperfect competition interact in rich, surprising, and potentially socially important ways. Yet despite these features, we are unaware of any systematic analysis of imperfect competition in selection markets. In this paper we try to fill this gap with a price-theoretic model that builds on existing literature on both topics and can be analyzed graphically to provide intuition.

We start by presenting a model of symmetric imperfect competition in selection markets. To abstract from a particular model of imperfect competition (such as Bertrand or Cournot), we use the *conduct parameter* approach pioneered by Bresnahan (1989) and further developed in Weyl and Fabinger (2013). Market power is indexed by a parameter θ that nests, as special cases, monopoly, perfect competition, versions of symmetric Cournot competition (with or without conjectural variations), and differentiated products Bertrand competition.

This one-dimensional conduct parameter approach is enabled by the assumption that consumers' willingness-to-pay is distributed symmetrically across products. To add selection to this model, we need to strengthen this notion of symmetry to account for variation in the cost of providing the product to different consumers. Following Rochet and Stole (2002) and White and Weyl (2016), we assume that, at symmetric prices, each firm receives a sample of consumers with costs that are representative of all consumers purchasing the product, and that a firm that cuts its price steals consumers with a similarly representative distribution of costs from its competitors.

These assumptions allow us to parameterize the degree of selection—arising from different correlations between willingness-to-pay and costs across markets—with a single parameter σ. Complete selection, which occurs when there is a perfect correlation between willingness-to-pay and costs is denoted by $\sigma = 1$, and no selection, which occurs when willingness-to-pay and costs are uncorrelated, is denoted by $\sigma = 0$. Reducing the degree of selection moves individuals towards having the costs of the average individual in the full population. A particularly useful feature of our model is that changes in σ isolate the effect of changes in the degree of selection, holding fixed the full population's average costs.

We use this model to examine the welfare effects of (i) changes in the degree of market power in industries with selection and (ii) changes to the degree of selection in industries with market power. Increasing the degree of market power always reduces the amount of the product supplied, so its social utility depends on whether supply under perfect competition is excessive or insufficient. Under adverse selection, there is undersupply with perfect competition, and market power is harmful because it exacerbates this distortion. Under advantageous selection, there is oversupply with perfect competition, and social surplus is inverse-U-shaped in market power because market power, up to a point, mitigates the excess supply created by advantageous selection.

The effect of changing the degree of selection is subtle and depends on the direction of selection, the prevailing quantity in the market, and the degree of market power. A monopoly firm, for example, internalizes the costs of its marginal consumers. When there is adverse selection and fewer consumers in the market, the marginal consumers will tend to be more costly than the population average, and reducing the degree of selection will thus lower costs and the price charged by the monopolist. When there are many consumers in the market, on the other hand, the marginal consumers will be less costly than the population average, and reducing the degree of selection will raise costs and the monopolist's price.

The welfare effects of changes in selection are even subtler, as they depend on the source of the change in selection. If selection is reduced by the implementation of risk adjustment, then welfare is determined entirely by whether quantity moves towards its socially optimal level under the non-risk-adjusted costs. When the underlying primitives change, shifts in the degree of selection directly impact welfare by affecting the social cost of providing the product.

We illustrate the implications of our results with three applications. The first application is to a canonical problem in competition policy: the merger of two symmetric competitors to monopoly. We show that several standard intuitions embodied in the latest revision of the United States Horizontal Merger Guidelines (United States Department of Justice and Federal Trade Commission, 2010, henceforth HMG) are reversed in selection markets. For example, advantageous selection can generate large values of "Upward Pricing Pressure" (UPP), a standard indicator used to assess a prospective merger's harm. However, since markets with advantageous selection can have too much competition, UPP can be large exactly in settings where additional market power can be socially beneficial.

Our other two applications flesh out the examples we used to motivate our analysis. First, we build a model of health plan choice and calibrate it to data and empirical estimates on the U.S. employer-sponsored insurance market (Dafny, Duggan and Ramanarayanan, 2012; Handel, Hendel and Whinston, 2015). For the baseline parameter values, and indeed most parameter values consistent with the empirical literature, risk adjustment has the unintended consequence of reducing surplus received by the employer and its workers, and often harms social welfare. To examine the effects of market power in consumer lending, we calibrate a model of subprime auto lending to the data in Einav, Jenkins and Levin (2012). We show that in this market, which has severe advantageous selection due to limited information on borrower creditworthiness, a realistic degree of competition generates a significant oversupply of loans, providing a cost subsidy to the marginal borrower of 41%. While these calibrations do not substitute for careful empirical analysis, they suggest that the forces we highlight may be quantitatively important in canonical empirical contexts.

While our analysis is general along some dimensions, its simplicity depends on a number of stylized assumptions. Most importantly, we assume that products are symmetric(ally differentiated) and that their (non-price) characteristics are exogenously determined. We thus rule out adjustment by either consumers or firms along the dimension of product quality. While some violations of symmetry across products, such as having a firm that is horizontally (and orthogonally to costs) viewed as more desirable, is probably innocuous, allowing adjustments along the quality dimension would likely change our results in important ways. In particular, Veiga and Weyl (2016) show that market power typically has additional social benefits in the presence of selection, as it limits firms' incentives to engage in socially wasteful design of products to "cream skim" from their rivals, a point originally emphasized in Rothschild and Stiglitz (1976).

Our paper is closely related to Einav, Finkelstein and Cullen (2010) and Einav and Finkelstein (2011), who conduct a general analysis of perfectly competitive selection markets that builds on the classical theory of a natural monopoly regulated to charge a price equal to average cost (Dupuit, 1849; Hotelling, 1938).[1] While this work has been influential, a constraint in applying the framework more broadly is that the assumption of perfect competition is questionable in many important selection markets.[2] Perhaps because of this, existing work on imperfect competition has relied more heavily on structural assumptions about firm and consumer behavior (e.g., Lustig, 2010; Starc, 2014). To provide a more general treatment, we extend the price-theoretic approach of Einav and Finkelstein, conveying our results whenever possible with simple graphs and verbal descriptions, with formal mathematical statements and proofs presented in the appendix.[3]

The remainder of the paper proceeds as follows: Section 2 presents the model and Section 3 presents the main results. Section 4 presents our applications to the Horizontal Merger Guidelines, health insurance, and consumer lending. Section 5 concludes.

2 Model

In this section, we describe a model of symmetric imperfect competition that nests monopoly, perfect competition and common models of imperfect competition including Cournot and differentiated products Bertrand competition. By placing these models in a common framework, we are able to develop results that are robust to the details of the industrial organization. Our model combines the model of selection markets proposed by Einav, Finkelstein and Cullen (2010, henceforth EFC) and Einav and Finkelstein (2011, henceforth EF) with the model of imperfect competition proposed by

[1] This is an application of Marshall (1890)'s observation that competitive industries with economies or diseconomies of scale that are external to an individual firm's production would operate identically to a monopolist regulated to charge a price at average cost.

[2] In their survey on empirical models of insurance markets, Einav, Finkelstein and Levin (2010) write that "there has been much less progress on empirical models of insurance market competition, or on empirical models of insurance contracting that incorporate realistic market frictions. One challenge is to develop an appropriate conceptual framework. Even in stylized models of insurance markets with asymmetric information, characterizing competitive equilibrium can be challenging, and the challenge is compounded if one wants to allow for realistic consumer heterogeneity and market imperfections." Similarly Chiappori et al. (2006) argue that "there is a crying need for . . . models . . . devoted to the interaction between imperfect competition and adverse selection on risk."

[3] See Weyl (2015) for a detailed discussion of price theory methodology more generally.

Weyl and Fabinger (2013, henceforth WF), with suitable modifications to each to accommodate the features of the other.

Consider an industry with symmetric firms that provide symmetric, though not necessarily identical, products.[4] When firms produce symmetric quantities, prices are given by $P(q)$, where $q \in [0,1]$ denotes the fraction of consumers served by the market. We do not specify the cardinality of the firms in the market to minimize the notational burden. For most of our analysis we assume, like EF, that individuals who do not purchase the product from the industry receive no product. However, as we discuss in some detail in Subsection 2.3, the outside option may in some cases be an alternative product, as emphasized by EFC.

As in EF, and as described more formally by Weyl and Veiga (Forthcoming), total costs for the industry are summarized by the aggregate cost function $C(q)$, given by the linear aggregation of the cost of all individuals served, and associated marginal and average cost functions $MC(q) \equiv C'(q)$ and $AC(q) \equiv \frac{C(q)}{q}$. These may be increasing or decreasing in aggregate quantity depending on whether selection is respectively "advantageous" or "adverse."[5]

We assume that firms have no internal economies or diseconomies of scale, and thus no fixed costs. At a symmetric equilibrium, firms supply segments of the market that are equivalent in terms of their distribution of costs and thus have average costs equal to $AC(q)$.

Industry profits are $qP(q) - C(q) = q\left[P(q) - AC(q)\right]$. A competitive equilibrium requires that firms earn zero profits and is characterized by $P(q) = AC(q)$. A monopolist or collusive cartel chooses q to maximize profit by equating marginal revenue to marginal cost:

$$P(q) + qP'(q) \equiv MR(q) = MC(q).$$

We also follow EF in assuming quasi-linear utility in price.[6] This allows us to define consumer surplus as $CS(q) = \int_0^q \left[P(x) - P(q)\right] dx$ and *marginal consumer surplus* is $MS(q) \equiv CS'(q) = -qP'(q)$. Social welfare is $CS(q) + qP(q) - C(q)$ and the first-order conditions for the maximization of social welfare are

$$-qP'(q) + qP'(q) + P(q) - MC(q) = 0 \iff P(q) = MC(q).$$

Thus, the socially optimal quantity (constrained as we are throughout the paper to uniform prices) is characterized by $P(q) = MC(q)$.

Panel (A) of Figure 1 shows the perfectly competitive equilibrium and monopoly and social optima in the case of "advantageous selection" where $AC'(q) > 0$ and the consumers with the highest willingness-to-pay are least costly. Panel (B) shows the same in the case of "adverse selection" where

[4]Some consumers may favor one product over another, but there must be an equal number of consumers who have the symmetric opposite preference.

[5]It is possible that these slopes have different signs over different ranges or that the two have slopes of different signs over a particular range. All of these cases do not fall cleanly into one category or the other and are not our focus in what follows. It would be interesting to extend our analysis to such cases.

[6]This assumption is literally validd, for instance, in the insurance application if individuals have constant absolute risk aversion (CARA) preferences and is an accurate approximation in most cases over the range of policy changes we consider (Willig, 1976).

Figure 1: Equilibrium Under Advantageous and Adverse Selection

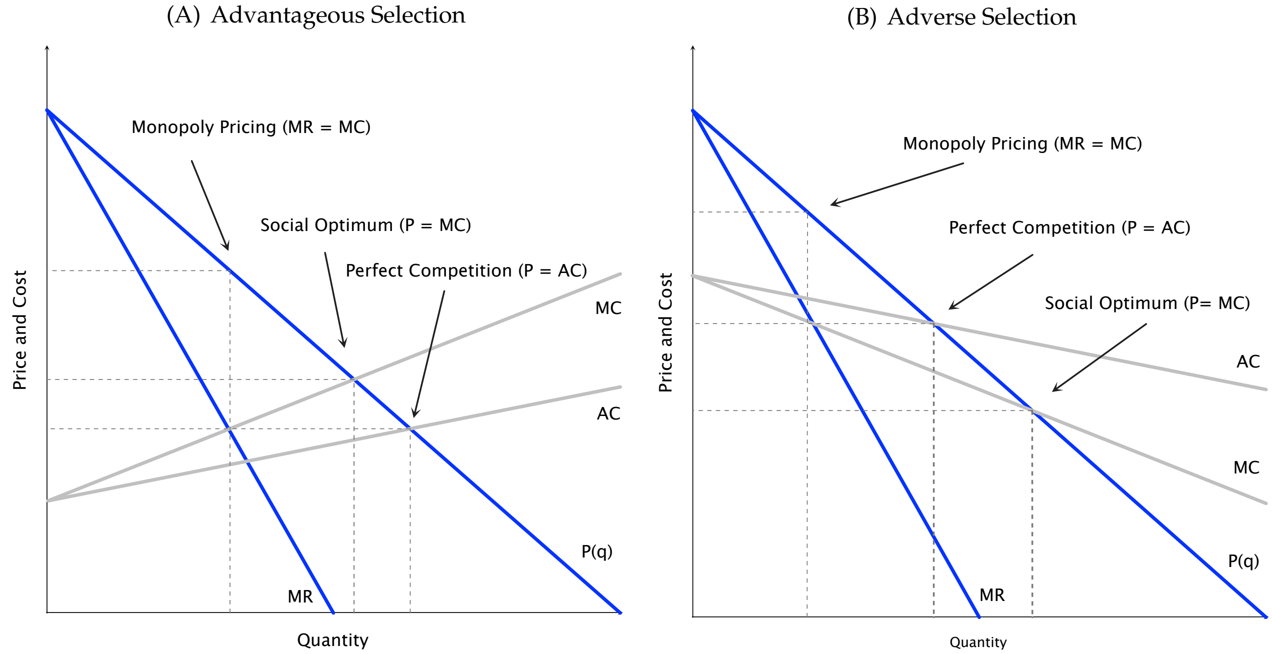

Note: This figure shows the perfectly competitive equilibrium and monopoly and social optima. Panel (A) shows these equilibria in the case of advantageous selection where average costs are upward slopping. Panel (B) shows these equilibria in the case of adverse selection where average costs slope downward.

$AC'(q) < 0$ and the consumers with the highest willingness-to-pay are most costly.[7]

2.1 Imperfect Competition (θ)

We can nest the monopoly optimization and competitive equilibrium conditions into a common framework by introducing a parameter $\theta \in [0,1]$. The parameter indexes the degree of competition in the market with $\theta = 0$ under perfect competition and $\theta = 1$ under monopoly. Equilibrium prices are given by

$$P(q) = \theta\left[MS(q) + MC(q)\right] + (1-\theta)AC(q). \quad (1)$$

Below we discuss how Equation 1 is a reduced-form representation of two canonical models of imperfect competition. Formal derivations of these representations appear in Appendix A.

1. *Cournot*: There are n symmetric firms that each choose a quantity $q_i > 0$, taking the quantity chosen by other firms as given. Price is set by Walrasian auction to clear the market so that the price is $P(q)$ where $q = \sum_i q_i$. If we assume that each firm gets a random sample of all consumers who purchase the product, then the equilibrium is characterized by Equation 1 with $\theta \equiv \frac{1}{n}$. Intuitively, just as in the standard Cournot model, firms internalize their impacts on

[7] We follow EF in defining the sign of selection in terms of the slope of the average cost curve as this determines the sign of the marginal distortion under perfect competition as $AC'(q) = \frac{MC(q)-AC(q)}{q}$.

aggregate market conditions proportional to their market share ($\frac{1}{n}$ at equilibrium) and otherwise act as price- and average cost-takers. This model can easily be extended to incorporate conjectural variations as in Bresnahan (1989); see WF for details.

2. *Differentiated Product Bertrand*: There are n single-product firms selling symmetrically differentiated products. Each firm chooses a price p_i taking as given the prices of all other firms. Consumers have a type that determines their utility for each product and their cost. The distribution of consumer types is symmetric in the interchange of any two products. In addition to these traditional assumptions of the symmetrically differentiated Bertrand model, we add two additional assumptions proposed by White and Weyl (2016) that imply our representation is valid. First, the distribution of costs is orthogonal to the distribution of preferences across products *given the highest utility a consumer can earn from any product*. Second, the distribution of utility among the *switching* consumers that definitely will buy one product but are just indifferent between any two products is identical to that among the set of all consumers who are currently purchasing a product. These two assumptions imply that the average cost of consumers that switch between firms in response to a small price change is the same as the average cost among all participating consumers.[8]

   In Appendix A we provide two micro-foundations for these assumptions. The first is a renormalized version of the Chen and Riordan (2007) "spokes" model that generalizes Hotelling (1929)'s linear city model in which the dimensions of consumer's type other than her spatial position are orthogonal to her spatial position as in Rochet and Stole (2002). The second is a discrete choice, random utility model in the spirit of Anderson, de Palma and Thisse (1992) in which, rather than utility draws being independent across products as in Perloff and Salop (1985), the *relative* utility of different products is independent of the draw of the *first-order statistic* of utilities, and the distribution of consumer costs is mean-independent of relative utilities conditional on the first-order statistic.

   In this case, again, our representation is valid if $\theta \equiv 1 - D$ where $D \equiv -\frac{\sum_{j \neq i} \partial Q_i / \partial p_j}{\partial Q_i / \partial p_i}$ is the *aggregate diversion ratio*, which, by symmetry, is independent of the i chosen at symmetric prices. Note that, unlike in the previous case, θ will not be constant in this case; it will typically increase in price and thus decline in quantity (WF).

2.2 Selection (σ)

We model a change in the degree of selection as a flattening or steepening of the industry average costs curve, because a completely flat average cost curve corresponds to a complete absence of selection. For this rotation to imply a *ceteris paribus* change in selection, it should leave some point on the average cost curve fixed. One possibility is to hold fixed average cost at the equilibrium quantity. However, under perfect competition, this rotation would leave price invariant to the degree of selection, in contrast to common intuition (Hendren, 2013). Moreover, a rotation around the equilibrium,

[8] Even if this assumption fails, so long as average switching consumers have a cost that is strictly between that of average exiting consumers and average purchasing consumers most of our results are left unchanged.

or any point other than $AC(1)$, would increase or decrease average population cost, a counterfactual that strikes us as conceptually separable from a change in the degree of selection. We therefore parameterize selection as a rotation of the industry average cost curve *holding population average costs,* $AC(1)$*, constant*.

We operationalize this concept by adding a parameter σ to the model. This parameter indexes the degree of selection with $\sigma = 0$ representing a situation in which costs are mean-independent of willingness-to-pay across individuals with $AC(q) = MC(q) = AC(1)$ and $\sigma = 1$ normalized to represent perfect correlation between costs and willingness-to-pay as in the standard uni-dimensional model of heterogeneity in Akerlof (1970) for adverse selection and de Meza and Webb (1987) for advantageous selection.

This parametrization maps to the type of regression approach taken to estimate the degree of selection in the empirical selection literature. Building upon work by Chiappori and Salanié (2000), a growing literature estimates the correlation between demand and marginal costs in range of selection markets (e.g., Finkelstein and Poterba, 2004; Bundorf, Levin and Mahoney, 2012). Consider a standard econometric model of product choice:

$$v = \widetilde{\beta}_0 + \widetilde{\beta}_1(c - \mu_c) + \epsilon.$$

Here willingness-to-pay v depends linearly on expected costs c, which are distributed normally in the population $c \sim \mathcal{N}(\mu_c, V_c)$, and a mean-zero idiosyncratic taste parameter ϵ, which is independent of costs and normally distributed $\epsilon \sim \mathcal{N}\left(0, V_v - \widetilde{\beta}_1^2 V_c\right)$. In this formulation, we parameterize the variance of v with V_v, rather than parameterizing the variance of ϵ, so that the correlation between c and v may be adjusted holding fixed the marginal distribution of v. Similarly, we normalize $\widetilde{\beta}_0$ and $\widetilde{\beta}_1$ so that changing $\widetilde{\beta}_1$ does not impact the mean of the marginal distribution of v.

In Appendix A we use standard calculations for the normal distribution to show that if we define $\sigma \equiv \left|\widetilde{\beta}_1\right| \sqrt{\frac{V_c}{V_v}}$, $MC(q) \equiv \sqrt{V_c}\Phi^{-1}(1-q) \pm \mu_c$ and

$$AC(q) \equiv \sqrt{V_c}\frac{e^{-\frac{\left[\Phi^{-1}(1-q)\right]^2}{2}}}{\sqrt{2\pi}q} \pm \mu_c,$$

we can write equilibrium conditions by replacing average costs with $\sigma AC(q) + (1-\sigma)AC(1)$ and marginal costs with $\sigma MC(q) + (1-\sigma)AC(1)$ in Equation 1. Collecting terms, this yields

$$P(q) = \theta MS(q) + \sigma\Big[\theta MC(q) + (1-\theta)AC(q)\Big] + (1-\sigma)AC(1). \tag{2}$$

Thus we have a representation of the first-order equilibrium condition where θ indexes the degree of market power and σ indexes the degree of selection in the market.

This linear interpolation between $AC(1)$ and $AC(q)$ or $MC(q)$ obviously relies on the joint normal structure of the example above. Another structure that yields the same results is if a fraction σ of the population is drawn from some arbitrary joint distribution of cost and willingness-to-pay while a fraction $1 - \sigma$ is drawn from the same marginal distributions of cost and willingness-to-pay but with the two independently distributed of one another. More generally, reductions in parameterizations of the dependence (i.e. "correlation") between cost and willingness-to-pay, holding fixed population average cost, often bring $AC(q)$ and $MC(q)$ towards $AC(1)$ at each point, though not necessarily linearly or proportionally. Given that all of the results in the next section depend only on this property of moving towards $AC(1)$ at each point, and not on the linear structure, our results apply more generally than these examples.

Nonetheless, we maintain this linear form in what follows both for expositional simplicity and because it conveniently represents one of the most commonly policies used to correct the effects of selection: risk adjustment. Medicare Advantage is a high-profile example. In the United States, elderly individuals with government health insurance can choose to opt out of the public Traditional Medicare (TM) program and purchase a private Medicare Advantage (MA) plan. For each enrollee, MA plans receive a payment from the government that is supposed to equal average costs under TM, partially risk adjusted to account for demographics and ex-ante health conditions.

We can use our framework with one additional modification to model changes in the degree of risk adjustment in this and other similar settings.[9] Let $1 - \sigma$ indicate the fraction of the difference between expected average and population average costs that is compensated for by risk adjustment. The average risk adjustment payments in this setting are $ARA(q) \equiv (1 - \sigma)\,[AC(q) - AC(1)]$ with $\sigma = 0$ indicating a setting where firms are fully compensated for any differential selection they receive and $\sigma = 1$ indicating a setting where firms receive no risk adjustment. Firms' perceived average costs are the difference between their actual average costs and the average risk adjustment payments:

$$\widehat{AC}(q) = AC(q) - ARA(q) = \sigma AC(q) + (1 - \sigma)AC(1).$$

Perceived industry marginal costs, as before, are the weighted average of marginal cost and $AC(1)$:

$$\widehat{MC}(q) = \sigma MC(q) + (1 - \sigma)AC(1).$$

The effects of risk adjustment on equilibrium price and quantity—and thus consumer and producer surplus—will be the same as a change in σ due to different correlations. However, the effect on social surplus will be different because implementing risk adjustment in this manner is not budget neutral. To reduce the degree of adverse selection, an exchange operator needs to make net payments to insurance plans, and will therefore run a deficit. To reduce the degree of advantageous selection, the exchange operator will run a surplus. As a result, social surplus depends only on whether quantity moves towards the socially optimal level under the *original, non-risk adjusted demand and cost curves*.

[9] For instance, Handel, Kolstad and Spinnewijn (2015) take an identical approach to modeling risk adjustment in their stylized model of an employer-sponsored health insurance market.

2.3 Interpretation of the Outside Option

Thus far we have focused on the case where a consumer who does not purchase from the industry receives no product. Much of the literature considers a more general case when consumers choose between two products of different quality levels and must choose one of the two. This has been formulated in several ways, some of which fit our model and others which do not.

The first setting, studied by EFC, is to view the product in the market as the *incremental* quality of a high-quality product, such as supplemental insurance coverage to "top up" a low-quality base plan. This model is fully equivalent to ours from a positive perspective. It is also equivalent from a normative perspective, so long as there are no externalities from the purchase of incremental quality on the cost of providing the low-quality base product.[10] A second, closely-related setting is when consumers choose between a high-quality product supplied according to our model and low-quality, base product provided at a fixed, administratively-set price (often zero). This is the approach used in our application to employer-sponsored health insurance in Subsection 4.2. Our model corresponds to this case if suppliers receive from the low-quality provider *baseline risk adjustment* to account for consumers' cost of service under the baseline plan.[11] This ensures that the low-quality provider is indifferent to how many customers she retains and allows for an exclusive focus on the market for the high-quality product.

To see what this baseline risk adjustment means, consider an example where costs under the low-quality baseline health insurance plan are the high-quality costs scaled down by $\lambda < 1$, as would occur with a linear actuarial rate in the absence of moral hazard. Let $\widetilde{P}(q)$ and $\widetilde{AC}(q)$ be the high-quality plan's price and average costs. Letting P_0 be the administratively-set price of the low-quality plan, the relevant price from the perspective of our model is $P = \widetilde{P} - P_0$, the net price for the high-quality plan. Similarly, letting $\lambda\widetilde{AC}(q)$ be the baseline risk adjustment payment, the relevant average cost is

$$AC(q) = \widetilde{AC}(q) - \lambda\widetilde{AC}(q) = (1-\lambda)\widetilde{AC}(q),$$

which is the average cost net of baseline risk adjustment.

Of course, baseline risk adjustment, which corresponds to $\sigma = 1$, is only one of many policies an employer or other risk adjuster might decide to pursue. A risk adjuster could, for example, make payments to fully account for costs under the high-quality plan or provide a flat subsidy and not risk adjust at all. Risk adjustment to fully cover costs under the high-quality plan would correspond to a subsidy of

$$\lambda\widetilde{AC}(q) + (1-\lambda)\left[\widetilde{AC}(q) - \widetilde{AC}(1)\right] = \text{Baseline Risk Adjustment} + \left[AC(q) - AC(1)\right],$$

[10] Such externalities could be caused by moral hazard in an insurance setting or by common-pool problems in a credit setting. For instance, Medigap supplemental insurance, which provides incremental insurance for the deductibles and coinsurance in the baseline Traditional Medicare, blunts patients' incentives to control utilization, thereby imposing an externality on baseline insurance provider (Cabral and Mahoney, 2014).

[11] Note that by the neutrality of the physical incidence of taxes it is equivalent if the risk adjustment subsidy is given as a voucher to the consumer or as a subsidy to the firm serving the consumer.

which is full risk adjustment ($\sigma = 0$) in our model. Providing a flat subsidy equal to the population average cost would correspond to

$$\lambda\widetilde{AC}(1) = \lambda\widetilde{AC}(q) + \left[\lambda\widetilde{AC}(1) - \widetilde{AC}(q)\right] = \text{Baseline Risk Adjustment} - \frac{\lambda}{1-\lambda}\left[AC(q) - AC(1)\right].$$

This can be thought about as *negative* risk adjustment in our model of an amount $\frac{\lambda}{1-\lambda}$ or $\sigma = 1 + \frac{\lambda}{1-\lambda} = \frac{1}{1-\lambda} > 1$.

A final approach, adopted by Cutler and Reber (1998) and Handel, Hendel and Whinston (2015), is to allow both the prices of the high-quality and baseline product to be endogenous. Extending this approach to imperfect competition is more analytically challenging, as it would require either an asymmetric treatment of the two plans or an equilibrium model where both plans are imperfectly competitively supplied. Such a model is an interesting direction for future research, but sufficiently different from our analysis here that we view it as beyond the scope of our work.[12]

2.4 Technical Notes

In the next section, we study equilibria characterized by Equation 2. To ensure a unique equilibrium exists, we impose global stability conditions that, while not necessary for our results, simplify the analysis. In particular we assume that $P' < \min\{AC', MC', 0\}$ and $MR' < \min\{MC', 0\}$. Under these conditions there is a unique equilibrium for a constant value of θ, the case we focus on below. While θ is not constant in the Bertrand case, all of our results below can be extended to the case of non-constant θ with appropriately generalized stability conditions at the cost of some notational complexity.

3 Results

In this section, we present results on the welfare effects of (i) market power in industries with selection and conversely (ii) selection in industries with market power. To do so, we build on the notation, equilibrium and stability conditions of the previous section. To ease the exposition, all propositions are stated verbally. When possible, the results are illustrated graphically assuming linear demand and costs, and often focusing on the extreme cases of monopoly and perfect competition. Formal statements and proofs of all results appear in Appendix B.

3.1 Imperfect Competition

Proposition 1. *Market power increases producer surplus and decreases consumer surplus*

As firms gain market power, they increasingly internalize the impact of their output decisions on equilibrium price and quantity. This leads them to raise their price so long as price slopes downward

[12]See Weyl and Veiga (Forthcoming) for a more detailed discussion of the relationship among these models under perfect competition and Veiga and Weyl (2016) for an alternative model of imperfect competition in selection markets with endogenous product quality.

more quickly than does average cost ($AC' > P'$), as implied by our stability assumptions. This internalization directly leads to higher producer surplus. The resulting higher price reduces consumer surplus by the logic of the envelope condition.

Proposition 2. *Under adverse selection, social surplus falls with market power. Anytime a market would collapse as a result of adverse selection no monopolist would choose to operate.*

With perfect competition, adverse selection leads to too little equilibrium quantity, as shown in Panel (B) of Figure 1. Since market power reduces quantity, market power only further reduces social surplus. An implication is that if the market collapses under perfect competition (Akerlof, 1970), and therefore the market generates no social surplus, no amount of market power will restore the market and enable it to contribute to aggregate welfare (Dupuit, 1844).

Thus, at least under adverse selection, standard intuitions about the undesirability of market power are confirmed. However, while these results are in this sense unsurprising, they contrast with intuitions in the contract theory literature that market power may be beneficial under adverse selection. For example, Rothschild and Stiglitz (1976) argue that imperfect competition may be necessary to sustain the existence of markets under adverse selection when non-price product characteristics are endogenous, and Veiga and Weyl (2016) show that imperfect competition can indeed restore the first-best, albeit in a stylized model. However, these analyses focus on the impacts of market power on product quality rather than on the fraction of individuals supplied. Our analysis indicates a trade-off between these quality benefits of market power and its quantity harms.[13]

Under advantageous selection our analysis more directly contradicts conventional intuitions on the impact of market power.

Proposition 3. *Under advantageous selection, there is a socially optimal degree of market power strictly between monopoly and perfect competition and social surplus is inverse-U-shaped in market power. The optimal degree of market power is increasing in the degree of advantageous selection.*

Perfect competition leads to excessive output under advantageous selection because, in an attempt to skim the cream from their rivals, competitive firms draw higher marginal cost consumers into the market (de Meza and Webb, 1987). On the other hand, a monopolist, who internalizes the industry cost and revenue curves, will produce too little. As a result, there is an intermediate degree of market power that leads to the optimal quantity being produced.

Figure 2 shows this result graphically. The monopoly equilibrium, determined by $MR = MC$, results in too little quantity. The perfectly competitive equilibrium, determined by $P = AC$, results in too much. An intermediate level of market power $\theta = \theta^*$, which leads to the equilibrium determined by $\theta^* MR + (1 - \theta^*) P = \theta^* MC + (1 - \theta^*) AC$, results in the same equilibrium level of quantity as the equilibrium achieved by setting $P = MC$ and is therefore socially optimal. Because advantageous

[13] In one extension of their baseline model, Veiga and Weyl consider a calibrated model that allows for both effects and find that an intermediate degree of market power is able to achieve welfare near the first-best and that even market power approaching monopoly leads to much higher welfare than does perfect competition. This suggests that, at least in some settings, the quality benefits may be more important than the quantity harms we emphasize here.

Figure 2: Optimal Market Power Under Advantageous Selection

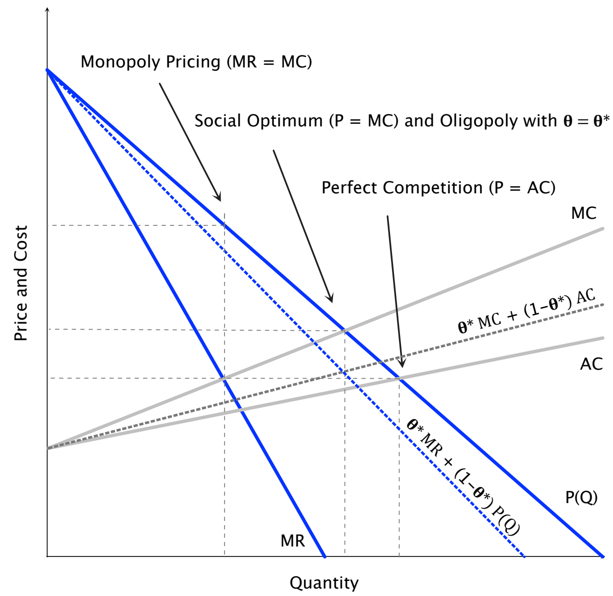

Note: This figure shows that under advantageous selection, there is a socially optimal degree of market power strictly between monopoly and perfect competition. The monopoly optimum ($MR = MC$) results in too little quantity, while perfect competition ($P = AC$) results in too much. There is intermediate level of market power θ^*, leading to an equilibrium $\theta^* MR + (1-\theta^*)P = \theta^* MC + (1-\theta^*)AC$, that results in the same equilibrium level of quantity as the socially optimum ($P = MC$).

selection always pushes firms towards excessive production, the degree of market power required to offset this selection and restore optimality increases with the extent of advantageous selection.

Appendix Table A1 summarizes our results, with Panel (A) presenting the results on market power in selection markets, discussed above.

3.2 Selection

We begin our analysis of selection by considering the impact of changes in the degree of correlation between willingness-to-pay and cost. Because the degree of correlation is a property of a market, and not the result of a policy intervention, these results apply most directly to comparative statics across markets rather than the impacts of policy interventions.[14] Our results are easiest to state verbally for the cases of monopoly and perfect competition. We thus confine our attention to these extreme cases. Results for intermediate cases are an interpolation between these extremes and are stated and proved in the formalization of these propositions in Appendix B.

Proposition 4. *Under monopoly, reducing the degree of adverse selection raises profits but can raise or lower consumer surplus. Less adverse selection harms consumers when demand is high ($q^\star > \underline{q}$). If demand is very*

[14] For example, Hendren (2013) compares outcomes in markets with different degrees of correlation under the assumption of perfect competition; our comparative statics with respect to σ would allow such analysis to be extended to imperfect competition. Hendren's analysis focuses on markets with very low quantities where the impact of reducing selection is similar, according to our results, to that under perfect competition. However, our results in Section 4.2 suggest that the presence of a realistic degree of market power could substantially alter comparative statics results like his in markets where quantity is sufficiently high.

Figure 3: Reducing Adverse Selection Under Monopoly

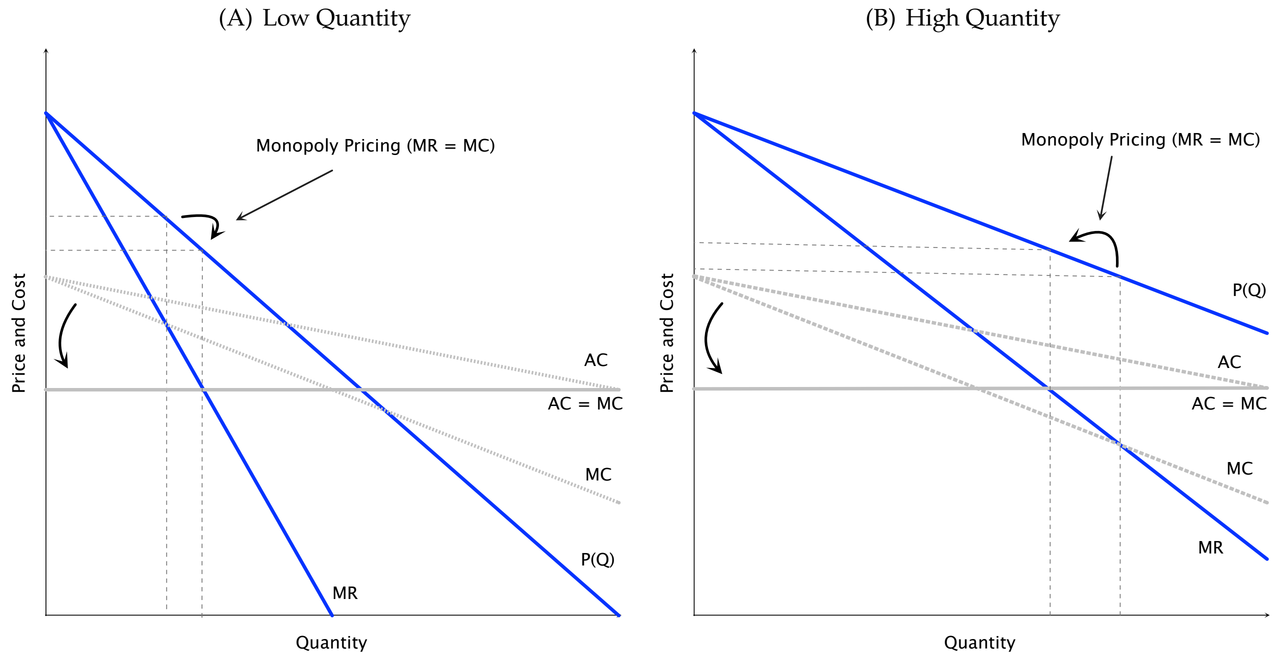

Note: This figure shows the effects of reducing the degree of adverse selection in a market served by a monopolist. Panels (A) considers a setting where the equilibrium quantity is low and reducing adverse selection lowers price and raises quantity. Panel (B) considers a setting where the equilibrium quantity is high and reducing adverse selection increases price and lowers quantity.

high ($q^\star > \overline{q} > \underline{q}$), and the monopolist's pass-through is bounded below by zero, less adverse selection lowers both consumer and social surplus.

Figure 3 shows the effect of reducing the degree of adverse selection in the market. Panel (A) shows the effect of this shift when the equilibrium quantity is low ($q^* < \underline{q}$) and Panel (B) shows the effect when the profit-maximizing quantity is high ($q^* > \underline{q}$).[15] When the equilibrium quantity is low, the reduction in selection lowers the cost of the average marginal consumer. This lowers the price and raises equilibrium quantity. When the equilibrium quantity is high, the reduction in selection raises the cost of the average marginal consumer, raising the price and lowering equilibrium quantity. In this setting with linear costs, reducing the degree of selection raises quantity whenever the profit-maximizing quantity is less than $\underline{q} = \frac{1}{2}$ because under this distribution the mean and median coincide; this result would occur for any distribution symmetric about its median, for example. More generally, reducing the degree of adverse selection reduces prices and increases quantity whenever the population average consumer has a cost lower than the average marginal consumer at the profit-maximizing level of quantity.

By the envelope theorem, we can determine the effect of a reduction in adverse selection on a monopolist's profits holding fixed the quantity the monopolist optimally chooses. Because a reduction

[15] Of course, anything that impacts equilibrium quantities must do so by shifting the demand or supply curve. The necessary thresholds for these effects, $\underline{q}$ and $\overline{q}$, can be defined as a function of the cost curves. We then interpret high and low quantities in terms of vertical shifts of the demand curve that thus vertically shift the marginal revenue curve without changing its shape.

in selection lowers average costs, as those participating in the market are selected adversely, producer surplus is necessarily increased. A reduction in the degree of adverse selection can lower welfare if the reduction in consumer surplus is large enough to offset the increase in firm profits. This only happens when profit-maximizing quantity is sufficiently high because in this case both the increase in marginal cost is large and the change in average cost is small, as the firm's average consumers are nearly representative of the whole population. The weight placed on the former effect relative to the latter effect in welfare terms is the monopolist's pass-through rate, so it must be bounded above zero at high quantities for the result to hold.

When there is advantageous selection, the conditions under which a decrease in the degree of selection raises consumer surplus are reversed.

Proposition 5. *Under monopoly, reducing the degree of advantageous selection lowers a monopolist's profits but can raise or lower consumer surplus. Less advantageous selection benefits consumers when demand is high ($q^\star > \underline{q}$). If demand is very high ($q^\star > \overline{q} > \underline{q}$) and the monopolist's pass-through is bounded away from zero, less advantageous selection raises both consumer and social surplus.*

We discuss the intuition and graphical illustration of this result in Appendix B.

Proposition 6. *Under perfect competition, reducing the degree of adverse selection raises consumer surplus and is socially beneficial. Reducing the degree of advantageous selection lowers consumer surplus and is socially harmful. Producer surplus is always zero under perfect competition.*

Under perfect competition, firms make no profits and thus the effect of selection on welfare is driven entirely by consumer surplus or, equivalently, prices. If consumers are adversely selected, consumers are always more costly than the population average, and therefore reducing the degree of selection always lowers average costs and prices, making consumers and society better off. If consumers are advantageously selected, then by the same logic, reducing the degree of selection raises average costs and prices, and reduces consumer and social surplus.

3.3 Risk Adjustment

We next consider the impact of risk adjustment, which, as discussed in the previous section, has the same positive implications as changing correlations but different normative implications.

Proposition 7. *Under monopoly and assuming demand is strictly log-concave and satisfies a weak regularity condition, using risk adjustment to eliminate adverse selection has effects that are defined by the thresholds q' and $\underline{q} > q'$, where $\underline{q}$ is defined exactly as in Proposition 4. The equilibrium quantity $q^\star$ is defined as its value after risk adjustment.*

1. *If $q^\star < q'$ then there is an interior optimal quantity of risk adjustment that achieves the socially optimal quantity. Social welfare is increasing in risk adjustment below this threshold and decreasing above it.*

2. *If $q' \leq q^\star < \underline{q}$ then welfare is monotonically increasing in risk adjustment and full risk adjustment achieves the socially optimal quantity if and only if $q^\star = q'$.*

Figure 4: Risk Adjustment of Adverse Selection Under Monopoly

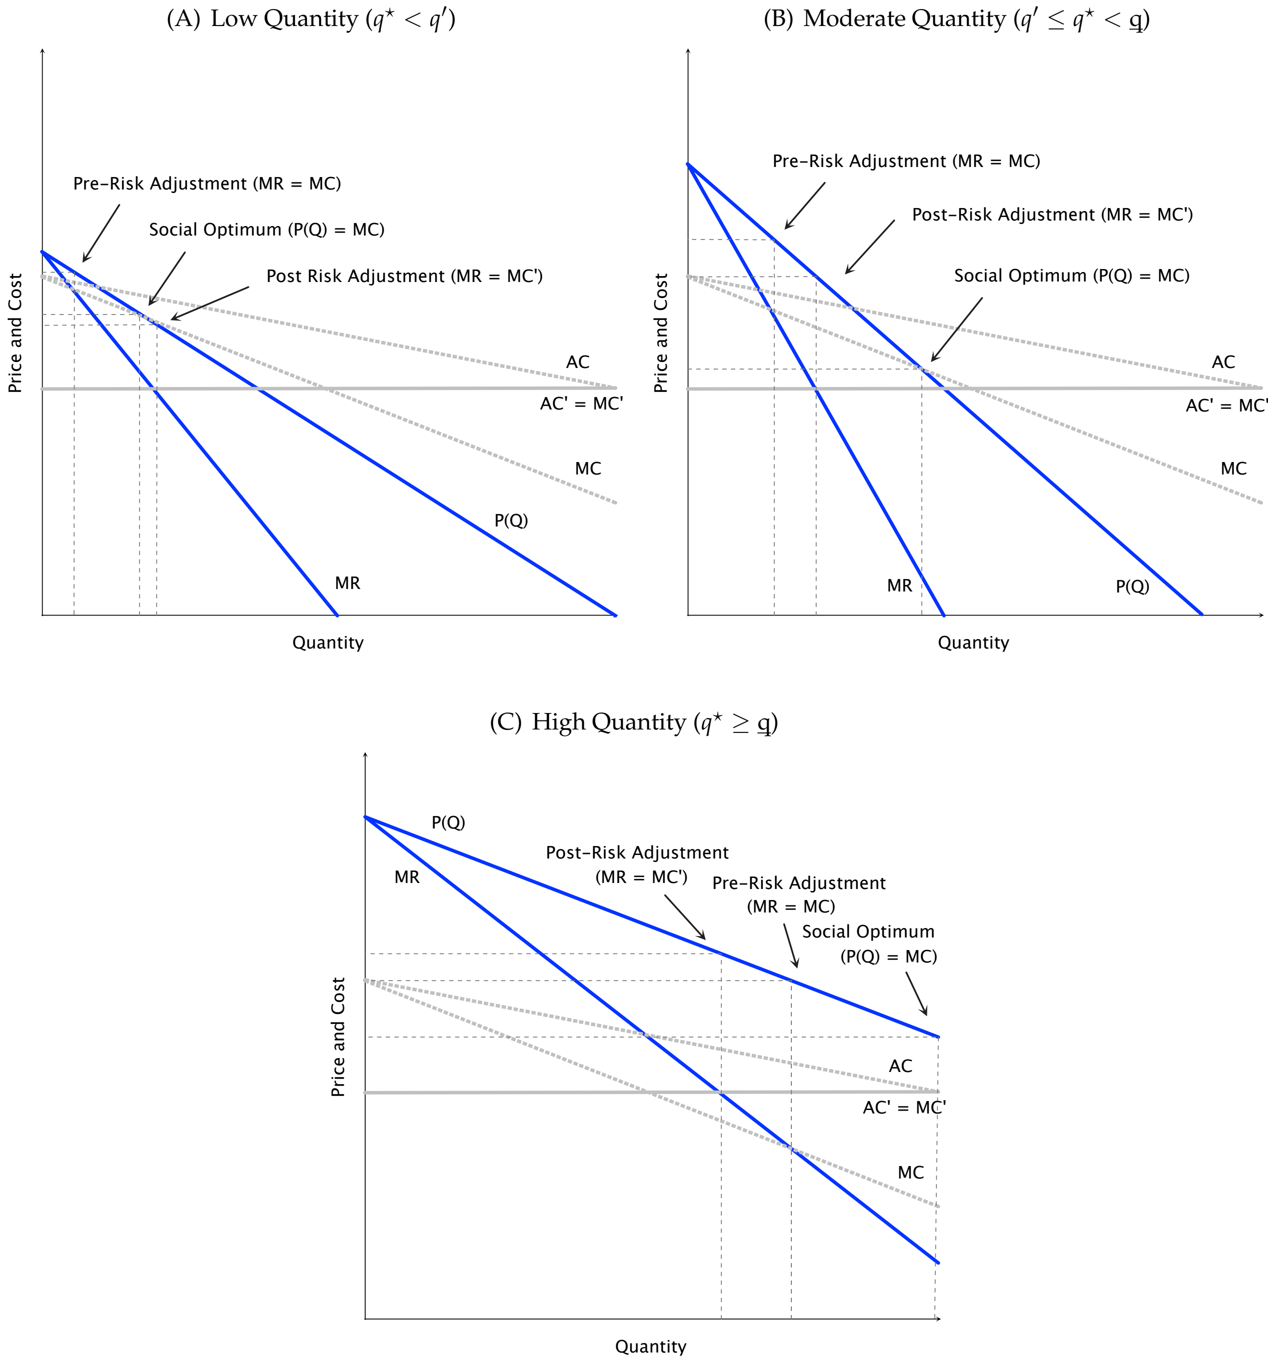

Note: This figure shows the effects of risk adjustment of adverse selection in a market served by a monopolist. Panel (A) shows a setting where full risk adjustment reduces price below the original marginal cost, leading to a socially excessive quantity. Panel (B) shows a setting where full risk adjustment is beneficial but insufficient to achieve the social optimal level of quantity. Panel (C) shows a setting where risk adjustment raises marginal costs perceived by the firm, lowering quantity and social welfare.

3. *If $q^\star \geq \underline{q}$ then risk adjustment is weakly socially harmful, and is strictly socially harmful if the inequality is strict.*

If demand is not log-concave (or violates the regularity condition) it may be that $q' = 0$ so that behavior 1) above is irrelevant or that there are multiple thresholds between 1) and 2), but one or the other always occurs when $q^\star < \underline{q}$.

Figure 4 graphically depicts these results for the different quantity ranges. The results are also summarized in Panel (A) of Appendix Figure A1. Social surplus depends on whether quantity is moved towards the socially optimal level under the *original, non-risk adjusted demand and cost curves*. Since monopoly results in too little quantity, risk adjustment that increases quantity is beneficial, so long as it does not increase quantity beyond the socially optimal level.

Panel (A) shows a setting where $q^\star < q'$. In this case, risk adjustment is initially beneficial, but full risk adjustment reduces price below the original marginal cost, leading to socially excess quantity. Intuitively, this occurs at low quantity because this is where (under log-concavity) the monopoly distortion $MS(q^*)$ is the smallest and where risk adjustment has the biggest effect on reducing perceived marginal costs. Panel (B) shows a setting where $q' \leq q^\star < \underline{q}$ and where full risk adjustment is always beneficial but insufficient to achieve the social optimal level of quantity. Indeed, in this setting it would be optimal for the exchange operator to make excess transfers to the firms. Panel (C) shows a setting where $q^\star \geq \underline{q}$, and risk adjustment raises marginal costs perceived by the firm, lowering quantity and thereby reducing social welfare.

Proposition 8. *Under monopoly and assuming that $MS' - MC'$ is globally signed, using risk adjustment to eliminate advantageous selection has effects that are defined by the thresholds $\underline{q}$ and $q'' > \underline{q}$, where $\underline{q}$ is defined as in Proposition 5. The equilibrium quantity $q^\star$ is defined as its value after risk adjustment.*

1. *If $q^\star \leq \underline{q}$ then risk adjustment is weakly harmful, and is strictly socially harmful if the inequality is strict.*

2. *If $\underline{q} < q^\star \leq q''$ then welfare is monotonically increasing in risk adjustment and full risk adjustment achieves the socially optimal quantity if and only if $q^\star = q''$.*

3. *If $q^\star > q''$ then there is an interior optimal quantity of risk adjustment that achieves the socially optimal quantity. Social welfare is increasing in risk adjustment below this threshold and decreasing above it.*

The threshold q'' may equal 1 in which case the last region irrelevant; this occurs if and only if $MC(1) < AC(1) + MS(1)$. If $MS' - MC'$ is not globally signed there may be back-and-forth between behaviors 2) and 3).

The results under advantageous selection are analogous to those under adverse selection with the regions reversed and are summarized in Panel (B) of Appendix Figure A1. When $q^\star \leq \underline{q}$, quantity is below the socially optimal level and risk adjustment further reduces quantity. When $\underline{q} < q^\star \leq q''$, risk adjustment increases quantity but is insufficient to achieve the socially optimal level. When $q^\star > q''$, there is an intermediate level of risk adjustment that increases quantity to the socially optimal level.

Under perfect competition, some risk adjustment is always beneficial although as before too much risk adjustment can sometimes be detrimental.

Proposition 9. *Under perfect competition and either adverse or advantageous selection:*

- *If $q^\star < \underline{q}$ then there is an interior optimal quantity of risk adjustment achieving the socially optimal quantity and social welfare increases in risk adjustment below and decreases in risk adjustment above this threshold.*
- *If $q^\star \geq \underline{q}$ then welfare is strictly increasing in the quantity of risk adjustment. Full risk adjustment achieves the socially optimal quantity if and only if $q^\star = \underline{q}$.*

Risk adjustment, at least initially, moves average cost towards marginal cost and thus moves quantity towards the social optimum. However, when $q^\star < \underline{q}$ it may overshoot. Under adverse selection, this occurs because $AC(1)$ is below $MC(q^\star)$. When $q^\star > \underline{q}$, even full risk adjustment is insufficient. Under adverse selection, this occurs because $AC(1)$ is above $MC(q^\star)$. When $q^\star = \underline{q}$ then full risk adjustment exactly achieves the socially optimal quantity.

In Subsection 2.3 we discussed a model where negative risk adjustment ($\sigma > 1$) is possible. Local to a small amount of risk adjustment, negative risk adjustment has precisely the opposite effect of positive risk adjustment. More globally there are various regions in terms of q where different effects may occur. In the interests of brevity, rather than cataloging these results, we here simply discuss an example that will prove relevant in our application in Subsection 4.2.

Consider negative risk adjustment in the case of monopoly when $q^\star > \underline{q}$, so that locally negative risk adjustment is beneficial. If the social optimum involves full coverage (as it will in our application), then any amount of negative risk adjustment is (weakly) beneficial as it always increases quantity. If the social optimum involves partial coverage then there is a socially optimal amount of negative risk adjustment ($\sigma^\star > 1$) that achieves the social optimum and further negative risk adjustment leads to excessive insurance coverage.

3.4 Other Forces Impacting Selection

Correlation and risk adjustment are only two of many forces that impact the extent of selection. Other commonly-discussed factors are changes in consumers' knowledge of their own costs (Handel and Kolstad, 2015) and changes in the permitted extent of risk-based pricing (Finkelstein and Poterba, 2006). Unlike the micro-foundations above, these interventions will *not only* result in a change in the cost curves but will also shift the demand curves. In the first case, this is because greater knowledge by consumers of their health risks will shift the distribution of willingness-to-pay for insurance. In the second case, characteristics that are used to price risk can also be used to price discriminate.

Because accounting for such effects requires a different analytical approach than the one we adopt here, we do not treat these forces generally. Instead, we consider specific examples that illustrate possible and plausible cases. First, in Appendix C, we show the discrimination allowed by risk-based pricing can offset or even reverse the results we derived above about the effects of selection under market power. Second, in Subsection 4.2, we use our calibrated model of the insurance market to study the impact of these changes. We find that allowing this price discrimination effect or consumer misinformation actually strengthens our main results, especially our most counterintuitive result that eliminating adverse selection may harm consumers.[16]

4 Applications

4.1 Merger Analysis

In this subsection we discuss how the results we developed above should change the welfare evaluation of mergers. In particular, we examine central principles articulated in the most recent revision of the United States Horizontal Merger Guidelines (United States Department of Justice and Federal Trade Commission, 2010) and show that many qualitative findings are altered or reversed in an industry with selection. To simplify the analysis, we focus on a symmetrically differentiated Bertrand industry in which a potential merger changes the industry from a duopoly to a monopoly. This is not intended to be a realistic applied merger model, but simply to illustrate our argument in the simplest case, which has also been emphasized in previous theoretical merger analysis (Werden, 1996; Farrell and Shapiro, 2010*a*).

1. *Price-raising incentives are harmful*: A basic principle of merger analysis is that the stronger are firms' incentives to raise prices as a result of a merger, the more antitrust authorities should suspect the merger. However, to the extent that the incentive to raise prices is stronger because of selection, mergers are likely to be more *beneficial* the stronger the incentive to raise prices.

   Consider the "first-order" incentive of a firm to raise prices after a merger (Farrell and Shapiro, 2010*a*; Jaffe and Weyl, 2013), or "Upward Pricing Pressure" (UPP), measured by the externality a firm imposes on its rivals when it increases its sales by one (infinitesimal) unit. When a firm increases its sales by one unit, it diverts D units from its rivals, where D is the aggregate diversion ratio. In a market without selection, the markup associated with this unit is $M = P - MC$ so that the sale exerts a negative externality on its rivals of $DM = D(P - MC)$.

   Suppose we naïvely calculate this object in a selection market. In a market with selection, the marginal cost perceived by an individual firm is

   $$\widehat{MC}(q) = \sigma\left(D(q)AC(q) + [1 - D(q)]\,MC(q)\right) + (1-\sigma)AC(1),$$

   so one would then compute

   $$DM = D(P - \widehat{MC}) = \left(P - \sigma\left[DAC + (1-D)\,MC\right] - (1-\sigma)AC(1)\right),$$

   where we have dropped arguments when possible for notational simplicity.

[16] However, as discussed in Subsection 4.2, price discrimination will typically increase social welfare (Cowan, 2016) and thus will not tend to generate the counter-intuitive social surplus results if one accounts for the payments made by the government for risk adjustment. Poorly informed consumers may reinforce or mitigate this effect depending on how welfare is evaluated.

In selection markets, this term does not capture the externality imposed by a firm on its rivals. Instead, our assumption that switching consumers are representative of all consumers and have costs given by AC means that the incremental profit from this unit is $P - \sigma AC - (1-\sigma)AC(1)$ and the sale creates a negative externality on rivals of $D\left[P - \sigma AC - (1-\sigma)AC(1)\right]$. As a result, the relevant UPP in selection markets is

$$\begin{aligned}\text{UPP in Selection Markets} &= D\left[P - \sigma AC - (1-\sigma)AC(1)\right] = \\ &= D\left(P - \sigma\left[DAC + (1-D)\,MC\right] - (1-\sigma)AC(1)\right) + \sigma D\,(1-D)\,(MC - AC) \\ &= \text{Standard UPP} + \sigma D\,(1-D)\,(MC-AC)\,,\end{aligned}$$

which is the standard measure plus an additional term $\sigma D(1-D)(MC - AC)$.

Increased advantageous selection (larger σ when $MC > AC$) creates more upward pricing pressure, yet is precisely the setting where market power can be a desirable check on cream-skimming externalities. Conversely, greater adverse selection (raising σ when $MC < AC$) reduces upward pricing pressure but, at the same time, is the setting where market power is most harmful because it further distorts the incentive to price above marginal cost. Thus, to the extent that it is selection rather than changes in D or M that generates upward pricing pressure, a merger is actually *most desirable* when pricing pressure is large rather than small. For the rest of this subsection, we assume $\sigma = 1$.

2. *Competition-reduction is harmful*: A second principle of merger analysis is that antitrust authorities should be suspect of mergers where the goods are close substitutes. However, in settings with *advantageous selection*, mergers between firms producing highly substitutable products are exactly the settings in which there may be too much competition and increases in market power may be beneficial. In standard analysis, a larger value of D suggests the merger is more problematic because it leads to a larger value of $UPP = D(P - MC)$. However, recall that $D = 1 - \theta$ and that under advantageous selection social surplus is inverse-U-shaped in market power. Thus if D is sufficiently small, and as a result $\theta = 1 - D$ is larger than the optimal level θ^*, the resulting merger will further increase θ above its optimal level, and always be harmful. And if D is very large, and as a result $\theta = 1 - D$ is smaller than θ^*, the resulting merger will reduce cream-skimming externalities, and may be desirable. Thus under advantageous selection mergers may be socially beneficial (absent other efficiencies) if and only if D is large enough.

3. *Marginal costs should be used to calculate markups*: A third principle is that a firm's marginal, not average, cost should be used to assess UPP. However, in selection markets, recall that the valid UPP is $D(P - AC)$ and not $D\left(P - [DAC + (1-D)MC]\right)$. Thus, if we want to use the simple formula suggested by Farrell and Shapiro (2010*a*) to calculate UPP, we should use *average cost not marginal cost* to calculate firms' markups.[17]

[17] This result is partly an artifact of our assumption that firms have additive costs across consumers, which is analogous to firms having linear (i.e., constant marginal) costs in a standard market. Accounting for firm-level (dis)economies of scale from forces other than selection would require an adjusted notion of marginal cost. Nonetheless, even in this case,

4. *Demand data is preferable to administrative data*: As a result of the focus on marginal costs, antitrust analysts often prefer demand side data to administrative data (Nevo, 2001). Marginal costs are hard to measure from firm administrative data (Laffont and Tirole, 1986), so it is standard to measure marginal costs by estimating demand and using the firm's first-order condition to recover marginal costs using the approach of Rosse (1970). However, in markets with selection, the demand-driven approach identifies the markup in

$$D\left(P - [DAC + (1 - D)MC]\right)$$

and not the relevant markup over average cost needed to calculate $D(P - AC)$. Indeed, in selection markets, demand data is insufficient, and it is necessary to have administrative data that reveals P and AC to calculate valid UPP. This implies that the administrative data obtained in recent studies of selection markets (cf. Einav, Finkelstein and Levin, 2010) are likely to be useful not only for the measurement of selection but also for antitrust policy.

In the discussion above, we considered only the price impact of a merger of a duopoly to monopoly, holding fixed all non-price product characteristics. This sort of analysis is very common, especially at the screening stage, of standard merger reviews. Enriching our analysis, for example, to allow for endogenous changes in product characteristics (called "product repositioning" in antitrust circles) would obviously complicate our conclusions, but actually seems likely to only reinforce our message that selection may fundamentally change the conclusions of standard merger analysis. In particular, in an accompanying policy piece where we discuss policy implications of this literature in greater detail (Mahoney, Veiga and Weyl, 2014), we argue that product repositioning typically makes at least some significant market power more desirable when selection is an important factor. Incorporating such repositioning explicitly into merger analysis in selection markets, and more broadly extending our results to deal with richer and more realistic settings with many asymmetric firms, is an exciting direction for future research.

4.2 Health Insurance

In this section, we examine the quantitative importance of our predictions in a benchmark health insurance example. Our model is designed to approximate basic features of an employer-sponsored health insurance setting and is similar to the models used in much of the recent literature (e.g., Einav, Finkelstein and Cullen, 2010; Handel, Hendel and Whinston, 2015; Handel, Kolstad and Spinnewijn, 2015). Below we briefly summarize the model. See Appendix D for more details.

We assume that consumers are expected utility maximizers with constant absolute risk aversion (CARA) preferences and heterogeneous absolute risk aversion, denoted α. We assume that consumers have private information about their health-type, denoted λ, which is correlated with their expected utilization. We calibrate the distribution of risk aversion using values from the lit-

firm-level marginal costs would be inappropriate for predicting UPP. And, if selection is the primary source of non-linear cost, average cost will be more accurate in predicting UPP than the standard notion of marginal cost.

erature. We calibrate the distribution of health types and medical spending using values from the 2009 Medical Expenditure Panel Survey (MEPS), under the assumption that consumers' knowledge of their future health costs is the same as that which can be predicted by standard risk adjustment software.[18] Appendix Table A2 provides more information on these parameters.

The choice set of health plans is meant to resemble those offered by a large employer. There are a number of symmetrically differentiated high-quality plans, such as Health Maintenance Organization (HMO) or Preferred Provider Organization (PPO) plans, which are supplied by private insurance companies, and have premiums determined by market forces. We define the outside option as a low-quality plan, which is provided by the employer and has a premium fixed at zero. This is a reasonable characterization of many High Deductible Health Plans (HDHP), which tend to be "self-insured," meaning that the employer bears the medical cost risk, and have administratively set premiums that are often zero or a nominal amount.[19]

As discussed in Subsection 2.3, we can fit our model to this setting by defining the products in the market as the movement from the low-quality to the high-quality plan. Let $P = \widetilde{P} - P_0$ be the incremental premium of the high quality plan, and let the functions $c_H = \kappa_H(c)$ and $c_L = \kappa_L(c)$ describe a consumer's out-of pocket costs under the high- and low-quality plans. The willingness-to-pay v for moving to the high-quality plan is the value that equates the consumer's expected utility with the high-quality plan to that with the outside option, defined implicitly by

$$\mathbb{E}_c[u(-\kappa_H(c) - v)|\alpha, \lambda] = \mathbb{E}_c[u(-\kappa_L(c))|\alpha, \lambda].$$

Consumers purchase the high-quality plan if and only if their willingness-to-pay is greater than the premium ($q = 1 \iff v \geq P$). The distribution of willingness-to-pay provides us with inverse demand and marginal revenue curves for the industry according to the standard identities.[20]

We define baseline risk adjustment as the case where the the employer compensates the private plans with each consumer's expected costs in the low-quality plan. Average costs perceived by the high-quality providers are then $AC(q) = \mathbb{E}_c\left[c - \kappa_H(c)|v \geq P(q)\right] - \mathbb{E}_c\left[c - \kappa_L(c)|v \geq P(q)\right]$ and marginal costs are $MC(q) \equiv AC'(q)q + AC(q)$. As shown in Section 2, equilibrium price for the high-quality plan is determined by Equation 2:

$$P(q) = \theta MS(q) + \Big[\theta MC(q) + (1-\theta)AC(q)\Big]$$

where θ indexes the degree of competition and we have normalized $\sigma = 1$ to the baseline degree of selection in our calibration.

We calibrate the level of market power to $\theta = 0.5$. This corresponds to Cournot competition with two high-quality plans. One justification for this choice is that in the EHBS, less than 1% of firms

[18]This assumption follows standard practice in the literature (Handel, 2012; Handel, Hendel and Whinston, 2015) and is supported by the finding from Bundorf, Levin and Mahoney (2012) of little private information conditional on an industry standard measure of predicted health risk.

[19]In the KFF EHBS, 56% of HDHP are self-insured. Twenty-three percent of HDHP are free, and 56% require employee payments of less than \$50 per month.

[20]Viz. $Q(p) = \mathbb{P}(v \geq p)$, $P(q) = Q^{-1}(p)$ and $MR(q) = P(q) + P'(q)q$.

have a choice set with more than two non-HDHP options. However, to the extent that firms bargain with insurers to be part of the choice set, it is more appropriate to define market power based on the number of insurance providers in the market. Thus, we show that our results are qualitatively robust to the market-wide concentration indices derived from Dafny, Duggan and Ramanarayanan (2012).

4.2.1 Results

A key decision for employers is how to set subsidies for the high-quality plans. Cutler and Reber (1998) argue that Harvard University's decision to provide a constant per-employee subsidy led to an "adverse selection death spiral" and the collapse of the high-quality plan. They propose that subsidies be risk adjusted to account for selection, and many employers—along with other health insurance exchanges—now implement risk adjustment schemes. However, Cutler and Reber's model, and other work that we are aware of on risk adjustment, assumes insurers are perfectly competitive, which contrasts with the findings of Dafny (2010) and Dafny, Duggan and Ramanarayanan (2012) on limited competition in employer-sponsored health insurance.[21]

Figure 5 examines the effects of risk adjustment on premiums and allocations in our calibrated model. Panel (A) shows the equilibrium with baseline risk adjustment ($\sigma = 1$). Panel (B) shows the equilibrium from an alternative calibration where we keep the demand curve unchanged and reduce the variation of health type λ, holding constant population average costs under the insurance contract.[22] Because demand is the same and there is less variation in costs, this exercise implements the same reduction in the degree of correlation between willingness-to-pay and costs that we explored theoretically. Panel (C) shows the equilibrium where an exchange operator implements full risk adjustment ($\sigma = 0$) so that consumers have constant marginal costs equal to the population average. Panel (D) shows an equilibrium with partial negative risk adjustment: in particular, the exchange operator risk adjusts subsidies by an equal and opposite amount to the full risk adjustment payments ($\sigma = 2$).[23]

With baseline risk adjustment, premiums are \$1,790 and 79.6% of the population purchases a high-quality plan. Because marginal costs are below average population costs at this equilibrium, reducing the degree of correlation increases the cost of the marginal consumer, raising premiums to \$1,820 and reducing quantity to 78.7%. Eliminating selection by means of perfect risk adjustment further raises the price and reduces the quality provided by the market. Negative risk adjustment, on the other hand, reduces premiums to \$1,682 and raises quantity to 84.1%.

Table 1 examines the normative implications of these counterfactuals. All values are presented as a percentage of the first best total surplus under the baseline scenario. Under the baseline scenario,

[21] Papers that assume perfect competition or a constant markup include Handel, Hendel and Whinston (2015), Bundorf, Levin and Mahoney (2012), Glazer and McGuire (2000), Pauly and Herring (2000), Feldman and Dowd (1982), and Carlin and Town (2009). In contrast, Dafny (2010) and Dafny, Duggan and Ramanarayanan (2012) show that not only is the insurance sector highly concentrated but that recent mergers have significantly raised premiums in the large-employer segment of the market.

[22] Because of the non-linearity of the insurance contract, holding constant population average costs under the insurance contracts requires us to adjust the mean population cost.

[23] Full negative risk-adjustment ($\sigma = 3$) has even more extreme effects in the same direction, but violates our stability conditions and thereby creates some unnecessary expositional challenges.

Figure 5: Reduced Adverse Selection in Health Insurance Model

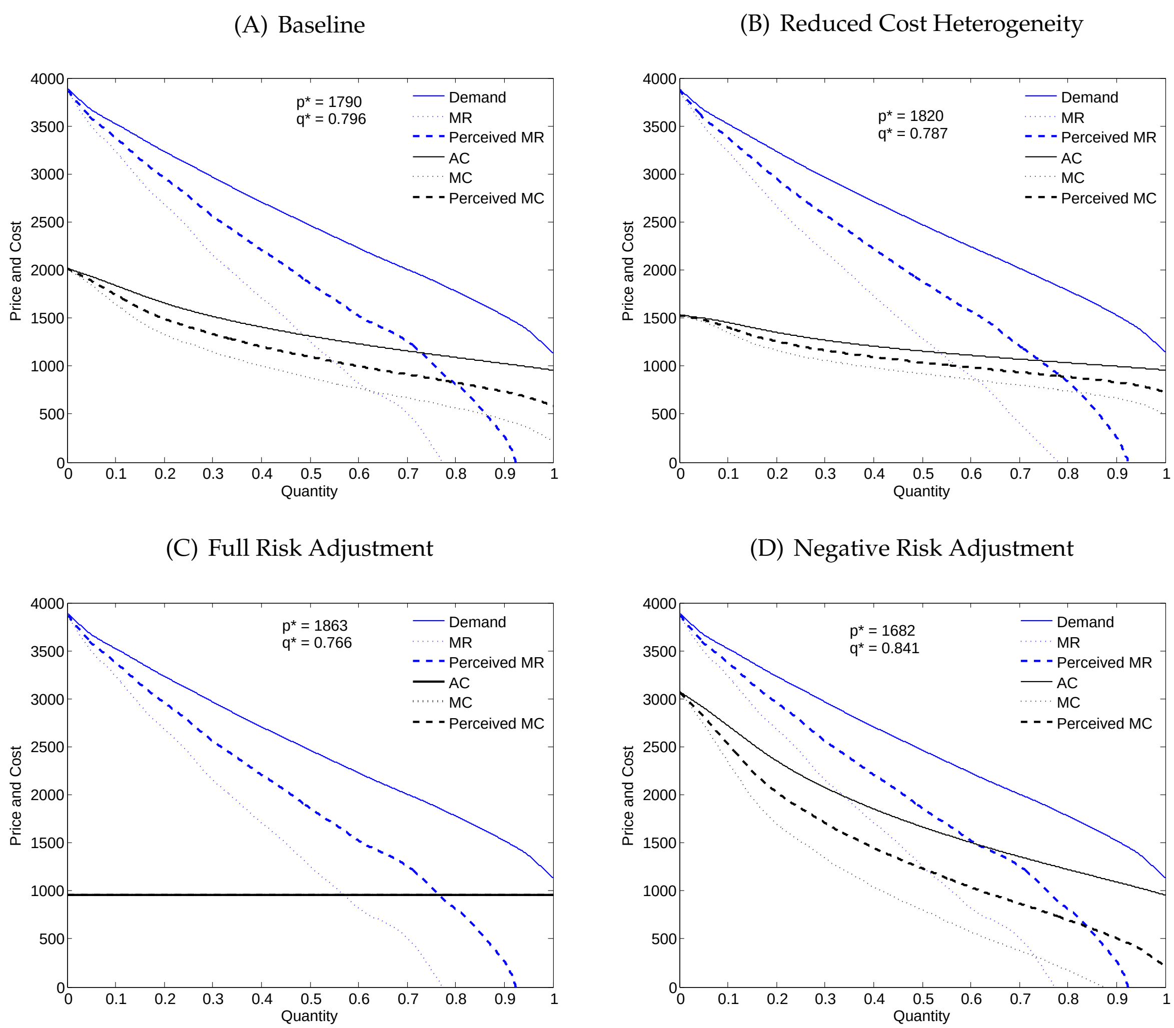

Note: This figure shows the effects of different amounts of adverse selection in the calibrated health insurance model. Panel (A) shows the baseline equilibrium ($\sigma = 1$). Panel (B) shows a scenario where the demand curve is unchanged but there is a lower correlation between willingness-to-pay and marginal costs. Panel (C) shows an equilibrium with full risk adjustment so that marginal costs are constant in the population ($\sigma = 0$). Panel (D) shows negative risk adjustment of an equal and opposite amount to the full risk adjustment payments ($\sigma = 2$).

shown in the first column, imperfect competition and selection combine to reduce total surplus to 85.7% of the first best level. Producers capture slightly less than half of this surplus, while employees capture the remainder. By raising prices, reduced correlations, shown in the second column, lowers employee surplus by 1.4 percentage points of the total surplus at the social optimum. Profits increase by 3.7 percentage points due to the lower costs of providing coverage, more than offsetting the decline in employee surplus and raising total surplus provided by the market. These results are consistent with Proposition 4 in the setting where optimal quantity takes a high, but not very high, value (i.e., $\overline{q} > q^{\star} > \underline{q}$).

Table 1: Welfare Effects of Reducing Adverse Selection

	Percent of First Best Total Surplus				
	Baseline	Reduced Cost Heterogeneity	Full Risk Adjustment	Negative Risk Adjustment	Segmented Market
Employee + Employer Surplus	49.6%	48.3%	37.1%	62.8%	32.1%
Employee Surplus	49.6%	48.3%	45.9%	55.4%	32.1%
Employer Surplus	0.0%	0.0%	-8.9%	7.4%	0.0%
Producer Surplus	36.0%	39.8%	45.0%	27.9%	67.2%
Total Surplus	85.7%	88.1%	82.1%	90.7%	99.3%

Note: This table shows the welfare effects of reducing the degree of adverse selection in the calibrated health insurance model. The first column shows welfare under the baseline equilibrium ($\sigma = 1$). The second column shows welfare in a scenario where the demand curve is unchanged but there is a lower correlation between willingness-to-pay and marginal costs. The third column shows welfare with full risk adjustment ($\sigma = 0$). The fourth column shows negative risk adjustment of an amount equal and opposite in sign to full risk adjustment ($\sigma = 2$). The fifth column shows welfare when the market is segmented into four quartiles based on consumer health type λ. All values are presented as a percentage of the first best total surplus in the baseline scenario.

Full risk adjustment, shown in column 3, exacerbates the effects of reducing correlations on employee surplus. Relative to the baseline scenario, full risk adjustment reduces employee surplus by 12.5 percentage points and increases profits by 9.0 percentage points of first best total surplus. Moreover, implementing full risk adjustment requires the employer to run a deficit equal to 8.9 percentage points of the optimized social surplus. Negative risk adjustment, shown in column 4, has the opposite effect, raising combined employee-employer surplus by 13.2 percentage points and reducing producer surplus by 8.1 percentage points relative to the baseline level. Thus, the calibrated results indicate that risk adjustment has the counterintuitive effect of reducing surplus for employees and surplus provided by the market, as described in Proposition 7 in settings where the optimal quantity is high (i.e., $q^{\star} > \underline{q}$).

Segmenting the market, shown in the fifth column, not only allows prices to reflect cost differences across employees but also allows the insurance companies to price-discriminate by charging different markups to different market segments. It, therefore, does not correspond cleanly to our pure cost-side parameter σ. To implement segmentation we partition the distribution of λ into quartiles and allow the firms to charge the profit-maximizing price to each market thus defined. Appendix Figure A3 shows plots which depict equilibrium price and quantity in each segment.

We find that the segmented markets have essentially no selection (a more-or-less flat cost curve) so that the results under segmentation reflect the elimination of selection as well as any price discriminatory effects. Segmentation reduces employee surplus by 17.5 percentage points of the optimized total surplus, which is more than the decline under full risk adjustment. The reduced selection combined with the ability to price discriminate raises profits by a substantial 31.2 percentage points of the optimized total surplus. Total surplus from the market is within 1 percentage point of first best level, but the incidence is significantly skewed, with producers capturing more than two-thirds of the surplus generated by the market. This suggests that employers' reluctance to adopt risk-based pricing may not only be due to legal restrictions and concerns about reclassification risk (Handel, Hendel and Whinston, 2015), but may also stem from the more familiar concern that allowing for price discrimination would transfer significant surplus from employees to insurance companies.

In Appendix E, we also use the model to examine the effect of a change in correlations that would result from a change in consumer perceptions about the distribution of risk they face. For instance, an insurance-choice decision-aide might reduce misperceptions of costs and therefore increase the degree of selection in the market. Reducing the correlation between perceived and actual health risk has the effect of decreasing the degree of selection in the market, as shown in Appendix Figure A4. This means that, similar to the results above, increased misperceptions raise price and reduce quantity in the market, even under the demand curves that result from perceived risk.[24] Employee surplus is even lower, and social surplus actually falls, under actual demand curves, since the misperceptions create an allocative inefficiency in who receives insurance coverage. This pushes against the argument made by Handel (2012) that in a perfectly competitive environment nudging (improving information) can hurt consumers by exacerbating the degree of selection. On the contrary, our result helps justify employer efforts to help employees optimize their health plan choice through engines provided by firms like Picwell.

4.2.2 Sensitivity and Caveats

These findings on risk adjustment and risk-based pricing are not universal. As discussed in Section 3, eliminating selection may raise or lower consumers and social surplus, and the same is famously true of the price discriminatory effects of market segmentation (Aguirre, Cowan and Vickers, 2010). To examine the robustness of these conclusions, consider Figure 6. The vertical axis represents the degree of market power θ. The horizontal axis represents the equilibrium fraction of individuals served in the market. The dots show simulated markets constructed using data on the distribution of market power based on concentration indices reported by Dafny, Duggan and Ramanarayanan (2012) and coverage rates in the EHBS.[25] The figure shows that a significant share of markets fall in the top-right region where risk adjustment reduces social surplus by lowering quantity. Furthermore, most markets fall at least into the central region, where risk adjustment reduces combined employer-

[24] Increasing misperceptions raises producer surplus, suggesting that policy efforts to de-bias consumers through decision aides may be opposed by the insurance industry.

[25] We use the reported HHI ratio to measure θ, which Cabral, Geruso and Mahoney (2014) find is an accurate proxy in a related setting.

Figure 6: Parameter Space Where Risk Adjustment is Harmful

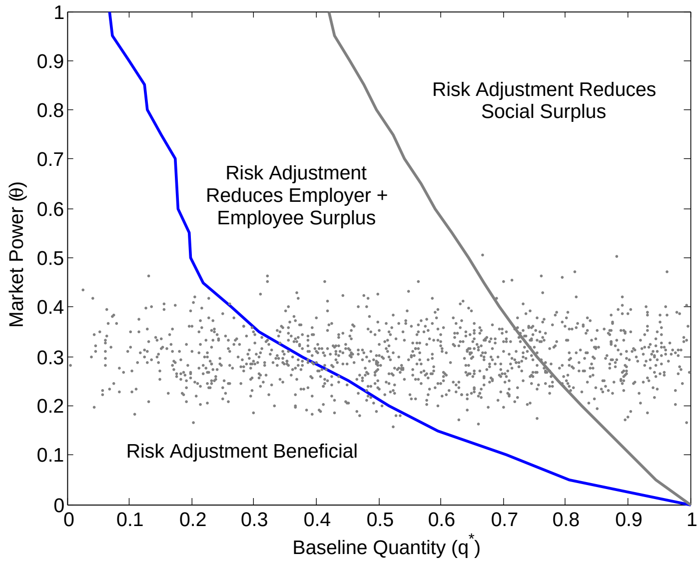

Note: This figure shows the effects of risk adjustment on welfare from the calibrated health insurance model. The vertical axis (θ) shows an index of market power. The horizontal axis ($q^{\star}$) shows shows the equilibrium fraction of consumers with coverage. Dots show simulated markets based on the distribution of market power reported in Dafny, Duggan and Ramanarayanan (2012) and coverage rates in the 2010 Employer Health Benefits Survey (EHBS). See text for details.

employee surplus. Our results are therefore fairly robust: Risk adjustment is, from the perspective of the firm and its workers, attractive in a relatively small part of the parameter space and in a minority of markets.

Our results should be taken with caution, however, because they depend heavily on our assumption that there are only two quality tiers in the market: high and baseline quality. Suppose, instead, that there were three quality tiers available, with very unhealthy individuals sorting into the highest-quality plan, intermediate-health individuals sorting into the medium-quality plan, and the healthiest individuals sorting into the baseline plan. It then seems likely that even the marginal customers of the highest quality plan would be costlier than the population average, while the average marginal customers of the medium quality plan would be a mix of the costly customers substituting from the high quality plan and the low-cost customers substituting from the baseline plan. Risk-adjustment would thus likely lower the price of the high quality plan and leave fixed the price of the medium quality plan and therefore might be beneficial in a fairly broad range of cases. Enriching our analysis to allow for such vertical differentiation in combination with market power is thus an important direction for future research.

4.3 Competition and Consumer Lending

We assess the potential for excess competition in consumer lending by using the Einav, Jenkins and Levin (2012, henceforth EJL) model of subprime auto lending, which they calibrate to proprietary data from a large firm. In their model, consumers have preferences over the down-payment d and monthly loan-payments required to payoff the total price p of the car. Costs to the firm depend on whether the consumer defaults on their loan, and conditional on default, how many payments are made prior to default and the recovery value of the car.

We apply our framework to this setting by modeling the down-payment as the "price" of the product, holding fixed the total size of the loan the consumer takes out and owes in the future. We think this choice is appropriate for two reasons. First, consumers are substantially more sensitive to down-payments than other characteristics of the loan.[26] Second, EJL consider a model with no savings so that the lender's future revenue from a borrower depends only on her type and the amount of her loan. Thus, holding fixed the loan size when changing the down-payment is the only formulation consistent, in their model, with our assumption that the cost of a consumer depends only on her type and not on the price she is charged. The market is adversely selected if lower down-payments, which increase quantity, decrease average default rates, and thereby lower costs ($AC'(q) < 0$). The market is advantageous selected if lower down-payments raise the average probability of default ($AC'(q) > 0$).

We recover the demand curve and degree of selection as *perceived by the firm* by considering the effects of small increase in the down-payment d holding fixed the size of the loan $l = p - d$ for the modal car in the data.[27] Demand is highly sensitive to the down-payment with a purchase elasticity of -0.63. There is considerable advantageous selection, with the marginal borrower with respect to the down-payment defaulting 79% of the time relative to a default rate of 59% among average borrowers.

We assess the potential for socially excess competition by calculating the *social markup* as a function of industry competition. The social markup is defined as the difference between the equilibrium price and the social marginal cost for the marginal borrower:

$$\text{Social Markup} = P - MC = \theta MS + (1-\theta)\left[AC - MC\right].$$

When the social markup is positive, there is too little equilibrium quantity and social surplus is increasing as the market becomes more competitive. When the social markup is negative, there is too much equilibrium quantity, and greater market power would improve social welfare.

The parameters of the social markup function can be recovered from the estimates of demand and selection perceived by the firm, which following the notation in Section 2 are indicated with a "hat." Because of our symmetry assumption, average costs for the industry are equal to those perceived by the firm: $AC = \widehat{AC}$. We can recover industry marginal costs for a given θ by rearranging the formula for perceived marginal costs to yield $MC = \frac{\widehat{MC}-(1-\theta)AC}{\theta}$.[28] We can similarly recover industry marginal surplus from perceived marginal surplus from the perspective of a single firm: $MS = \frac{1}{\theta}\widehat{MS}$, where $\widehat{MS} = \frac{p}{\epsilon}$ and ϵ is the absolute value of the lender's residual demand elasticity.

Figure 7 plots the social markup (y-axis) as a function of the market power parameter θ (x-axis).[29]

[26] Using the same data, Adams, Einav and Levin (2009) find that consumers are indifferent between a \$100 increase in the down-payment and a \$3,000 increase in the total amount borrowed.

[27] The modal car has a down-payment of \$1,000 and a loan size of \$10,000.

[28] The formula for perceived marginal costs is $\widehat{MC} = \theta MC + (1-\theta)AC$.

[29] We only consider variation in θ and not other parameters because, given our symmetry assumptions, all other param-

Figure 7: Social Markup

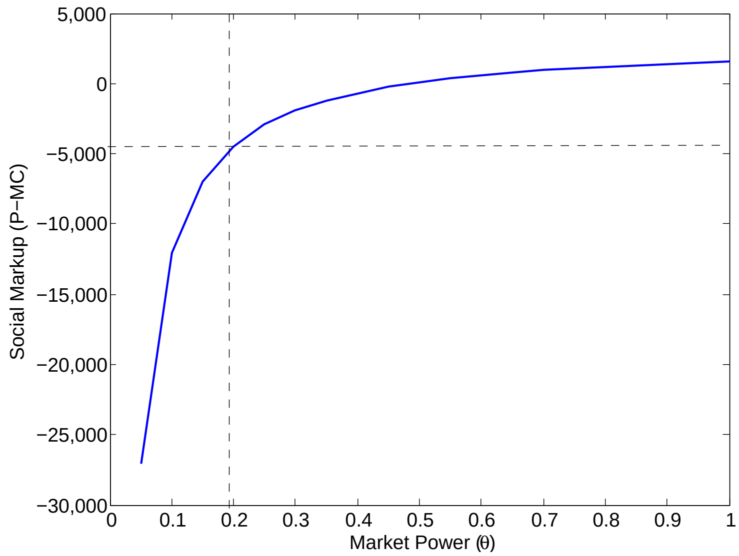

Note: This figure shows the social markup $P - MC$ (y-axis) of the down-payment on a $10,000 car loan as a function of the degree of market power (x-axis). The social markup on the down-payment is defined as the difference between the required down-payment amount and the cost for the marginal borrower, with a negative value indicating a subsidy to borrowers. The values are calibrated using the model and estimated parameters from Einav, Jenkins and Levin (2012)'s study of subprime auto lending.

The value $\theta = 0.2$ is a useful benchmark—with symmetric firms in Cournot competition it corresponds to an HHI of 2,000, just above the threshold the Department of Justice used to define markets as highly concentrated during this period. For $\theta = 0.2$, the marginal borrower is subsidized by $4,462 or 41% of the price of the car. Indeed, the marginal borrower receives a subsidy for all $\theta < 0.5$, or symmetric Cournot duopoly, indicating that high levels of concentration may be desirable.

The exercise above restricts attention to a single dimension of competition. However, as EJL highlight, competition is potentially very multidimensional. To take the simplest example, there is no clear reason that down-payments need to move in lock-step with the total price for the car and thus leave the debt burden fixed. Veiga and Weyl (2016) analyze the EJL data in a model allowing for endogenous multidimensional contract terms and find that our basic conclusion, that increased market power may be socially desirable, is actually reinforced by allowing for endogenous multidimensional contracts. Competition leads to cream-skimming, distorting the nature of and not just the price of contracts. However, they find that in the EJL data, this would lead to excessively high, not excessively low, down-payment requirements. Whether lending is overall too lax or too tight thus depends on whether the cream-skimming effect they emphasize or the advantageous selection effect we focus on is of greater magnitude. Their quantitative analysis suggests the later, but a richer empirical equilibrium analysis would be necessary to precisely resolve this ambiguity.

Thus, as we discuss in greater detail in an accompanying policy piece (Mahoney, Veiga and Weyl, 2014), while our conclusion about the surprising potential social benefits of market power seems

eters are identified by EJL's model.

fairly robust, the implications for financial regulation should be approached with caution.

5 Conclusion

This paper makes three contributions. First, we propose a simple but general model nesting a variety of forms of imperfect competition in selection markets. Second, we derive from this model several basic, yet often counter-intuitive, comparative statics. Third, we show the empirical and policy relevance of these comparative statics by applying them to merger policy and calibrated models of health insurance and subprime auto lending.

Our work here suggests several directions for future research. We have shown calibrated and empirical examples where the counter-intuitive comparative statics we derived are relevant. However, it is not clear how prevalent such examples are or the extent to which the issues we raise are first-order in determining optimal competition or selection policy. Further empirical research is required to investigate these questions. We have also focused on a small number of policy instruments: merger policy, risk adjustment, risk-based pricing and consumer information campaigns. While these may be the most canonical policies for addressing selection and market power, many other policies, such as price controls and restraints on exclusive dealing, play an important role. Studying these policies in imperfectly competitive selection markets would be informative.

As we repeatedly emphasize above, a central limitation of our analysis is our assumption of symmetry across products and that non-price product characteristics are exogenous. These together imply that uptake of the product varies only along the extensive margin. However, in many selection markets, such as insurance markets, the intensive margin is at least as important as the extensive margin, especially with the increasing prevalence of mandates for basic coverage. As a result, the literature on selection markets has increasingly turned to studying the quality of products offered (Veiga and Weyl, 2016; Azevedo and Gottlieb, 2016). While Veiga and Weyl (2016) analyze the impact of market power in detail, the richest work in this literature (Azevedo and Gottlieb, 2016) focuses on perfect competition. We hope future research will study more richly the interaction between market power and the intensive margin of product uptake.

References

Adams, William, Liran Einav, and Jonathan Levin. 2009. "Liquidity Constraints and Imperfect Information in Subprime Lending." *American Economic Review*, 99(1): 49–84.

Aguirre, Iñaki, Simon George Cowan, and John Vickers. 2010. "Monopoly Price Discrimination and Demand Curvature." *American Economic Review*, 100(4): 1601–1615.

Akerlof, George A. 1970. "The Market for "Lemons": Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84(3): 488–500.

Anderson, Simon P., André de Palma, and Jacques-François Thisse. 1992. *Discrete Choice Theory of Product Differentiation.* Cambridge, MA: MIT Press.

Azevedo, Eduardo M., and Daniel Gottlieb. 2016. "Perfect Competition in Markets with Adverse Selection." http://ssrn.com/abstract=2491867.

Bresnahan, Timothy F. 1989. "Empirical Studies of Industries with Market Power." In *Handbook of Industrial Organization*. Vol. 2, , ed. Richard Schmalensee and Robert Willig, 1011–1057. North-Holland B.V.: Amsterdam, Holland.

Brown, J., M. Duggan, I. Kuziemko, and W. Woolston. 2014. "How Does Risk Selection Respond to Risk Adjustment? New evidence from the Medicare Advantage Program." *American Economic Review*, 104(10): 3335–3364.

Bulow, Jeremy I., and Paul Pfleiderer. 1983. "A Note on the Effect of Cost Changes on Prices." *Journal of Political Economy*, 91(1): 182–185.

Bundorf, Kate, Jonathan Levin, and Neale Mahoney. 2012. "Pricing and Welfare in Health Plan Choice." *American Economic Review*, 102(7): 3214–3248.

Cabral, Marika, and Neale Mahoney. 2014. "Externalities and Taxation of Supplemental Insurance: A Study of Medicare and Medigap." http://ssrn.com/abstract=2372163.

Cabral, Marika, Michael Geruso, and Neale Mahoney. 2014. "Does Privatized Health Insurance Benefit Patients or Producers? Evidence from Medicare Advantage." *NBER Working Paper*.

Carlin, Caroline, and Robert Town. 2009. "Adverse Selection, Welfare and Optimal Pricing of Employer-Sponsored Health Plans." http://goo.gl/KVdRJf.

Chen, Yongmin, and Michael H. Riordan. 2007. "Price and Variety in the Spokes Model." *Economic Journal*, 117(522): 897–921.

Chiappori, Pierre-André, and Bernard Salanié. 2000. "Testing for Asymmetric Information in Insurance Markets." *Journal of Political Economy*, 108(1): 56–78.

Chiappori, Pierre-André, Bruno Jullien, Bernard Salanié, and François Salanié. 2006. "Asymmetric Information in Insurance: General Testable Implications." *RAND Journal of Economics*, 37(4): 783–798.

Cohen, Alma, and Liran Einav. 2007. "Estimating Risk Preferences from Deductible Choice." *American Economic Review*, 97(3): 745–788.

Cowan, Simon. 2016. "Welfare-Increasing Third-Degree Price Discrimination." *RAND Journal of Economics*, 47(2): 326–340.

Cutler, David M., and Sarah J. Reber. 1998. "Paying for Health Insurance: the Trade-Off between Competition and Adverse Selection." *Quarterly Journal of Economics*, 113(2): 433–466.

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

Dafny, Leemore S. 2010. "Are Health Insurance Markets Competitive?" *American Economic Review*, 100(4): 1399–1431.

de Meza, David, and David C. Webb. 1987. "Too Much Investment: A Problem of Asymmetric Information." *Quarterly Journal of Economics*, 102(2): 281–292.

Dupuit, Arsène Jules Étienne Juvénal. 1844. *De la Mesure de L'utilité des Travaux Publics.* Paris.

Dupuit, Jules. 1849. "De L'Influence des Péages sur l'Utilitée des Voies de Communication."

Einav, Liran, Amy Finkelstein, and Jonathan Levin. 2010. "Beyond Testing: Empirical Models of Insurance Markets." *Annual Review of Economics*, 2(1): 311–336.

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

Einav, Liran, and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Empirics in Pictures." *Journal of Economic Perspectives*, 25(1): 115–138.

Einav, Liran, Mark Jenkins, and Jonathan Levin. 2012. "Contract Pricing in Consumer Credit Markets." *Econometrica*, 80(4): 1387–1432.

Fabinger, Michal, and E. Glen Weyl. 2014. "A Tractable Approach to Pass-Through Patterns with Applications to International Trade." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2194855.

Farrell, Joseph, and Carl Shapiro. 2010*a*. "Antitrust Evaluation of Horizontal Mergers: An Economic Alternative to Market Definition." *Berkeley Electronic Journal of Theoretical Economics: Policies and Perspectives*, 10(1).

Farrell, Joseph, and Carl Shapiro. 2010*b*. "Recapture, Pass-Through and Market Definition." *Antitrust Law Journal*, 76(3): 585–604.

Feldman, Roger D, and Bryan E Dowd. 1982. "Simulation of a Health Insurance Market with Adverse Selection." *Operations Research*, 30(6): 1027–1042.

Finkelstein, Amy, and James Poterba. 2004. "Adverse selection in insurance markets: Policyholder evidence from the UK annuity market." *Journal of Political Economy*, 112(1): 183–208.

Finkelstein, Amy, and James Poterba. 2006. "Testing for Adverse Selection with "Unused Observables"." http://economics.mit.edu/files/699.

Glazer, Jacob, and Thomas G McGuire. 2000. "Optimal Risk Adjustment in Markets with Adverse Selection: An Application to Managed Care." *American Economic Review*, 1055–1071.

Handel, Benjamin. 2012. "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts." *American Economic Review*, 103(7): 2643–2682.

Handel, Benjamin, Jonathan T. Kolstad, and Johannes Spinnewijn. 2015. "Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets." *Working Paper*.

Handel, Benjamin R., and Jonathan T. Kolstad. 2015. "Health Insurance for "Humans": Information Frictions, Plan Choice, and Consumer Welfare." *American Economic Review*, 105(8): 2449–2500.

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

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

Hotelling, Harold. 1929. "Stability in Competition." *Economic Journal*, 39(153): 41–57.

Hotelling, Harold. 1938. "The General Welfare in Relation to Problems of Taxation and of Railway and Utility Rates." *Econometrica*, 6(3): 242–269.

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

Laffont, Jean-Jacques, and Jean Tirole. 1986. "Using Cost Observation to Regulate Firms." *Journal of Political Economy*, 94(3): 614–641.

Lustig, Joshua. 2010. "Measuring Welfare Losses from Adverse Selection and Imperfect Competition in Privatized Medicare." http://neumann.hec.ca/pages/pierre-thomas.leger/josh.pdf.

Mahoney, Neale, André Veiga, and E. Glen Weyl. 2014. "Competition Policy in Selection Markets." *CPI Antitrust Chronicle*, 10(1).

Marshall, Alfred. 1890. *Principles of Economics.* London: Macmillan and Co.

Nevo, Aviv. 2001. "Measuring Market Power in the Ready-to-Eat Cereal Industry." *Econometrica*, 69(2): 307–342.

Pauly, Mark V, and Bradley J Herring. 2000. "An Efficient Employer Strategy for Dealing with Adverse Selection in Multiple-Plan Offerings: An MSA Example." *Journal of Health Economics*, 19(4): 513–528.

Perloff, Jeffrey M., and Steven C. Salop. 1985. "Equilibrium with Product Differentiation." *Review of Economic Studies*, 52(1): 107–120.

Rochet, Jean-Charles, and Lars A. Stole. 2002. "Nonlinear Pricing with Random Participation." *Review of Economic Studies*, 69(1): 277–311.

Rosse, James N. 1970. "Estimating Cost Function Parameters Without Using Cost Data: Illustrated Methodology." *Econometrica*, 38(2): 256–275.

Rothschild, Michael, and Joseph E. Stiglitz. 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, 90(4): 629–649.

Starc, Amanda. 2014. "Insurer Pricing and Consumer Welfare: Evidence from Medigap." *RAND Journal of Economics*, 45(1): 198–220.

United States Department of Justice, and Federal Trade Commission. 2010. "Horizontal Merger Guidelines." http://www.justice.gov/atr/public/guidelines/hmg-2010.html.

Veiga, André, and E. Glen Weyl. 2016. "Product Design in Selection Markets." *Quarterly Journal of Economics*, 126(2): 1007–1056.

Werden, Gregory J. 1996. "A Robust Test for Consumer Welfare Enhancing Mergers Among Sellers of Differentiated Products." *Journal of Industrial Economics*, 44(4): 409–413.

Weyl, E. Glen. 2015. "Price Theory." http://goo.gl/dzHgQw.

Weyl, E. Glen, and André Veiga. Forthcoming. "Pricing Institutions and the Welfare Cost of Adverse Selection." *American Economic Journal: Microeconomics*.

Weyl, E. Glen, and Michal Fabinger. 2013. "Pass-Through as an Economic Tool: Principles of Incidence under Imperfect Competition." *Journal of Political Economy*, 121(3): 528–583.

White, Alexander, and E. Glen Weyl. 2016. "Insulated Platform Competition." http://papers.ssrn.com/sol3/papers.cfm?abstract_id=1694317.

Willig, Robert D. 1976. "Consumer's Surplus Without Apology." *American Economic Review*, 66(4): 589–597.

Appendix

A Model

This appendix provides formal micro-foundations for the representations in the text.

A.1 Cournot model

Potential consumers of a homogeneous service are described by a multi-dimensional type $t = (t_1, \dots t_T)$ drawn from a smooth and non-atomic distribution function $f(t)$ with full support on a hyper-box $(\underline{t}_1, \bar{t}_1) \times \cdots (\underline{t}_T, \bar{t}_T) \subseteq \mathbb{R}^T$. Consumers receive a quasi-linear utility $u(t) - p$ if they purchase the service for price p. When the prevailing price is p, therefore, the set of consumers purchasing the service is $T(p) = \{t : u(t) \geq p\}$ and the number of purchasers $Q(p) = \int_{T(p)} f(t)dt$. $T(p)$ is clearly decreasing in p in the strong set order so that by our assumption of full support $Q(p)$ is strictly decreasing. Thus we can define the *inverse demand function* $P(q)$ as the inverse of $Q(p)$.

Each consumer also carries with her a cost of service, $c(t) > 0$ that must be incurred to supply the service to her by any supplier. Thus the average cost of all individuals served when the aggregate quantity is q is

$$AC(q) \equiv \frac{\int_{T(P(q))} c(t)f(t)dt}{Q\left(P(q)\right)}.$$

There are n firms that can each choose a quantity q_i of the service to supply non-cooperatively. If $q \equiv \sum_i q_i < 1$ then the prevailing market price is set by by market clearing as $P(q)$. If $q > 1$ then price is 0. Clearly no equilibrium can involve $q > 1$ as all firms would make losses. Firms receive a uniform random sample of all customers who are in the market at the prevailing prices and thus earn profits $q_i \left[P(q) - AC(q)\right]$. Thus, to maximize profits non-cooperatively they must satisfy

$$P(q) - AC(q) + P'(q)q_i - \frac{MC(q) - AC(q)}{q} q_i = 0.$$

At a symmetric equilibrium where $q_i = \frac{q}{n}$ for all i this becomes

$$P(q) - \left(1 - \frac{1}{n}\right) AC(q) - \frac{MS(q)}{n} - \frac{MC(q)}{n} = 0$$

as claimed in the text.

A.2 Differentiated Bertrand model

There are n firms $i = 1, \dots n$ each selling a single service. Consumers are described by two types, each possibly multidimensional, (t, ϵ). t is drawn as in the Cournot case. ϵ consists of two components: $\epsilon = (l, e)$ where l is an integer between 1 and L, with each value of l having equal probability, and e is drawn from a real hyper rectangle in E dimensions. The distribution of e is atomless, symmetric in all coordinates, independent of the value of l and given by the distribution function g. The distributions of t and ϵ are independent.

Consumers may consume at most a single service and receive a quasi-linear utility from consuming the service of firm i, $u_i(t, \epsilon) - p_i$, where p_i is the price charged for service i. Let the *first order statistic* of utility $u^\star(t, \epsilon) \equiv \max_i u_i(t, \epsilon)$. We assume (without loss of generality yet) that $u^\star(t, \epsilon) = u^\star(t)$;

that is that the value of the first-order statistic depends only on t and not on ϵ. Second, and this does entail a loss of generality, we make the following assumption.

Assumption 1. *$u_i = u^\star(t) + \hat{u}_i(\epsilon)$ so that all valuations shift up uniformly with a shift in $u^\star$ induced by changes in t.*

This implies that the relative utility of services other than the one the individual most prefers, compared to that which she most prefers, are determined purely by ϵ and not t. Third we assume, with only a modest loss of generality, that $u^\star(t)$ is smooth in t and that $\partial u^\star / \partial t_T > k > 0$ for some constant k. This implies that raising t_T sufficiently causes $u^\star > u$ for any fixed u and lowering it sufficiently causes the reverse to be true.

Services are symmetrically differentiated in the sense that distribution of $u(t,\epsilon) = (u_1(t,\epsilon), \ldots u_n(t,\epsilon))$ induced by the distribution of (t,ϵ) is symmetric in permutations of coordinates. The set of individuals purchasing service i is

$$T_i(p) = \{(t,\epsilon) : u_i(t,\epsilon) \geq p_i \wedge i \in \text{argmax}_i u_i(t,\epsilon) - p_i\}$$

and the demand for good i is thus $Q_i(p) = \int_{T_i(p)} f(t,\epsilon) d(t,\epsilon)$.

As in the Cournot example, the cost of serving a consumer depends on her type. However, we make the substantive assumption now that cost depends only on t and not on ϵ.

Assumption 2. *The cost of serving a consumer of type (t,ϵ) is $c(t)$ and thus the total cost faced by firm i is $C_i(p) = \int_{T_i(p)} c(t) f(t,\epsilon) d(t,\epsilon)$.*

This assumption states that only the determinants of the highest possible utility a consumer can achieve, and not of her relative preferences across services, may directly determine her cost to firms. Given the independence of t and ϵ, this assumption implies a clean separation between determinants of relative "horizontal" preferences across services and "vertical" utility for the most preferred service that also determines the cost of service. Absent this assumption it is possible that the consumers that firms attract from their rivals when lowering their price are very different in terms of cost from the average consumers of the service more broadly.

Let $1 \equiv (1, \ldots, 1)$. Then by symmetry $Q_i(p1) = Q_j(p1)\, \forall i,j$ and similarly for C_i and C_j. Let the *aggregate demand* $Q(p) \equiv nQ_i(p1)$ for any i and similarly for aggregate cost. Then we define the inverse demand function $P(q)$ as the inverse of the aggregate demand. Average cost is then $AC(q) \equiv \frac{C(P(q))}{q}$ and marginal cost $MC(q) \equiv C'(P(q)) P'(q)$.

We now describe two particular models satisfying these assumptions and show how they yield the reduced-form representation we use in the text. Any other micro-foundation of these assumptions should also yield our representation, but the notation required to encompass different cases is sufficiently abstract and not relevant enough to any results we derive. We thus omit it here and focus on specific micro-foundations.

First consider a random utility model in the spirit of Anderson, de Palma and Thisse (1992) proposed by White and Weyl (2016) in the context of heterogeneity of preferences for non-price product characteristics. $L = n$ and the value of l represents which product is the individual's favorite. $e = (e_1, \ldots, e_E)$ and $E \geq n - 1$. We assume that

$$u_i(u^\star(t), l, e)$$

is increasing in $e_{i^\star}$ where $i^\star$ is i if $i < l$ and is $i - 1$ if $i > l$ and that it is constant in all other e_i where $i \leq n - 1$ and not $i^\star$. We also assume that u_i is smooth in its arguments other than l, bounded and that and that $\lim_{e_{i^\star} \to \overline{e_{i^\star}}} u_i(u^\star(t), l, e) = u^\star(t)$ and $\lim_{e_{i^\star} \to \underline{e_{i^\star}}} u_i(u^\star(t), l, e) = 0$ for any value of the other entries $u^\star$, l and $e_{-i^\star}$ where $\underline{e_i}$ and $\overline{e_i}$ are respectively the lowest and highest values of e_i. This implies that raising $e_{i^\star}$ sufficiently for any i while holding fixed the other components of e makes (in the limit) service i equally desirable to the most desirable service for the individual and lowering it makes it always uncompetitive with the best service regardless of the price differential.

An individual firm i's profits are $p_i Q_i\left(p_1, \ldots p_i, \ldots p_n\right) - C_i\left(p_1, \ldots p_i, \ldots p_n\right)$. Thus the first-order condition for the optimization of any firm i is

$$p_i \frac{\partial Q_i}{\partial p_i} + Q_i = \frac{\partial C_i}{\partial p_i}. \tag{3}$$

Because price does not appear in the interior of the integrals defining Q_i and C_i, the derivatives of these with respect to p_i is, by the Leibniz Rule applied to multidimensional integrals (Weyl and Veiga, Forthcoming), given by the sum of the effects of the extensive margin effects from the change in the boundaries of integration. There are many such boundaries, so we use a shorthand notation for them. $\partial T_i^X(p) \equiv \{(t,\epsilon) \in T_i(p) : u^\star(t) = p_i\}$ denotes the set of *exiting* consumers from product i who are just indifferent between buying service i and no service. $\partial T_{ij}^S(p) \equiv \left\{(t,\epsilon) \in T_i(p) \cap T_j(p)\right\}$ denotes the set of *switching* consumers between services i and j who are just indifferent between the two services, but prefer purchasing one over purchasing nothing. To formally define the density of consumers on such boundaries it is useful to express the multidimensional integrals representing Q_i and C_i more explicitly.

At symmetric prices p, every individual i with t_T above this threshold buys from her most preferred services l and any individual below this threshold buys no service. If a single price p_i is elevated to $p_i + \delta$ then all individuals with $l \neq i$ continue to buy their preferred product as at symmetry. However, individuals with $l = i$ and $u^\star(t) \in (p, p+\delta)$ will stop consuming any service and those with $l = i$ and $e_{j^\star}$ sufficiently close to $\overline{e_{j^\star}}$ will switch to purchasing service j. Let $t_T^\star(p; t_{-T})$ be defined implicitly by $u^\star\left(t_{-T}, t_T^\star(p; t_{-T})\right) = p$ and let $e_{j^\star}^\star\left(\Delta; u^\star(t), e_{-\{1,\ldots,n-1\}}\right)$ be implicitly defined for positive Δ by

$$u^\star(t) - u_j\left(u^\star(t), e_{j^\star}^\star\left(\Delta; u^\star(t), e_{-\{1,\ldots,n-1\}}\right) e_{-\{1,\ldots,n-1\}}\right) = \Delta$$

where $e_{-\{1,\ldots,n-1\}}$ is all components of e other than the first $n-1$ and the dependence of u_j on the other components of e is dropped as these do not impact u_j.

Then when prices are symmetric except for price p_i being above the other prices, we can write

$$Q_i\left(p, \ldots, p_i, \ldots, p\right) =$$

$$\frac{1}{n}\int_{t_{-T}}\int_{e_{-\{1,\ldots,n-1\}}}\int_{t_T^\star(p_i;t_{-T})}^{\overline{t_T}}\int_{\underline{e_1}}^{e_1^\star\left(p_i - p; u^\star(t), e_{-\{1,\ldots,n-1\}}\right)}\cdots\int_{\underline{e_n}}^{e_n^\star\left(p_i - p; u^\star(t), e_{-\{1,\ldots,n-1\}}\right)} f(t)g(e)d\left(t,e\right)$$

and similarly

$$C_i\left(p, \ldots, p_i, \ldots, p\right) =$$

$$\frac{1}{n}\int_{t_{-T}}\int_{e_{-\{1,\ldots,n-1\}}}\int_{t_T^\star(p_i;t_{-T})}^{\overline{t_T}}\int_{\underline{e_1}}^{e_1^\star\left(p_i - p; u^\star(t), e_{-\{1,\ldots,n-1\}}\right)}\cdots\int_{\underline{e_n}}^{e_n^\star\left(p_i - p; u^\star(t), e_{-\{1,\ldots,n-1\}}\right)} c(t)f(t)g(e)d\left(t,e\right).$$

To fill in the first-order condition (Equation 3), we need to differentiate these using the Leibniz rule.

$$\frac{\partial Q_i}{\partial p_i}\left(p1\right) =$$

$$-\frac{1}{n}\left[\int_{\partial T_i^X(p1)}\frac{f(t_{-T},t_T^\star(p;t_{-T}))g(e)}{\partial u^\star/\partial t_T\left(t_{-T},t_T^\star(p;t_{-T})\right)}d\left(t_{-T},e\right)+(n-1)\int_{\partial T_{ij}^S(p1)}\frac{f(t)g\left(e_{-j^\star},\overline{e_{j^\star}}\right)}{\partial u_j/\partial e_j\left(t,e_{-j^\star},\overline{e_{j^\star}}\right)}d\left(t,e_{-j^\star}\right)\right] \tag{4}$$

for any $j \neq i$ by symmetry and similarly

$$\frac{\partial C_i}{\partial p_i}(p1) =$$

$$-\frac{1}{n}\left[\int_{\partial T_i^X(p1)}\frac{c\left(t_{-T},t_T^\star(p;t_{-T})\right)f\left(t_{-T},t_T^\star(p;t_{-T})\right)g(e)}{\partial u^\star/\partial t_T\left(t_{-T},t_T^\star(p;t_{-T})\right)}d(t_{-T},e)+(n-1)\int_{\partial T_{ij}^S(p1)}\frac{c(t)f(t)g\left(e_{-j^\star},\overline{e_{j^\star}}\right)}{\partial u_j/\partial e_j\left(t,e_{-j^\star},\overline{e_{j^\star}}\right)}d\left(t,e_{-j^\star}\right)\right]. \tag{5}$$

By contrast and following the same logic

$$\frac{dQ_i}{dp}(p1) = -\frac{1}{n}\int_{\partial T_i^X(p1)}\frac{f\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)g\left(e\right)}{\partial u^\star/\partial t_T\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)}d\left(t_{-T},e\right)$$

and

$$\frac{dC_i}{dp}(p1) = -\frac{1}{n}\int_{\partial T_i^X(p1)}\frac{c\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)f\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)g\left(e\right)}{\partial u^\star/\partial t_T\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)}d\left(t_{-T},e\right).$$

Thus by symmetry

$$Q'(p) = -\int_{\partial T_i^X(p1)}\frac{f\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)g\left(e\right)}{\partial u^\star/\partial t_T\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)}d\left(t_{-T},e\right)$$

and

$$C'(p) = -\int_{\partial T_i^X(p1)}\frac{c\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)f\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)g\left(e\right)}{\partial u^\star/\partial t_T\left(t_{-T},t_T^\star\left(p;t_{-T}\right)\right)}d\left(t_{-T},e\right),$$

so that

$$MC\left(Q(p)\right) = \frac{\int_{\partial T_i^X(p1)}\frac{c\left(t_{-T},t_T^\star(p;t_{-T})\right)f\left(t_{-T},t_T^\star(p;t_{-T})\right)g(e)}{\partial u^\star/\partial t_T\left(t_{-T},t_T^\star(p;t_{-T})\right)}d\left(t_{-T},e\right)}{\int_{\partial T_i^X(p1)}\frac{f\left(t_{-T},t_T^\star(p;t_{-T})\right)g(e)}{\partial u^\star/\partial t_T\left(t_{-T},t_T^\star(p;t_{-T})\right)}d\left(t_{-T},e\right)}.$$

Furthermore

$$\int_{\partial T_{ij}^S(p1)}\frac{c\left(t\right)f\left(t\right)g\left(e_{-j^\star},\overline{e_{j^\star}}\right)}{\partial u_j/\partial e_j\left(t,e_{-j^\star},\overline{e_{j^\star}}\right)}d\left(t,e_{-j^\star}\right) = n\int_{T_i(p)}c(t)f(t)dt\int_{e_{-j^\star}}\frac{g\left(e_{-j^\star},\overline{e_{j^\star}}\right)}{\partial u_j/\partial e_j\left(t,e_{-j^\star},\overline{e_{j^\star}}\right)}d\left(t,e_{-j^\star}\right) =$$

$$AC\left(Q(p)\right)Q(p)\int_{e_{-j^\star}}\frac{g\left(e_{-j^\star},\overline{e_{j^\star}}\right)}{\partial u_j/\partial e_j\left(t,e_{-j^\star},\overline{e_{j^\star}}\right)}d\left(t,e_{-j^\star}\right) \equiv -AC\left(Q(p)\right)s(p)$$

where $s(p)$ is the density of consumers diverted to a rival from a small increase in one firms price starting from symmetric prices p. Thus we can rewrite Expression 4 as

$$\frac{Q'(p)-(n-1)s(p)}{n} = Q'(p)\frac{1}{n\left[1-D\left(Q(p)\right)\right]},$$

where $D(q) \equiv -\frac{(n-1)s(P(q))}{Q'(P(q))-(n-1)s(P(q))}$ is the *aggregate diversion ratio* (Farrell and Shapiro, 2010*b*), the fraction of consumers lost to a small increase in prices by a single first that go to rivals rather than the outside good. We can also rewrite expression 5 as

$$Q'(p)\frac{MC\left(Q(p)\right)+\frac{D(Q(p))}{1-D(Q(p)}AC\left(Q(p)\right)}{n}.$$

Then Equation 3 becomes, at symmetric prices

$$Q'(p)\frac{p}{n\left[1-D\left(Q(p)\right)\right]}+\frac{Q(p)}{n}=Q'(p)\frac{MC\left(Q(p)\right)+\frac{D(Q(p))}{1-D(Q(p)}AC\left(Q(p)\right)}{n}\implies$$

that at any symmetric equilibrium

$$\frac{P(q)}{1-D(q)}-MS(q)=MC(q)+\frac{D(q)}{1-D(q)}AC(q)$$

because $MS(q)=\frac{Q(P(q))}{Q'(P(q))}$. Letting $\theta(q)\equiv 1-D(q)$ this becomes

$$P(q)-\theta(q)MS(q)=\theta(q)MC(q)+\left[1-\theta(q)\right]AC(q)$$

as reported in the text.

A second model that delivers our form builds on the Chen and Riordan (2007) "spokes" extension of the Hotelling linear city model, combining it with modifications from Rochet and Stole (2002). There are n firms $i=1,\dots,n$. For every pair of firms, (i,j) with $i<j$ there is a line segment of unit length of potential consumers who will only consider purchasing either service i or service j. Thus there are $\frac{n(n-1)}{2}$ such segments and we denote the segment (i,j) by the integer $\frac{i(i-1)}{2}+(j \mod i)$. $\epsilon=(l,e)$ where l is the integer representing the line segment on which the consumer lives and $e\in(0,1)$ is the distance of the consumer from i or $1-$ her distance from j. In particular let $i(l)\equiv\max_{i\in\mathbb{Z}:\frac{i(i-1)}{2}<l}\frac{i(i-1)}{2}$ and let $j(l)\equiv l \mod i(l)$; then e is the distance of the consumer from $i(l)$. There are an equal number of consumers on each segment so $\frac{2}{n(n-1)}$ of the consumers are on each segment.

In addition to maintaining our assumptions about t and ϵ, we make two modifications to the set-up of Chen and Riordan:

1. We modify the exact form of consumer utility. In particular, $u^\star(t)$ is the utility a consumer earns from service $i(l)$ if $e\leq\frac{1}{2}$ and from good $j(l)$ if $e\geq\frac{1}{2}$ *regardless of the other details of her position*. This contrasts with the standard Chen and Riordan, and Hotelling (1929), model because it implies no transport cost to an individual's most preferred service.

2. Consumers' highest possible utility is not constant across consumers but instead follows a distribution $u^\star(t)$.

3. The gross utility a consumer derives from purchasing from $j(l)$ if $e<\frac{1}{2}$ is $u^\star(t)-(1-2e)t$, where t is a transportation cost parameter absent in the Chen and Riordan model. If $e>\frac{1}{2}$ the consumer derives gross utility of $u^\star(t)-(2e-1)t$ from purchasing from $i(l)$.

4. We allow arbitrary smooth and symmetric-about-$\frac{1}{2}$ distributions of e on the unit interval, as long as this distribution is the same for all l.

Calculations to derive the representation in the text are tedious and extremely similar to those in our modified Anderson, de Palma and Thisse model above. We therefore omit these calculations and simply explain why there results are the same. At symmetric prices, every consumer purchases from her most preferred firm, $i(l)$ if $e\leq\frac{1}{2}$ and $j(l)$ if $e>\frac{1}{2}$. All consumers with the same t make the same purchase decision at this price because only $u^\star$ impacts their total utility. Consumers with $e=\frac{1}{2}$ are "switchers" between a pair of firms (if $u^\star(t)\geq p$) and have the same distribution of t as all purchasers by the independence of ϵ and t. Thus switchers will be representative of the full population of consumers and exiters everywhere will be on average identical. This is precisely what gave rise to our structure above.

A.3 Parameterizing selection with Gaussian heterogeneity

Consumers purchase the product if and only if their willingness-to-pay is greater than the price:

$$q = 1 \iff v > p \iff \widetilde{\beta}_0 + \widetilde{\beta}_1 c + \epsilon > p.$$

If we divide through by the standard deviation of the taste parameter $\sqrt{V_\epsilon} = \sqrt{V_v - \widetilde{\beta}_1^2 V_c}$ and define $\beta_2 = 1/\sqrt{V_v - \widetilde{\beta}_1^2 V_c}$ and the coefficients $\beta_i = \beta_2 \widetilde{\beta}_i$ for $i = 0, 1$, the model can be estimated by a Probit regression of product choice on expected costs and premiums, assuming we have a source of exogenous variation in premiums:

$$\Pr(q = 1 | c, p) = \Phi(\beta_0 + \beta_1 c - \beta_2 p),$$

and the parameters μ_c and V_c can be estimated directly from the data: $\widetilde{\beta}_1 = \beta_1/\beta_2$ and $V_v = 1/\beta_2^2 + \widetilde{\beta}_1^2 V_c$.

Standard properties of the normal distribution and some algebra yield that

$$\widehat{MC}(q) = \mathbb{E}[c | v = P(q)] = \widetilde{\beta}_1 \sqrt{\frac{V_c}{V_v}} \left[\sqrt{V_c} \Phi^{-1}(1-q) \pm \mu_c \right] + \left(1 - \left| \widetilde{\beta}_1 \right| \sqrt{\frac{V_c}{V_v}} \right) \mu_c$$

and average cost is

$$\widehat{AC}(q) = \mathbb{E}[c | v \geq P(q)] = \widetilde{\beta}_1 \sqrt{\frac{V_c}{V_v}} \left[\sqrt{V_c} \frac{e^{-\frac{\left[\Phi^{-1}(1-q)\right]^2}{2}}}{\sqrt{2\pi} q} \pm \mu_c \right] + \left(1 - \left| \widetilde{\beta}_1 \right| \sqrt{\frac{V_c}{V_v}} \right) \mu_c$$

where the $\pm$ has the sign of $\widetilde{\beta}_1$. To fit our domain of $\sigma \in (0, 1)$, we define $\sigma \equiv \left| \widetilde{\beta}_1 \right| \sqrt{\frac{V_c}{V_v}}$, which is always between 0 and 1 because it is the absolute value of the correlation between v and c. Then letting $MC(q) \equiv \sqrt{V_c} \Phi^{-1}(1-q) \pm \mu_c$ and

$$AC(q) \equiv \sqrt{V_c} \frac{e^{-\frac{\left[\Phi^{-1}(1-q)\right]^2}{2}}}{\sqrt{2\pi} q} \pm \mu_c,$$

we can write equilibrium conditions by replacing average costs with $\sigma AC(q) + (1-\sigma) AC(1)$ and marginal costs with $\sigma MC(q) + (1-\sigma) AC(1)$ in Equation 1. Collecting terms this yields

$$P(q) = \theta MS(q) + \sigma \Big[\theta MC(q) + (1-\theta) AC(q) \Big] + (1-\sigma) AC(1). \tag{6}$$

B Proofs

Throughout we assume that $\theta, \sigma \in [0, 1]$, that selection is either globally adverse or advantageous (either $AC', MC' > 0$ or $AC', MC' < 0$ for all q) and impose a global equilibrium stability condition: $P' < \min\{AC', MC', 0\}$ and $MR' < \min\{MC', 0\}$. Most of the results may be obtained absent these global monotonicity assumptions, but the additional expositional complexities add little insight. We also assume that θ and σ are constant parameters, independent of q; all results can be extended to the case when this fails, but again, the additional notation is cumbersome.

Lemma 1. *Let* $F(q) \equiv P(q) - \sigma\left(\theta MC(q) + (1-\theta)\, AC(q)\right) + (1-\sigma)AC(1) - \theta MS(q)$. *Then* $F' < 0$.

Proof. The derivative of the expression is

$$P' - \sigma\theta MC' - \sigma(1-\theta)AC' - \theta MS' =$$

$$\sigma\left[\theta\left(MR' - MC'\right) + (1-\theta)(P' - AC')\right] + [1-\sigma]\left[\theta MR' + (1-\theta)P'\right] < 0.$$

by our monotonicity assumptions. □

Proposition (Formal) 1. *For* $\theta \in (0,1)$, $\frac{\partial PS}{\partial \theta} \geq 0 \geq \frac{\partial CS}{\partial \theta}$, *with strict inequality if* $q^\star > 0$.

Proof. By the implicit function theorem,

$$F'\frac{\partial q^\star}{\partial \theta} - \sigma\left[MC\left(q^\star\right) - AC\left(q^\star\right)\right] - MS\left(q^\star\right) = 0 \implies \frac{\partial q^\star}{\partial \theta} = \frac{MS\left(q^\star\right) + \sigma\left[MC\left(q^\star\right) - AC\left(q^\star\right)\right]}{F'}.$$

Focusing on the numerator

$$MS + MC - AC = -P'q + \sigma AC'q = q\left[\sigma\left(AC' - P'\right) - (1-\sigma)P'\right] > 0$$

by our monotonicity assumptions. Thus by Lemma 1, $\frac{\partial q^\star}{\partial \theta} < 0$ if $q \neq 0$ and weakly if $q = 0$. This immediately implies that price rises in θ by monotonicity and thus that CS falls. Producer surplus is

$$PS(q) = q\left[P(q) - \sigma AC(q) - (1-\sigma)AC(1)\right]$$

so

$$PS'(q) = P(q) - \sigma AC(q) - (1-\sigma)AC(1) + q\left[P'(q) - \sigma AC'(q)\right] =$$
$$P(q) - MS(q) - \sigma MC(q) - (1-\sigma)AC(1) < F'(q)$$

as $MS + MC - AC, MS > 0$ by the argument above so long as $\theta < 1$. Thus at $q^\star$ for any $\theta < 1$, $PS' < 0$. □

Proposition (Formal) 2. *If* $AC' < 0$ *and* $\theta \in (0,1)$, $\frac{\partial SS}{\partial \theta} \leq 0$, *strictly if* $q^\star > 0$.

Proof. $SS(q) = \int_0^q \left(P(q) - \left[\sigma MC(q) + (1-\sigma)AC(1)\right]\right) dq$ so $SS'(q) = P(q) - \sigma MC(q) - (1-\sigma)AC(1)$. Thus

$$SS'\left(q^\star\right) = \sigma\left(1-\theta\right)\left[AC\left(q^\star\right) - MC\left(q^\star\right)\right] + \theta MS\left(q^\star\right) > 0$$

because $MS > 0$ and $AC'(q) = \frac{MC(q) - AC(q)}{q} < 0$. Thus the result follows from the chain rule and the fact that $\frac{\partial q^\star}{\partial \theta} < 0$ as shown in the proof of the previous proposition. □

Proposition (Formal) 3. *If* $AC' > 0$ *and* $q^\star > 0$ *for every* $(\theta, \sigma) \in (0,1)^2$, $\exists \theta^\star \in (0,1)$ *such that* $\frac{\partial SS}{\partial \theta} > (< / =) 0$ *if* $\theta < (> / =) \theta^\star$. $\frac{\partial \theta^\star}{\partial \sigma} > 0$ *if* $\sigma \in (0,1)$.

Proof. By the logic of the previous proof, $SS'\left(q\right) = P(q) - \sigma MC(q) - (1-\sigma)AC(1)$ so

$$SS''(q) = P'(q) - \sigma MC'(q) = (1-\sigma)P'(q) + \sigma\left[P'(q) - MC'(q)\right] < 0.$$

Thus social surplus is concave in quantity. Quantity is below its optimal level at $\theta = 1$ by the standard monopoly argument and quantity is above its optimal level at $\theta = 0$ by the argument in the proof of the previous proposition. Thus the result follows from the fact, shown in the proof of Proposition 1, that $\frac{\partial q^\star}{\partial \theta} < 0$. $\square$

Proposition (Formal) 4. *$\frac{\partial q^\star}{\partial \sigma}$ has the same sign as $AC(1) - \theta MC(q^\star) - (1-\theta) AC(q^\star)$. If $AC' < 0$ then providing a specific subsidy to the industry can only cause the sign of $\frac{\partial q^\star}{\partial \sigma}$ to move from being negative to being positive; a sufficiently large such subsidy guarantees this sign is positive. If $\theta = 1$ then $\frac{\partial PS}{\partial \sigma}$ has the same sign as AC'. Again if $\theta = 1$, $AC' < 0$ and if the pass-through rate, $\rho(t) \equiv \frac{dP(q^\star)}{dt} > M > 0$ for some M and all t such that $q^\star \in (0,1)$ then $\frac{\partial PS}{\partial \sigma} + \frac{\partial CS}{\partial \sigma} > 0$ starting from a sufficiently large subsidy $-t$ such that $q^\star < 1$.*

Proof. With a specific tax (negative specific taxes are specific subsidies), the equilibrium condition is

$$P(q) - \sigma\left(\theta MC(q) + (1-\theta)\, AC(q)\right) + (1-\sigma) AC(1) - \theta MS(q) - t = 0.$$

Thus by the Implicit Function Theorem

$$F'(q^\star)\frac{\partial q^\star}{\partial \sigma} - \left[\theta MC(q^\star) + (1-\theta)\, AC(q^\star) - AC(1)\right] = 0 \implies$$

$$\frac{\partial q^\star}{\partial \sigma} = \frac{\theta MC(q^\star) + (1-\theta)\, AC(q^\star) - AC(1)}{F'(q^\star)}. \tag{7}$$

Because $F' < 0$ by Lemma 1, this has the same sign as $AC(1) - \theta MC(q^\star) - (1-\theta)\, AC(q^\star)$ *regardless of the degree of tax or subsidy*. By the same arguments as above, $\frac{\partial q^\star}{\partial t} < 0$. Thus if $AC', MC' < (>) 0$ the sign of this expression can only move, with an increase in tax, from being positive to being negative (from being negative to being positive).

Furthermore as t (a sufficiently large subsidy) becomes arbitrarily negative (a sufficiently large subsidy is give), $q^\star \to 1$. Thus the denominator on the right-hand side of Equation 7 must approach $\theta MC(1) - AC(1)$, which is negative if $AC' < 0$ and thus $\frac{\partial q^\star}{\partial \sigma} > 0$ eventually.

As before $PS(q) = q\left[P(q) - \sigma AC(q) - (1-\sigma) AC(1)\right]$. When $\theta = 1$ profits are maximized over q so by the envelope theorem we can calculate $\frac{\partial PS}{\partial \sigma}$ while holding fixed $q^\star$ yielding $AC(1) - AC(q^\star)$. For $q^\star < 1$ this clearly has the same sign as AC'. By the envelope theorem for consumers and this result for producers

$$\frac{\partial CS}{\partial \sigma} = -q^\star P'(q^\star)\frac{\partial q^\star}{\partial \sigma} = -\rho(t) q^\star \left[MC(q^\star) - AC(1)\right],$$

where the second equality uses the fact that when $\theta = 1$,

$$P'(q^\star)\frac{\partial q^\star}{\partial \sigma} = P'(q^\star)\frac{MC(q^\star) - AC(1)}{F'(q^\star)} = \rho(t)\left[MC(q^\star) - AC(1)\right]$$

as the tax enters linearly into the expression for F. Thus

$$\frac{\partial CS}{\partial \sigma} + \frac{\partial PS}{\partial \sigma} = -\rho(t) q^\star \left[MC(q^\star) - AC(1)\right] + AC(1) - AC(q^\star).$$

As t becomes sufficiently negative, $q^\star \to 1$ so that the second term vanishes and the first term is bounded away from 0 as $MC(q^\star) - AC(1)$ grows in absolute value (becomes more negative) monotonically in $q^\star$ and $\rho(t) > M > 0$ by hypothesis. $\square$

Note that this result is formulated in terms of subsidies, but these are equivalent to upward demand shifts of we confine attention to in-market quantities and thus ignore the impact on the subsidy provider, as we do here.

Proposition (Formal) 5. *If $AC' > 0$ then giving a specific subsidy to the industry can only cause the sign of $\frac{\partial q^\star}{\partial \sigma}$ to move from being positive to being negative; a sufficiently large such subsidy guarantees this sign is negative. If $\theta = 1$, $AC' > 0$ and if the pass-through rate, $\rho(t) \equiv \frac{dP(q^\star)}{dt} > M > 0$ for some M and all t such that $q^\star \in (0,1)$ then $\frac{\partial PS}{\partial \sigma} + \frac{\partial CS}{\partial \sigma} < 0$ starting from a sufficiently large subsidy $-t$ such that $q^\star < 1$.*

Proof. This follows exactly from the logic of the proof of Proposition (Formal) 4. □

The graphs for this scenario are analogous to those for adverse selection and are shown in Appendix Figure A2. Reducing the degree of advantageous selection rotates the average cost curve around AC(1) in a clockwise direction. When the profit-maximizing quantity is low ($q^\star < \underline{q}$), this rotation increases the cost of the average marginal consumer, raising prices and lowering equilibrium quantity. When the profit-maximizing quantity is high ($q^\star > \underline{q}$), the reduction in the degree of selection lowers the cost of the average marginal consumer, lowering prices and increasing quantity. A reduction in advantageous selection lowers industry profits by the same envelope logic discussed above. Reduced advantageous selection lowers welfare except when quantity is sufficiently high, in which case the increase in consumer surplus outweighs the decrease in firm profits.

Panels (B) and (C) of Appendix Table A1 summarize these results on the effects of selection in settings with market power. The results under adverse and advantageous selection can be understood together by noticing that a reduction in the degree of selection lowers the component of cost heterogeneity in the population correlated to willingness-to-pay, moving average individuals at any willingness-to-pay quantile q towards the population average cost. Because the monopolist internalizes the costs of the average marginal consumer, reducing selection will reduce this marginal cost exactly when the average marginal consumer is more costly than the population average consumer. Under adverse selection the average marginal consumer has higher cost at lower quantity and under advantageous selection the average marginal consumer has higher cost at higher quantity. Therefore, the benefits from reducing selection occur at low equilibrium quantities under adverse selection and high equilibrium quantities under advantageous selection.

Proposition (Formal) 6. *If $\theta = 0$ then $\frac{\partial SS}{\partial \sigma}$ has the same signs as AC'.*

Proof. At $\theta = 0$ there is no producer surplus so only the impact on consumer surplus is relevant. Because $P(q) = \sigma AC(q) + (1-\sigma)AC(1)$, $\frac{\partial P(q^\star)}{\partial \sigma}$ has the same sign as $AC(q) - AC(1)$ (given that $P' > AC'$ by our stability assumptions) which is opposite to that of AC'. By the envelope theorem, $\frac{dCS}{dP} = -q$. Thus the impact of σ on consumer and thus social surplus has the same sign as AC'. □

For the following results, σ represents risk-adjustment rather than correlation. Again we use a tax or subsidy to shift the demand curve and measure welfare now with respect to the primitive demand and supply curves, including the tax/subsidy, excluding any impacts of the tax/subsidy on the government budget and ignoring the risk adjustment (as this is just a transfer) except through its impacts on equilibrium quantity. We let $q^{\star\star}$ denote the socially optimal quantity. In what follows we treat the "specific tax" t a simply a uniform inverse demand/cost shifter and thus irrelevant to welfare quantities.

Proposition (Formal) 7. *Let $q_0 = q^\star$ when $\sigma = 0$ (full risk-adjustment). If $\theta = 1$, $AC' < 0$ and $MS' > MC'$ then there exist thresholds $q' < \underline{q}$ that are invariant to the level of a specific tax t such that*

1. *If $q_0 < q'$ then $\frac{\partial q^\star}{\partial \sigma} < 0$ and there exists $\sigma^\star \in (0,1)$ such that at $\sigma^\star$, $q^\star = q^{\star\star}$.*
2. *If $q_0 = q'$ then $\frac{\partial q^\star}{\partial \sigma} < 0$ and $q^\star = q^{\star\star}$ when $\sigma = 0$.*

3. *If* $q' < q_0 < \underline{q}$ *then* $\frac{\partial q^\star}{\partial \sigma} < 0$ *and* $q^\star < q^{\star\star}$ *even when* $\sigma = 0$.

4. *If* $q_0 = \underline{q}$ *then* $\frac{\partial q^\star}{\partial \sigma} = 0$ *and* $q^\star < q^{\star\star}$.

5. *If* $q_0 > \underline{q}$ *then* $\frac{\partial q^\star}{\partial \sigma} > 0$ *and* $q^\star < q^{\star\star}$

$q' > 0$ *if* $\lim_{q\to 0} P'(q)q = 0$. q_0 *ranges between* 0 *and* 1 *as a sufficiently large tax or subsidy is imposed.*

The additional discussion about the direction of welfare in the text follows from this result and the observation that under our stability assumptions welfare is strictly concave in quantity.

This proposition imposes two additional conditions not discussed previously: that $MS' > AC'$ and that $\lim_{q\to 0} P'(q)q = 0$. Log-concavity of direct demand is sufficient, but not necessary, for the first condition assuming that $MC' < 0$, as it implies $MS' > 0$ (Weyl and Fabinger, 2013) and thus clearly $> MC'$. We have typically assumed that when $AC' < 0$, $MC' < 0$ as well. The second condition is neither necessary nor sufficient for log-concavity but is true of every log-concave demand function we are aware of, as shown by Fabinger and Weyl (2014). It also implies that demand is log-concave at sufficiently high prices as, letting $\overline{p} \equiv \lim_{q\to 0} P(q)$ and Q be the direct demand,

$$\lim_{q\to 0} P'(q)q = 0 \iff \lim_{p\to\overline{p}} \frac{Q(p)}{Q'(p)}.$$

Bulow and Pfleiderer (1983) show that the sign of the derivative of $\frac{Q}{Q'}$ positive if and only if Q is locally log-concave; clearly $Q' < 0$ so $\lim_{p\to\overline{p}} \frac{Q(p)}{Q'(p)}$ only if in the limit this quantity is increasing (towards 0). Thus it is closely connected to log-concavity and $\lim_{q\to 0} P'(q)q = 0$ are closely-allied concepts and thus we view the gap between them as being a "regularity" condition, as quoted in the text.

If the first condition fails it is possible that there are other points of switching between regimes 1) and 3) from the proposition; this does not lead to qualitatively different behavior, but would be more complex to state and thus we omitted discussing it in the text. If the second condition fails, then, as discussed in the text, it is possible (though not necessary) that even at very low q_0 full risk-adjustment is still insufficient.

Proof. First note that risk-adjustment payments are pure transfers and thus social surplus is invariant to them except in how they impact quantity. Second, note that their impact on quantity is precisely as in Proposition (formal) 4 as the equilibrium equations are identical to there. This establishes the claims about $\frac{\partial q^\star}{\partial \sigma}$. Point 5) follows because quantity is always too low when $AC' > 0$ and becomes only lower with risk-adjustment; it may never cross $\underline{q}$ because there $\frac{\partial q^\star}{\partial \sigma} = 0$. Point 4) follows directly from this observation: if $q^\star = \underline{q}$ if $\sigma = 1$, $q^\star$ is invariant to σ. All of this is, as claimed, invariant to the value of t by the same logic in the proof of Proposition (Formal) 4.

On the other hand when $q^\star < \underline{q}$ if $\sigma = 1$ then $\frac{\partial q^\star}{\partial \sigma} < 0$ but by the same logic $q^\star < \underline{q}$ for all $\sigma \in [0,1]$. By the logic in the proof of Proposition (Formal) 3, social welfare is concave in quantity and quantity is too low when $\sigma = 1$. Thus either social surplus monotonically increases as q falls or social surplus reaches a peak and then declines beyond some point if q becomes too low. Which occurs is determined by the sign of $SS'(q_0)$ as, by concavity and monotonicity of $q^\star$ in σ, $SS'(q^\star) < SS'(q_0)$ for all $\sigma > 0$.

$$SS'(q_0) = P(q_0) - MC(q_0) = MS(q_0) + AC(1) - MC(q_0).$$

Thus if $MS(q_0) > MC(q_0) - AC(1)$ then $q^\star \le q_0 < q^{\star\star}$ while if $MS(q_0) = MC(q_0) - AC(1)$ then $q^\star \le q_0 = q^{\star\star}$ and if $MS(q_0) < MC(q_0) - AC(1)$ then there is an interior optimal σ as there was an interior optimal $\theta^\star$ in Proposition (Formal) 3 and as described in point 1). Note that this is all invariant to t as this has no impact on either MS or $AC(1) - MC$ as it shifts the latter two in parallel.

By definition of $\underline{q}$, $q_0 < \underline{q}$ implies that $MC(q_0) > AC(1)$. Thus if $\lim_{q \to 0} MS(q) = 0$ then for sufficiently small q_0 the second case holds. Conversely $MS(q_0) > 0$ for all $q_0 > 0$ and as $q_0 \to \underline{q}$, again by definition of $\underline{q}$, $MC(q_0) \to AC(1)$ and thus the first case holds.

q' is then simply defined as the threshold between these regimes, which exists by the assumption that $MS' > MC'$ and thus $MS' - MC' + AC(1)$ has a single crossing of 0 (from below to above). The range claim on q_0 as a function of t follows from the fact that $F' < 0$.

□

Proposition (Formal) 8. *Let q_0 be defined as in Proposition (Formal) 7. If $\theta = 1$, $AC' > 0$ and $MC' - MS'$ is signed globally, there exist thresholds $q'' > \underline{q}$ that are invariant to the level of a specific tax t, with $\underline{q}$ being identical to its value in Proposition (Formal) 7, such that*

1. *If $q_0 < \underline{q}$ then $\frac{\partial q^\star}{\partial \sigma} > 0$ and $q^\star < q^{\star\star}$ even at $\sigma = 1$.*
2. *If $q_0 = \underline{q}$ then $\frac{\partial q^\star}{\partial \sigma} = 0$ and $q^\star < q^{\star\star}$ for all $\sigma \in [0, 1]$.*
3. *If $\underline{q} < q_0 < q''$ then $\frac{\partial q^\star}{\partial \sigma} < 0$ and $q^\star < q^{\star\star}$ even at $\sigma = 0$.*
4. *If $q_0 = q''$ then $\frac{\partial q^\star}{\partial \sigma} < 0$ and $q^\star = q^{\star\star}$ when $\sigma = 0$.*
5. *If $q_0 > q''$ then $\frac{\partial q^\star}{\partial \sigma} < 0$ and there exists a $\sigma^\star \in (0, 1)$ such that at $\sigma^\star$, $q^\star = q^{\star\star}$.*

$q'' < 1$ if and only if $MC(1) - AC(1) < MS(1)$ and, as in Proposition (Formal) 7, adjusting t traces out the full possible range of q_0.

The additional conditions in this result have less intuitive content than those in the previous proposition. Again $MC' - MS'$ being signed is necessary to ensure a simple structure on the regions of potential outcomes. A sufficient condition for this is log-convexity of demand as in this case $MC' > 0 > MS'$, assuming that MC' has the same sign as AC'. But $MS' > MC' > 0$ would also satisfy the condition and would have demand being very log-concave.

The second condition has little intuitive content, but is only possible in the case when $MC' > MS'$. It states that the downward distortion from market power is smaller than the upward distortion from advantageous selection (that would occur under perfect competition) when equilibrium quantity is sufficiently high.

Proof. The proof follows precisely the logic of Proposition (Formal) 7, *mutatis mutandis* for the differences between the adverse and advantageous cases. □

Proposition (Formal) 9. *If $\theta = 0$ then $\frac{\partial q^\star}{\partial \sigma}$ has the same sign as $-AC'$. If $q_0 \leq \underline{q}$ then there exists a $\sigma^\star \in [0, 1)$ such that at $\sigma^\star$, $q^\star = q^{\star\star}$. If $q_0 > \underline{q}$ then $SS'(q^\star) \frac{\partial q^\star}{\partial \sigma} > 0$ and $q^\star - q^{\star\star}$ has the sign of AC' for $\sigma \in [0, 1]$.*

Proof. The first claim follows directly from the logic of Proposition (Formal) 6 and the fact that the equilibrium conditions with $\theta = 0$ are the same for a given σ under the two models as above.

The second claim comes from a logic similar to the preceding two propositions. Social surplus is still concave for the same reasons. At $\sigma = 1$ is it is always declining in σ because for $AC' > 0$ quantity is too high and for $AC' < 0$ quantity is too low. It is thus sufficient to verify whether this sign is maintained or not at q_0. We just consider one of the four cases; the other three are analogous.

Suppose that $q_0 < \underline{q}$ and that $AC' < 0$. Then by definition of $\underline{q}$, $MC(q_0) > AC(1)$.

$$SS'(q_0) = P(q_0) - MC(q_0) = AC(1) - MC(q_0) < 0,$$

reversing the sign compared to $SS'(q^\star)$ when $\sigma = 1$ and implying an interior optimum by the reasoning in the proof of the previous two propositions. □

C Example with Large Demand-Driven Effects of Risk-Based Pricing

One intervention commonly applied in selection markets is cost-based pricing. In Subsection 4.2 we showed that these discriminatory effects may reinforce the cost-based effects of selection. In this appendix we discuss how price discriminatory effects of cost-based pricing may instead reverse the results we established about the impact of changing the degree of selection in Section 3.

Proposition 5 states that increasing advantageous selection increases monopoly profits. However, clearly allowing cost-based pricing may never hurt a monopolist as she may maintain uniform pricing. It will generically aid the monopolist. Thus her gains from price discrimination swamp the effects we highlight.

To see that impacts on consumers may also be reversed by price discriminatory effects, consider our result (Proposition 4 and 5) that decreasing adverse and advantageous selection may both benefit consumers, depending on the equilibrium quantity. This may be true of cost-based pricing, but not in one simple, extreme case. Suppose that there is only a single dimension of heterogeneity determining both cost and valuation and that we move from uniform pricing to full cost-based pricing. This operates as perfect, first-degree price discrimination, extracting all surplus from consumers regardless of the equilibrium quantity and thus contradicting the natural extrapolation of our result.

Thus cost-based pricing *cannot* cleanly be interpreted as an example of increasing selection in our framework; price discrimination may be more important in some cases than are cost-based effects. However, in the leading counter-intuitive case we emphasize, the two effects reinforce one another to lower consumer surplus.

D Health Insurance Model

We build a model of health insurance choice that matches key features of the U.S. employer-sponsored health insurance market, focusing in particular on individual health plans with a duration of one year. We assume that consumers are expected utility maximizers with constant absolute risk aversion (CARA) preferences. Consumers are heterogeneous in their absolute risk aversion, denoted α, and their health-type, denoted λ, which we assume are jointly log-normally distributed according to

$$\begin{array}{c}\ln \alpha \\ \ln \lambda\end{array} \sim \mathcal{N}\left(\left[\begin{array}{c}\mu_\alpha \\ \mu_\lambda\end{array}\right], \left[\begin{array}{cc}V_\alpha & \rho_{\alpha,\lambda}\sqrt{V_\alpha V_\lambda} \\ \rho_{\alpha,\lambda}\sqrt{V_\alpha V_\lambda} & \sqrt{V_\lambda}\end{array}\right]\right).$$

Consumers with health-type λ are exposed to a distribution of shocks with realized values c. We assume that consumers' health type and health outcomes are jointly log-normally distributed according to

$$\begin{array}{c}\ln \lambda \\ \ln c\end{array} \sim N\left(\left[\begin{array}{c}\mu_\lambda \\ \mu_c\end{array}\right], \left[\begin{array}{cc}\sqrt{V_\lambda} & \rho_{\lambda,c}\sqrt{\sigma_\lambda \sigma_c} \\ \rho_{\lambda,c}\sqrt{V_\lambda V_c} & \sqrt{V_c}\end{array}\right]\right).$$

This implies that a consumer's realized health risk, conditional on their health-type, is distributed according to

$$\ln c|\ln\lambda \sim \mathcal{N}\left(\mu_c + \sqrt{\frac{V_c}{V_\lambda}}\rho_{\lambda,c}\left[\ln\lambda - \mu_\lambda\right], \sqrt{1-\rho^2_{\lambda,c}V_c}\right).$$

We calibrate the distributions of risk aversion using values from the literature and the distribution of health types and medical spending using values from the 2009 Medical Expenditure Panel Survey (MEPS). Table A2 summarizes the exact calibrated variables. Below we discuss the calibrated values in more detail.

- *Risk aversion (α).* We calibrate the distribution of absolute risk aversion to the values estimated by Handel, Hendel and Whinston (2015), which are identified using over-time variation in the choice set of health insurance plans offered to employees at a large firm. These values are similar to those estimated by Cohen and Einav (2007). The mean value of $\alpha = 0.000439$ implies indifference between a 50-50 gamble for $\{\$100, -\$96\}$ and \$0 with certainty.

- *Realized costs (c).* We calibrate the distribution of realized medical costs c to match the population mean and standard deviation of medical spending for non-elderly individuals in the 2009 MEPS, excluding individuals with coverage from a public program such as Medicaid. The mean level of spending for this sample is \$3,139 and the standard deviation in \$10,126.

- *Health-type (λ).* To calibrate the degree of private information, we assume that consumers' knowledge of their future health costs is the same as that which can be predicted by standard risk adjustment software.[30] The 2009 MEPS provides information on individual's Relative Risk Scores, which is calculated using the Hierarchical Clinical Classification (HCC) model that is also used to risk adjust Medicare Advantage payments.

- *Correlation between risk aversion and health-type ($\rho_{\alpha,\lambda}$).* We assume that risk aversion and health risk are uncorrelated in the population. This is probably a reasonable assumption given the diverging estimates of the sign of this correlation in the literature.

- *Correlation between realized costs and health-type ($\rho_{\lambda,c}$).* Following our model, we estimate the correlation $\rho_{\lambda,c}$ with a regression of log realized health costs on the log Relative Risk Score, where both variables are normalized by subtracting the mean and dividing by the standard deviation. We estimate a coefficient of $\rho_{\lambda,c} = 0.498$. This estimate, combined with information on the mean and standard deviation of the Relative Risk Scores and realized costs, allows us to simulate the joint distributions of λ and c.

- *Cost-sharing ($\kappa_H(c)$ and $\kappa_L(c)$).* We calibrate the cost-sharing of the high-quality plan to cover 90% of the cost of medical care in the population on average, known as a 90% actuarial value (AV) plan. This is the level of coverage provided by a "platinum" plan on an Affordable Care Act (ACA) Health Insurance Marketplace. We calibrate the low-quality plan to have an 60% AV, which is typical for an HDHP, and would qualify as a "bronze" plan on an ACA Marketplace. The AV 90% plan has no deductible, 10% co-insurance, and an \$8,000 out-of-pocket maximum.[31] The AV 60% plan has a \$500 deductible, 40% coinsurance, and an \$8,000 out-of-pocket maximum.

[30] This assumption follows standard practice in the literature (Handel, 2012; Handel, Hendel and Whinston, 2015) and is supported by the finding from Bundorf, Levin and Mahoney (2012) of little private information conditional on an industry standard measure of predicted health risk.

[31] Because it has an out-of-pocket maximum, in addition to the 10% co-insurance, the AV 90% actually covers 90.42% of costs. We round to 90% for simplicity.

- *Market power (θ).* We calibrate the level of market power to $\theta = 0.5$. This corresponds to Cournot competition with two high-quality plans. While this is significantly higher than the representative values we discussed above, we focus on this value for two reasons. First, in the EHBS, less than 1% of firms offer more than two non-HDHP options; thus the market-wide concentration indices derived from Dafny, Duggan and Ramanarayanan (2012)'s data may understate effective market power. Second, with greater market power our results are more visible. However, as we showed above for the case of risk adjustment, they are qualitatively true even with lower market power.

E Misperceptions of Health Risk

We also use the model to examine the effect of a change in correlations that would result from a change in consumer perceptions about the distribution of risk they face. For instance, an insurance-choice decision-aide might reduce misperceptions of costs and therefore increase the degree of selection in the market. We model this potential misperception by generating a perceived health-type $\hat{\lambda}$ that is jointly log-normally distributed with a consumer's actual health-type λ with correlation $\rho_{\lambda,\hat{\lambda}}$. We calculate equilibria where willingness-to-pay is determined by the consumer's perceived health type while costs are determined by the consumer's actual health type. We calculate surplus under perceived demand, which might be relevant if consumers are never de-biased of their misperceptions, and under the actual demand curve. Appendix Figure A4 plots the equilibrium allocations generated by the perceived demand and marginal cost curves for different values of $\rho_{\lambda,\hat{\lambda}}$.

Table A3 shows the results of this exercise. The first column shows the setting where perceptions are fully accurate ($\rho_{\lambda,\hat{\lambda}} = 1$). The second column shows a setting where consumers have perceptions that are partially correlated with true health risk ($\rho_{\lambda,\hat{\lambda}} = 0.5$). The third column shows a setting where perceived health risk is completely uncorrelated with the truth ($\rho_{\lambda,\hat{\lambda}} = 0$). As above, the welfare values are shows as a percent of the optimized total surplus.

Reducing the correlation between perceived and actual health risk has the effect of decreasing the degree of selection in the market, as shown in Appendix Figure A4. This means that, similar to the results above, increased misperceptions raise price and reduce quantity in the market, even under the demand curves that result from perceived risk.[32] Employee surplus is even lower, and social surplus actually falls, under actual demand curves, since the misperceptions create an allocative inefficiency in who receives insurance coverage. This pushes against the argument, made in a perfectly competitive environment, that nudging (improving information) can hurt consumers by exacerbating the degree of selection (Handel, 2012) and helps justify employer efforts to help employees optimize their health plan choice.

[32] Increasing misperceptions raises producer surplus, suggesting that policy efforts to de-bias consumers through decision aides may be opposed by the insurance industry.

Figure A1: Summary of Welfare Effects of Risk Adjustment Under Monopoly

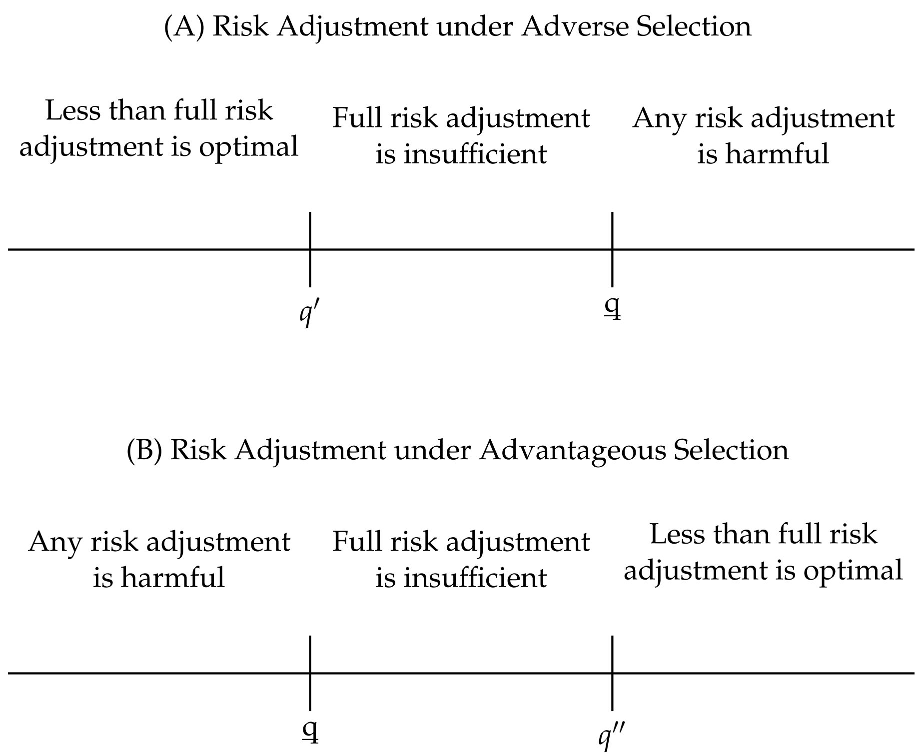

Note: This figure summarizes the effects of risk adjustment for different ranges of the post-risk adjustment equilibrium level of quantity $q^{\star}$. Panel (A) shows these ranges in an adversely selected market. Panel (B) shows these ranges in an advantageously selected market.

Figure A2: Reducing Advantageous Selection under Monopoly

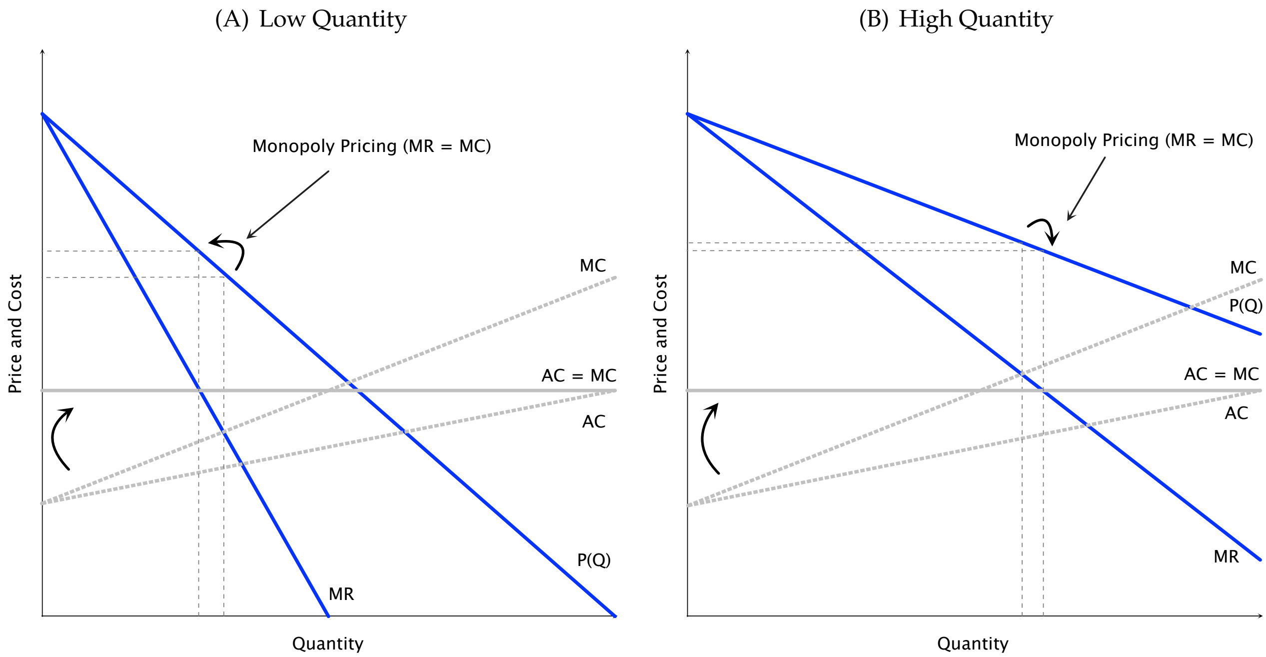

Note: This figure shows the effect of reducing the degree of advantageous selection in a market served by a monopolist. Panel (A) considers a setting where the equilibrium quantity is low and reducing advantageous selection raises price and lowers quantity. Panel (B) considers a setting where the equilibrium quantity is high and reducing advantageous selection lowers price and increases quantity.

Figure A3: Risk-Based Pricing: Segmented Market

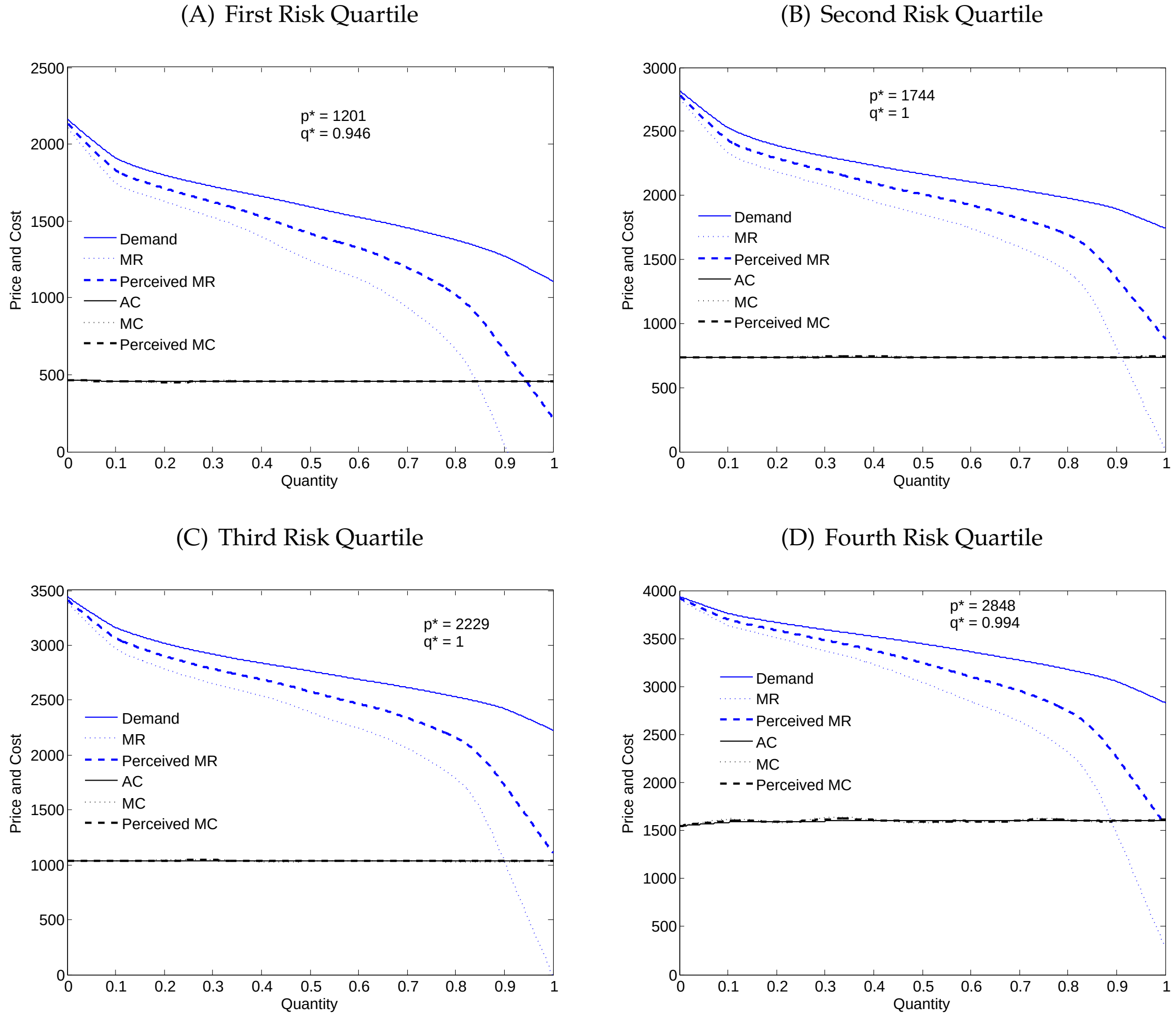

Note: This figure shows the effects of risk-based pricing, achieved by segmenting the market into quartiles using the risk-type parameter λ. The first risk quartile corresponds to the set of consumers with the lowest expected costs and the fourth risk quartile corresponds to the consumers with the highest expected costs in the market. See Subsection 4.2 for more details.

Figure A4: Misperceptions: Imperfect Correlation Between Perceived and Actual Risk

(A) Fully Accurate Perceptions: $\rho_{\lambda,\hat{\lambda}} = 1$

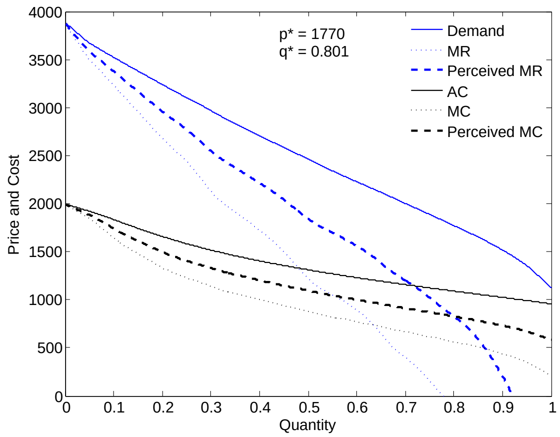

(B) Partial Misperceptions: $\rho_{\lambda,\hat{\lambda}} = 0.5$

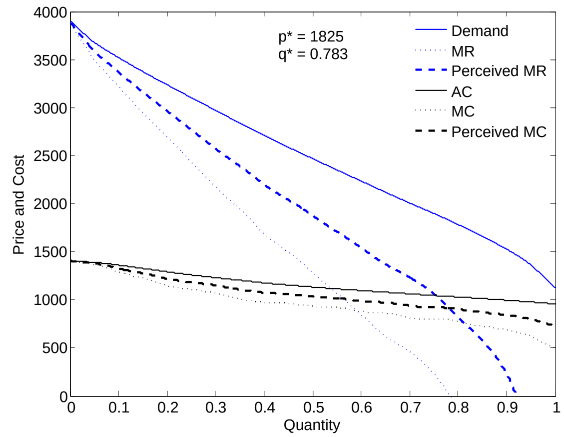

(C) Full Misperceptions: $\rho_{\lambda,\hat{\lambda}} = 0$

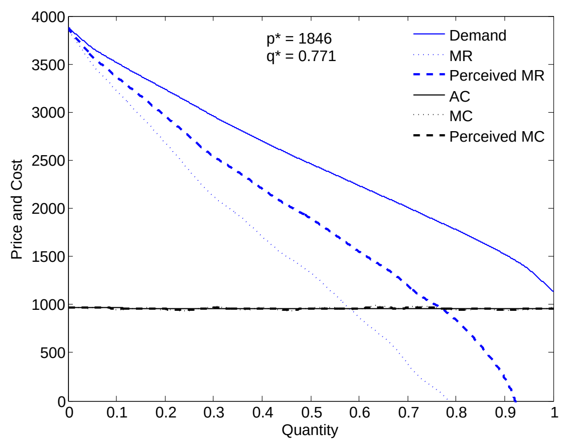

Note: This figure shows the effects of consumer misperceptions about health risk, modeled by allowing consumers' perceived health type $\hat{\lambda}$ and actual health type λ to be jointly log-normally distribution with correlation parameter $\rho_{\lambda,\hat{\lambda}}$. See Subsection 4.2 for more details.

Table A1: Summary of Results

Panel (A): Greater Market Power		
	Adverse Selection	Advantageous Selection
Producer Surplus	Higher	Higher
Consumer Surplus	Lower	Lower
Social Surplus	Lower	Inverse-U shaped
Panel (B): Less Adverse Selection		
	Perfect Competition	Monopoly
Producer Surplus	Always zero	Higher
Consumer Surplus	Higher	Lower $\iff q^\star > \underline{q}$
Social Surplus	Higher	Lower $\iff q^\star > \overline{q} > \underline{q}$
Panel (C): Less Advantageous Selection		
	Perfect Competition	Monopoly
Producer Surplus	Always zero	Lower
Consumer Surplus	Lower	Higher $\iff q^\star > \underline{q}$
Social Surplus	Lower	Higher $\iff q^\star > \overline{q} > \underline{q}$

Note: This table summarizes the main results. Panel (A) shows the effects of increasing market power in industries with adverse and advantageous selection. Panels (B) and (C) show the effects of reducing the degree of adverse and advantageous selection, respectively, under perfect competition and monopoly market power.

Table A2: Calibration Values for Health Insurance Model.

Parameter	Description	Mean	Std. Dev.	Note
α	Absolute risk aversion	4.39×10^{-4}	6.63×10^{-5}	Estimates of absolute risk aversion from Table 3 of Handel, Hendel and Whinston (2015).
λ	Privately known health type	0.979	1.378	Values for Relative Risk Score (HCC, Private) in the 2009 MEPS.
c	Realized medical spending	\$3,139	\$10,126	Realized medical spending for the non-elderly population without public insurasnce in the 2009 MEPS.
ρ	Correlation of $\ln\lambda$ and $\ln c$	0.498		Estimated from a regression of normalized log realized medical spending on normalized log Relative Risk Scores in the 2009 MEPS.

Note: This table lists the calibrated values used in the health insurance model and their sources.

Table A3: Welfare Effects of Misperception of Health Risk

	Correlation Between Perceived and Actual Risk		
	$\rho_{\lambda,\hat{\lambda}} = 1$	$\rho_{\lambda,\hat{\lambda}} = 0.5$	$\rho_{\lambda,\hat{\lambda}} = 0$
Price	1,754	1,825	1,846
Quantity	81.3%	66.4%	66.1%
Employee Surplus			
Perceived Risk	51.5%	2.2%	1.6%
Actual Risk	51.5%	-4.8%	-11.9%
Producer Surplus	35.5%	40.2%	44.4%
Total Surplus			
Perceived Risk	87.0%	42.4%	45.9%
Actual Risk	87.0%	35.4%	32.5%

Note: This table shows the effect of a reduction in correlation that would result from consumer misperceptions' of their health risks, modeled by allowing consumers' perceived and actual health type to be jointly log-normally distribution with correlation parameter $\rho_{\lambda,\hat{\lambda}}$. The first column shows the setting where perceptions are fully accurate ($\rho_{\lambda,\hat{\lambda}} = 1$), the second column where perceptions are partially correlated ($\rho_{\lambda,\hat{\lambda}} = 0.5$), and the third column where perceived health risk is completely uncorrelated with the truth ($\rho_{\lambda,\hat{\lambda}} = 0$). We show employee and total surplus under the demand curves that result from the perceived health risk and actual health risk distributions. All values are presented as a percentage of the first best total surplus in the baseline scenario where perceptions are fully accurate.