

Adverse Selection Pricing and Unraveling of Competition in Insurance Markets *

Edward Kong
Harvard University

Timothy J. Layton
Harvard and NBER

Mark Shepard
Harvard and NBER

October 6, 2022

Abstract

Market-based health insurance programs rely on robust insurer participation to function well. We develop a model of a health insurance market with adverse selection where insurer participation is endogenous. We show that when low-cost consumers also tend to be highly price sensitive — the key feature of adverse selection — insurers often lose money at their profit-maximizing prices, even without fixed costs. Adverse selection (like fixed costs) generates strong incentives for insurers to undercut each other’s prices, leading to equilibria with few surviving competitors — and in the extreme, to natural monopoly. Using data from Massachusetts’ health insurance exchange, we leverage exogenous subsidy-driven variation in health plan prices to estimate high consumer price sensitivity and strong adverse selection. Using an estimated structural model, we find price sensitivity and adverse selection strong enough to lead to only monopoly or duopoly in equilibrium, absent corrective policies. We show how corrective policies, including price floors and risk adjustment, can restore equilibria with greater insurer participation and higher consumer welfare.

1 Introduction

Over the last two decades, social health insurance programs in the United States have shifted away from offering a single publicly managed health insurance plan for all beneficiaries and toward regulated and carefully crafted marketplaces in which private insurers compete for enrollees (Gruber, 2017). This shift has taken place in the Medicare program and the Medicaid program and undergirded a primary mechanism for coverage expansion under Affordable Care Act: the subsidized state Health Insurance Marketplaces. These types of marketplaces are also used in social health insurance programs in the Netherlands, Switzerland, Germany, Chile, and Israel. A growing literature has studied the economics of these marketplaces, including subsidy design, adverse selection, and competition.

*Contact: Kong (edk571@g.harvard.edu), Layton (Layton@hcp.med.harvard.edu), Shepard (Mark.Shepard@hks.harvard.edu). We are grateful to Liran Einav, Kate Ho, Lee Lockwood, Neale Mahoney, Amanda Starc, Pietro Tebaldi, Nick Tillipman, and seminar participants at Duke University, Northwestern University, the University of Chicago, University of Virginia, Johns Hopkins University, the University of Illinois - Chicago, Erasmus University Rotterdam, the University of Pennsylvania, Stanford University, the American Economic Association meetings, and the Annual Meeting of the American Society of Health Economists for comments. Ilana Salant provided superb research assistance.

While the goal of this type of market-based provision of social health insurance benefits is to improve efficiency beyond what might be achieved by a single publicly managed program, there are well-documented forces working to impede the achievement of this goal. Imperfect competition and limited insurer entry has been shown to lead to higher prices for consumers and higher costs for the government. Asymmetric information has been shown to lead to important price distortions that cause consumers to sort inefficiently across insured and uninsured states and across the various levels of generosity available in the market. Further, asymmetric information has been shown to cause insurers to offer inefficient contracts distorted toward the preferences of healthy, profitable consumers and away from the preferences of sick, unprofitable consumers.

In this paper, we study an additional feature of these markets that pushes against efficiency. The feature we study is a firm strategic pricing response that is related to, but distinct from, classic notions of adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976). In the public economics literature, adverse selection typically plays out in a vertical setting, where sick consumers have a stronger preference for more generous insurance than healthy consumers. Under a typical notion of perfect competition, this preference causes the prices of more generous insurance plans to be “too high” and sometimes causes those plans to “death spiral” altogether (Einav et al., 2010a; Cutler and Reber, 1998). We consider a different type of setting and consumer behavior, one where the preferences of sick and healthy consumers still differ in important ways, but where plans are more generally (horizontally) differentiated products that compete strategically on price. In this type of setting, there may be a different type of selection, characterized by strong correlations between price sensitivity and risk, where healthy consumers will often be more price sensitive than sick ones.

We show that in this type of setting, it is important to consider firms’ incentives to price strategically. Rather than taking plan demand and cost curves as market primitives, we bring insights from the industrial organization literature that consider these curves to be equilibrium objects, where demand and costs for a particular plan depend not only on that plan’s price but also on the prices of the other plans in the market. We show that strong correlations between price sensitivity and risk don’t just influence plan prices or contract design; they also cause plans to have strong price undercutting incentives that can destabilize markets via downward price spirals. In contrast to the prior literature on vertical selection, where it is well known that competitive equilibria are often characterized by distorted prices and distorted contracts, we show that strategic firm pricing in the presence of horizontal selection can lead to such strong price cutting incentives that multiple plans cannot break even in equilibrium.

To fix ideas, consider a large urban market where there are two large hospital systems. One system dominates the western part of the city while the other system dominates the eastern part of the city, though both systems have some presence throughout the city. While two *vertically* differentiated plans might consist of one plan that includes both hospital systems and one plan that includes only one, two *horizontally* differentiated plans would each cover one system. Consumers have heterogeneous preferences, with consumers living in the west on average preferring the plan

that covers the western hospital system and consumers living in the east on average preferring the plan that covers the eastern hospital system. Selection may seem like a minor concern in such a market, as preferences are correlated with some dimension (geography) that is unlikely to be highly correlated with underlying health risk. However, geography is not the only dimension of preference heterogeneity. In such a market, it is likely that sicker consumers care deeply about which hospital system is in their plan’s network – they interact with the healthcare system often, so they want to be sure that their preferred (i.e., nearby) providers are in the network. Healthier consumers, on the other hand, may care little about which hospital system is included in their plan’s network – they rarely interact with the healthcare system, making them much more sensitive to price than to the network.

While this type of horizontal differentiation would typically result in low price sensitivity and an ability for plans to charge high mark-ups, in this setting the highly price sensitive healthy consumers prevent this from happening. If the western plan charges a slightly higher price than the eastern plan, the western plan will lose many of its healthy price sensitive consumers, driving up its average cost and causing it to lose money. The western plan thus has a strong incentive to *undercut* the eastern plan and take the healthy consumers, thus disciplining the plan’s desire to charge high mark-ups, a result that has been pointed out in the prior literature [Starc \(2014\)](#); [Mahoney and Weyl \(2017a\)](#); [Einav et al. \(2018\)](#). But, the key new insight here is that the eastern plan has the same incentive—after the western plan lowers its price, the eastern plan still faces a strong incentive to drop its price even lower.

If the correlation between price sensitivity and risk is strong enough, this incentive remains even after the plans are pricing below the market-wide average cost, as either plan can do better by pricing just below their competitor and taking all of the healthy enrollees in the market, driving their plan-specific average cost below the market-wide average. There is thus a “race to the bottom” in price. Importantly, prices may ultimately settle below average cost, resulting in no pure-strategy Nash equilibria where both plans are offered and earn non-negative profits. Thus, instead of disciplining the mark-up, selection leads to the elimination of equilibria with multiple firms participating in the market. In these types of settings, the only valid equilibria may be ones where only one of the two plans is offered, resulting in losses to consumer surplus due to (1) the loss of one of the two plan options and (2) higher mark-ups that can now be charged by the remaining (monopolist) plan. Indeed, everyone in this market would be better off if the two plans could coordinate on price and split the market, but such an equilibrium is not stable, as both plans have a strong incentive to undercut the other in order to attract the entire healthy group to themselves.

We formalize this idea by specifying a general model of an insurance market where multiple firms decide whether to offer fixed contracts (on which they then compete on prices) or to refrain from entering the market. We use the model to draw a novel connection between the economics of firm behavior under fixed costs and the economics of firm behavior under adverse selection. It is widely known that fixed costs lead to increasing returns to scale via a downward-sloping average cost

curve. In such settings, firms have an incentive to undercut their competitors in order to increase their market share and spread the fixed costs across a larger set of consumers. This undercutting incentive may be difficult for the market to sustain many competitors due to the need to spread the firm-specific fixed costs of production across as many purchasers as possible. In the extreme, the market produces a natural monopoly. We show that markets with strong horizontal selection exhibit similar properties. It is widely known that markets with vertical selection also feature downward-sloping average cost curves at the market level (Einav et al., 2010b). But horizontal selection can also lead to *firm-level* downward-sloping average cost curves that, like in the case of fixed costs, can make it difficult for a market to sustain many competitors.

To assess the empirical importance of this conceptual possibility, we turn to data from Massachusetts’ Commonwealth Care (CommCare) insurance exchange, the state’s pre-ACA health insurance marketplace for low-income people. The market consisted of 4-5 insurers that each offered a single plan with standardized cost sharing provisions but that differed in provider networks.

We start by presenting a few case studies of instances where one plan undercut another, showing the consequences of these price changes for the market shares and average costs of both the *undercutting* and the *under-cut* plans. These case studies reveal enormous increases in market share and enormous decreases in average cost for the under-cutting plans, with the opposite occurring for the under-cut plan. This illustrates the undercutting incentive implied by our model by showing that undercutting one’s competitors on price can simultaneously lead to large gains in market share *and* large increases in the under-cutting plan’s average profit margins. We then move beyond our case studies and estimate summary measures of price sensitivity and average cost in a difference-in-differences design, leveraging the fact that individuals below 100% of FPL are fully subsidized and are thus unaffected by price changes, while individuals above 100% of FPL experience changes in price from one year to the next based on changes in plan bids year-over-year. These more general quasi-experimental estimates show that the case study results are indeed general, with price decreases generally leading to large increases in market share and large decreases in average cost. Importantly, we estimate that the average slope of the average cost curve at observed prices is around 1, raising serious concerns about existence of Nash equilibria with the observed number of firms.

We then study the effects of these incentives on market equilibrium and how equilibria vary with a variety of corrective policies. To do so, we estimate a full structural model of demand for and costs of the insurance plans in the Massachusetts market. When estimating the model, we leverage similar variation in prices to what we used in the difference-in-difference design. We show that our structural model fits our reduced form results quite well, with predictions from the model replicating our event studies.

We then use the model to show that in the absence of corrective policies, there are no pure strategy Nash pricing equilibria where all firms participate in the Massachusetts market. We then consider a two-stage game where in stage 1 insurers decide whether to enter the market and in stage

2 insurers set Nash-Bertrand prices to maximize profits. We posit an equilibrium concept where a combination of firms and prices is only an equilibrium if no additional firm can unilaterally enter and earn positive profits in the stage 2 Nash equilibrium that results when said firm enters. Under this equilibrium definition, we show that in the absence of corrective policies, the undercutting incentives are so strong that the only surviving equilibria are ones with a single monopoly firm. The issues raised by our theoretical model thus appear to be empirically relevant in this setting.

We then use the model to perform counterfactual simulations exploring the effects of various corrective policies on market equilibrium. The CommCare market featured two key corrective policies: risk adjustment and price floors.¹ Risk adjustment involves the regulator forcing transfers from firms that enroll healthier-than-average individuals to firms that enroll sicker-than-average individuals. Such transfers weaken the correlation between price sensitivity and risk and flatten the firm-specific average cost curve, weakening undercutting incentives. However, conditional on participation, risk adjustment also limits incentives for a firm to lower its price, potentially limiting gains in consumer surplus (Mahoney and Weyl, 2017a). The second policy is a price floor. During much of its existence, the Connector imposed both price ceilings and floors, which were technically applied via rate regulation to ensure actuarial soundness. If set correctly, a price floor could cause equilibria with multiple participating firms to exist, by effectively allowing firms to coordinate on a price at the floor and splitting the group of price-sensitive healthy consumers. Counterintuitively, a price floor could result in *lower* prices, as the alternative may be a monopoly equilibrium with high prices. However, if set too high, a price floor might encourage entry but push up premiums far above average costs, reducing consumer surplus.

Our simulations show that these corrective policies can indeed increase firm participation. In the absence of risk adjustment, there is no equilibrium where more than one firm participates. With moderate risk adjustment, equilibria with two participating firms survive, and perfect risk adjustment allows equilibria with four participating firms. Not surprisingly, consumer welfare improves when going from no risk adjustment to moderate risk adjustment, as the market goes from a monopoly to a duopoly, and prices drop below an imposed regulatory price ceiling. When moving from moderate risk adjustment to perfect risk adjustment, we find that welfare may improve but also that prices are higher. This is consistent with prior work by (Mahoney and Weyl, 2017a) showing that under imperfect competition, risk adjustment can raise prices, because it weakens the downward price pressure exerted by selection. Overall, our results suggest that moderate levels of risk adjustment increase insurer participation and reduce prices below monopoly levels; however, too much risk adjustment can lead to higher prices.

Finally, we simulate various price floors. First, we show that the optimal price floor is usually slightly above the average cost across all consumers in the market, and is usually non-zero. Price floors are more beneficial in the presence of fixed costs and in the absence of risk adjustment, but they almost always increase consumer welfare. Similar to risk adjustment, the optimal price floor

¹By contrast, the ACA exchanges today use only risk adjustment, not price floors.

usually delivers *lower* prices, but may sometimes result in higher prices.

We conclude that undercutting incentives caused by horizontal selection have important implications for market stability in the Massachusetts exchange. Our simulations suggest that without policies such as risk adjustment and price floors, this market would be a natural monopoly characterized by high premiums and limited choice. Moderate risk adjustment goes a long way toward improving the situation, and price floors are often welfare improving.

These results raise questions about the ability of insurance markets to sustain competition. Indeed, many have raised concerns over the current state of ACA individual insurance exchanges. As of 2021, 54% of U.S. counties (covering 22% of consumers) had just one or two participating insurers.² Our results suggest that high adverse selection and price sensitivity may be at least partially responsible for the low participation in these markets. However, our results also suggest that corrective policies such as risk adjustment and price floors could address these problems and allow the market to support more competition. This is consistent with many calls for these marketplaces to engage in “active purchasing,” wherein market regulators could use price floors, ceilings, and coordination to achieve desired market outcomes.

Related Literature Our paper contributes to several literatures. First, we contribute to work on adverse selection in insurance markets. It has long been recognized that adverse selection can distort prices and contracts (Rothschild and Stiglitz, 1976; Einav et al., 2010b; Bundorf et al., 2012; Azevedo and Gottlieb, 2017). Previous work has also shown that selection can result in no trade at all in insurance markets (Akerlof, 1970; Hendren, 2013). Our paper shows that even when subsidies or mandates ensure that trade occurs, adverse selection can still limit the ability of the market to support multiple competing firms, with important implications for consumer welfare. Our work also builds on previous work studying the interaction of imperfect competition and selection (Starc, 2014; Mahoney and Weyl, 2017a). That work showed that selection can reduce price markups in settings with imperfect competition, implying that policies such as risk adjustment can reduce consumer welfare. Our analysis points out that selection may impact the *level* of competition, potentially outweighing the impacts of markups conditional on participation. When market structure is endogenous, corrective policies such as risk adjustment can sustain higher participation and therefore lower markups.

Our paper also contributes to a growing literature studying policies used to combat selection, such as risk adjustment, subsidies, and contract and price regulation (see Geruso and Layton (2017)). This literature has shown that in some cases these policies can correct price distortions. Some work has also established a variety of unintended consequences of these policies (Geruso and Layton, 2018; Geruso et al., 2021). Our paper introduces an additional benefit of these policies: They can improve consumer welfare by allowing the market to support more competitors. We also consider a new policy to combat selection problems — price floors.

²See data from the Kaiser Family Foundation at <https://www.kff.org/private-insurance/issue-brief/insurer-participation-on-the-aca-marketplaces-2014-2021/>

Finally, our paper contributes to the literature on firm entry in industrial organization (see [Berry and Reiss \(2007\)](#) for a review). This work has focused on fixed and sunk costs and the nature of competition as explanations for limited entry. Our work shows the role adverse selection can play in shaping entry outcomes, including leading to limited entry in settings without fixed costs.

2 Model

We start by presenting a model of insurance market competition to illustrate the paper’s key adverse selection pricing mechanism and its implications for market outcomes. We integrate adverse selection into a standard two-stage entry model from the Industrial Organization literature, in which differentiated firms choose whether to enter a market and then compete on prices to maximize profits. Our goal is to show that even with regulated plan quality, *strategic pricing* becomes a tool for “cherry picking” low-risk enrollees. If sufficiently strong, the result can be a pricing “death spiral” that leads to non-existence of multi-firm Nash equilibrium where all firms earn non-negative profits. We draw parallels between our analysis and the classic idea of natural monopoly driven by large fixed costs.

Section [2.1](#) sets up a general model and discusses the nature of adverse selection. Section [2.2](#) derives general conditions under which the market cannot support a given set of competitors in Nash pricing equilibrium. Finally, Section [2.3](#) analyzes a simple toy model example to illustrate the key economic forces and outcomes and build intuition.

2.1 Model Setup

Consider an insurance market where a set of potential firms $j \in \{1, \dots, J\}$ each has the ability to offer a single contract (or “plan”). Plans differ on a vector of non-price attributes, X_j , which following much of the previous literature, we treat as fixed and determined outside the model. However, in contrast with much of the selection literature, we allow for *general differentiation* on a (multi-dimensional) set of attributes in X_j – including provider networks, cost sharing, and brand preferences – which allows consumers to vary in preference rankings across plans, making variety inherently beneficial. This contrasts with the standard assumption of *pure vertical differentiation* – where there is a clear (single-dimensional) ranking of plan “generosity” (see [Einav et al. \(2010a\)](#) and [Marone and Sabety \(2022\)](#) for examples).

Although the model is general, we focus much of our discussion on cases where plans where plans are roughly “symmetric” in terms of their cost structure and consumer demand for their attributes, X_j . As a leading example—based closely on our empirical setting—consider a market where two plans (“ E ” and “ W ”) differ on their covered hospital networks, with the W plan network better on the west side of the area and the E plan network better on the east side. As a result, there is meaningful differentiation: at equal prices, east-side consumers prefer the E plan that covers their

nearby hospitals, and west-side consumers prefer the W plan. But the plans are symmetric in that neither is higher-cost for an average patient, and consumer risk (though it varies) is uncorrelated with location.³

Our goal is to show the relevance of adverse selection for this type of “symmetric” or “horizontal” differentiation setting where its role has not been previously appreciated. Even when product attributes are fixed (e.g., by regulation), adverse selection dynamics arise through *strategic price competition*. While prior work has suggested that adverse selection problems may be less severe in settings where contracts are horizontally differentiated (Bundorf et al., 2012) and that selection can even lead to lower markups and more consumer surplus (Starc, 2014; Mahoney and Weyl, 2017a), we show that adverse selection pricing dynamics can actually be severe enough to cause competition to unravel entirely.

Insurer Game Insurers compete in a simple two-stage entry game, following a standard setup in the IO literature (Berry and Reiss, 2007). In stage 1, each firm j (with attributes X_j) decides whether to enter the market, with its entry decision denoted $E_j \in \{0, 1\}$. Entering involves incurring fixed costs $F_j \geq 0$ but gives a firm the opportunity to sell its product and earn profits in stage 2.⁴ Non-entering firms earn zero profits. In stage 2, firms observe the set of entrants (denoted E), and participating firms compete on prices (P_j) to maximize profits in standard Nash-Bertrand equilibrium. Regulators require that each participating insurer charge a single price (no price discrimination) and cannot deny coverage to any willing buyer (guaranteed issue), policies that are standard in health insurance markets around the world.

Consumer Cost and Demand Consumers (i) vary in both their demand and costs, with the interaction of the two creating risk selection. Costs vary because of heterogeneous medical risk (R_i) and plan effects on costs, $\delta_j = \delta(X_j)$, which captures a plan’s costliness for an average consumer (e.g., due to varying networks). The expected insurer cost for consumer i in plan j equals

$$C_{ij} = R_i \cdot \delta_j. \quad (1)$$

Consumers also vary in their demand for different plans. We assume that choice utility of consumer i for plan j (at price P_j) takes the following standard form:

$$U_{ij}(P_j) = -\alpha_i P_j + f(X_j, \zeta_i) + \varepsilon_{ij} \quad (2)$$

where $\alpha_i > 0$ is the consumer’s price-sensitivity coefficient; $f(X_j, \zeta_i)$ is predictable utility based on plan (X_j) and consumer attributes (ζ_i , which may include risk); and ε_{ij} is an i.i.d. idiosyncratic

³This example can be thought of as a simple extension of the classic Hotelling line or Salop circle model to incorporate multiple consumer risk types. We work out this model in detail in Section 2.3.

⁴Fixed costs are a standard feature of entry models to help rationalize limited competition. This will be true in our model as well, but we will show that adverse selection can generate limited competition even without fixed costs.

preference shifter. Because this is a discrete choice problem and demand takes a linear-in-price form, we can re-normalize utility by multiplying by $\beta_i \equiv \frac{1}{\alpha_i}$ to yield an equivalent model:

$$\tilde{U}_{ij}(P_j) = -P_j + \beta_i \cdot \underbrace{[f(X_j, \zeta_i) + \varepsilon_{ij}]}_{\text{“Match Quality” } (Q_{ij})} \quad (3)$$

where the price coefficient is now -1 for all consumers, and $V_{ij} \equiv \beta_i \cdot [f(X_j, \zeta_i) + \varepsilon_{ij}] = \beta_i Q_{ij}$ is the consumer’s willingness to pay (WTP) for plan j . WTP is in turn a product of $Q_{ij} \equiv f(X_j, \zeta_i) + \varepsilon_{ij}$, which we call the consumer’s “*match quality*” for plan j , and β_i , the consumer’s “*value for quality*” relative to price, which maps match quality into dollar values. High- β_i types care a lot and will choose their preferred plan (e.g., the one that best covers their preferred doctors) even if it costs a lot more; low- β_i types care only a little bit and will switch to a less-ideal plan if its price is modestly lower. Thus, β_i relates to the *dispersion* of WTP across plans for a given consumer, or the degree to which consumers value plans as highly differentiated (high β_i) vs. little differentiated (low β_i).

The specification in (3) is intentionally parallel to the classic vertical demand model in IO, with the difference that consumers may disagree about plan quality rankings in Q_{ij} . In the special case where consumers agree ($Q_{ij} = Q_j \forall i$), it collapses to a vertical model. We now draw a parallel between the forces driving adverse selection in the two models.

Heterogeneity Underlying Adverse Selection We say that the market features adverse selection if *higher-risk consumers place higher value on plan quality differences*; that is, if:

$$\text{Corr}(\beta_i, R_i) > 0. \quad (4)$$

While consumers no longer agree on plan rankings (captured in each consumer’s Q_{ij}), we can still think about how much different consumers will pay for *their* preferred plan relative to other available options. Adverse selection occurs when this value for match quality, β_i , is positively correlated with consumer risk, and therefore cost. Note that this WTP-cost correlation is precisely the driver of adverse selection in the standard vertical selection model, often motivated via an underlying model of risk-averse utility and varying risk distributions.⁵

Recall that β_i is the inverse of the price coefficient, α_i , from (2). Therefore, the adverse selection condition suggests that $\text{Corr}(\alpha_i, R_i) < 0$, or that *sicker consumers are less price-sensitive* (and healthy consumers more price-sensitive) in their demand. This is a natural and testable condition that we examine in our empirical work. It is also the key force underlying the “adverse selection pricing” incentives that we lay out below.

While we can (and do) test whether condition (4) holds in our empirical setting, there is reason

⁵In the standard vertical model, adverse selection occurs if high-risk consumers have higher demand for generous insurance. Intuitively, sicker people tend to be willing to pay more to upgrade to a generous plan (high- β_i), typically because these plans have lower cost sharing that covers their high expected medical expenditures. Healthier types have lower WTP and are much more price-sensitive (low- β_i , high- α_i).

to think that this condition will often hold in health insurance more generally. This is because the value of most insurance attributes depends on *how frequently* one expects to use care. Consider again the example of the E and W plans that vary in provider networks. A consumer’s extra WTP for the plan that covers their nearby hospitals depends on how often they expect to use care. A healthy east-side consumer who rarely uses care may have low WTP and be willing to switch to the less ideal W plan for modest price savings. A cancer patient who gets weekly treatments at their nearby hospital will pay a lot to ensure it is in-network.⁶ An analogous reasoning applies to a wide variety of insurance attributes, including: financial coverage, drug formularies, utilization limits (e.g., prior authorization rules), and customer service quality.⁷ In all these cases, sicker people are likely to place higher value on a plan with good match quality for their needs. This creates the conditions for adverse selection unless a third factor driving WTP (e.g., income) has a strong enough *inverse* correlation with risk to offset this direct relationship (Cutler et al., 2008).

Insurer Profits and Adverse Selection Pricing We now set up the insurer’s profit function. Consumers choose the plan that maximizes utility in (3), implying demand function $D_{ij}(P) = 1\{V_{ij} - P_j \geq V_{ik} - P_k \ \forall k \in E\}$, which is a function of all entrants’ prices $P \equiv \{P_k\}_{k \in E}$ and (implicitly) the set of entrants, E . (In what follows, all demand, cost, and profit functions depend on E ; we suppress this for notational conciseness.) Summed up across consumers, total demand for plan j is $D_j(P) \equiv \sum_i D_{ij}(P)$.

The profits of a participating firm j equal:

$$\pi_j(P) = [P_j - AC_j(P)] \cdot D_j(P) - F_j \quad (5)$$

where

$$AC_j(P) = \frac{1}{D_j(P)} \sum_i [C_{ij} \cdot D_{ij}(P)] \quad (6)$$

is the insurer’s average costs, aside from fixed costs. Importantly, average costs depend on (all plans’) prices because they affect which types of (cost-varying) consumers choose each plan. This is a key feature of insurance and other “selection markets” (Einav et al., 2021). We say that plan j (at prices P) faces *adverse selection pricing* incentives if when it cuts its price, it attracts relatively healthy consumers so lowers its average costs. Mathematically:

$$\text{Adverse Selection Pricing:} \quad \frac{\partial AC_j(P)}{\partial P_j} > 0 \quad (7)$$

⁶This relationship between demand and risk is built into the classic “option demand” model for networks of Capps et al. (2003). Mapping that model to our setup, Q_{ij} is a measure of network utility conditional on needing care (from a hospital demand system), and $\beta_i = \gamma_i \cdot \text{Freq}_i$ where γ_i is a taste parameter and Freq_i is the frequency of needing care. Because β_i is a direct function of Freq_i , the two will tend to be positively correlated unless the taste parameter γ_i is inversely correlated with risk.

⁷The few attributes where the healthy are likely to value quality more—e.g., discounted gym memberships—are notable exceptions but less common. Moreover, if all plans cover these benefits, they are no longer *differentiating* attributes.

Because quantity and price move inversely, this is equivalent to a “downward sloping” average cost curve in quantity. It is also equivalent to a firm facing a “marginal cost” – the cost of its marginal consumers attracted when it cuts price – below its average cost. Mathematically, defining marginal cost as $MC_j(P) \equiv \frac{1}{\partial D_j / \partial P_j} \sum_i \left[C_{ij} \cdot \frac{\partial D_{ij}}{\partial P_j} \right]$, adverse selection pricing is equivalent to:⁸

$$\iff AC_j(P) - MC_j(P) > 0 \quad (8)$$

Although conditions (7) and (8) are local conditions that vary across plans and depend on P , they relate naturally to the underlying heterogeneity driving adverse selection discussed above. We noted that this heterogeneity suggests low-risk consumers will be more price-sensitive (higher α_i) in their demand. This suggests that the healthy will often comprise a disproportionate share of marginal consumers who switch plans in response to price changes.

When plans are symmetric and charge roughly similar prices – so there is no one cheap plan that has already attracted all the healthiest people – all firms may have adverse selection pricing incentives at once, since by cutting price they can differentially attract the healthiest consumers who would (at equal prices) prefer a competitor’s plan. This is a key feature for the incentives we discuss below. It also distinguishes our model from the two-plan vertical model, in which where only the H plan faces adverse selection pricing, whereas L faces advantageous selection on its pricing margin.

Parallel between Adverse Selection and Fixed Costs It is also useful to define average *total* costs, $ATC_j(P)$, which includes fixed costs:

$$ATC_j(P) = AC_j(P) + \frac{F_j}{D_j(P)} \quad (9)$$

which lets us write profits as $\pi_j(P) = [P_j - ATC_j(P)] \cdot D_j(P)$. A firm is profitable only if $P_j - ATC_j(P) > 0$.

Equation (9) points to a key parallel we draw between adverse selection and fixed costs for competitive incentives. As just discussed, a firm faces adverse selection in pricing if $\partial AC_j / \partial P_j > 0$, or equivalently AC_j is “downward sloping” in quantity. Equation shows (9) that fixed costs generate a similar downward sloping relationship with average *total* costs, ATC_j , since cutting prices to boost demand lets the firm spread fixed costs across more consumers. Both adverse selection and fixed costs contribute to a steeper pricing slope of the average cost curve. To see this, note that by

⁸To see this equivalence, use the product rule to differentiate $AC_j(P)$ in (6). After rearranging terms, this yields $\partial AC_j / \partial P_j = \left(\frac{-1}{D_j} \frac{\partial D_j}{\partial P_j} \right) (AC_j(P) - MC_j(P))$. Because $\frac{-1}{D_j} \frac{\partial D_j}{\partial P_j} > 0$ by the law of demand, $\frac{\partial AC_j}{\partial P_j} > 0$ if and only if $AC_j(P) - MC_j(P) > 0$.

differentiating $ATC_j(P)$ and grouping terms, we can write:

$$\frac{\partial ATC_j}{\partial P_j} = \underbrace{\eta_{j,P_j}}_{\text{Price sensitivity}} \times \left[\underbrace{(AC_j - MC_j)}_{\text{Adverse selection}} + \underbrace{\frac{F_j}{D_j}}_{\text{Fixed costs}} \right] \quad (10)$$

where $\eta_{j,P_j} \equiv -\frac{\partial D_j/D_j}{\partial P_j} > 0$ is the firm's own price semi-elasticity of demand. This expression shows that the price slope of average total costs – a key statistic in our model, as we show below – relates to three features of demand and cost:

1. **Consumer price sensitivity:** Average cost is steeper when consumers are highly price-sensitive, as captured by a high semi-elasticity of demand (η_{j,P_j}).
2. **Adverse selection:** Average cost is steeper if a firm faces greater adverse selection, captured by the gap in costs between its average and marginal consumer.
3. **Fixed costs per consumer:** Average cost is steeper when fixed costs are large relative to demand.

Equation (10) shows that adverse selection and fixed costs affect the slope of average total costs in a very parallel way. Moreover, both of these forces are magnified when consumers are highly price elastic. Further, when these forces are sufficiently strong on the margin, it is possible for $\frac{\partial ATC_j}{\partial P_j}$ to grow large and even exceed 1. As we discuss next, this creates the conditions for unstable price competition that make it difficult to sustain robust competition.

2.2 Adverse Selection Pricing and Unraveling of Competition

When will firms find it profitable to compete in the insurance market, and how does this relate to adverse selection? To understand this, we return to the two-stage game setup above and (proceeding by backward induction) consider the conditions for Nash-Bertrand pricing equilibrium *conditional* on a set of competitors E . Firm j 's first-order condition for profit maximization is: $\frac{\partial \pi_j}{\partial P_j} = \left(1 - \frac{\partial ATC_j}{\partial P_j}\right) D_j(P) + [P_j - ATC_j(P)] \cdot \frac{\partial D_j}{\partial P_j} = 0$, which after rearranging implies:

$$\begin{aligned} \underbrace{P_j - ATC_j}_{\text{Profit margin}} &= \frac{1}{\eta_{j,P_j}} \times \left(1 - \frac{\partial ATC_j}{\partial P_j}\right) \\ &= \underbrace{\frac{1}{\eta_{j,P_j}}}_{\text{Lerner Markup}} - \underbrace{(AC_j - MC_j)}_{\text{Adverse selection}} - \underbrace{\frac{F_j}{D_j}}_{\text{Fixed costs}} \end{aligned} \quad (11)$$

where η_{j,P_j} is the firm's own-price semi-elasticity of demand, and the second line follows from plugging in $\frac{\partial ATC_j}{\partial P_j}$ from equation (10). The firm's profit margin equals the standard Lerner markup over marginal costs (the inverse of the semi-elasticity of demand) times a term that accounts for

adverse selection and fixed costs, $(1 - \frac{\partial ATC_j}{\partial P_j})$. Without adverse selection or fixed costs, this term is 1; as both forces become stronger, $\frac{\partial ATC_j}{\partial P_j}$ grows in magnitude, reducing the effective markup over average costs. The second line makes clear how adverse selection and fixed costs (per consumer) both reduce markups in a parallel fashion.

Prior work on the interaction of adverse selection and imperfect competition has pointed out the selection-pricing relationship in (11) (Starc, 2014; Mahoney and Weyl, 2017a). This work shows that conditional on a set of competitors, adverse selection *disciplines market power* by reducing optimal markups. Firms realize that they must price low to attract healthier consumers, so they sacrifice some markups to do so. This creates a “theory of the second best” style tradeoff between mitigating adverse selection vs. encouraging price competition. Importantly, however, this prior work makes two assumptions:

1. It treats as *fixed* the set of competitors in the market; and
2. It assumes *existence* of a (stable) Nash equilibrium characterized by the FOC in (11).

Our basic conceptual argument is that when adverse selection is strong enough (or fixed costs are large enough), neither assumption can be taken for granted. Strong adverse selection may unravel the existence of profitable Nash equilibrium with a robust number of competitors.

To see why, start by considering what happens as the equilibrium magnitude of $\frac{\partial ATC_j}{\partial P_j}$ grows, corresponding to stronger adverse selection or larger fixed costs. Initially, price markups decline. But suppose that selection or fixed costs grow large enough that $\frac{\partial ATC_j}{\partial P_j} > 1$, or equivalently, $(AC_j - MC_j) + \frac{F_j}{D_j} > \frac{1}{\eta_{j,P_j}}$. In this case, the FOC in (11) cannot characterize an equilibrium. If it did, a firm’s profits would be *negative* even at its optimal price.⁹ Anticipating this, it would not wish to enter the market in stage 1 of the game, and the set of entrants E would not be an equilibrium.

In fact, it is straightforward to see that if a firm at prices P is *not* losing money and faces a steep enough average cost curve that $\frac{\partial ATC_j}{\partial P_j}(P) > 1$, it will not be at an optimum. Instead, the firm will want to cut its price, which will increase its demand and raise its per-consumer profit margin. The lesson is the following: strong adverse selection (or fixed costs) tends to generate strong incentives to further cut prices, implying that a candidate price P is not an equilibrium. We call this the “*price undercutting incentives*” created by strong adverse selection and fixed costs. The following proposition generalizes this logic.

Proposition 1. *Consider a candidate Nash equilibrium price vector $P = (P_j, P_{-j})$. P is not an equilibrium if there exists an “undercutting deviation” price $\tilde{P}_j < P_j$ for any participating firm j such that:*

$$\frac{\Delta ATC_j}{\Delta P_j} > 1 - \tilde{\eta}_{j,P_j} \cdot [P_j - ATC_j(P)] \quad (12)$$

⁹In the case without fixed costs, the firm could raise its price to ∞ and get zero demand and profits, implying that while P_j^* might be a local profit maximum, it cannot be the global maximum. With $F_j > 0$, however, it is possible that (11) characterizes the global maximum and the firm simply cannot earn enough profits to cover its fixed costs.

where $\tilde{P} \equiv (\tilde{P}_j, P_{-j})$, $\Delta P \equiv P_j - \tilde{P}_j$, $\Delta ATC_j \equiv ATC_j(P) - ATC_j(\tilde{P})$, and $\tilde{\eta}_{j,P_j} \equiv \frac{-\Delta D_j / \Delta P_j}{D_j(\tilde{P})}$ with $\Delta D_j \equiv D_j(P) - D_j(\tilde{P})$.

Proof: For P_j to be firm j 's profit maximizing price, it must be the case that $\pi_j(P_j, P_{-j}) \geq \pi_j(\tilde{P}_j, P_{-j})$ for all \tilde{P}_j , or using the definition of $\pi_j(P)$, $[P_j - ATC_j(P)] \cdot D_j(P) \geq [\tilde{P}_j - ATC_j(\tilde{P})] \cdot D_j(\tilde{P})$. Rearranging terms, this is equivalent to $\frac{\Delta ATC_j}{\Delta P_j} \leq 1 - \tilde{\eta}_{j,P_j} [P_j - ATC_j(P)]$, which is the reverse inequality of (12). Therefore, if (12) holds for any \tilde{P}_j , P cannot be an equilibrium. ■

Corollary 1. *If P is a Nash pricing equilibrium where all firms earn non-negative profits, there cannot be an undercutting deviation such that $\frac{\Delta ATC_j}{\Delta P_j} > 1$, or equivalently, $\frac{\Delta AC_j}{\Delta P_j} > 1 - \tilde{\eta}_{j,P_j}(\frac{F_j}{D_j})$.*

When are the conditions of Proposition 1 likely to apply? A natural case is an adverse selection market where multiple “horizontally” differentiated firms (with similar cost structures) compete and charge similar prices (a symmetric candidate equilibrium). In that case, each firm j has a strong incentive to undercut all others’ prices slightly. By doing so, it can attract a sizable group of price-sensitive and low-cost consumers, who are primarily interested in choosing the lowest-price plan in the market. For this undercutting deviation, the effective price semi-elasticity will be high and the marginal consumers’ costs low – precisely the conditions for a steep average cost curve slope (see (10) above). We explore this case further in the simple example in Section 2.3 below.

One way to think about the analysis so far is in terms of the classic Bertrand pricing paradox. When products are identical (so consumers infinitely price elastic between the two), two competing firms will always tend to compete price down to marginal cost. But with adverse selection or fixed costs, average total costs exceed marginal cost — and precisely by the factor $(AC_j - MC_j) + F_j/D_j$ in (11). Therefore, “perfect” competition will lead to prices below average total costs, making it impossible for plans to break even. Differentiated products allows for firms to have some market power, weakening the Bertrand paradox. But the presence of a subset of price-sensitive and healthy consumers — who raise the demand elasticity and are a disproportionate share of marginal consumers — can make it difficult for even differentiated plans to survive.

Taken on its own, Proposition 1 suggests only that a given P will not be an equilibrium if any individual firm j wishes to further cut prices to \tilde{P}_j . But suppose firm j does so. Will the resulting price vector $\tilde{P} = (\tilde{P}_j, P_{-j})$ be an equilibrium? There are two more ingredients of selection markets that makes it challenging for this to be sustainable. The first is the existence of *cost spillovers* between firms. When firm j undercuts its competitors, it reduces its average costs by “stealing” healthy consumers from other firms, raising rivals’ average costs (Equation (13)). The second key ingredient is the *strategic cost interaction* between firm prices and competitors’ average cost slope. By j undercutting and cherry picking healthy types, firm k now has a stronger incentive to “undercut back” and try to retake these healthy consumers. This effect implies that a cut in P_j tends to *steepen* the average cost curve slope of $k \neq j$ (Equation 14).

These two ingredients can be expressed mathematically, for $k \neq j$, as:

$$\text{Cost spillovers: } -\frac{\partial ATC_k}{\partial P_j} > 0 \quad (13)$$

$$\text{Strategic cost interactions: } -\frac{\partial}{\partial P_j} \left(\frac{\partial ATC_k}{\partial P_k} \right) > 0 \quad (14)$$

These strategic interactions are what can create an unstable “race-to-the bottom” in prices. They tend to increase strategic complementarity between firms’ prices: when firm j cuts its price, other firms both lose money (the cost spillover) and have an increased incentive to undercut back (the strategic cost interaction). The undercutting logic of Proposition 1 applies iteratively between firms. This cycle can easily continue until multiple firms are unable to break even, resulting in a non-existence or “unraveling” of pure strategy Nash equilibrium with multiple firms. Below, we show a simple example where this occurs, and we also find these forces to be highly relevant in our empirical work.

Implications for Competition Suppose this unraveling occurs. How should we think about the implication for market outcomes? One possibility is that some firms exit the market—or foreseeing these dynamics, choose not to enter in the first place. Importantly, this exit (or non-entry) may stabilize the market by reducing price elasticity. With fewer firms, any given firm has more inframarginal consumers, making its demand less price-elastic (lower η_{j,P_j}), and raising its Lerner markup in (11). This larger markup allows the market to support a greater degree of adverse selection or fixed costs.

To see this, consider the case with firms that are fully symmetric and price in a symmetric equilibrium at P , with no fixed costs ($F_j = 0$).¹⁰ By applying the requirement that $\frac{\partial AC_j}{\partial P_j} < 1$, one can show that the maximum number of symmetric firms N that the market can support is:

$$N < \frac{1}{(AC_j - MC_j) \cdot \left(-\frac{1}{D_{tot}} \frac{\partial D_j}{\partial P_j} \right)} \quad (15)$$

where D_{tot} is the total number of consumers in the market, and $N < 2$ implies natural monopoly.¹¹ Note that multiple firms could survive by *coordinating* to avoid the race-to-the-bottom pricing described above.¹² Indeed, as we argue later, this may be a desirable setting for price regulation.

¹⁰In symmetric equilibrium, $D_j = D_{tot}/N$.

¹¹To consider an example (based on our empirical work), suppose that by cutting price by \$10 below P , firm j can attract 10% of all consumers to switch to its plan, or $\left(-\frac{1}{D_{tot}} \frac{\partial D_j}{\partial P_j} \right) = 0.10/\$10 = 0.01$. If these marginal consumers cost \$50 less than average (about 13% below average in our data), then the denominator in this formula is 0.50, implying that the RHS is 2. Therefore, the market can support $N < 2$ firms, which implies natural monopoly.

¹²This coordination could possibly occur through dynamic price cycling, perhaps facilitated by inertial consumers (as in “invest-then-harvest” pricing). Importantly, however, if the conditions of Proposition 1 apply, any firm would have a static incentive to “cheat” and undercut its competitors in a year when it is supposed to price high.

2.3 A Simple Example

To illustrate the forces just described, we adapt the [Salop \(1979\)](#) model of monopolistic competition among differentiated firms to allow for adverse selection. This is a classic model for understanding how fixed costs affect entry; it is therefore natural to use it to understand the relevance of adverse selection. In the model, a population of consumers reside uniformly around a unit-circumference circle. A set of N firms (which we will solve for) locate equidistantly around the circle. Firms incur fixed cost F to enter the market. They sell a homogeneous product of equal value V to all consumers (i.e., there is no vertical differentiation), but consumers dislike travel so prefer nearby firms. Location, therefore, captures horizontal differentiation – which for health insurance might include features like local provider network coverage, insurer reputation, and past consumer experiences.

The standard model includes a single type of consumer with fixed marginal cost to firms and disutility of travel. We enrich this setup by allowing for two (unobserved) types of consumers: (1) healthier type L consumers, who comprise share θ_L of the population, and (2) sicker type H consumers, who comprise share $\theta_H = 1 - \theta_L$. Type H incurs higher medical costs C_H to the insurer and also has a higher travel cost t_H , which implies a higher value for firm location (the horizontal dimension of differentiation). Type L has lower medical costs $C_L < C_H$ and lower travel costs $t_L < t_H$.¹³ Consumers of type $i \in \{L, H\}$ have utility for firm j of

$$U_{ij} = -P_j + (V - t_i \cdot \|\ell_i - \ell_j\|) \quad (16)$$

We assume for simplicity that all consumers buy exactly one good, and there is no outside option of not buying. It will also be convenient to consider price-sensitivity $\alpha_i \equiv 1/t_i$, which is the coefficient on price in (re-scaled) utility if all consumers care equally about travel distance.

In equilibrium, each of the N firms competes with its two adjacent neighbors for consumers living in between them. Share demanded among type- i consumers for firm j (with neighbors $j - 1$ and $j + 1$) equals $D_{ij}(P) = \frac{1}{N} - \frac{1}{2}\alpha_i((P_j - P_{j-1}) + (P_j - P_{j+1}))$, which is a linear demand curve with slope $\frac{\partial D_{ij}}{\partial P_j} = -\alpha_i$, where recall $\alpha_i = 1/t_i$. Total demand is $D_j(P) = \sum_i \theta_i D_{ij}(P)$, and overall profits equal $\pi_j(P) = \sum_i (P_j - C_i) \theta_i D_{ij}(P) - F$. In symmetric equilibrium, all N firms charge the same price ($P_j = P^*$ for all j) and split the demand overall and for each type ($D_{ij}(P) = D_j(P) = \frac{1}{N}$). Solving for a firm's pricing FOC and imposing symmetry yields:

$$P^* = \underbrace{\frac{\sum_i (\theta_i \alpha_i) \cdot C_i}{\sum_i \theta_i \alpha_i}}_{=MC_j} + \underbrace{\frac{1/N}{\sum_i \theta_i \alpha_i}}_{\text{Lerner Markup}} \quad (17)$$

where the second is the Lerner markup ($= 1/\eta_{j,P_j}$) and the first term is marginal costs. Marginal costs equal a weighted average of type-specific costs C_i , with weights proportional to population

¹³The model yields identical insights if we instead allow medical and travel costs to be positively (but imperfectly) correlated. We focus on the simpler case here for expositional simplicity; our empirical model allows for flexible heterogeneity.

shares (θ_i) times type-specific price sensitivity ($\frac{\partial D_{ij}}{\partial P_j} = -\alpha_i$). We denote this marginal consumer share as $s_i^{MC} \equiv \frac{\theta_i \alpha_i}{\sum_k \theta_k \alpha_k}$.

A key point in this model is that adverse selection plays a major role in shaping equilibrium outcomes — despite firms being symmetric and attracting equal shares of healthy and sick consumers in equilibrium.¹⁴ To see this, note that by symmetry, $AC_j = \sum_i \theta_i C_i$, and

$$AC_j - MC_j = \sum_i \theta_i C_i - \sum_i s_i^{MC} C_i > 0 \quad (18)$$

Because $\alpha_L > \alpha_H$ (i.e., healthy type L consumers are more price-sensitive), $s_L^{MC} > \theta_L$ and $s_H^{MC} < \theta_H$. In words, healthy consumers comprise a larger share of marginal than average consumers (and inversely for sicker consumers), implying that average costs exceed marginal costs, the key feature of adverse selection.

Adverse selection, in turn, has implications for the number of firms that can survive in equilibrium. For instance, with $F = 0$, equation (11) implies that profits will only be positive if $\frac{1}{\eta_{j,P_j}} - (AC_j - MC_j) > 0$, which in this model simplifies to:

$$N < \frac{1}{(AC_j - MC_j) \cdot \left(-\frac{\partial D_j}{\partial P_j}\right)} \quad (19)$$

where note that AC_j , MC_j , and the demand slope $\frac{\partial D_j}{\partial P_j} = \sum_i \theta_i \alpha_i$ are all constants determined by primitives in this model. Thus, the maximum number of firms that can be sustained decreases with the degree of adverse selection and consumer price sensitivity. Fixed costs further reduce the number of firms that can be sustained, though the formula becomes more complicated.

Calibrated Outcomes How does this outcome play out empirically? To understand this, Figure 1 presents results from a simple calibrated version of this model, with parameters calibrated based on estimates from our Massachusetts exchange data.¹⁵ We plot the maximum number of firms surviving (panel A) and equilibrium price (panel B) across varying degree of adverse selection and fixed costs. The x-axis captures the degree of adverse selection, captured by the ratio C_H/C_L , which varies from 1 (all enrollees have equal cost; no selection) up to 3 (the sicker enrollees have costs 3x that of healthier). We use $\theta_L = 0.5$, so the L (H) types represent people with below-

¹⁴This implies, for instance, that a positive correlation test (Chiappori and Salanié, 2000) would not detect adverse selection, despite its importance.

¹⁵Based on our data, we set the overall market average cost at \$375 per month. We set $\theta_L = 0.5$ and plot outcomes at varying levels of adverse selection – captured by the ratio C_H/C_L , or the extent to which H types are more expensive (but holding overall average cost fixed at \$375). Based on model estimates, we set the overall price semi-elasticity of demand to be 2.5% per \$1 of price increase and assume $t_H = 2t_L$, which generates price sensitivity twice as high for healthy L types as sicker H types. Finally, we consider fixed costs ranging from \$0 up to \$30 per enrollee-month, where the latter is roughly equal to insurers total administrative costs reported on financial reports to the regulator. This is a plausible upper bound on fixed costs, since some administrative expenses are variable costs.

(above-)median price sensitivity. The different series on each graph are four levels of fixed costs ranging from $F = \$0$ up to $F = \$30$ per enrollee in the market (relative to average medical costs of \$375 per enrollee), where \$30 is a rough upper bound based on what the average insurer reports for *total* administrative expenses on exchange financial reports.

Figure 1A shows the number of competing firms. Competition declines with stronger adverse selection for any level of fixed costs, and with fixed costs for any degree of adverse selection. With $F = \$0$, while (in theory) infinite firms can survive without adverse selection ($C_H/C_L = 1$), this quickly declines to just four firms if sick types are just twice as expensive as healthy. Even moderate fixed costs of $F = \$5$ per enrollee-month mean that only $N = 6$ firms survive without adverse selection, and this declines to $N = 3$ firms with $C_H/C_L = 2$. With high-end fixed costs of $F = \$30$, only a monopolist can survive as long as adverse selection is strong enough that $C_H/C_L \geq 2$.

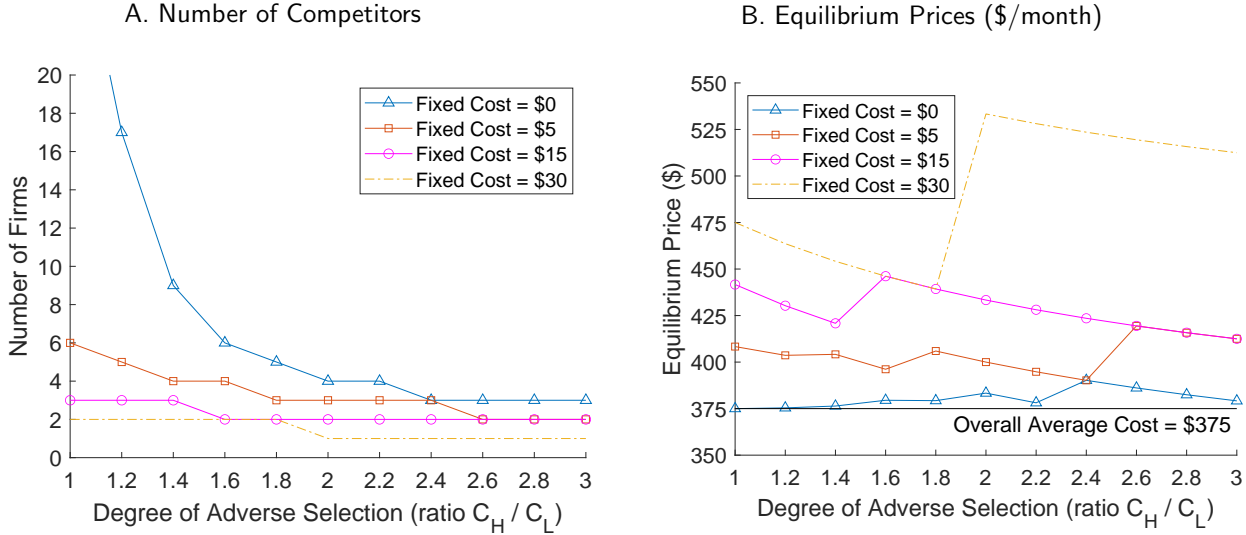
Figure 1B shows equilibrium prices (P^*). In all cases, prices exceed the market average costs of \$375 plus fixed costs, which is the minimum required for firms to break even. But in many cases, prices are substantially higher because of the lack of competition. For instance, in the case with $F = \$15$, the minimum sustainable price if insurers could coordinate is $\$375 + \$15 = \$390$, but actual prices range from \$413 to \$446 (or 6-15% higher). Prices get particularly high when only a monopolist can survive (with $F = 30$), reaching \$533, more than 30% above the minimum coordination price of $AC + F = \$405$.

Another important point in Figure 1B is that prices are a *non-monotonic* function of the degree of adverse selection. For segments where the number of competitors is constant, stronger adverse selection leads to *lower prices*, consistent with the results of [Mahoney and Weyl \(2017b\)](#). But when adverse selection crosses a threshold where a firm exits, prices jump upward because of the weaker competition.

3 Empirical Setting and Data

To investigate the empirical importance of adverse selection on insurer participation, we turn to data from the Massachusetts Health Connector, the state’s precursor to the Affordable Care Act (ACA) Health Insurance Marketplaces. We use this setting to provide reduced form estimates of the key parameters highlighted by our model: the firm-specific price semi-elasticity of demand and the slope of the firm-specific average cost curve. We estimate these parameters using two natural experiments, described in detail below. We then estimate a full structural model of consumer demand and plan costs and use that model to run counterfactual simulations illustrating the effects of policies such as risk adjustment, price floors, price smoothing, etc. on insurer participation, prices, and consumer surplus.

Figure 1. Equilibrium Number of Competitors and Prices in Simple Model



3.1 Setting: Subsidized Massachusetts Exchange (CommCare)

We study the pre-ACA subsidized Massachusetts health insurance exchange, a program called Commonwealth Care (or “CommCare”). CommCare provided subsidized health insurance for Massachusetts residents with incomes below 300% of the federal poverty level (FPL) who did not qualify for Medicaid, Medicare, or job-based coverage. Consumers chose between four to five competing health insurers, with each insurer offering a single highly regulated plan that followed standardized cost sharing rules. The primary way plans differed was on their provider networks. Insurers were primarily Medicaid-based insurers offering limited networks similar to those of their Medicaid managed care plans. In the 2007-2011 period we study, three participating insurers (BMC, NHP, and Network Health) had comparably broad but differentiated networks, covering 75-85% of hospitals (see Appendix Figure A7 for a graph of network size over time). One plan (Fallon) was a regional carrier offering coverage only in central Massachusetts, and a final plan (CeltiCare) was a new entrant in 2010 offering a much narrower network.

The determination of premiums in the CommCare market proceeded in two steps. First, insurers set *pre-subsidy* premiums each year, which determined their base revenue for a standard enrollee. These premiums could only vary based on specific factors: at the income group and region level (2007-09), regional level (2010), or statewide (2011+). Premiums could not vary with other factors including age or medical status.¹⁶

Second, premiums were then *subsidized* to ensure affordability. CommCare’s subsidies were more complex than in the ACA exchanges today, but we explain the details because they matter

¹⁶Technically, during 2007-09 firms could set base prices separately for certain age-sex groups. However, this variation did not affect subsidized enrollee premiums. Instead, the state first calculated a “composite” average premium by income-region cell, and applied subsidies to this composite to determine enrollee premiums.

for the economics of the market and provide our key source of identifying variation. Unlike the ACA exchanges, CommCare’s subsidies were not a flat amount (which reduce the prices of all plans equally) but followed a progressive formula that affected both premium levels and *differences* across plans. Specifically:

- “Below-poverty” enrollees (0-100% of FPL) were *fully subsidized*; all available plans were \$0.
- “Above-poverty” enrollees (100-300% of FPL) were *partly subsidized*. The cheapest plan cost an income-varying “affordable amount,” which rose from \$0 for 100-150% FPL to \$116 per month at 250-300% FPL. Higher-price plans cost more, following a progressive formula where price differences passed through to enrollees at a rate that rose with income.

As an example, consider the subsidy schedule in 2009. For below-poverty enrollees all available plans were \$0. For enrollees 100-150% of FPL, the cheapest plan cost \$0, and each higher-price plan cost 50% of the pre-subsidy price gap between it and the cheapest plan (a 50% pass-through). For higher-income enrollees, the price of the cheapest plan was either \$39 (150-200% FPL), \$77 (200-250% FPL), or \$116 (250-300% FPL) per month, and pre-subsidy price differences were fully passed through to enrollees.

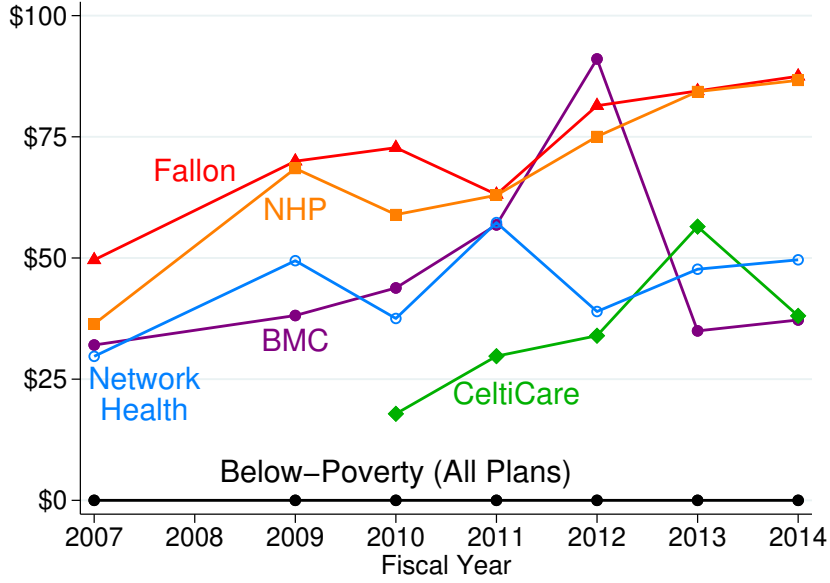
This subsidy structure, while complex, has two important implications for our analysis. First, they create useful identifying variation in the premiums different consumers pay *for the same plan choice set*. In particular, below-poverty enrollees, who are insulated from prices, serve as a sort of “control group” for estimating the demand impact of premium changes. By comparing demand responses to plan price changes for above- and below-poverty groups, we can infer the impact of premiums separately from any unobserved changes in plan quality. We also get additional variation from the heterogeneous pass-through rate among above-poverty enrollee groups. We use precisely this identification strategy in both our reduced form analysis and for our structural demand model.

Second, this structure implied that CommCare’s subsidies played an important role in softening insurer price competition, a key force in our model. Because of subsidies below-poverty enrollees, who comprised about half of the market, were completely inelastic to firm price changes. This substantially lowers firms’ effective price elasticity of demand as is relevant for the undercutting incentives highlighted in our model.

Figure 2 plots average net-of-subsidy premiums for each plan as paid by above-poverty enrollees. The black line at \$0 indicates that all plans were free for below-poverty enrollees, even as premiums varied across plans for above-poverty enrollees. There is substantial variation across plans and over time, including in the identity of the cheapest plan, which we make use of in our empirical analysis.

In addition to subsidies, CommCare used two other policies that are relevant for the key forces in our model. First, starting in 2010, it applied risk adjustment to insurer revenues based on their enrollees’ measured health risk. For an insurer that set a base price of P_j and attracted enrollees with average risk score $\bar{\varphi}$, the plan received revenues of $\bar{\varphi} \cdot P_j$. Risk adjustment, which is also used in the ACA, can mitigate the degree of adverse selection by reducing differences in risk-adjusted costs

Figure 2. Variation in Enrollee Premiums in CommCare



Note: The graph shows post-subsidy enrollee premiums for each insurer’s plan in the CommCare market, by fiscal year and income group. Enrollment-weighted average premiums for above-poverty enrollees (100-300% of poverty) are shown in different colors by plan and labeled. As shown, these vary substantially across plans and over time. For below-poverty enrollees, subsidized premiums are \$0 for *all plans* in all years. We use this subsidy-driven, within-plan premium variation as the key source of identification for our DD estimates and for our structural demand model.

across enrollees. However, risk adjustment is known to be often imperfect, and there is evidence of imperfect risk adjustment in CommCare specifically (Shepard, 2022).

Second, the exchange directly regulated plan using price ceilings and floors. Price ceilings were intended to actively control costs — especially given the presence of price-inelastic below-poverty consumers — and were gradually tightened so that they became binding for about half of plan prices set during 2011-13. Price floors were imposed under rules requiring that premiums be “actuarially sound,” that is no lower than a minimum level defined by an independent actuary. These floors were often binding, especially for BMC, CeltiCare and Network Health. As a result it was quite common that 2+ plans tied for the lowest premium in a region. Although not explicitly intended to ensure participation, these floors may have had the effect of constraining a race-to-the-bottom in prices. Neither price floors nor ceilings are explicitly used in ACA markets today.

In sum, CommCare’s regulator oversaw the market using a highly active form of the “managed competition” model envisioned by Enthoven (1993) — much more so than in ACA markets today. This active management may help explain the ability of the market to sustain robust and stable set of competitors over time, despite the forces we highlight in our model. We use the CommCare market (and its rich premium variation) as a laboratory to infer key demand and cost elasticities relevant for the theory. This lets us assess the strength of these forces and assess their counterfactual

relevance in a setting (like the ACA) that does much less to soften and regulate price competition.

3.2 Enrollment and Claims Data

We acquired enrollment records and full medical and prescription drug claims data for the universe of CommCare enrollees. The enrollment records provide demographic and geographic information for each enrollee as well as monthly enrollment information so we can observe when the individual first enrolled in a CommCare plan, which plan she enrolled in, and if she ever switched plans or left the market and returned again later. We also observe each enrollee’s income and geographic market, allowing us to identify the net-of-subsidy prices of each plan in the enrollee’s choice set.

In addition to enrollment information, we also have full claims data for each enrollee for all months that they are enrolled in a CommCare plan. This data allows us to construct measures of healthcare utilization and spending for each person, including total insurer claims costs. This data also allows us to construct diagnosis-based risk scores, similar to the ones used by the CommCare in its risk adjustment program, where plans that enrolled healthier-than-average (according to the risk scores) enrollees transferred money to plans that enrolled sicker-than average enrollees. We use these risk scores in counterfactual simulations below.

4 Descriptive Evidence

We start by presenting descriptive evidence of the type of strategic pricing in response to adverse selection implied by our model in Section 2. Specifically, we present 3 ‘case studies’ of plans in the Massachusetts Connector undercutting one another. We show that this type of undercutting on price has significant effects on the market share and average cost of the under-*cutting* plan and the under-*cut* plan, with the undercutting plan gaining substantial market share and seeing a large reduction in its average cost, while the undercut plan loses market share and sees an increase in its average cost.

Following these cases studies, We then leverage exogenous variation in plan prices over time to estimate summary parameters describing overall levels of price sensitivity and adverse selection. Specifically, to estimate these summary parameters we leverage changes in plan prices over time in a difference-in-differences design, comparing consumers who face premiums to consumers who are fully subsidized and whose plan choices are thus unaffected by year-to-year price changes. Again, we find strong evidence of high levels of price sensitivity and adverse selection.

4.1 Case Studies

To illustrate the consequences of undercutting for a plan’s market share and average cost, we identify three cases where one plan undercuts another on price. The primary case focuses on Network Health and BMC in plan-years 2012 and 2013. In 2013, BMC dropped its bid from just under \$450 to just under \$350 to undercut Network Health, previously the cheapest plan in the

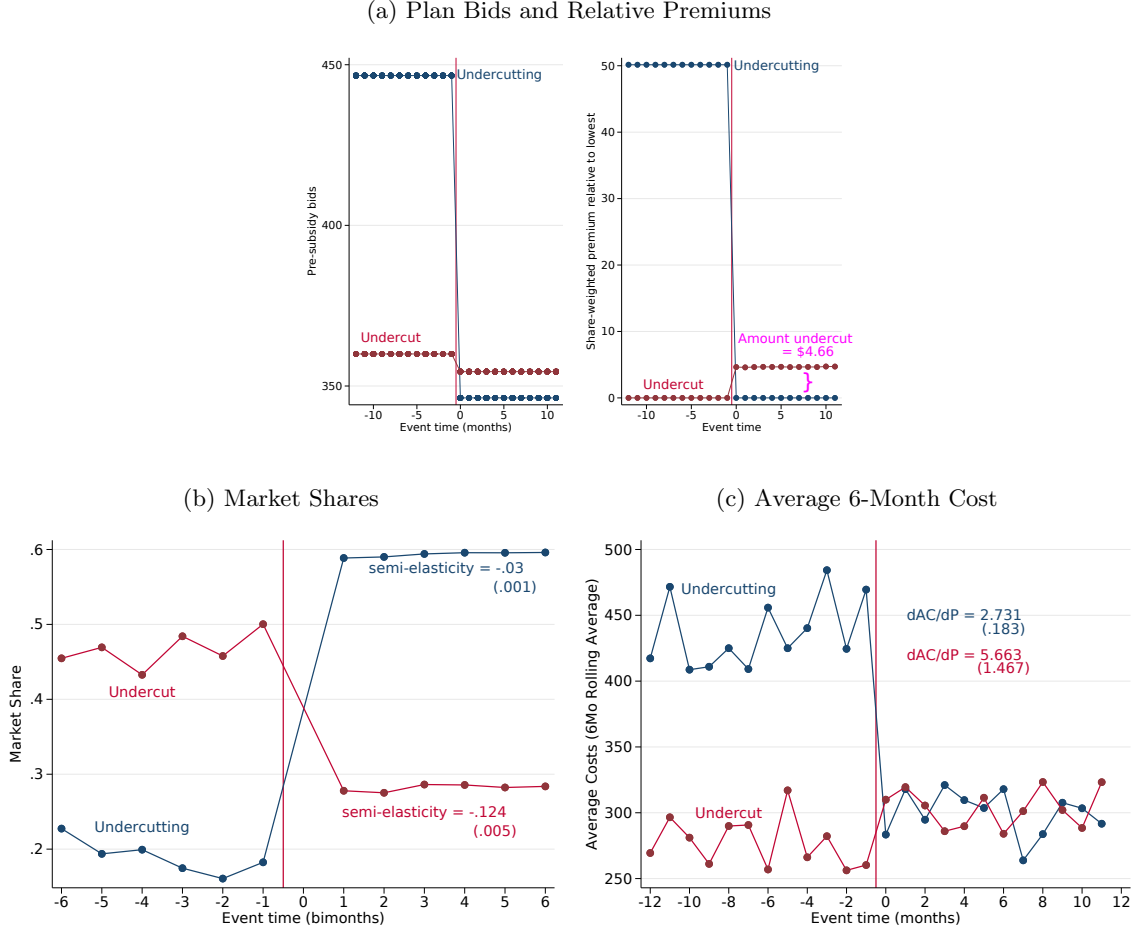
market. This can be seen in Panel (a) of Figure 3. In 2013, the after-subsidy price gap between these two plans was just under \$5 per month on average, ranging from \$3 for the lowest-income group to \$8 for the highest-income group. Despite this small difference in 2013 prices between Network Health and BMC, Network Health’s (the *under-cut* plan) market share plummeted in 2013, dropping from around 50% to around 30% (Panel (b)). BMC (the *under-cutting* plan), on the other hand, saw its market share spike from around 20% to around 60%. These shifts in market share correspond with price semi-elasticities of -0.124 and -0.03, respectively.

Panel (c) of Figure 3 shows how the average cost of the BMC and Network Health enrollees shifted around the time of the price change. BMC saw an enormous drop in average cost from around \$450 to around \$300 at the time of the price change. Given that the drop in BMC’s price was around \$100, this shift in average cost makes it clear why BMC would want to undercut Network Health in this way — With a \$100 drop in price, BMC simultaneously increased its market share by 300% and *increased* its profit average margin. This is exactly the type of adverse selection pricing incentive implied by our model in Section 2. Network Health’s average cost likewise increased, though only by around \$50. While this change in average cost is much smaller than BMC’s change, Network Health’s (relative) price only moved slightly between 2012 and 2013, increasing by only around \$5. This shift in average cost thus also implies a steep own-firm average cost curve for Network Health. It also implies that *both* plans were adversely selected on price.

The remaining two case studies are presented in the appendix. Appendix Figure A1 shows a case where Network Health was priced around \$18 per month above Celticare in 2011 but cut its price to tie Celticare in 2012. In response, Network Health saw a large increase in market share (around 20 percentage points, or more than 60%) and an enormous decrease in average cost (around \$100, or around 30%). Appendix Figure A2 shows where Celticare undercut Network Health in the prior year, going from effectively being tied with Network Health in 2010 to being priced around \$16 below Network Health in 2011. Relative shifts in market share were again substantial, while shifts in average cost were noisy but with point estimates still suggesting very steep own-firm average cost curves.

Overall, these case studies illustrate the strong under-cutting incentives we discuss in Section 2. In all cases, the undercutting plan simultaneously increased its market share *and* increased its profit margin per enrollee by dropping its price below the price of its competitor. The under-cut plans simultaneously saw decreases in market share and decreases in their profit margins. These case studies thus suggest that plans in this market have strong incentives to undercut their competitors. However, these cases were just a small (potentially cherry-picked) subset of price changes in this market. In the next section, we test whether these large shifts in market share and average cost are restricted to the cases we chose or more general. We also address the possibility that other plan (networks, quality) or market factors were changing simultaneously with price, implying that the shifts in cost and market share were not due to the change in price but to these other factors.

Figure 3. Case Study: BMC and Network Health in 2012-2013



Note: Figure shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over time. In all of the figures x-axis is time in months or bimonths relative to month 1 in 2013.

4.2 Difference-in-Differences

We now estimate summary parameters describing overall average levels of firm-specific price sensitivity and adverse selection. Specifically, we set out to estimate the average price semi-elasticity of demand across all plan-years ($E_J[\frac{\partial D_i}{\partial P_j}]$) and the average slope of the firm-specific average cost curve at the observed prices across all plan-years ($E_J[\frac{\partial AC_i}{\partial P_j}]$). While these parameters are clearly equilibrium objects rather than market primitives, it is still useful to understand their values at observed market prices to provide some sense of whether this market is generally characterized by high versus low levels of price sensitivity and strong versus weak adverse selection. This exercise

also helps validate and describe our method for estimating the key parameters of a structural model of demand in Section 5.

To estimate these parameters, we leverage variation in prices over time and across income groups in a difference-in-differences design. Specifically, we leverage the fact that changes in plan bids from one year to the next affect different income groups differently, as described in Section 3. Changes in bids have no effect on the incremental net-of-subsidy prices paid by those with incomes less than 100% of FPL, because for that group all plans are free. Changes in bids *do*, however, affect the incremental net-of-subsidy prices paid by those with incomes greater than 100% of FPL, with those net-of-subsidy prices determined by the ordering and relative levels of the plan bids in the market. While prices shift in January of each year, enrollment occurs continuously throughout the year. Thus, we use individuals with incomes below 100% of FPL as a control group to capture any shifts in market share across plans due to time-varying factors other than the change in premiums, such as changes in the composition of Connector enrollees, changes in plan networks, or other changes in plan benefits. We note, however, that plans and the composition of enrollees were generally fairly consistent over time, resulting in little change in plan market shares among the control group between consecutive years.

We combine many price changes across many markets and several years in a “stacked” difference-in-differences design. Define an experiment e at the plan-region-consecutive year pair level. Basically, we consider each year-to-year price change for each plan in each market, which affects people in different income groups differently, as a single experiment. We restrict each experiment to the 12 months before and the 12 months after the price change. Within the experiment, we effectively compare changes in market share and average cost for the premium-paying groups to changes in market share and average cost for the non-premium-paying control group. Formally, each experiment contains multiple income groups $g \in \{0, 1, \dots, 4\}$, which correspond to the five income groups in the market (0-100% FPL (control) and 100-150%, 150-200%, 200-250%, and 250-300% FPL (treatment)).

To establish the validity of the difference-in-differences design, we start by presenting event study plots. Initially, we stay as close as possible to the raw data by dividing experiments into two groups: price increases E^{incr} and price decreases E^{decr} . We estimate the effects of price changes on market shares and average cost separately for each of these groups to show symmetry. We leverage all plan-market-year experiments to estimate the average effect of the price change across experiments. To do so, we stack all experiments and estimate the following event-study regression specification:¹⁷

$$Y_{egt} = \tilde{\alpha}_{et} + \tilde{\gamma}_{eg} + \sum_{k \in \{-T, T\} \setminus \{-2\}} \tilde{\delta}_t \times 1\{g > 0, t = k\} + \tilde{\varepsilon}_{egt} \quad (20)$$

¹⁷The following specification assumes a homogeneous treatment effect across all income groups, and estimates a single coefficient across all groups.

Each outcome Y_{egt} is measured at the experiment-income group-month level. The regression specification includes experiment-by-event time fixed effects $\tilde{\alpha}_{et}$ to ensure that we only use *within-experiment* variation in prices to identify the firm-specific demand response to the change in prices. We also include experiment-by-income group fixed effects to net out average differences in costs or preferences across income groups — if the higher-income treated groups have lower costs in general, these fixed effects ensure that the pre-period difference in costs is netted out, such that the treatment effects $\tilde{\delta}_t$ only represent post-period differences in the outcome variable (i.e., changes due to the premium change). The coefficients of interest are the $\tilde{\delta}_t$ s. These are the event study coefficients, and they reflect how the gap in Y_{egt} between the control group and the other groups changes over time, relative to the gap 2 months prior to the price change.¹⁸ We estimate this regression for three key outcomes: relative premiums p , (log) market shares (\log_sh), and average costs (AC).¹⁹

Figure 4 presents the event study plots for the positive and negative price change experiments. Panel (a) shows the average price change for each group, around \$20 for both price increases and price decreases. Panel (b) shows the changes in log market share. The event study plot shows that differential market share trends between the treatment and control groups are steady throughout the pre-period for both the positive and negative price change experiments, suggesting that the parallel trends assumption is likely to be satisfied here. At time $t = 0$, however, market shares diverge between the treatment and control groups. For the price increases, market shares decrease by around 20%. For the price decreases, market shares increase by slightly more than 20%. The effects of price changes on market shares thus appear to be symmetric. This provides evidence of the credibility of our empirical strategy, as spurious trends in market share are unlikely to be positively correlated with price decreases and negatively correlated with price increases. The effects are also quite strong, with a \$20 (115%) increase in the premium causing a 20% shift in market share.

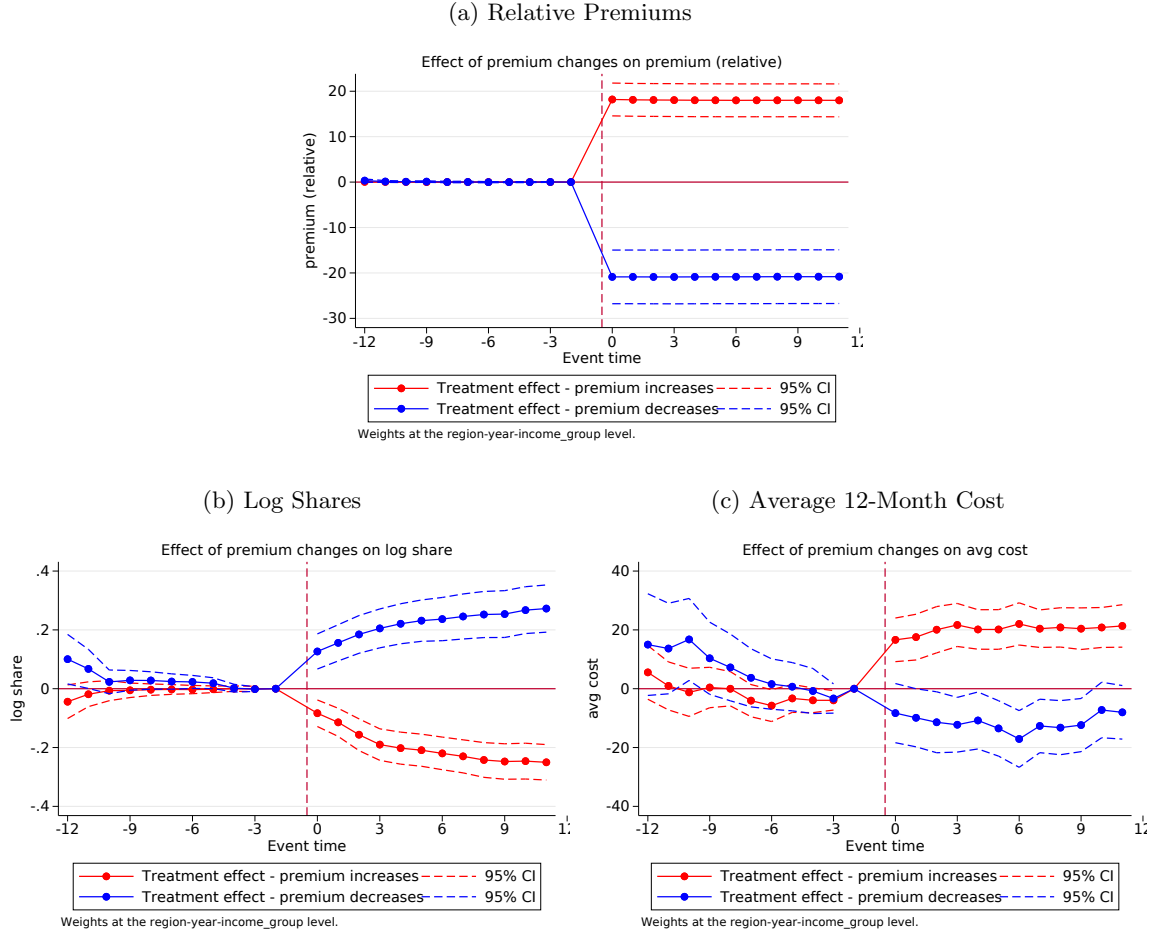
Panel (c) shows the changes in the average cost of plan enrollees that correspond to the shifts in market share documented in Panel (b). Here, estimates are noisier, and there is some evidence of a pre-trend for the premium decreases. However, the plot suggests that when prices increase, average costs increase and when prices decrease, average costs decrease. These results are consistent with adverse selection. The magnitudes are also large: A \$20 price increase leads to an increase in the average cost of a plan’s enrollees of around \$20, suggesting that the slope of the firm-specific average cost curve is close to one. As discussed in Section 2, this raises concerns about the ability of this market to support multiple competing plans in the absence of policies used to combat selection, such as price floors and risk adjustment.²⁰

¹⁸We normalize to month -2 and exclude month -1 because prices for the following year were publicized one month prior to the open enrollment period. We have some evidence that these start to affect demand and costs in month -1.

¹⁹Costs for each enrollee-month are defined as the average monthly cost of the following 12 months, or until we no longer observe the enrollee in the data.

²⁰Note that many of these corrective policies (risk adjustment, price floors, etc.) were in place in the Connector during our sample period, explaining why this market was able to sustain multiple competing plans during this period.

Figure 4. Event Study Estimates for All Enrollees

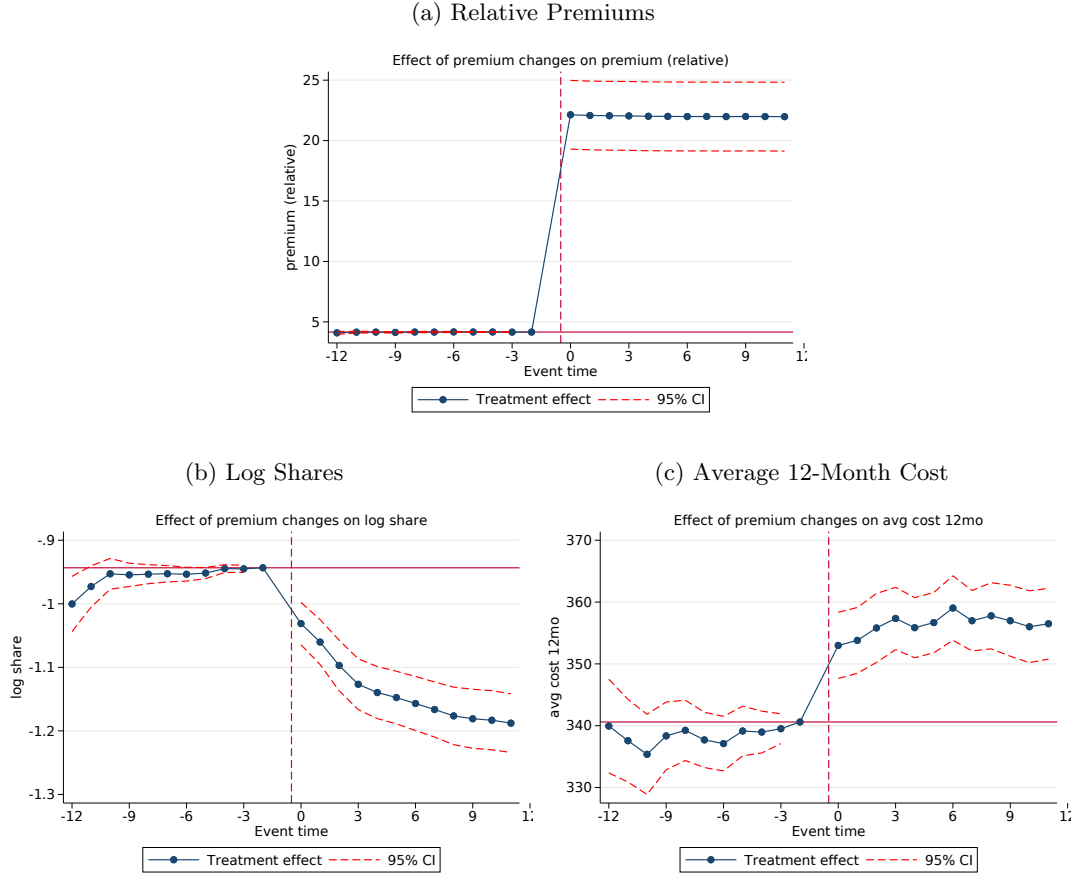


Note: Figure shows event study estimates of the impact of premium increases and decreases on relative premiums. Panel (a) shows results for relative premiums, Panel (b) shows results for log plan shares, and Panel (c) shows results for average costs.

Next, we leverage all experiments in a single unified regression in order to maximize power. Specifically, we multiply all outcomes from negative price change experiments by -1 , stack *all* experiments (both positive and negative price changes), and re-run the specification described in Equation 20. Figure 5 presents the event studies for these “pooled” regressions. Again, Panel (a) shows the average price change (now across both positive and negative price change experiments), just over \$20. Panel (b) shows the effect of that \$20 price increase on log market share. Again, the event study shows that treatment and control market shares trend similarly prior to the price change. After the price change, however, treatment and control market shares diverge, with the \$20 price increase resulting in a 25% decrease in market share.

In Section 6 we perform counterfactual simulations to show the importance of these policies for generating equilibria with multiple competing plans.

Figure 5. Pooled Event Study Estimates for All Enrollees

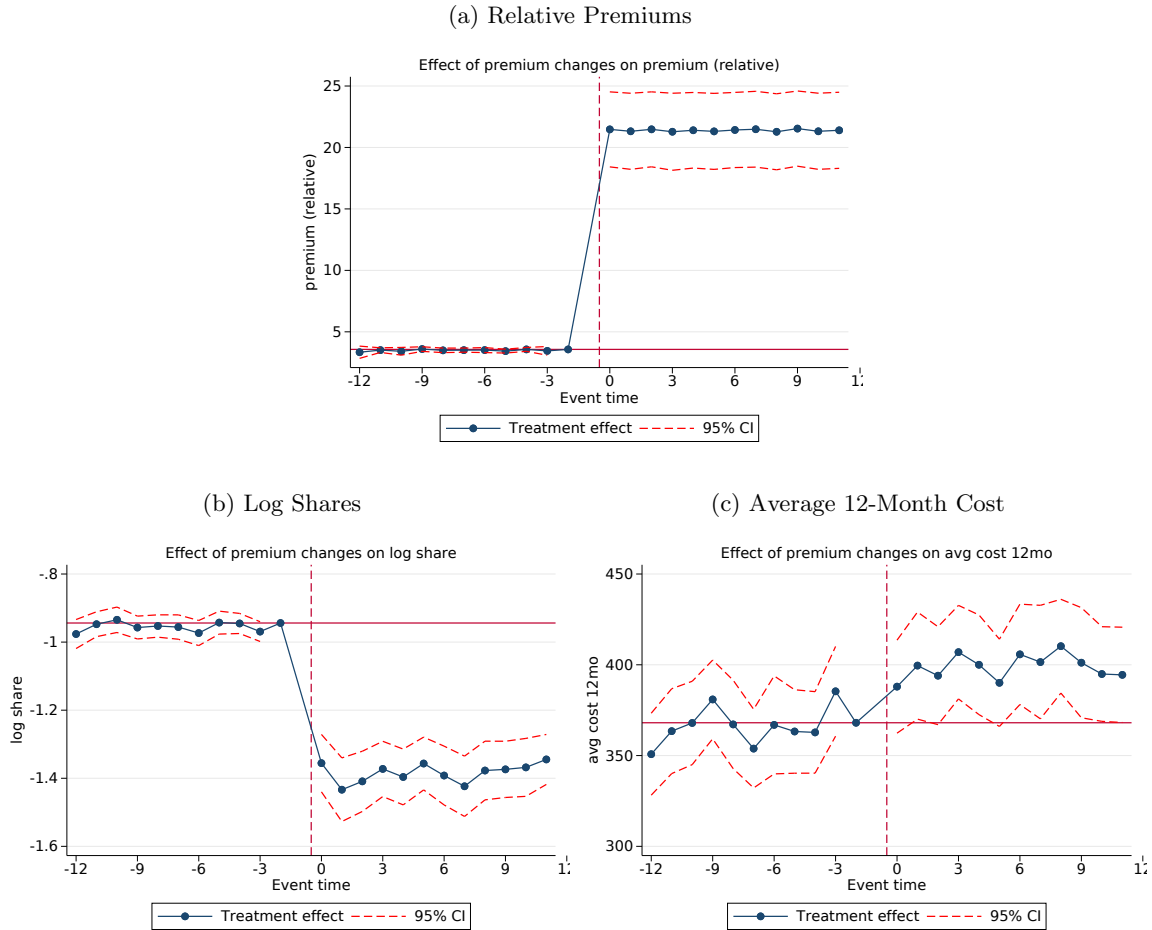


Note: Figure shows pooled event study estimates of the impact of premium increases and decreases, where all outcomes are multiplied by -1 for premium decreases. Panel (a) shows results for relative premiums, Panel (b) shows results for log plan shares, and Panel (c) shows results for average costs.

Here, a dynamic effect is clear, with the effect growing over time. Such a dynamic may seem odd, but it is due to enrollment churn. In the first months of the new year, a large share of the enrollees are incumbent enrollees subject to inertia in their plan choices. As the new year goes on, however, more and more of the inertial enrollees drop out of the market, leaving a larger share of enrollees who entered the market under the new prices. To illustrate this, in Figure 6 we replicate Figure 5, restricting only to new enrollees rather than estimating the effects of prices on both incumbent and new enrollees combined. Again, the change in prices is a little over \$20. The effects on market share, however, are much larger: The same \$20 increase in prices decreases a plan's market share by around 40% among new enrollees compared to the 20% decrease across all enrollees, revealing that inertia is strong in this setting. Further, Panel (b) of Figure 6 shows that among new enrollees, the treatment effect is not dynamic. Instead, as soon as the price change, the market share among new

enrollees drops by 40% and continues at that level in perpetuity. These results combine to show that we would expect the effect of a price change on market share to increase over time as fewer of the consumers in the market are inertial.

Figure 6. Pooled Event Study Estimates for New Enrollees Only



Note: Figure shows pooled event study estimates of the impact of premium increases and decreases, where all outcomes are multiplied by -1 for premium decreases. Sample is limited to new enrollees only. Panel (a) shows results for relative premiums, Panel (b) shows results for log plan shares, and Panel (c) shows results for average costs.

Panel (c) of Figure 5 shows the effect of the price increase on the average cost of plan's enrollees. Here, after pooling across the price increases and decreases, results are less noisy than before. Now we see that treatment and control groups have similar average cost trends prior to the price changes. And, as before, we find that the average cost of a plan's enrollees increases markedly following the price increase, though now the estimates are much cleaner. Specifically, we again estimate that a \$20 increase in the premium results in an increase in the average cost of a plan's enrollees of around \$20. Panel (c) of Figure 6 shows that selection is also stronger among new enrollees, with

a \$20 price increase producing an increase of around \$40 in the average cost of a plan’s enrollees. Thus, the overall slope of the firm-specific average cost curve appears to be around 1, while the slope of the firm-specific average cost curve for new enrollees appears to be nearly twice that level, around 2. Recall that according to our model a slope greater than 1 is sufficient to induce plans to prefer not to enter (and earn negative profits at the profit maximizing price). This market seems to meet that condition overall, and go well beyond that condition when it comes to new enrollees, suggesting (1) serious concerns about this market to sustain multiple competitors in the absence of corrective policies and (2) that in markets with high levels of enrollment “churn” (i.e. many more new enrollees than incumbent enrollees) it may be nearly impossible to support multiple competing plans without serious regulation.

In Table 1, we report the difference-in-differences coefficients for a variety of outcomes and specifications. In Panel (a) we present price elasticity estimates. In Panel (b) we present estimates of the slope of the average cost curve. Column (1) corresponds to Figure 5 and Column (2) corresponds to Figure 6. These estimates again suggest high levels of price sensitivity and strong adverse selection. Columns (3–5) are estimated using all enrollees, but show results for different sets of price changes. Column (3) presents estimates for a specification where we restrict to experiments where there was a change in the cheapest plan. For these cases, it appears that demand elasticities are slightly larger with similar amounts of selection, suggesting that this market may include a group of relatively healthy “choose-the-cheapest plan” consumers. In Column (4), we restrict to experiments with price changes below the median price change magnitude, ultimately finding more selection but lower price sensitivity. Finally, in Column (5) we restrict to experiments with price changes below the median price change *and* where there was a change in the cheapest plan. We find that even small changes in premiums (of about \$5 on average) lead to very large changes in shares (of 14 percentage points). Results for cost selection retain the expected sign but are no longer statistically significant due to smaller samples.

In Appendix B we also present supplemental evidence on the high degree of price sensitivity via a regression discontinuity design, leveraging the sharp difference in prices for consumers with incomes just below 100% of FPL (for whom the price of all plans is \$0) and consumers with incomes just above 100% of FPL (for whom the price of the lowest-priced plan is \$0 but all other plans have positive prices). In Panel (a) of Figure A4, we present the RD plot for the market share of the lowest-priced plan (red) and the market share of the combination of all other plans (blue). The shift in market share at 100% of FPL is striking. The market share of the lowest-priced plan increases from around 20 percentage points to around 40 percentage points, a relative increase of 100%. The increase in the differential price that produces this shift in market share is only \$11.27 per month, or around \$135 per year. Price sensitivity is even stronger when restricting to new enrollees. Ultimately, these estimates suggest that the average implied firm-specific own-price elasticity of demand ranges from -0.015 across all enrollees to -0.020 for new enrollees, again

Table 1. Difference-in-Differences Results

	All Enrollees	New Enrollees	Cheapest Plan Changes	Price Changes < Median	Cheapest plan & Price Changes < Median
<i>Panel (a): Price Sensitivity</i>					
Premium	17.87*** (0.287)	17.90*** (0.318)	21.06*** (0.611)	7.135*** (0.113)	5.329*** (0.829)
Market Share	-0.181*** (0.00522)	-0.429*** (0.0103)	-0.336*** (0.0100)	-0.0459*** (0.00514)	-0.141*** (0.0311)
Demand Elasticity	-0.0101	-0.0240	-0.0159	-0.00643	-0.0264
<i>Panel (b): Adverse Selection</i>					
Premium	17.87*** (0.287)	17.90*** (0.318)	21.06*** (0.611)	7.135*** (0.113)	5.329*** (0.829)
Average Cost	17.81*** (0.758)	32.48*** (3.949)	20.10*** (1.424)	9.689*** (0.919)	3.479 (6.572)
Slope of Avg Cost Curve	0.997	1.815	0.954	1.358	0.653
N	5888	4922	2323	2967	667

Standard Errors reported in parentheses.

*** $p < 0.01$, ** $p < 0.05$, and * $p < 0.10$.

Notes: Table shows estimates from difference-in-difference specifications where price increases and decreases are pooled by multiplying outcomes for premium decreases by -1 . Panel (a) presents price elasticity estimates. Panel (b) presents estimates of adverse selection. Each row corresponds to a different outcome variable. Cells contain coefficient estimates and standard errors from separate regressions of the row outcome variable on premium increases, relative to the below-poverty control group. We compute demand elasticities by dividing market share coefficients by premium coefficients and compute the slope of the average cost curve by dividing the average cost coefficients by the premium coefficients.

indicating extremely high levels of price sensitivity.²¹

Ultimately, the results from the case studies and the diff-in-diff and RD designs combine to provide strong evidence of high levels of price sensitivity and strong adverse selection in this market. Indeed, many of our estimates of price elasticities and slopes of the average cost curve reach levels at which our model suggests firms have strong incentives to undercut their competitors and that there may not exist a set of prices with the observed number of firms that constitute a Nash equilibrium. These results raise major concerns about the ability of this market to sustain multiple competing plans, absent corrective policies.

That said, this market did in practice sustain multiple competitors. How could this be so? First, we note that the relevant elasticities and slopes are equilibrium objects, not fixed market primitives. Our estimates reflect the values of these equilibrium objects at observed market prices, but to understand these elasticities and slopes more generally, we need estimates of the full distribution of consumer types which requires us to estimate a structural model of demand. Second, we note

²¹We consider the RD results descriptive in nature, as there is some concern about whether the assumptions necessary for a valid RD design are met here. This concern is stronger for results around the level of adverse selection, which is why we do not use this design to estimate that parameter. See Appendix B for details.

that this market included a variety of policies meant to correct adverse selection. Such policies may have made it viable for multiple firms to participate.

In the next section, we thus leverage this same quasi-experimental variation in prices to estimate a full structural model of consumer demand and plan costs in this market. We then use that model to map out the elasticities and slopes and to perform counterfactual simulations and assess the roles of various corrective policies in achieving an equilibrium with multiple competing plans.

5 Structural Model and Estimation

In this section, we describe and estimate our structural model of insurance plan choice (demand) and enrollee-level insurer costs. In Section 6 we combine these estimates with a model of equilibrium entry and pricing to study the implications of adverse selection (and corrective policies) for insurer participation, prices, and consumer welfare. Our model follows closely the approaches of [Shepard \(2022\)](#) and [Jaffe and Shepard \(2020\)](#), who also study the CommCare market. We therefore discuss the model briefly and refer readers to the original papers for further details.²²

5.1 Insurance Demand Model

We model enrollees’ plan choices at the start of each enrollment spell as a function of net-of-subsidy premiums and prior plan choices, all interacted with enrollee characteristics. To construct our demand estimation sample, we restrict the data to “choice instances,” defined as one of three times when consumers can choose/switch plans: (1) when an enrollee newly enrolls in the market, (2) when an enrollee re-enrolls after a break in coverage, and (3) when continuing enrollees have the option to switch plans during annual open enrollment.²³ A single enrollee may have multiple choice instances; we index unique enrollee-choice instance pairs by (i, t) .

We estimate a multinomial logit choice model where enrollees choose among one of five CommCare health insurance plans (or the subset available to them in their area-year). We specify the utility enrollee i receives from enrolling in plan j at time t as:

$$u_{ijt} = -\alpha(Z_{it}) \cdot P_{ijt} + f(X_{jt}, Z_{it}; \beta) + \xi_j(W_{it}) + \varepsilon_{ijt}, \quad j = 1, \dots, J \quad (21)$$

where P_{ijt} is plan j ’s post-subsidy premium for consumer i in year t (based on their income group and region). Following [Shepard \(2022\)](#), price sensitivity $\alpha(Z_{it}) = Z_{it}\alpha$ is allowed to vary by income bins, medical diagnoses (chronic disease, cancer, or neither), demographics (age-sex bins), medical (HCC) risk scores, and immigrant status. Relative to [Shepard \(2022\)](#), we add one key covariate to the Z_{it} on which price sensitivity can vary: an estimate of enrollee’s unobserved health risk, based on residuals from a regression of enrollee costs on medical observables (see Appendix Section C.1).

²²We will also include these details in an appendix, which we have not yet written for this draft.

²³The open enrollment period for plan year 2009 was three months long. We code this period as one choice instance, where the final plan chosen during open enrollment is back coded to the first month of plan year 2009.

We bin these residuals into deciles and include them in Z_{it} . We find that this additional covariate is helpful in matching the empirical patterns of adverse selection in response to price changes.

The function $f(X_{jt}, Z_{it}; \beta)$ includes interactions of observed plan and consumer factors that affect demand. These include terms capturing the quality of a plan’s hospital network for a given enrollee, derived from a hospital demand system and from existing relationships with physicians and hospitals (see [Shepard \(2022\)](#) for details). We also include dummy variables for the immediate prior choice of continuing enrollees’ (capturing switching costs and other drivers of state dependence) as well as the interactions of these variables with income, age-sex bins, and risk scores. Finally, we capture unobserved plan quality with $\xi_j(W_{it})$, which are plan-region-year, plan-region-income, plan-risk score, plan-age-sex, and plan-immigrant-status fixed effects. These allow plan quality to vary flexibly to capture changes in plan quality across different areas and years and differences in plan quality for enrollees who differ in their health status. Additionally, these fixed effects ensure that the price coefficient, α , is identified only off of the exogenous subsidy-driven premium variation discussed in Section 3.

Demand Estimates. We estimate the plan choice model using maximum likelihood. The price-sensitivity parameters are identified by within-plan premium variation created by the exchange subsidy rules, as discussed in Section 4 above. Below-poverty enrollees pay no premiums for any plan, whereas higher income groups pay more for more expensive plans. The subsidy rules also result in additional premium variation over time. Plans that increase their premiums over time become more expensive for higher income groups but remain free for below-poverty enrollees. As in our reduced-form analyses, the rich set of plan dummies limits our identifying variation to these differential changes in premiums across income groups. The lack of pre-trends in the reduced-form analyses of shares and average cost lend strong support to our identification strategy.

Table 2 reports the implied own-price semi-elasticities and own-price average cost slopes dAC/dP for each plan and by enrollee group (averaged across plans). These are calculated using the demand model evaluated at observed prices in the data and are based on a \$1 change in a plan’s *post-subsidy* enrollee premium. It also uses the cost model (discussed below). The analysis excludes below-poverty enrollees, who never pay premiums so for whom we cannot estimate price coefficients.

The model implies that a \$1 increase in a plan’s monthly premium (a 0.25% increase as a share of average pre-subsidy prices) lowers its demand by an average of 1.6 percent. These estimates reproduce semi-elasticities from [Jaffe and Shepard \(2020\)](#) implying that CommCare enrollees are very price elastic, with some variation in price sensitivity across plans. In addition, these estimates imply a high degree of adverse selection. Absent risk adjustment, a \$1 *decrease* in monthly premium would lower a plan’s average cost by \$0.899. These estimates imply margins of 21% above marginal cost before risk adjustment. As a result of adverse selection, the implied margins above *average* cost are much smaller, at 1.8%.

For new enrollees, we find that a \$1 increase in monthly premium lowers demand by 2.9 percent.

Without risk adjustment, a \$1 *decrease* in monthly premium would lower average cost by \$1.685. These estimates imply margins of 12% above marginal cost before risk adjustment, if the market were to consist of only new enrollees. As a result of adverse selection, the implied margins above *average* costs are again negative, at -7%.

Table 2. Implied own-price semi-elasticities and dAC/dP

	<i>Panel (a) All Enrollees</i>		<i>Panel (b) New Enrollees</i>	
	Semi Elasticity	dAC/dP	Semi Elasticity	dAC/dP
<i>By Plan</i>				
BMC	-0.013	0.647	-0.024	1.107
Celticare	-0.037	0.891	-0.041	1.057
NHP	-0.021	1.376	-0.037	3.060
Network	-0.015	0.881	-0.028	1.682
<i>By Income Group</i>				
100–150% poverty	-0.022	0.843	-0.043	1.763
150–200% poverty	-0.013	0.778	-0.022	1.335
200–250% poverty	-0.012	0.960	-0.019	1.365
250–300% poverty	-0.010	0.843	-0.016	1.242
Overall	-0.016	0.899	-0.029	1.685

Notes: Table reports own-price semi-elasticities and average cost price derivatives (dAC_j/dP_j) by plan, by income group, and for the market as a whole. Own-price semi-elasticities are computed for each plan using the formula $\eta_j = \sum_i (ds_{ij}/dp_j) / (\sum_i s_{ij})$. The average cost price derivatives are computed assuming no risk adjustment, using the formula $\eta_j(AC_j - MC_j)$ introduced in Equation 10. The “Overall” row corresponds to the averages of the plan-specific values, weighted by plan market shares. Results for Fallon are omitted because they enroll few enrollees and because we omit Fallon in our simulation results in Section 6

5.2 Cost Model

To compute the degree of adverse selection predicted by the demand model and to simulate equilibrium plan prices and participation, we need to model each insurer’s expected cost of covering each consumer. Our approach to doing so closely follows the approach of [Jaffe and Shepard \(2020\)](#). We assume a model where observed costs for insurer j on enrollee i at time t is the product of enrollee risk (R_{it}) times a factor capturing plan effects on costs ($\delta_{j,r}$) which we allow to vary by region r :

$$C_{ijt}^{obs} = R_{it} \times \delta_{j,r(i)}. \quad (22)$$

We then proceed in two steps. First, we estimate $\delta_{j,r}$. To do so, we leverage cases where the same individual enrolls in the market in two separate spells in which they choose different plans.²⁴ This

²⁴We also explored using individuals who switch plans within a given spell (i.e., at open enrollment). However, we found that this sample was small and non-representative, likely due to the large role of inertia. Moreover, we found cost pre-trends for this analysis, suggesting that year-to-year plan switching is affected by unobserved health shocks.

lets us estimate a model of plan effects after controlling for both time-varying enrollee observables and also individual fixed effects.

Our estimation sample has observations at the enrollee x enrollment spell level, and we limit to individuals observed in at least two spells, separated by a gap in CommCare enrollment.²⁵ We estimate the following Poisson regression specification:

$$E(C_{ijt}^{obs} | Z_{it}) = \exp(\alpha_i + \beta_t + Z_{it}\gamma + \lambda_{j,r}) \quad (23)$$

This specification controls for individual fixed effects (α_i), year fixed effects (β_t), and time-varying enrollee observables Z_{it} (age-sex bins, a spline in risk score, income group, and enrollee location). The $\lambda_{j,r}$ coefficients represent the plan-specific cost effects, which we allow to vary across regions r to account for differential cost structures based on a plan’s regional provider network. The estimated multiplicative plan cost effect of interest is $\hat{\delta}_{j,r} = \exp(\hat{\lambda}_{j,r})$. We normalize the scale of these fixed effects so that $\hat{\delta}_{j,r}$ has an (enrollment-weighted) mean of 1.0 across all plans.

The model above assumes that plan cost effects are constant over time, though they can vary by region. This is reasonable only if the determinants of costs — in our setting, primarily networks — are stable. To facilitate this, we limit the estimation sample to 2007-2011 data, a period over which plan networks are relatively stable (and prior to a major network change that occurs in 2012).

Having estimating $\hat{\delta}_{j,r}$, our second step is to predict enrollees’ costs in counterfactual plans. To do so, we simply follow the specification in (22) to estimate enrollee risk as $\hat{R}_{it} = C_{ijt}^{obs} / \hat{\delta}_{j,r}$. The cost model’s prediction for enrollee i ’s cost in a counterfactual plan k , therefore, simply equals their observed costs times the ratio of the two plan effects, $\hat{\delta}_{k,r} / \hat{\delta}_{j,r}$.

Two points are worth noting about this approach. First, it assumes that plan cost effects take a constant multiplicative form for all enrollees (though they can vary by region), which lets us extrapolate the estimates of $\hat{\delta}_{j,r}$ to the full sample. We think this captures the first-order impacts that seem most relevant for our analysis, but it does miss richer enrollee-level heterogeneous effects that may be relevant for certain issues (e.g., “selection on moral hazard”; see Einav et al. (2013)). Second, the risk estimate, \hat{R}_{it} , should be thought of a *realized* enrollee risk, rather than *ex-ante* risk. In our analysis, we will always average cost outcomes over large groups of enrollees (e.g., all enrollees in a plan), which should generate a measure of expected costs that averages out any (additive) idiosyncratic shock.

Plan Cost Effect Estimates Table 3 shows estimates of cost heterogeneity.²⁶ As expected, CeltiCare has the lowest cost effect, with costs that are 27% lower than average. On the other hand, NHP has costs that are 11% higher than average. The estimated cost effects for each plan

This appears much less true for enrollees who actively choose different plans across two separate spells.

²⁵We also drop a small number of individuals enrolled in Fallon in the Boston and Southern regions to avoid fitting parameters on small cells.

²⁶These are obtained by generating predicted costs with all controls set to their omitted categories and re-normalizing such that the average plan effect is 1.0.

are broadly similar across regions. These estimates imply meaningful heterogeneity in costs across plans, consistent with prior work focusing on cost heterogeneity across Medicaid plans in New York City (Geruso et al., 2020). On the other hand, BMC and Network—which are the two largest plans empirically—have relatively similar cost structures, as in the “horizontal” differentiation case we highlighted in the theory in Section 2.

Table 3. Plan Cost Heterogeneity Estimates

Region	BMC	Celticare	NHP	Network
Boston	1.12	0.70	1.17	1.09
Central	0.83	0.61	1.19	0.92
North	0.84	0.76	1.04	1.01
South	0.95	0.73	1.09	0.87
West	0.97	1.03	1.01	0.90
Average	0.97	0.73	1.11	0.98


Notes: Table shows cost heterogeneity estimates from the Poisson regression model with fixed effects and controls. Reported coefficients describe the multiplicative effect of each plan on costs relative to the average plan, separately by region and on average. Results for Fallon are omitted because they enroll few enrollees and because we omit Fallon in our simulation results in Section 6

5.3 Model Validation and Analysis.

Comparison to Reduced Form In order to test the validity of our demand and cost model estimates, we compare our model to the actual data. Figure 7 below shows that our estimated demand model is capable of reproducing the extreme price sensitivity evident in our reduced form results, as well as the large response of average costs to premium changes.

To conduct this comparison, we start from our demand estimation sample at the choice instance by plan level, assigning predicted choice probabilities and plan-specific costs to each observation. We then extend each choice instance up to the month of the subsequent choice instance or when the enrollee exits the data. Next, we collapse the resulting data set to the plan-region-income group-month level, weighting each observation by either the observed plan choice $y_{ij} \in 0, 1$ (observed share and average cost) or the predicted share s_{ij} derived from the demand model (predicted share and average cost). Finally, we run event studies following the reduced form analysis, described previously in Section 4.2, for the observed vs. predicted shares and costs.

Figure 7. Comparing Model-Predicted Shares and Costs with Reduced Form Results

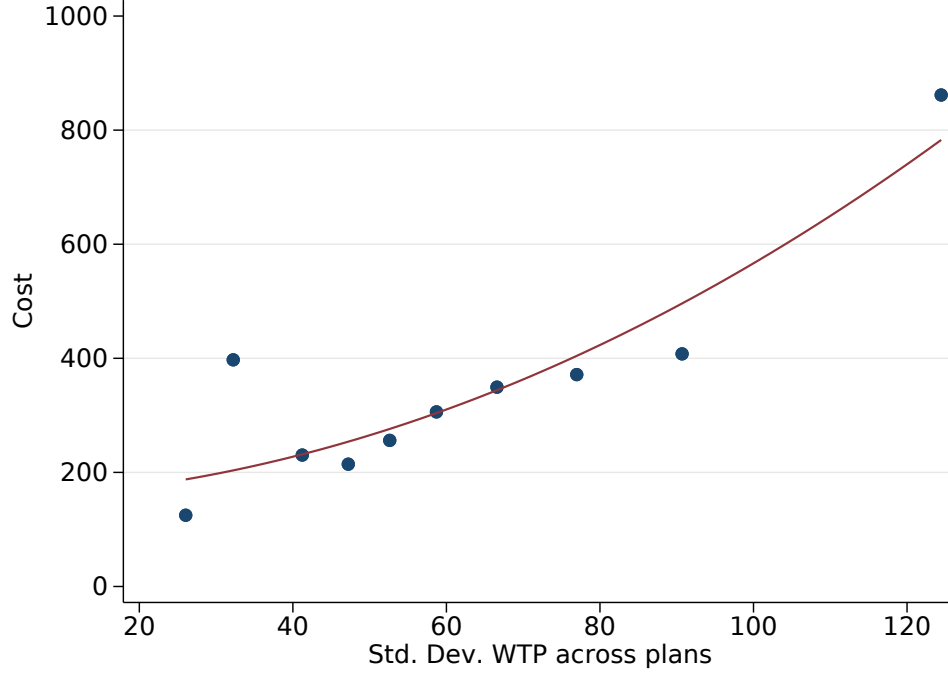


Notes: Figure shows how shares and average costs respond to increases and decreases in post-subsidy premiums, comparing actual shares and average costs for each spell with predictions using the demand estimates from Section 5. Panel (a) shows results for shares; Panel (b) shows results for average costs over each enrollment spell.

Analysis of heterogeneity generating adverse selection The theory in Section 2 argued that adverse selection would be generated from a correlation between enrollee value for plan differences (captured in the parameter β_i there) and enrollee risk or cost. We use our model to test this relationship empirically. Specifically, we use the estimated demand model in (21) to calculate “predictable” WTP for each plan (ignoring ε_{ijt}) as $V_{ijt} = \alpha(Z_{it})^{-1} \cdot [f(X_{jt}, Z_{it}; \beta) + \xi_j(W_{it})]$. We then calculate the standard deviation of V_{ijt} across plans j for a given (i, t) consumer-year choice instance, which we can compare with observed consumer costs (adjusted for plan effects).

Figure 8 shows a binned scatter plot of this relationship. Individual costs are strongly positively correlated with willingness to pay, confirming that adverse selection is a strong feature of our market. In particular, individuals in the top decile of WTP-variance across plans have a monthly average cost of about \$850, compared to a monthly average cost of about \$150 for the lowest decile of wTP-variance.

Figure 8. Enrollee Cost vs. their Dispersion of WTP across Plan Options (both in \$/month)



Notes: Figure shows a binned scatter plot of individual-level costs (for a given year t , and adjusted for plan effects) vs. deciles of the standard deviation of willingness to pay (V_{ijt}) across plans (j) in the individual's choice set. WTP is derived from the structural demand estimates of (21), as described in the text. Sample includes all new enrollees in the market from 2007-2014.

6 Counterfactual Simulations

In this Section, we use our demand estimates from Section 5 to calibrate a model of insurer entry and pricing under a variety of policy settings. As in Section 2, we model the insurance market in two stages. In the first stage, single-plan insurers choose whether to enter the market. In the second stage, each insurer compete on prices in Nash-Bertrand equilibrium. Conditional on a set of entrants and prices, consumers choose plans and incur health care costs, which determines insurer profits.

Our baseline simulations are conducted on new enrollees in the market (to avoid dynamic considerations with continuing enrollees) based on demand parameters for a single year (2011), and they do not impose risk adjustment.²⁷ We also conduct simulations under perfect risk adjustment, various degrees of partial risk adjustment, with different levels of fixed costs, and with inertial current enrollees included (but without modeling pricing dynamics). For specification, we assess the impact of imposing price floors on average insurer premiums, total insurer profits, and consumer

²⁷To increase computational speed, simulations in this draft are based on a random 10% sample of $N = 5,155$ enrollees in 2011.

welfare. We assume that the four statewide CommCare insurers are potential entrants in all simulations. For simplicity, we exclude one insurer (Fallon) that only operates in a handful of regions. Lastly, as in the estimation, we assume that all enrollees must choose a plan—i.e., we do not model substitution to the outside option of uninsurance. While uninsurance is quite relevant in the ACA today, price-linked subsidies ensure that *enrollee premiums* for the cheapest option(s) are fixed regardless of the prices insurers set. This minimizes the degree that insurer-set prices lead to substitution and adverse selection on the extensive margin (Geruso et al., 2019).

6.1 Equilibrium Definition and Simulation Assumptions

In our model, insurers first decide whether to enter the market. Second, they set Nash-Bertrand prices to maximize profits. We define a “valid pricing equilibrium” as a combination of entrants and a corresponding vector of pre-subsidy premiums that satisfies the following conditions corresponding to each of these stages (solving backwards).

In Stage 2, conditional on a set of entrants E , insurers set Nash-Bertrand equilibrium prices (subject to any policies restricting prices, e.g., price floors). As discussed below, we ensure that each firm’s price P_j is a *global* optimum best-response to competitors’ prices, P_{-j} . In Stage 1, insurers decide whether to enter, fully anticipating outcomes in later stages. An equilibrium is defined as a set of entrants and prices where:

1. Stage 2 has a pricing equilibrium where all firms make positive profits net of any fixed costs.
2. No non-entering firm can unilaterally enter and earn positive profits in the Stage 2 Nash equilibrium that results when said firm enters.²⁸

Because we have assumed that all enrollees must choose a plan, monopoly firms face no constraint on their prices, other than constraints imposed by the regulator. In all simulations, we impose a price ceiling of \$475 and assume that all monopoly firms price at the ceiling, due to the lack of an extensive margin response to price from consumers.²⁹

In general, there may multiple valid equilibria corresponding to different combinations of firms that could profitably enter. In these cases, we report all valid equilibria.³⁰ Often, these correspond to different combinations of the same *number* of entering firms. In graphs, we often summarize results by grouping cases these cases together and reporting the range of outcomes. See Appendix D for additional details on the method of searching for equilibrium.

²⁸Note that this allows other insurers to respond to the new entrant by adjusting prices. This is both realistic given the structure of regulated insurance markets (where prices are rebid annually after observing participants) and standard in two-stage entry models in IO. This also embeds a notion of requiring entry to be a “safe” best response, as in the equilibrium notion of Riley (1979).

²⁹We base this roughly on an expected markup that would result if we use the extensive margin elasticity estimated for CommCare by Finkelstein et al. (2019), which is 25% per \$40 monthly premium increase, or 0.00625. This implies a Lerner markup of $1/0.00625 = \$160$ over average costs (about \$375 per month), or \$535. We reduce this down to \$475 to account for adverse selection on the extensive margin.

³⁰We find that multiplicity arises only in the set of entrants; conditional on entrants, we do not find cases where multiple price vectors satisfy Nash-Bertrand equilibrium conditions.

We make several further assumptions in our simulations. For the enrollee population, we include only CommCare enrollees with incomes 100-300% of poverty, which is the price-paying population for whom we can estimate demand responses to premiums. This also better matches our simulations to the ACA exchanges' population, since enrollees with incomes below 138% of poverty are covered by Medicaid in most states. Also following ACA policy, we set subsidies as a flat amount for all plans, which preserves pre-subsidy price differences. We do not model CommCare's policy of full or incremental subsidies that narrow price differences (though we expect to do so in a future draft). Because there is no extensive margin participation decision, the subsidy amount is arbitrary; we set it to \$350 per month (based on average subsidies in CommCare). For each outcome, we calculate "consumer welfare" as enrollee surplus (which accounts for both consumer premiums and plan utility, using the standard inclusive value formula) minus the government's subsidy spending. Because subsidies are fixed, consumer welfare moves one-for-one with enrollee surplus.

Finally, we specify the following policy for risk adjustment. In our baseline simulations, we have no risk adjustment. We then simulate various levels of risk adjustment strength by calculating risk scores, φ_i , for each enrollee and having insurers receive $\varphi_i * P_j$ for covering person i . We set φ_i as a scaled function of enrollee costs relative to the mean:

$$\varphi_i = (C_{it}/\bar{C}_t)^\lambda \quad (24)$$

where \bar{C}_t is the overall mean cost at time t and $\lambda \in [0, 1]$ is a factor that scales the strength of risk adjustment from none ($\lambda = 0$) up to perfect ($\lambda = 1$, which implies that risk scores perfectly align with costs). [Curto et al. \(2021\)](#) shows that equilibrium with this style of risk adjustment is defined by standard Nash-Bertrand conditions, but replacing raw enrollee costs, C_{ij} , with risk-adjusted costs, $C_{ij}^{RA} = C_{ij}/\varphi_i$, and raw demand D_{ij} with risk-scaled demand, $D_{ij}^{RA} = \varphi_i D_{ij}$.

6.2 Solving for Equilibria

We solve the model backwards, starting with price competition (step 2). Given a set of insurer entrants, pre-subsidy premiums are determined by a Nash-Bertrand pricing assumption, where each insurer sets its pre-subsidy premium to maximize its profit, subject to the prices of other firms. The maximization problem for insurer j is identical to that in Section 2 and its first-order condition is given by Equation (11).

Due to adverse selection, not all solutions to the FOCs will be global optima for the insurers' maximization problem. Indeed, some solutions to the FOCs are local minima for certain firms. To surmount this issue, we search for equilibrium using a grid search approach. For each possible combination of firms, we test all possible price vectors in a grid of plausible prices from \$350 to \$500 to identify candidate price vectors that are close to satisfying the equilibrium conditions outlined above—including being both a local and global profit maximizing price (see Appendix D for details). We then search within a local region of each candidate price vector to obtain exact

equilibrium prices at which all firms' FOCs are satisfied.

6.3 Baseline Simulation Results

Our simulations demonstrate how undercutting incentives can cause market instability by eliminating possible equilibria. First, we show baseline simulation results that demonstrate how adverse selection reduces the set of possible equilibria relative to the case with perfect risk adjustment. Then, we visualize the undercutting phenomenon using the best response curves of a pair of insurers. Lastly, we show how price floors can recover the existence of equilibria, allow the market to support a larger number of firms, increase consumer welfare, and even sometimes lead to *lower* equilibrium prices. We show that fixed costs exacerbate instability and demonstrate how risk adjustment increases entry at the expense of higher markups.

Panel (a) of Table 4 shows the set of possible equilibria that exist in the absence of risk adjustment, in the population of new enrollees, with no fixed costs. Our baseline simulations focus on new enrollees to abstract away from consumer inertia and dynamic considerations. Due to high price sensitivity and adverse selection, the market can only support one firm in equilibrium without risk adjustment. This result holds with or without fixed costs. The single monopoly firm prices at the ceiling, \$475. This is clearly an unfavorable outcome for the consumers in this market, and it is consistent with the reduced form results in Section 4 that indicated that the degree of adverse selection in this market was in the range under which the model in Section 2 implied that supporting multiple competing firms would be difficult.

Panel (b) shows equilibria in a market with moderate risk adjustment ($\lambda = 0.5$). In the absence of fixed costs, risk adjustment clearly matters for firm participation, allowing the market to support two firms instead of one. With low fixed costs, risk adjustment still results in equilibria with two firms, though not all two firm equilibria are supported. Not surprisingly, with high fixed costs the market only supports one firm with or without risk adjustment. In the two-firm equilibria generated by moderate risk adjustment, prices are significantly lower, reflecting duopoly instead of monopoly. Consumer welfare is significantly higher, reflecting both the lower prices and the additional choice.

Panel (c) shows that perfect risk adjustment ($\lambda = 1$) goes even further, allowing the market to support all four firms. Consumer welfare is often, but not always, higher with perfect risk adjustment. This reflects the fact that risk adjustment leads to more firm participation (welfare-increasing) but can sometimes also generate higher prices (consumer welfare-decreasing). This is due to the logic from Mahoney and Weyl (2017a) that risk adjustment limits firms' disincentive to charge high mark-ups by limiting the adverse selection caused by price increases. Clearly, our results show that this logic is not universal, however, with moderate risk adjustment leading to lower prices due to the effects on firm participation.

Table 4. Baseline Simulated Equilibria

(1)	(2)	(3)	(4)	(5)	(6)	(7)
Possible Entry Combinations	Prices	Shares	Avg Price	Welfare	Remains with: $F = 10$	$F = 30$
<i>Panel (a) No Risk Adjustment</i>						
Monopoly [BMC]	[475]	[1]	475	27	Yes	Yes
Monopoly [Celticare]	[475]	[1]	475	0	Yes	Yes
Monopoly [NHP]	[475]	[1]	475	29	Yes	Yes
Monopoly [Network]	[475]	[1]	475	30	Yes	Yes
<i>Panel (b) Moderate Risk Adjustment ($\lambda = 0.50$)</i>						
Two firms [BMC,NHP]	[410, 432]	[0.67, 0.33]	417	103	–	–
Two firms [BMC,Network]	[381, 385]	[0.51, 0.49]	383	140	Yes	–
<i>Panel (c) Perfect Risk Adjustment</i>						
Four firms [BMC,Celticare,NHP,Network]	[415, 367, 452, 408]	[0.18, 0.52, 0.05, 0.25]	390	131	Yes	–

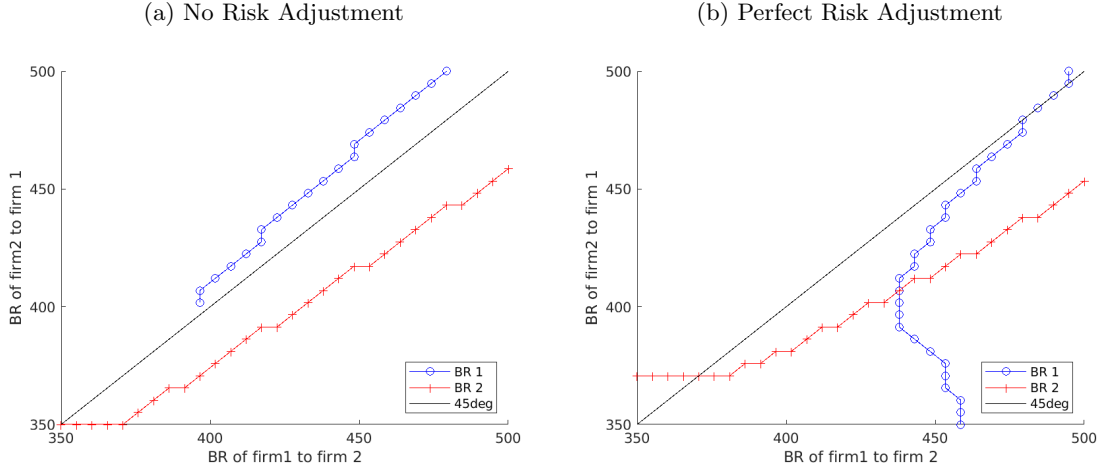
Notes: Table shows baseline simulated equilibria with no risk adjustment (Panel (a)), moderate risk adjustment (Panel (b)), and perfect risk adjustment (Panel (c)). Each panel shows the possible combinations of entrants and prices that constitute valid equilibria (Columns 1 and 2) and resulting market shares (Column 3). Column 4 reports consumer welfare per enrollee-month, relative to the equilibrium where Celticare is the monopoly insurer. Columns 5-6 report whether each equilibrium remains when firms are assumed to have fixed costs equal to F per enrollee-month, divided equally among the four potential entrants.

6.4 Mechanisms Underlying Undercutting: Best Response Curves

To illustrate how adverse selection causes certain equilibria with more firm participation to be eliminated, we plot best response curves that illustrate pricing incentives for a two-firm case. For a given price of firm 2, p_2 , firm 1's best response is defined as the value of p_1 that maximizes firm 1's profit conditional on p_2 . Plotting firm 1's best responses to all possible prices of firm 2 traces out firm 1's best response curve. Figure 9 plots the best response curve of BMC (firm 1) to Celticare (firm 2) on the Y-axis, and the best response curve of Celticare to BMC on the X-axis. Panel (a) shows the case with no risk adjustment, and Panel (b) shows the case with perfect risk adjustment. We exclude the sections of each best response curve corresponding to negative profits (i.e., conditional on firm 2's price, firm 1 is unable to make positive profits at any p_1). Points where the best response curves intersect represent valid equilibria. If the curves do not intersect, then there is no valid equilibrium for the given combination of firms.

In panel (b), we see that BMC and Celticare have a (unique) valid equilibrium under perfect risk adjustment, but not in the case with no risk adjustment (panel (a)). In both cases, if the other firm sets a high price (above about \$450), the best response of the other firm is to set a lower price. In this region, prices are strategic complements regardless of risk adjustment. That is, if one firm were to lower its price, then the best response of the competing firm would also be to lower its price. We also notice that the two firms do not price symmetrically: Celticare tends to set lower prices in response to BMC. This is reflected in Celticare's curve being below the 45 degree line.

Figure 9. Best Response Curves for BMC and Celticare



Notes: Panel (a) shows best response curves for BMC and Celticare in the case with no risk adjustment. Panel (b) shows the case with perfect risk adjustment. In both panels, the red curve labeled BR 2 shows the optimal pre-subsidy premium of Celticare on the Y-axis, given BMC's pre-subsidy premium on the X-axis. The curve labeled BR 1 shows the optimal pre-subsidy premium of BMC on the X-axis, given Celticare's pre-subsidy premium on the Y-axis. For each firm's best response curve, we exclude points where the firm makes negative profits.

This reflects Celticare's lower costs and lower estimated plan quality.

Without risk adjustment, prices continue to be strategic complements even at low prices (in the range of market average costs, which are about \$375). Adverse selection leads to undercutting incentives that persist at low prices: this is because price reductions not only increase market share but also decrease average costs. Undercutting occurs to the point where BMC's profits become negative. Celticare, by virtue of attracting lower-cost individuals, remains profitable and could undercut BMC even at prices significantly below market average costs. Thus, undercutting incentives prevent the best response curves from intersecting, resulting in no valid equilibrium.³¹

However, under perfect risk adjustment, prices are no longer strategic complements when they are low (below about \$425). In this region, both best response curves flatten out: firms no longer respond to their competitor's price cuts by cutting their own price. This allows the best response curves to intersect, yielding a valid equilibrium. Moving from Panel (a) to Panel (b) thus illustrates how risk adjustment can reduce market instability caused by undercutting.

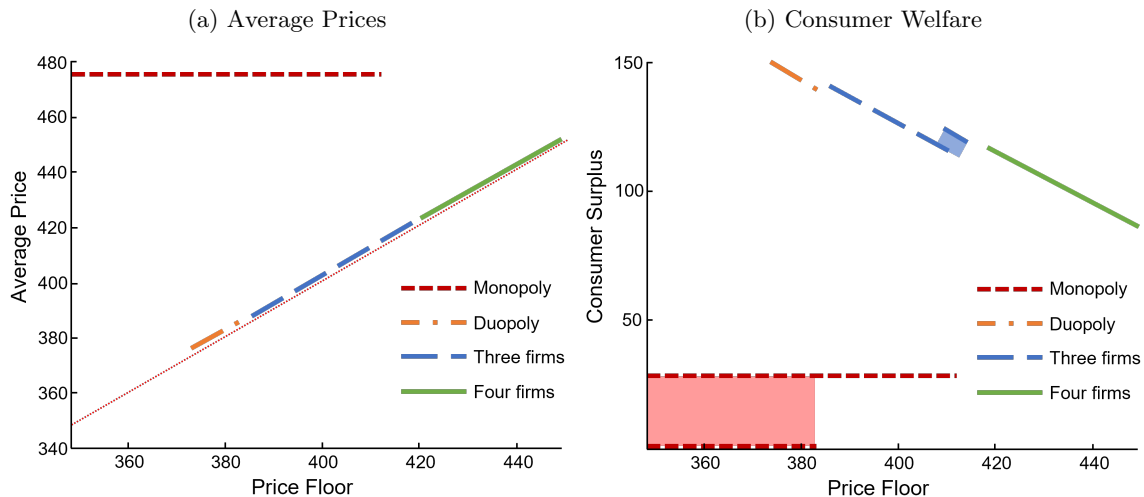
6.5 Price Floors

Our results and our model highlight that while adverse selection can sharpen price competition and reduce markups, it also can also generate pricing externalities (a firm that lowers its price

³¹Extending the curves of both BMC and Celticare into lower price ranges would eventually yield an intersection, but at that point profits for both firms would be negative, and both firms' best response would be to exit the market.

increases costs for other firms while reducing its own). We have shown that in the absence of strong risk adjustment these undercutting incentives can eliminate possible equilibria and reduce entry, potentially leading to higher prices and reduced plan choice. Risk adjustment can address these problems to some extent. In this section, we demonstrate how price floors can also be used to address undercutting and compare their effectiveness. We find that both risk adjustment and price floors can recover equilibria with more entrants and can increase consumer welfare. Importantly, we find that the optimal price floor are typically non-zero for most levels of risk adjustment.

Figure 10. Impact of Price Floors in Simulations with No Risk Adjustment



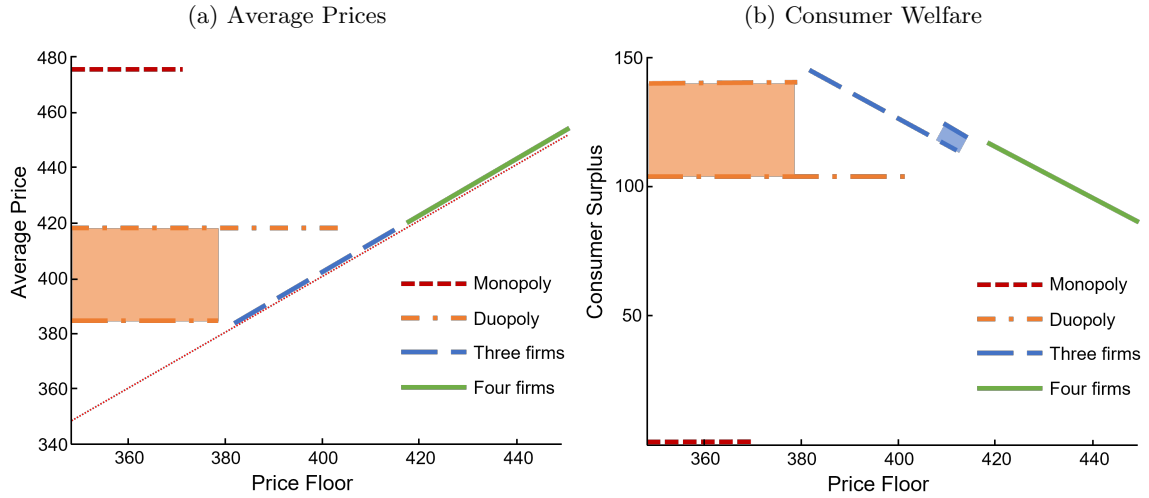
Notes: Figure shows equilibria as a function of the price floor, with no fixed costs and no risk adjustment for new enrollees. Panel (a) shows share-weighted average pre-subsidy premiums and Panel (b) shows consumer surplus per enrollee-month, normalized such that \$0 corresponds to the lowest-welfare monopoly case (with CeliCare as a monopolist).

Figure 10 plots all possible market outcomes against increasing levels of price floors in the case with no fixed costs and no risk adjustment ($\lambda = 0$). The figure enumerates all possible equilibria that satisfy the conditions introduced above. (Equilibria are grouped by number of surviving firms, with shading showing the range of possible outcomes where there are multiple specific-firm combinations.) Without a price floor, or with very low price floors, only equilibria where one firm participates survive. In these equilibria, prices are high (at the \$475 ceiling) and consumer welfare is low. But at a certain point — when they reach approximately the market average costs, or around \$375 — price floors can stabilize equilibria where 2+ firms participate. These equilibria actually involve *lower* prices than what would be charged in the absence of a price floor. Consumer welfare is also significantly higher, both due to the lower prices and due to the additional plan choice available to consumers.

Figure 11 mimics Figure 10 but for the case with moderate risk adjustment ($\lambda = 0.5$). Here, we see that in the absence of a price floor, or for very low floors, most equilibria that survive involve

two firms participating, with average prices ranging from \$385 to \$420. As the price floor increases and becomes binding for certain firms, three-firm equilibrium become feasible, with average prices that are actually *lower* than one of the possible duopoly equilibria (BMC + NHP) and similar to the other (BMC + Network). Panel (b) shows that consumer welfare is strictly higher because of the similar prices and increased plan choice available. A somewhat higher price floor of about \$420 (or 12% above market average costs) allows all four firms to enter. However, this results in higher prices and therefore somewhat lower consumer welfare than the better no-floor duopoly equilibrium (though higher welfare than the worse one).

Figure 11. Impact of Price Floors in Baseline Simulations



Notes: Figure shows equilibria from baseline simulations as a function of the price floor, with no fixed costs and moderate risk adjustment ($\lambda = 0.50$). Panel (a) shows share-weighted average pre-subsidy premiums and Panel (b) shows consumer surplus per enrollee-month, normalized such that \$0 corresponds to the lowest-welfare monopoly case (with CultiCare as a monopolist).

Table 5 shows how the consumer-welfare optimal price floor affects entry and welfare in a wide range of cases.³² Panel (a) compares results between the sample of new enrollees and all enrollees for the case with moderate risk adjustment ($= 0.5$). The average cost among all enrollees is \$368, hence the optimal price floor is 4% above average cost in this case. Given an average cost of \$375 among new enrollees, we find that the optimal price floor is just above this level.

The remainder of Table 5 shows that the optimal price floor is rarely \$0. Indeed, only in the cases of perfect risk adjustment (without fixed costs) does a price floor not improve welfare. In most cases, a price floor just above average cost maximizes welfare. This is even true for cases where there are large fixed costs, suggesting that price floors may be effective even if the source of reduced entry is fixed costs rather than adverse selection. Panel (d) further shows that price floors

³²We define the “optimal” price floor as the price floor level that maximizes consumer welfare. In cases where a given price floor permits multiple equilibria, we consider the equilibrium that gives the highest consumer welfare.

improve welfare when firms' costs are assumed to be homogeneous.

The welfare improvements from the optimal price floor vary according to market primitives and other policies. The largest welfare gains are achieved in cases with limited risk adjustment and large fixed costs. Welfare gains are positive, but smaller for cases with moderate risk adjustment and small fixed costs. Ultimately, however, price floors appear to be a useful policy in markets with extreme levels of adverse selection. In the majority of simulated policy environments, we find that these floors can actually lead to *lower* prices for consumers, as well as a wider set of plan options to choose from.

Table 5. Price Floors Recover Equilibria and Increase Welfare

	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Optimal Price Floor	Firms	Avg. Price	Profits	Welfare	Firms Without Price Floor	Welfare Gain from Price Floor
<i>Panel (a) New Enrollees vs. All Enrollees</i>							
New Enrollees (baseline)	382	Three firms [1 2 4]	382	19	144	Two firms [1 4]	5
All Enrollees	383	Four firms [1 2 3 4]	412	42	174	Three firms [1 3 4]	11
<i>Panel (b) Impact of Risk Adjustment</i>							
No Risk Adj	375	Two firms [1 4]	375	6	149	Monopoly [4]	119
$\lambda = 0.25$	375	Two firms [1 4]	375	6	149	Monopoly [4]	119
$\lambda = 0.50$ (baseline)	382	Three firms [1 2 4]	382	19	144	Two firms [1 4]	5
$\lambda = 0.75$	367	Three firms [1 2 4]	385	36	136	Three firms [1 2 4]	1
Perfect ($\lambda = 1$)	None	Four firms [1 2 3 4]	390	45	131	Four firms [1 2 3 4]	0
<i>Panel (c) Impact of Fixed Costs</i>							
F = 0 (baseline)	382	Three firms [1 2 4]	382	19	144	Two firms [1 4]	5
F = 10	386	Three firms [1 2 4]	386	23	141	Two firms [1 4]	1
F = 30	403	Three firms [1 2 4]	403	40	124	Monopoly [4]	94
<i>Panel (d) Impact of Cost Heterogeneity</i>							
With cost het (baseline)	382	Three firms [1 2 4]	382	19	144	Two firms [1 4]	5
Without cost het	386	Three firms [1 3 4]	386	12	148	Two firms [3 4]	13

Notes: Table shows market outcomes with optimal price floors (reported in Column 1), where “optimal” is defined as the price floor that yields the highest consumer welfare. In cases where a given price floor permits multiple equilibria, we report the equilibrium that gives the highest consumer welfare. Column 2 shows the set of entrants that maximizes welfare given the optimal price floor. Firms are numbered 1=BMC, 2=Celticare, 3=NHP, and 4=Network. Columns 3, 4, and 5 show the corresponding average pre-subsidy monthly premium, profits per enrollee-month, and consumer welfare per enrollee-month relative to the baseline case of where Celticare is the monopoly insurer. Column 6 reports the set of entrants in the equilibrium without price floors. As before, in cases where multiple equilibria are possible, we report the equilibrium that gives the highest consumer welfare. Column 7 reports the welfare gain from imposing price floors, relative to the equilibrium in Column 6. Unless otherwise noted, results are for no fixed costs, moderate risk adjustment ($\lambda = 0.50$), new enrollees only, and include cost heterogeneity.

7 Conclusion

Adverse selection has been shown to cause many problems in insurance markets. The prior literature has focused on two key sets of problems: price distortions and contract distortions. In this paper, we show that selection can cause a third problem that may be even more important: It can limit the number of firms that the market can support. Indeed, in the extreme case it can cause a market to become a natural monopoly. We show this via a general model of an insurance market that highlights the effects of selection in a market where firms are horizontally, rather than vertically, differentiated. We also show that the natural monopoly result is not just theoretical—it is actually the outcome predicted by our empirical estimates of the individual health insurance market in Massachusetts, in the absence of corrective policies. Fortunately, our counterfactual simulations reveal that this outcome can be reversed via risk adjustment, a common policy in health insurance markets, or by price floors.

Ultimately, these findings have important implications for health insurance markets. These types of individual markets are now highly prevalent in US social health insurance programs and around the world. Our results show just how fragile these markets are and just how much they rely on corrective policies such as risk adjustment to succeed. Our results also suggest an additional policy, price floors, could improve outcomes in many settings.

Our findings also generalize to other markets with downward-sloping average cost curves, such as markets with significant fixed costs (e.g., pharmaceuticals). We show how market competition is stymied by undercutting incentives, which endogenously determine the number of firms that can exist in a market. In both insurance markets and pharmaceuticals, firms have an additional incentive to undercut their rivals: not only does the undercutting firm acquire larger market shares, but they also move down their own average cost curve. As others have observed, undercutting can accentuate competition by reducing markups. However, we show that undercutting can lead to lower welfare overall by limiting the number of entrants and can even lead to higher prices. We show that two types of regulation can limit these problems. First, in some markets regulators are able to directly manipulate the average cost curve (e.g., with risk adjustment). However, this is not possible in cases where average cost curves slope downwards because of fixed costs. For these cases, regulators could ensure sufficient entry by setting price floors high enough to cover fixed costs for a desired number of firms.

In recent years, the individual health insurance Marketplaces created by the Affordable Care Act have struggled to achieve robust levels of competition. Indeed, in 2021 fewer than 50% of counties had more than two insurers competing in their local market. Low levels of competition have correlated with high prices. Many have suggested that political factors are responsible for this lack of participation. Our results suggest that the lack of competition may instead be a natural product of extreme levels of price sensitivity and adverse selection in these markets. Thus, counterintuitively, the best policies to improve competition in these markets may be policies that target adverse selection rather than competition policy.


References

- Akerlof, George A.** 1970. "The Market for "Lemons": Quality Uncertainty and the Market Mechanism." *Quarterly Journal of Economics*, 84(3): 488–500.
- Azevedo, Eduardo M., and Daniel Gottlieb.** 2017. "Perfect Competition in Markets With Adverse Selection." *Econometrica*, 85(1): 67–105.
- Berry, Steven, and Peter Reiss.** 2007. "Empirical models of entry and market structure." *Handbook of industrial organization*, 3: 1845–1886.
- Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney.** 2012. "Pricing and Welfare in Health Plan Choice." *American Economic Review*, 102(7): 3214–48.
- Capps, Cory, David Dranove, and Mark Satterthwaite.** 2003. "Competition and market power in option demand markets." *RAND Journal of Economics*, 737–763.
- Chiappori, Pierre-André, and Bernard Salanié.** 2000. "Testing for asymmetric information in insurance markets." *Journal of political Economy*, 108(1): 56–78.
- Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya.** 2021. "Can health insurance competition work? evidence from medicare advantage." *Journal of Political Economy*, 129(2): 570–606.
- Cutler, David M, Amy Finkelstein, and Kathleen McGarry.** 2008. "Preference heterogeneity and insurance markets: Explaining a puzzle of insurance." *American Economic Review*, 98(2): 157–62.
- Cutler, David M, and Sarah J Reber.** 1998. "Paying for health insurance: the trade-off between competition and adverse selection." *The Quarterly Journal of Economics*, 113(2): 433–466.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen.** 2010*a*. "Estimating Welfare in Insurance Markets Using Variation in Prices." *The Quarterly Journal of Economics*, 125(3): 877–921.
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen.** 2010*b*. "Estimating Welfare in Insurance Markets Using Variation in Prices." *Quarterly Journal of Economics*, 125(3): 877–921.
- Einav, Liran, Amy Finkelstein, and Neale Mahoney.** 2021. "The IO of selection markets." In *Handbook of Industrial Organization*. Vol. 5, 389–426. Elsevier.
- Einav, Liran, Amy Finkelstein, and Pietro Tebaldi.** 2018. "Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies."
- Einav, Liran, Amy Finkelstein, Stephen P Ryan, Paul Schrimpf, and Mark R Cullen.** 2013. "Selection on moral hazard in health insurance." *The American economic review*, 103(1): 178–219.
- Enthoven, Alain C.** 1993. "The history and principles of managed competition." *Health affairs*, 12(suppl 1): 24–48.
- Geruso, Michael, and Timothy Layton.** 2017. "Selection in Health Insurance Markets and Its Policy Remedies." *Journal of Economic Perspectives*, 31(4): 23–50.


- Geruso, Michael, and Timothy Layton.** 2018. "Upcoding: Evidence from Medicare on Squishy Risk Adjustment." *Journal of Political Economy*, forthcoming.
- Geruso, Michael, Timothy J. Layton, and Jacob Wallace.** 2020. "Are All Managed Care Plans Created Equal? Evidence from Random Plan Assignment in New York Medicaid Managed Care." National Bureau of Economic Research Working Paper 27762.
- Geruso, Michael, Timothy J Layton, Grace McCormack, and Mark Shepard.** 2019. "The two margin problem in insurance markets." *The Review of Economics and Statistics*, 1–46.
- Geruso, Michael, Timothy Layton, Grace McCormack, and Mark Shepard.** 2021. "The Two Margin Problem in Insurance Markets." *Review of Economics and Statistics*, Forthcoming.
- Gruber, Jonathan.** 2017. "Delivering Public Health Insurance Through Private Plan Choice in the United States." *Journal of Economic Perspectives*, 31(4): 3–22.
- Hendren, Nathaniel.** 2013. "Private Information and Insurance Rejections." *Econometrica*, 81(5): 1713–1762.
- Jaffe, Sonia, and Mark Shepard.** 2020. "Price-linked subsidies and imperfect competition in health insurance." *American Economic Journal: Economic Policy*, 12(3): 279–311.
- Mahoney, Neale, and E Glen Weyl.** 2017*a*. "Imperfect competition in selection markets." *Review of Economics and Statistics*, 99(4): 637–651.
- Mahoney, Neale, and E. Glen Weyl.** 2017*b*. "Imperfect Competition in Selection Markets." *The Review of Economics and Statistics*, 99(4): 637–651.
- Marone, Victoria, and Adrienne Sabety.** 2022. "When Should There Be Vertical Choice in Health Insurance Markets." *American Economic Review*, 112(1): 304–42.
- Rothschild, Michael, and Joseph Stiglitz.** 1976. "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information." *Quarterly Journal of Economics*, 90(4): 629–649.
- Salop, Steven C.** 1979. "Monopolistic competition with outside goods." *The Bell Journal of Economics*, 141–156.
- Shepard, Mark.** 2022. "Hospital network competition and adverse selection: evidence from the Massachusetts health insurance exchange." *American Economic Review*, 112(2): 578–615.
- Starc, Amanda.** 2014. "Insurer pricing and consumer welfare: Evidence from Medigap." *RAND Journal of Economics*, 45(1): 198–220.

Appendix Figure A1. Case Study: CeltiCare and Network Health in 2011-2012


(a) Plan Bids and Relative Premiums



(b) Market Shares



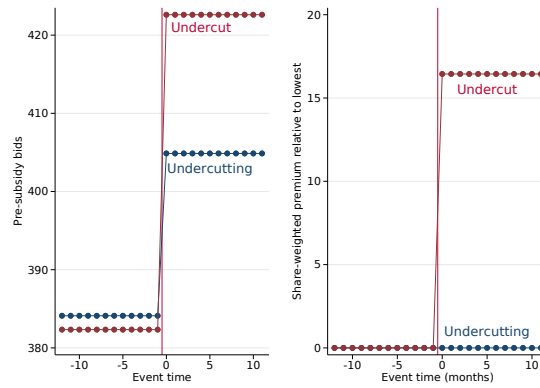
(c) Average 6-Month Cost




Note: Figure shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over time. In all of the figures x-axis is time in months or bimonths relative to month 1 2012.

Appendix Figure A2. Case Study: CeltiCare and Network Health in 2010-2011


(a) Plan Bids and Relative Premiums



(b) Market Shares



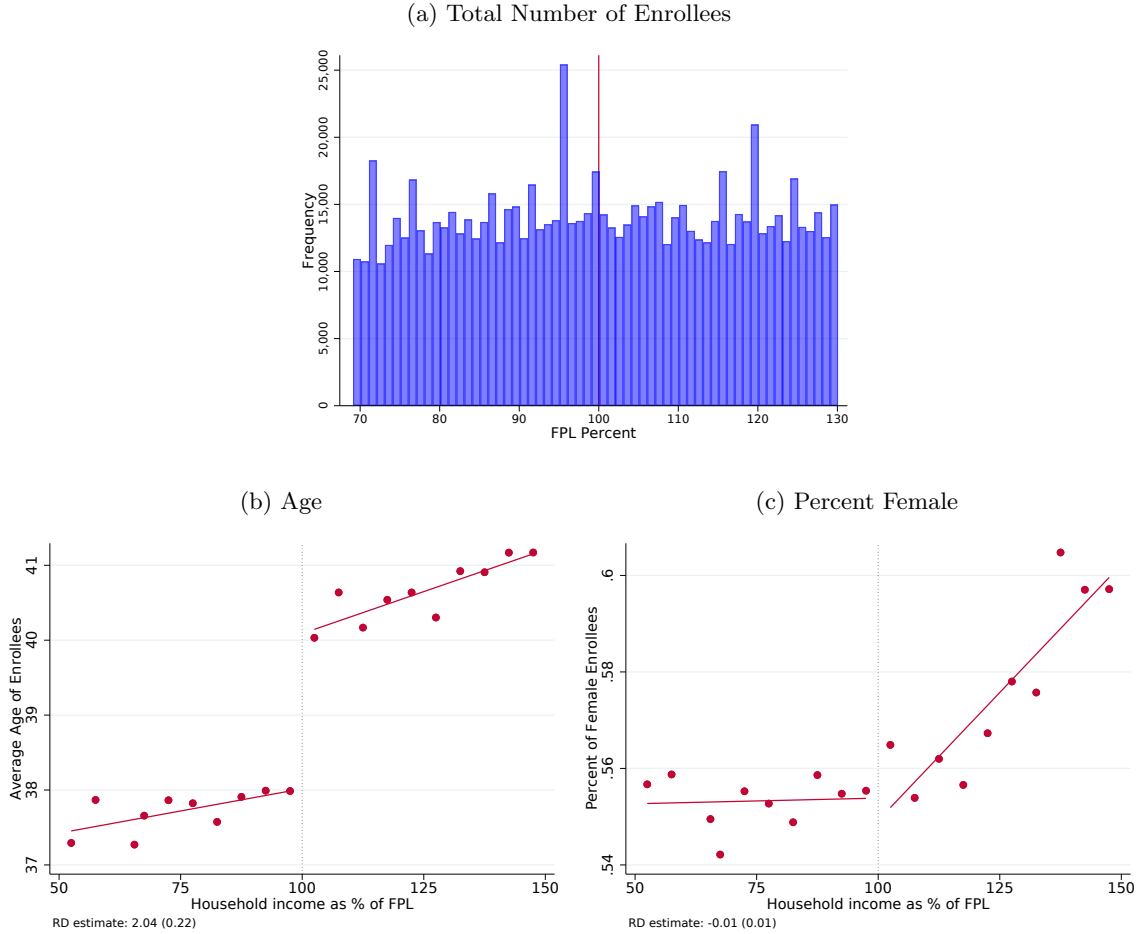
(c) Average 6-Month Cost



Note: Figure shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over time. In all of the figures x-axis is time in months or bimonths relative to month 1 of 2011.

A Regression Discontinuity: Further Analysis

Appendix Figure A3. RD's



B Regression Discontinuity Design

As discussed in Section 3, consumers with incomes just below 100% of FPL face no variation in premiums across plans (all plans are free) while consumers with incomes just above 100% of FPL face modest variation in premiums. We leverage this discontinuous change in premiums to estimate price elasticities and the effect of price on average cost using a regression discontinuity (RD) design.

Unfortunately, balance tests reveal that our setting is not ideal for an RD design. Specifically, a McCrary density test reveals a discontinuity in the density of enrollees around the 100% FPL cutoff. Panel (a) of Appendix Figure A3 also shows a possible decrease in the total number of enrollees on either side of the discontinuity. Further, as shown in Panel (b) of Appendix Figure A3, we find imbalance in a key , age, on either side of the discontinuity. Panel (c) shows that we do

not find such a discontinuity for gender, but there is a clear shift in slopes for this characteristic at 100% of FPL. These results combine to suggest that the key assumption for a valid RD design — that individuals are as good as randomly assigned to one side of the discontinuity versus the other — is violated in this setting.

Because of these violations of the key identifying assumption, we interpret all RD results as descriptive — revealing patterns that are suggestive of strong price sensitivity but not cleanly identifying the key parameter of the firm-specific price semi-elasticity of demand. We also do not present RD results related to selection outcomes, as we believe these outcomes to be more vulnerable to the biases introduced by compositional differences on either side of the discontinuity.

With those caveats, we implement the RD design both graphically and via a local linear regression. In both cases, we restrict to individuals with incomes between 50% and 150% of FPL (the income level at which premiums again change). Specifically, we use the following regression specification:

$$Y_i = \beta_0 + \beta_1 \text{Above100}_i + \beta_2 \text{FPL}_i + \beta_3 \text{FPL}_i \times \text{Above100}_i + \epsilon_i \quad (25)$$

β_2 and β_3 control for linear trends to the left and the right of the cutoff, respectively. β_1 estimates the change in the outcome at the cutoff, and represents the causal effect of exposure to price variation on the outcome Y_i .

We estimate the effects of prices on the market share of the lowest-priced plan (i.e., $Y_i = 1$ if i is enrolled in the lowest-priced plan and 0 otherwise) and the combined market share of all other plans. We then divide β_1 by the enrollee-weighted average premium of the plans other than the lowest-priced plan to recover the firm-specific price semi-elasticity of demand.

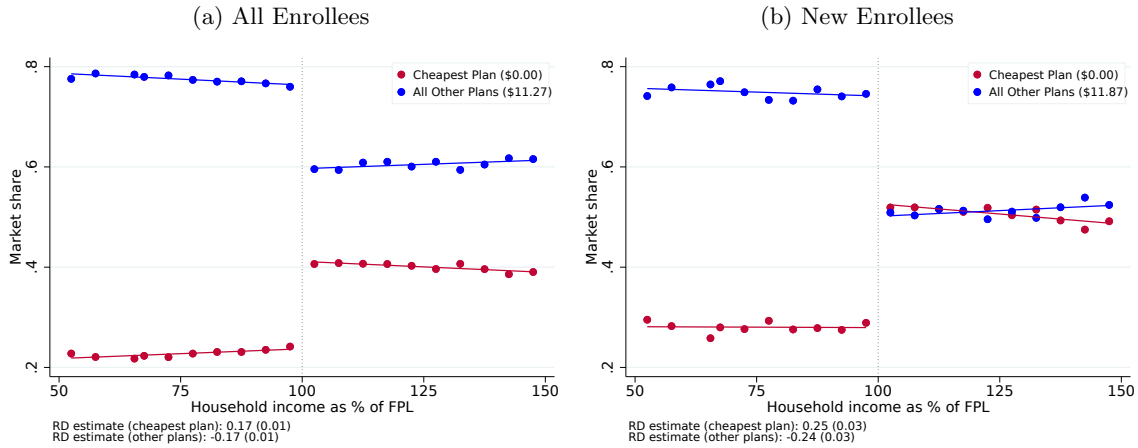
In Panel (a) of Figure A4, we present the RD plot for the market share of the lowest-priced plan (red) and the market share of the combination of all other plans (blue). The shift in market share at 100% of FPL is striking. The market share of the lowest-priced plan increases from around 20 percentage points to around 40 percentage points, a relative increase of 100%. The increase in the differential price that produces this shift in market share is only \$11.27 per month, or around \$135 per year. As noted above, we cannot fully attribute this shift in market share to the effect of the price because the enrollees just above 100% of FPL are slightly older than the enrollees just below 100% of FPL. However, we would probably expect older enrollees to be *less* likely to choose the cheapest plan (due to stronger preferences for the more generous, higher-priced plans), not more, suggesting that our estimate of the shift in market share may be an *under*-estimate of the true shift caused by the change in price.

In Panel (b) of Figure A4, we present the same RD plot but only for new enrollees (dropping incumbent enrollees). For these non-inertial consumers, the effects of prices are even more striking. When all plans are free, only around 25% of enrollees choose the cheapest plan. But when there is an average price gap of \$11.87 between the cheapest plan and the other plans, the cheapest plan enrolls a full 50% of the market.

In Table A1, we present the RD coefficient estimates. Columns 1 and 2 present estimates corresponding to Panels (a) and (b) of Figure A4. Column 3 presents the coefficient estimate for all enrollees, focusing only on market-years where the premium gap was less than \$10. This coefficient estimate is similar to the overall estimate from Column 1, revealing that price sensitivity is strong, even when the price gap is small, suggesting the presence of “choose-the-cheapest-plan” consumers. Column 4 focuses only on 2009-2010. In 2011, a discount insurer, Celticare, entered and was the cheapest plan in all markets. To ensure that we are estimating a general price-elasticity rather than a Celticare-specific price elasticity, we restrict to the years prior to Celticare entry. We find that shifts in market share are similar when restricting to these years.

These results combine to provide suggestive evidence of strong price sensitivity in this market. The implied firm-specific own-price elasticity of demand ranges from -0.015 across all enrollees to -0.020 for new enrollees. These price elasticities are high and raise concerns about under-cutting incentives. We now turn to the diff-in-diff analysis to provide precise estimates of price elasticities as well as estimates of the slope of the firm-specific average cost curve.

Appendix Figure A4. RD Estimates - Market Shares



Note: Figure shows regression discontinuity plots for the market share of the lowest-priced plan (red) and the combination of all other plans (blue). Panel (a) shows results for all enrollees, Panel(b) shows results for new enrollees.

Appendix Table A1. RD Main Estimates

	All Enrollees	New Enrollees	All Enrollees, < \$10 change	All Enrollees, 2009-2010 only
Market share	-0.17*** (0.01)	-0.24*** (0.03)	-0.15*** (0.02)	-0.17*** (0.02)
Premium Change	11.27	11.87	6.55	11.04
Elasticity	-.0354*** (.0025)	-.042*** (.0051)	-.0659*** (.008)	-.0312*** (.0033)

Standard errors in parentheses

* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Notes: Table shows estimates related to the RD specification. Each row corresponds to a different outcome variable and each column corresponds to a different sample. Row 1 contains coefficient estimates and standard errors from the RD with market shares as the dependent variable. Row two shows the enrollee-weighted average (using the enrollees in the 105-125 FPL bins) of the premiums of all but the 0 premium plan (for the IIA group). Row three shows semi-elasticities, which are computed using arc-elasticities i.e. dividing market share coefficients by the midpoint of the values to the left and right of the RD and then dividing by the premium change.

C Demand Estimation

C.1 Cost residual model

We model individual health costs using the set of covariates in the demand model, and use the residuals from this model as an additional input in our demand model. This allows us to capture variation in demand that may be correlated with unobserved factors that are correlated with health costs.

We start with choice-instance-level observations of individuals' monthly health costs, averaged over each choice instance. We then adjust for plan-specific cost heterogeneity using the estimates from Equation 23 of $\eta_j(Z_i)$, the plan-specific cost multiplier. For each individual i , the adjusted plan cost is equal to $c_{i,adj} = c_i/\eta_j(Z_i)$ for the chosen plan j and represents the individual's cost if they had been in the average plan.

Then, we estimate a Poisson regression of $c_{i,adj}$ on the other covariates in the demand model (specifically: indicator variables for income group, age-sex groups, immigrant status, HCC risk score quantiles, diagnoses for chronic disease, diagnoses for cancer, region, and year. We weight this regression by the number of monthly observations in each choice instance. We compute the "cost residual" for each individual as the residual from this regression (computed as the ratio of $c_{adj}/\hat{c}_{i,adj}$ where $\hat{c}_{i,adj}$ is the predicted cost for individual i . In the demand model, we allow individuals' price sensitivity to vary by deciles of these cost residuals.

D Counterfactual Simulations: Detailed Methods

D.1 Solving for Equilibria

For a given set of insurer entrants, we adopt a step-by-step approach to solve for price equilibria. For a price vector to permit a valid equilibrium, it must be a Nash equilibrium (no firm can deviate to another price and achieve higher profits) and all firms must have positive profits net of their fixed costs. Firms' prices may be restricted by price floors. For some combinations of entrants, there may be a price vector that satisfies the Nash equilibrium conditions but does not satisfy the positive profits condition. No valid equilibrium exists in those cases.

To solve for the insurer combinations that deliver valid equilibria (as defined above in Section 6.1), we search for valid price equilibria for all possible combinations of entrants. We adopt the following grid search approach.


For each possible combination of plans, we evaluate first order conditions and profits for a grid of all possible price vectors, where each price takes one of 30 evenly spaced values from \$350 to \$500. For 4 firms, there are $30^4 = 810,000$ possible price vectors. We then identify candidate price vectors that satisfy the following four conditions. First, all firms' prices fall between the price floor and \$500. Second, all firms make positive profits. Third, the FOCs are satisfied within a pre-defined tolerance³³ Fourth, no firm can deviate to a higher or lower price (within the bounds of the price floor and \$500) and make higher profits. To account for imprecision from the grid, we allow price vectors in a ± 1 grid point box around each of these candidate vectors. For each of the resulting candidate price vectors, we solve for the exact equilibrium prices using the `fmincon` function in Matlab to solve the system of FOCs within a box of ± 2 grid points around the candidate price vector. Finally, we check whether this exact price vector continues to satisfy the four conditions above. This procedure delivers up to one candidate equilibrium price vector for each possible combination of firms.³⁴

Among these possible equilibria, we exclude combinations where adding an additional firm would result in a valid equilibrium. In some cases, this procedure predicts a unique equilibrium (a key example is where all four firms comprise an equilibrium). However, there are cases where multiple equilibria may exist. For example, if the four-firm equilibrium is not possible, multiple different combinations of three firms may constitute equilibria. We do not take a stand on equilibrium selection in this context, but we report all possible plan combinations when such cases arise. The following section reports market outcomes for the equilibria described above.

³³The tolerance we use is $N/J * .20$, where N is the total number of individuals and J is the total number of plans. We also consider corner solutions where a firm's first order condition is negative at a grid point near the price floor (the firm would like to price lower but cannot) or positive at the \$500 grid point (the firm would like to price higher but cannot).

³⁴Although we do not formally prove uniqueness of the equilibria we identify using this procedure, in practice, conditional on a given set of entrants, we never observe cases where multiple different price vectors satisfy all of the conditions and are thus equilibria. Any multiplicity of equilibria occur when multiple different sets of entering firms satisfy the conditions.

Appendix Figure A5. CommCare Plans Pre-Subsidy Prices




Note: The graphs show average pre-subsidy insurer prices for each insurer's plan in the CommCare market, by fiscal year. The five plans are shown in different colors and labeled. Values shown are averages for the plan's actual enrollees; underlying premiums and (in some years) prices vary by income group and region. There are no data points for 2008 because prices were not re-bid that year but instead mechanically carried over from 2007.

Appendix Figure A6. Premium Variation Example: Network Health (Boston region), 2010-13



Note: The graphs shows the example of Network Health’s (post-subsidy) enrollee premiums by income group over the 2010-2013 CommCare years. “FPL” refers to the federal poverty level. Pre-subsidy prices (and enrollee premiums) vary at the regional level in 2010, and the graph shows premiums specifically for the Boston region. Both are constant statewide in 2011-2013. Panel A shows the level of the premium for Network Health in dollars per month. Panel B shows the plan’s “relative” premium, equal to the difference between its premium and the premium of the cheapest plan. The graph shows that different subsidies by income group translate a single pre-subsidy price into variation across income groups in the plan’s post-subsidy relative premium

Appendix Figure A7. Hospital Coverage in Massachusetts Exchange Plans



Note: The graph shows the shares of Massachusetts hospitals covered by each CommCare plan, where shares are weighted by hospital bed size in 2011. Fallon's hospital coverage share is much lower than other plans largely because it mainly operates in central Massachusetts and therefore does not have a statewide network.