Adverse Selection and (un)Natural Monopoly in Insurance Markets *

Edward Kong
Harvard University

Timothy J. Layton
Harvard and NBER

Mark Shepard
Harvard and NBER

July 1, 2023

Abstract

Adverse selection is a classic market failure known to limit or "unravel" trade in insurance markets and many other settings. We show that even when subsidies or mandates ensure trade, adverse selection also tends to unravel *competition* among differentiated firms — leading to fewer surviving competitors and in the extreme, what we call "un-natural monopoly." Like fixed costs in standard natural monopoly, adverse selection creates a wedge between marginal and average costs, as firms compete aggressively on price to attract (or "cherry-pick") price-sensitive low-risk consumers. This wedge must be covered by sufficiently large markups, which limits how many firms can profitably survive. Unlike fixed costs, the underlying problem is a coordination failure that can be addressed via (careful) price regulation — a policy often used in practice but which existing models have difficulty motivating. We show the empirical relevance of strong adverse selection on price using subsidy-driven price variation and a structural model of competition in Massachusetts' health insurance exchange. Our analysis suggests a new rationale for policies mitigating adverse selection: Without them, the market devolves to monopoly; with them, the market can sustain robust insurer competition.

*Contact: Kong (edk571@g.harvard.edu), Layton (Layton@hcp.med.harvard.edu), Shepard (Mark_Shepard@hks.harvard.edu). We are grateful to Chris Avery, Peter Blair, Liran Einav, Amy Finkelstein, Nathan Hendren, Kate Ho, Lee Lockwood, Neale Mahoney, Joe Newhouse, Jim Poterba, Amanda Starc, Pietro Tebaldi, Nick Tillipman, Richard Zeckhauser for helpful comments and discussions. We especially thank Richard for suggesting the term "un-natural monopoly" used in our title. We also thank seminar participants at MIT, Congressional Budget Office, 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, the Annual Meeting of the American Society of Health Economists, and the NBER Insurance meetings for comments. Ilana Salant provided superb research assistance.

1 Introduction

Adverse selection is a classic market failure in economics. Dating back to seminal insights by Akerlof (1970) and Rothschild and Stiglitz (1976), a growing body of theory and empirics shows how welfare-improving trade can break down when either buyers or sellers have private information relevant to the other party's payoffs. The insight that adverse selection limits or "unravels" trade (Hendren, 2013) is relevant in a wide variety of settings, from used cars (Akerlof, 1970) to labor markets (Greenwald, 1986; Stantcheva, 2014) to corporate finance (Michaely and Shaw, 1994). But perhaps the most prominent application of adverse selection is insurance, and notably health insurance. While there is debate about the ideal role of public vs. private insurance, essentially all high-income nations have meaningful health insurance markets, whether for basic coverage (as in the Netherlands, Switzerland, Germany, Chile, Israel, and the U.S.) or for supplementary insurance (as in Canada, France, and the U.K.). Addressing adverse selection —- and more generally, "making insurance markets work" —- is therefore a critical policy objective.

While policymakers have developed many tools to address selection, most policies —- like subsidies, mandates, and quality regulation —- focus on ensuring *trade* in decent quality insurance, since these are the key problems in standard models of adverse selection. However, as trade has increased (e.g., as health uninsurance declines in the U.S.), there is growing recognition of another concern: *limited competition*. By conventional antitrust standards, health insurance markets are often highly concentrated.[1] Nowhere is this more evident than in the insurance exchanges created by the Affordable Care Act (ACA) where as of 2021, one-fifth of enrollees lived in areas with just 1-2 participating insurers and 24 whole states had three or fewer competitors (McDermott et al., 2020). Given these facts, there is growing interest in understanding the industrial organization of "selection markets" and the implications for insurance market policy (Einav et al., 2021).

In this paper, we ask whether there may be a deep link between the economics of adverse selection and limited competition (or "natural monopoly") in insurance markets. We show that adverse selection, through its impact on *strategic price competition*, creates a barrier to robust firm entry – operating much like (and in tandem with) fixed costs in standard models of natural monopoly. In short, adverse selection may unravel not just trade (as in standard models) but also unravel *competition* among differentiated firms. We develop these insights using a simple model of imperfectly competitive insurance markets, which we then estimate using administrative data from a health insurance exchange. Our analysis suggests an important insight: that price (not just quality) competition may be a problem in selection markets. Addressing adverse selection, therefore, may involve policy that regulates or softens price competition, which (counterintuitively) may lower prices and increase consumer welfare by stimulating firm entry.

To make these points, we start by developing a general model of a health insurance market,

[1] For instance, over 70% of Medicare Advantage markets are "highly concentrated" by antitrust standards (HHI > 2500), with the typical market having just 2.5 competitors (Frank and McGuire, 2019). The Medicare supplemental insurance (Medigap) market is dominated by two firms with three-quarters of the overall market share (Starc, 2014).

where differentiated firms first decide whether to offer a fixed insurance contract and then select a price for that contract. In a typical market without selection, consumer price sensitivity constrains markups—when a firm raises its price by \$1, it earns more profits on each inframarginal consumer but it also loses marginal consumers. Fixed costs and consumer price sensitivity limit the number of participating firms—as more firms enter, markups are lower and firms spread fixed costs across fewer consumers. Eventually, markups available to potential entrants fall below fixed costs, such that additional firms prefer not to enter (Salop, 1979). Price sensitivity and fixed costs thus impose a cap on the number of firms a market can support.

We show that when there is adverse selection on price, i.e. when low-risk (healthy) consumers are more price sensitive than high-risk (sick) consumers, markups are even more constrained (a result first pointed out by Starc (2014) and Mahoney and Weyl (2017*a*)). The intuition behind this result mimics the logic behind firm markups in standard markets, but with a twist—in a market characterized by adverse selection on price, when a firm raises its price by \$1, it doesn't just lose marginal consumers, it loses its *lowest-risk* consumers, driving up its average cost and further reducing profit margins. With adverse selection, firms thus set lower markups than in settings without selection. With these lower markups, fixed costs exceed markups at even lower levels of entry, resulting in even more limited competition than in the absence of selection.

Importantly, we show that the previous result (that adverse selection lowers prices) relies critically on the assumption that market structure is fixed. When entry is endogenous, selection lowers prices only up to the point at which it decreases the number of firms the market can support. As the number of firms decreases, prices *increase* instead of decrease due to lower consumer price sensitivity stemming from the more limited number of choices. In extreme conditions, markets may not be able to support more than one firm, resulting in high (monopoly) prices and limited consumer choice. We also show that in our model, prices are not constrained by the *threat* of entry, as entry would result in a set of equilibrium prices where the entering firm loses money.

Finally, we show that the effects of selection on firm participation are not dependent on the existence of fixed costs. Indeed, adverse selection can unravel market participation in the complete absence of fixed costs. The logic behind this result stems from a simple observation—consider two firms offering identical (non-differentiated) health plans and competing on price. If both firms price at the average cost across the entire market, each firm faces a strong incentive to undercut its competitor, charge a slightly lower price, enroll all of its competitor's low-risk enrollees, and earn substantial profits. This incentive persists far below the average cost in the market, resulting in either a pure strategy Nash pricing equilibrium where one or both firms earn negative profits or no pure strategy pricing equilibrium conditional on the given number of firms participating.[2] At least one of the firms would thus prefer not to have entered in the first stage of the entry game.

To assess whether adverse selection could actually lead to limited firm participation in the real world, we turn to data from Massachusetts' Commonwealth Care (CommCare) insurance exchange,

[2]In cases with no pure strategy pricing, there is always a mixed strategy pricing equilibrium (Nash, 1951). In these cases, at least one of the firms will lose money in expectation under the mixed strategy.

the state's pre-ACA health insurance marketplace for low-income individuals. The market consisted of 4-5 insurers, each offering a single plan with standardized cost sharing provisions but differing provider networks.

We start by presenting a few case studies where one plan undercut another, showing the consequences of these price changes for the market shares and average costs of both the under-*cutting* and the under-*cut* plans. These case studies illustrate the undercutting incentive implied by our model by showing that undercutting one's competitors can lead to both large gains in market share *and* large increases in average profit margins. We then estimate summary measures of price sensitivity and the effects of price changes on average plan costs in a difference-in-differences design based on year-to-year plan price changes, leveraging a control group of individuals below 100% of the federal poverty line (FPL), whose premiums are fully subsidized in the CommCare market. These 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 at observed prices, a price decrease of \$1 lowers a plan's average claims cost by around \$1 — a critical condition that our model suggests will lead to unraveling of firm participation in the market, even without fixed costs.

We then use the same quasi-experimental variation in prices to estimate a full structural model of demand and costs for insurance plans in the Massachusetts market. We show that our structural model fits our reduced form results quite well, with predictions from the model replicating our event studies. We use the model to simulate equilibria with and without various corrective policies. The Massachusetts market used strong corrective policies to soften price competition, aggressive risk adjustment, binding price floors, and incremental subsidies for higher-price plans. We show that these corrective policies were critical for achieving the modest levels of participation observed in this market — In the absence of corrective policies, the undercutting incentives are so strong that the only surviving equilibrium is one with a single monopoly firm. In the absence of these policies, choice would thus be limited and prices would be high.[3] The issues raised by our theoretical model thus appear to be empirically relevant in this setting.

We use the estimated structural model to perform counterfactual simulations exploring the effects of various corrective policies on market equilibrium. Specifically, we study two key corrective policies featured in the CommCare market: risk adjustment and price floors.[4] Under risk adjustment, the regulator enforces 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 weakens the downward pressure adverse selection exerts on firm mark-ups, potentially limiting gains in consumer surplus (Mahoney and Weyl, 2017*a*). The second policy is a price floor. During much of its existence, the

[3] Again, we note that in our model, monopoly prices are not constrained by the *threat* of entry, as that threat is not credible because the entering plan would lose money at the new equilibrium price.

[4] By contrast, the ACA exchanges today use only risk adjustment, not price floors.

Connector imposed both price ceilings and floors, which were implemented via rate regulation to ensure "actuarial soundness." If set correctly, a price floor could restore equilibria with multiple participating firms, by effectively allowing firms to coordinate on a price at the floor and split 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, in equilibrium only a single firm participates and charges the monopoly price. 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 (2017*a*) 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 in the absence of risk adjustment, a modest price floor just above the average cost across all consumers in the market can induce entry and lead to much lower prices. Indeed, across a range of simulations, the optimal price floor is often 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 usually delivers *lower* prices, but may sometimes result in higher prices.

We conclude that undercutting incentives caused by adverse 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. The results from our simulations suggest that strong adverse selection and price sensitivity may be at least partially responsible for the low participation in these markets. In the last section of the paper, we implement a test of this hypothesis. Specifically, we leverage the end of the reinsurance program in the ACA marketplaces in 2017. This program reimbursed insurers for enrollee costs exceeding a threshold, effectively cushioning insurers against adverse selection in this market. We use differential exposure to reinsurance across markets to test whether the removal of this program led to a decrease in participation. Consistent with our model, we show that the end of the reinsurance program led to decreased insurer participation, accounting for 21% of the overall decline in the number of insurers per county from 2014-16 to 2017-18.

Overall our results illustrate the fragility of health insurance markets. Indeed, they indicate that in some cases, without the "managed" part of the type of "managed competition" called for by Enthoven (1993), there may not be *any* competition. Our results are thus consistent with suggestions for these marketplaces to engage in "active purchasing," wherein market regulators could use price floors, ceilings, and coordination to achieve desired market outcomes (Shepard and Forsgren, 2022). They also warn against designing markets with too narrow a focus on price competition, especially in markets where consumers are highly price sensitive.

In markets with low consumer price sensitivity, the concerns we lay out in this paper may not be first order. For example, price sensitivity is fairly low in Medicare Advantage and Medicare Part D due to inertia and passive choice. However, in the ACA Marketplaces, where there are high rates of enrollee churn and a large portion of consumers are new buyers making active choices, consumer price sensitivity is likely to be very high, and concerns about adverse selection leading to limited firm participation may be very real. Indeed, our analyses of the pre-ACA Massachusetts individual market and the end of the ACA reinsurance program indicate that these effects are highly empirically relevant in these types of markets. Policymakers may thus wish to be cautious when pushing policies that strengthen price competition in these settings.

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., 2010; 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, 2017*a*). That work showed that selection can reduce price mark-ups in settings with imperfect competition, implying that policies such as risk adjustment can reduce consumer welfare. Our analysis points out that selection may reduce the number of competing firms, potentially outweighing the impacts of mark-ups conditional on participation. When market structure is endogenous, corrective policies such as risk adjustment can sustain higher participation and therefore *lower* mark-ups.

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 these policies can sometimes be beneficial. 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 propose a new policy to combat selection problems — price floors.

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

In this section, we present a simple model of firm pricing and market participation under adverse selection. Our goal is to show how adverse selection works in a similar way as fixed costs in classic theories of natural monopoly to shape the *market structure* of competing firms that can be supported in equilibrium. The key force is *strategic pricing* – specifically, pricing to undercut competitors in order to "cherry pick" low-risk consumers, who tend to be highly price-sensitive and comprise a large share of marginal consumers who are responsive to price cuts. This force makes it challenging for multiple differentiated firms to sustain prices above average costs, since each firm has an incentive to undercut and cherry pick low-risk types. Unless firms can find a way to coordinate – or regulators intervene with policies to soften price competition – competition itself tends to unravel.

Section 2.1 sets up the basic model and shows the main conceptual point. Section 2.2 analyzes an example that incorporates adverse selection into the classic Salop (1979) model of monopolistic competition.

2.1 Basic Theory

Consider an insurance market where a set of potential firms $j \in \{1, ..., J\}$ each has the ability to offer a single contract (or "plan"). Contracts 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. Consumers (i) vary in a vector of characteristics ζ_i that affect their utility for different plans, $U(X_j; \zeta_i)$, and may also include risk attributes $R_i \subset \zeta_i$ that affect their expected costs, $C_{ij} = C(X_j; R_i)$. For instance, in our empirical setting, plans differ in their networks of covered hospitals and doctors, and consumers differ in their preferences for certain providers (e.g., based on where they live and their existing relationships). This preference heterogeneity means that there is real value in having a *variety* of different plans available in the market.

Firms have full information about demand and costs and compete in a simple two-stage entry game, following a standard setup in the IO literature (Berry and Reiss, 2007):

1. **Entry:** In stage 1, firms simultaneously choose whether to participate in (or "enter") the

market, which involves incurring fixed cost $F \geq 0$. Non-entrants earn zero profits and incur no fixed cost.

2. **Competition:** In stage 2, the set of entrants, $E \subset \{1, ..., J\}$, compete on prices (P_j) to maximize profits $\pi_j(P)$ in standard Nash-Bertrand equilibrium.

Although costs C_{ij} can vary across individuals, we assume firms cannot price discriminate against high-cost consumers. Indeed, health insurers are usually forbidden to do so by law. The market, therefore, may feature cost-relevant asymmetric information.[5]

Given a set of entrants E, the (variable) profit function governing price competition in stage 2 is:

$$\pi_j(P) = [P_j - AC_j(P)] \cdot D_j(P) \tag{1}$$

where $P = \{P_j\}_{j \in E}$ is the prices of competing firms, $D_j(P)$ is firm j's demand, and $AC_j(P)$ is its average (variable) cost at these prices.[6] Average costs equal:

$$AC_j(P) = \tfrac{1}{D_j(P)} \sum_i [C_{ij} \cdot D_{ij}(P)] \tag{2}$$

where C_{ij} are consumer-specific costs and $D_{ij}(P)$ is consumer-specific demand for firm j.[7] Importantly, average costs are a function of (all firms') prices because consumers vary in their costs and prices affect *which types of consumers* select into each plan. This is a defining feature of insurance and other "selection markets" (Einav et al., 2021). The fact that average costs may vary with prices, and specifically that $\frac{\partial AC_j}{\partial P_j} \neq 0$, is the key driver of selection's impact on competition in our analysis.

We consider (subgame perfect) Nash equilibria of this game, which can be solved by backward induction. For any set of entrants E, stage-2 Bertrand-Nash prices P_E^* occur when all firms maximize profits such that $\frac{\partial \pi_j(P_E^*)}{\partial P_j} = 0$ given competitors' prices $P_{-j,E}^*$.[8] An equilibrium in the overall game is a set of entrants E^* and Nash-Bertrand prices $P_{E^*}^*$ such that: (1) all entrants $j \in E^*$ earn non-negative profits net of fixed costs $(\pi_j(P^*) - F \geq 0)$, and (2) no non-entrant $j' \notin E^*$ can unilaterally enter and earn profits (net of fixed costs) at the stage-2 Nash equilibrium prices among

[5] In some markets, insurers can vary prices on a limited set of factors (e.g., age) but not detailed health risk. Regulators may also use policy tools like risk adjustment transfers to try to offset this cost heterogeneity and limit risk selection (Geruso and Layton, 2017). We abstract from these details in our simple theory here but analyze risk adjustment in our structural model simulations. All of our math below carries through if (unadjusted) average costs are replaced with "risk-adjusted" average costs and demand with risk-scaled demand, a point noted by Curto et al. (2021).

[6] Throughout our model, demand, costs, and profits are implicitly functions of the set of entrants E. We suppress this in the notation for readability.

[7] We think of $D_{ij}(P)$ as a continuous function of prices, which lets us differentiate it in the math below. This can be motivated either by stochastic demand or by i representing a type of consumers (rather than a specific person).

[8] This is the condition for a pure-strategy Nash pricing equilibrium. As we discuss later, when adverse selection "unravels" price competition among a set of firms E, pure strategy equilibrium may not exist. In these cases, we consider mixed strategy equilibrium. Although in principle there could be multiple Nash pricing equilibria, we do not find this to arise in our empirical work.

firms $E^* \cup j'$.

This setup is a classic way to model the degree of competition – or the *market structure* – that emerges in an imperfectly (or "monopolistically") competitive market. More entry benefits consumers both via greater product variety and via lower markups, as firms compete prices down closer to marginal costs. But firms are only willing to enter if their expected variable profits (in stage 2) are sufficient to cover their fixed costs, F. When an additional entrant j would compete down markups so that stage-2 profits fall short of F, entry stops and the level of competition equilibrates. Therefore, in this classic model of monopolistic competition, fixed costs are the key limiting factor on competition. As $F \to 0$, more and more firms can enter and survive in equilibrium.

Our model's central point is that *adverse selection*, through its influence on price competition, works alongside fixed costs as a force that limits entry and competition. For any given F, fewer competitors can survive if the market is more adversely selected, and even with $F = 0$, entry may be quite limited. To see how this occurs, we proceed backward through the game, starting with price competition in stage 2, then analyzing the implications for entry in stage 1.

Price Competition with Adverse Selection

Consider how price competition proceeds among a given set of entrants E. As we noted above, a key feature of selection markets is that average costs may vary with prices, and in particular $\frac{\partial AC_j}{\partial P_j} \neq 0$. At a given set of prices P, we say that a firm faces *adverse selection* in its pricing incentives (on the margin) if when it raises its price, its average costs also rise:

$$\textbf{Adverse Selection in Pricing:} \qquad \frac{\partial AC_j(P)}{\partial P_j} > 0 \tag{3}$$

Because price and quantity move inversely, adverse selection in pricing corresponds to the familiar condition of a "downward-sloping" average cost curve in quantity (Einav and Finkelstein, 2011). By contrast, a firm faces *advantageous selection* in pricing if $\frac{\partial AC_j(P)}{\partial P_j} < 0$ and *no selection* if $\frac{\partial AC_j(P)}{\partial P_j} = 0$.

When is adverse selection in pricing likely to be relevant? In the next subsection, we present a model that micro-founds it and shows its relevance in a model with general (or "horizontal") differentiation – that is, where there are no clear "vertical" quality rankings among plans. The key idea is the following. Even though consumers differ in their preferences, high-risk (sicker) consumers are willing to pay more for a plan with a greater "match quality" for them, while low-risk (healthy) consumers are more willing to choose a less-ideal plan to save money. In other words, *low-risk consumers are more price-sensitive* in their demand. This is a natural and testable condition for which there is much evidence, both in prior work and our data.[9] It generalizes the classic notion of adverse selection in a vertical model to a setting with more general contract differentiation.[10]

[9] For empirical evidence, see e.g., Finkelstein et al. (2019), Saltzman (2017), and Tebaldi (2022).

[10] To see this, consider a vertical model where plans vary in quality Q_j, and consumer utility is $U_{ij} = \beta_i Q_j - P_j$, where $\beta_i > 0$ is a consumer's WTP for a unit of quality. Adverse selection occurs if high-risk (R_i) consumers have

Another way to think about adverse selection is that it implies a gap or "wedge" between a firm's average and marginal costs, where "marginal costs" refers to the consumers a firm attracts when it cuts its price (or loses when it raises its price). Because low-risk types are more price sensitive, they form a disproportionate share of people who switch plans in response to price changes. Defining this marginal cost formally as $MC_j(P) \equiv \frac{1}{\partial D_j/\partial P_j} \sum_i \left[C_{ij} \cdot \frac{\partial D_{ij}}{\partial P_j} \right]$, we note that:[11]

$$\frac{\partial AC_j(P)}{\partial P_j} = \underbrace{\eta_{j,P_j}}_{\text{Price sensitivity}} \times \underbrace{[AC_j(P) - MC_j(P)]}_{\text{Degree of adverse selection}} \tag{4}$$

where $\eta_{j,P_j} \equiv -\frac{\partial \log D_j}{\partial P_j} > 0$ is the firm's (own price) semi-elasticity of demand. Because $\eta_{j,P_j} > 0$ (by the law of demand), adverse selection in pricing ($\frac{\partial AC_j}{\partial P_j} > 0$) is equivalent to $AC_j(P) - MC_j(P) > 0$. Moreover, equation (4) shows that the steepness of the firm's average cost curve is a product of two factors:

1. **Price semi-elasticity of demand** (η_{j,P_j}): How *many* marginal consumers a firm can attract with a \$1 price cut, and

2. **Degree of adverse selection** ($AC_j - MC_j$): How *much cheaper* those marginal consumers are than the firm's average consumer.

These features of adverse selection – downward sloping average costs and $AC_j > MC_j$ – are familiar from textbook models of adverse selection (Einav and Finkelstein, 2011). However, in our model, the relevant curves are *firm-specific* cost curves, which reflect how consumers substitute across firms. Importantly, this cross-firm substitution usually becomes stronger with more competitors in the market: in the model, adding firms results in a larger price semi-elasticity of demand (η_{j,P_j}). Thus, cross-firm adverse selection tends to become more intense with more competitors (a point noted by Lustig (2010)).

Adverse selection in pricing has implications for price competition. To see this, note that the firm j's FOC for profit maximization – a necessary condition for standard Nash equilibrium in prices – is $\frac{\partial \pi_j}{\partial P_j} = \left(1 - \frac{\partial AC_j}{\partial P_j}\right) D_j(P) + [P_j - AC_j(P)] \cdot \frac{\partial D_j}{\partial P_j} = 0$. After rearranging terms, this can

higher WTP for quality: $Corr(R_i, \beta_i) > 0$. Equivalently, this model can be stated as $U_{ij} = Q_j - \alpha_i P_j$ where $\alpha_i = 1/\beta_i$ is a consumer's price sensitivity and $Corr(R_i, \alpha_i) < 0$. In our model, consumers have heterogeneous preferences for "match quality," which we can denote as $Q_{ij} = U(X_j; \zeta_i)$. If we assume this match quality has a similar "scale" for all consumers and write $U_{ij} = Q_{ij} - \alpha_i P_j$, our notion of adverse selection again corresponds to $Corr(R_i, \alpha_i) < 0$, that is higher-risk consumers are less price sensitive.

[11] To show this, differentiate $AC_j(P)$ from (2) using the product rule and rearrange terms to get (4). The formula can also be derived using the fact that $MC_j(P) = \frac{\partial TC_j/\partial P_j}{\partial D_j/\partial P_j}$ where $TC_j(P) \equiv D_j(P) \cdot AC_j(P)$ is total variable costs.

be written:

$$\begin{aligned}\underbrace{P_j - AC_j(P)}_{\text{Profit Margin}} &= \frac{1}{\eta_{j,P_j}} \times \left(1 - \frac{\partial AC_j}{\partial P_j}\right) \\ &= \underbrace{\frac{1}{\eta_{j,P_j}}}_{\text{Lerner Markup}} - \underbrace{[AC_j(P) - MC_j(P)]}_{\text{Degree of adverse selection}}\end{aligned} \tag{5}$$

where the second line follows from plugging in $\frac{\partial AC_j}{\partial P_j}$ from equation (4). The firm's profit margin equals the standard Lerner markup over marginal costs ($= \frac{1}{\eta_{j,P_j}} > 0$) minus the "wedge" between average and marginal costs ($AC_j - MC_j$). As in non-selection markets, the Lerner term is still the correct markup over *marginal* cost, i.e. $P_j - MC_j(P) = \frac{1}{\eta_{j,P_j}} > 0$. But adverse selection creates a "wedge" between average and marginal costs ($AC_j - MC_j$) that pushes firms to cut prices in order to attract, or "cherry pick," low-cost marginal consumers from other firms.

Prior work on the interaction of adverse selection and imperfect competition has pointed out the selection-pricing relationship in (5) (Starc, 2014; Mahoney and Weyl, 2017*a*). This prior work shows that (conditional on a set of competitors) adverse selection *disciplines market power* by reducing 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) assumes that there *exists* a Nash pricing equilibrium where *all competitors are profitable* net of fixed costs. Our basic conceptual argument, which we turn to next, is that when adverse selection is strong enough, neither assumption can be taken for granted. Adverse selection may constrain how many firms can profitably compete in a market, and this constraint may become quite tight when selection and price competition are strong.

Implications for Firm Entry

We now move backward to the firm entry decision in stage 1 of the game. At an equilibrium, all entrants $j \in E$ must expect to earn enough profits at the resulting stage-2 Nash-Bertrand prices P^* to cover their fixed costs, or $\pi_j(P^*) = \left[P_j^* - AC_j(P^*)\right] \cdot D_j(P^*) \geq F$. Plugging in (5), all entrants must have:

$$P_j^* - AC_j(P^*) = \underbrace{\frac{1}{\eta_{j,P_j}}}_{\text{Lerner Markup}} - \underbrace{[AC_j(P^*) - MC_j(P^*)]}_{\text{Degree of adverse selection}} \geq \underbrace{\frac{F}{D_j(P^*)}}_{\text{Fixed cost per consumer}} \tag{6}$$

Further, any non-entrant $j' \notin E$ must expect to lose money net of fixed costs if they enter, given the resulting Nash-Bertrand prices $P^*_{E \cup j'}$ among competitors $E \cup j'$ (i.e., $\pi_{j'}\left(P^*_{E \cup j'}\right) < F$).

These conditions have important implications for the interaction between adverse selection

Figure 1. How Adverse Selection Affects Equilibrium Competition

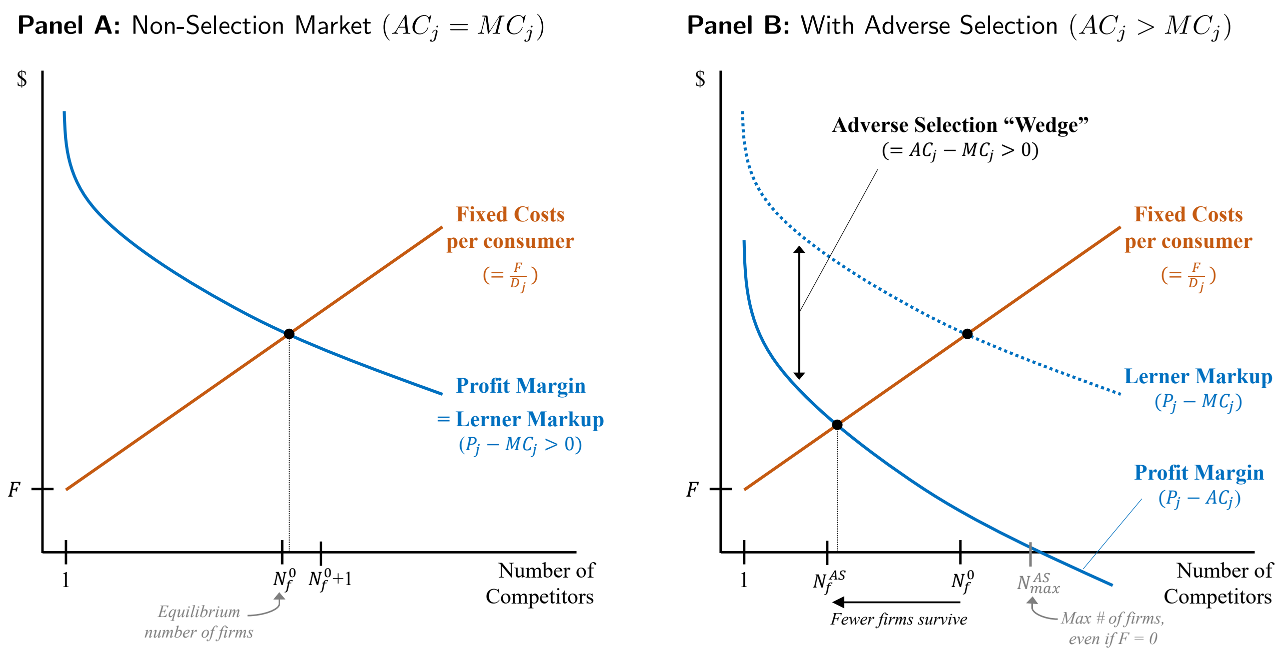

Note: The figure shows the implication of adverse selection for equilibrium firm entry into a market, based on the condition in equation (6). See the body text for a detailed discussion. An equilibrium occurs at the maximum integer number of firms, N_f, at which the profit margin ($P_j - AC_j(P)$) exceeds fixed costs per consumer ($F/D_j(P)$). Panel A shows this for a non-selection market, where average and marginal costs are equal, and the profit margin equals the (positive) Lerner markup over marginal costs. Panel B shows that adverse selection (which implies that $AC_j > MC_j$) drives a wedge between profit margins and the Lerner markup, reducing the number of firms that can survive. Moreover, there may be a maximum number of firms that can enter, even with $F = 0$, which occurs where the profit margin curve crosses into negative territory.

and competition, which we visualize in Figure 1. The figure shows how condition (6) can define equilibrium entry both in a non-selection market (Panel A) and a market with adverse selection (Panel B). The x-axis is the number of competing firms, N_f, and the graphs show two sets of curves: (1) the variable profit margins for competing firms (the LHS of (6)) in blue and (2) each firm's fixed costs per consumer (the RHS of (6)) in orange. The plots shown are based on a simple example (fleshed out in the next section) with symmetric firms and a consumer population size normalized to 1, implying that $D_j(P^*) = 1/N_f$ and $\frac{F}{D_j(P^*)} = F \cdot N_f$, but the logic underlying the graph is general.

Panel A shows the standard case of a *non-selection market*, in which $AC_j = MC_j$. As a result, variable profit margins are identical to the Lerner markup, which is always positive: $P_j^* - AC_j(P^*) = P_j^* - MC_j(P^*) = \frac{1}{\eta_{j,P_j}} > 0$. As more firms enter the market, price competition strengthens (higher η_{j,P_j}), so markups decline as shown in the downward-sloping blue curve. Meanwhile, firms further split the fixed consumer population, implying that the orange curve (fixed costs per consumer) rises. An equilibrium number of firms (N_f^0) occurs at the largest integer to the left of where the Lerner markup and fixed cost curves cross. Notably, as $F \to 0$ (corresponding to lower fixed costs or a growing market size) the number of firms that can enter and survive also grows in an unlimited way.

Panel B shows how the economics of entry differ in a market with *adverse selection*, in which $AC_j > MC_j$. As a result, the profits margin curve $(P_j - AC_j)$ is now shifted downward relative to the Lerner markup curve $(P_j - MC_j)$, with the gap equal to what we have called the "adverse selection wedge," $AC_j(P^*) - MC_j(P^*)$. This downward shift has two consequences. First, the equilibrium number of competitors, N_f^{AS}, is weakly lower than without selection. Second, there is no longer a guarantee that profit margins will be positive. Indeed, as more firms enter, we might expect cross-firm adverse selection to become stronger (a point noted by Lustig (2010)) while the Lerner markup declines, and eventually, variable profits are likely to become negative. This suggests a limit on the number of firms that can profitably survive in an adverse selection market, even as $F \to 0$.

Our analysis also implies that policies like risk adjustment that mitigate adverse selection may be able to promote more entry and competition, and the associated product variety it brings. We return to this implication below.

A Testable Condition Our theory implies a testable condition for when adverse selection and price competition have become "too strong" to sustain a given market structure. Rearranging condition (6) and plugging in the expression for $\frac{\partial AC_j}{\partial P_j}$ from (4), it must be true that:

$$\frac{\partial AC_j(P^*)}{\partial P_j} \leq 1 - \eta_{j,P_j} \cdot \frac{F}{D_j(P^*)} \tag{7}$$

In words, the price-slope of a firm's average cost curve must be less than a threshold that is 1.0 if $F = 0$ and strictly less than 1.0 if $F > 0$. When the average cost curve has a slope exceeding one, a given market structure cannot be sustained without corrective government policies. We examine this condition in our empirical work, showing that absent corrective policies like risk adjustment, it is very plausibly violated in our empirical setting.

Parallel between Adverse Selection and Fixed Costs We now note another way of seeing the very similar implications of fixed costs and adverse selection as limits to entry: both forces imply a *steeper average total cost curve*. To see this, we first define a firm's average *total* costs, $ATC_j(P)$, which includes fixed costs as: $ATC_j(P) = AC_j(P) + \frac{F}{D_j(P)}$. This lets us write net profits as $\pi_j^{Net}(P) = [P_j - ATC_j(P)] \cdot D_j(P)$, where $P_j - ATC_j(P)$ is the *net* profit margin. Differentiating ATC_j and rearranging terms yields:

$$\frac{\partial ATC_j}{\partial P_j} = \underbrace{\eta_{j,P_j}}_{\text{Price sensitivity of demand}} \times \left[\underbrace{(AC_j(P) - MC_j(P))}_{\text{Adverse selection}} + \underbrace{\frac{F}{D_j(P)}}_{\text{Fixed costs}} \right] \tag{8}$$

Equation (8) shows that the degree of adverse selection and fixed costs per consumer both lead to a steeper average total cost curve in a very parallel way. Equivalently, both drive a wedge between marginal costs and average total costs. Indeed, both features – a downward-sloping average cost curve and a wedge between marginal and average costs – are standard in textbook treatments of natural monopoly (e.g., Tirole, 1988) and of adverse selection (Einav and Finkelstein, 2011).[12]

2.2 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

$$U_{ij} = -P_j + (V - t_i \cdot \|\ell_i - \ell_j\|) \tag{9}$$

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 $D_{ij}(P) = \frac{1}{N} - \frac{1}{2}\alpha_i\left((P_j - P_{j-1}) + (P_j - P_{j+1})\right)$, which is a linear demand

[12] Indeed, plugging in the formula for $\frac{\partial ATC_j}{\partial P_j}$ into the condition for profitable equilibrium in (7) shows that it collapses to $\frac{\partial ATC_j}{\partial P_j} \leq 1$. Equivalently, in price-quantity space, the average total cost curve must be weakly less steep than demand. The point where $\frac{\partial ATC_j}{\partial P_j} = \frac{\partial ATC_j/\partial D_j}{\partial P_j/\partial D_j} = 1$ occurs when demand and average cost have the same slope, that is when ATC_j is tangent to inverse demand in price-quantity space. This is the familiar condition for equilibrium in monopolistic competition models.

[13] The 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.

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:

$$P^* = \underbrace{\frac{\sum_i (\theta_i \alpha_i) \cdot C_i}{\sum_i \theta_i \alpha_i}}_{=MC_j} + \underbrace{\frac{1/N}{\sum_i \theta_i \alpha_i}}_{\text{Lerner Markup}} \tag{10}$$

where the 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 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

$$AC_j - MC_j = \sum_i \theta_i C_i - \sum_i s_i^{MC} C_i > 0 \tag{11}$$

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 (5) 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:

$$N < \frac{1}{(AC_j - MC_j) \cdot \left(-\frac{\partial D_j}{\partial P_j}\right)} \tag{12}$$

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 2 presents results from a simple calibrated version of this model, with parameters calibrated based

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

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- (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 2A 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 2B 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 2B 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 (2017*b*). But when adverse selection crosses a threshold where a firm exits, prices jump upward because of the weaker competition.

[15] Based 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.

Figure 2. Equilibrium Number of Competitors and Prices in Simple Model

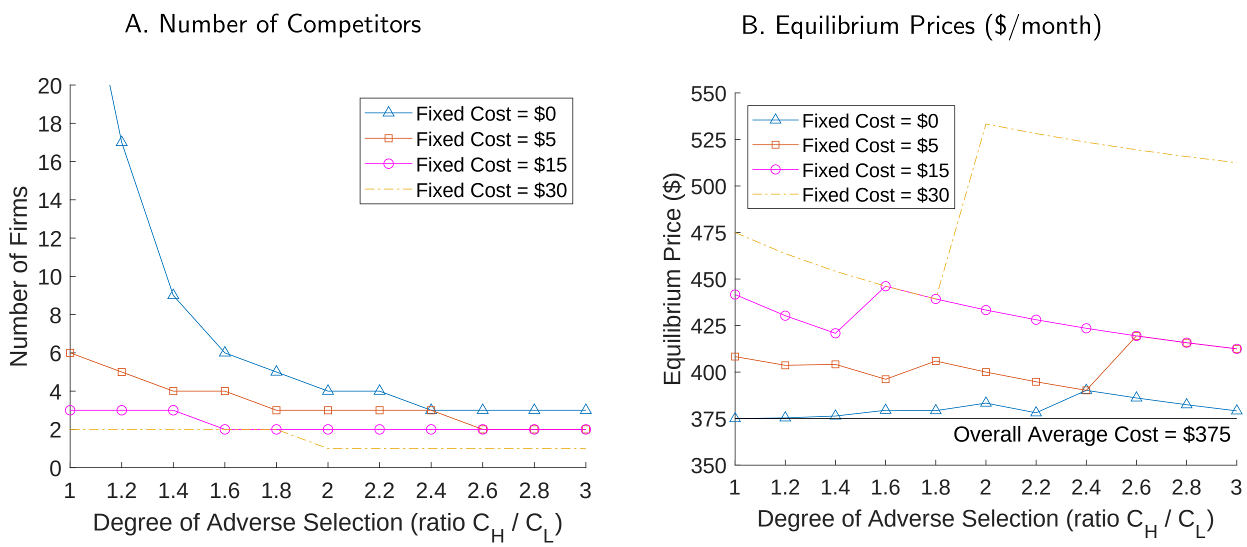

3 Empirical Setting and Data

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.

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) without access to Medicaid, Medicare, or job-based coverage. Because of its ACA-like structure, rich policy variation, and comprehensive administrative data, the Massachusetts exchange has been a fruitful setting for research on health insurance markets.[16]

The market featured four to five competing health insurers, with each insurer offering a single

[16]Prior work on CommCare includes Chandra et al. (2014); Finkelstein et al. (2019); Jaffe and Shepard (2020); Shepard (2022); McIntyre et al. (2021); Shepard and Forsgren (2022); Shepard and Wagner (2021). Other work has studied the pre-ACA unsubsidized Massachusetts health insurance exchange, a program known as "CommChoice" (Ericson and Starc, 2015*a*,*b*, 2016).

highly regulated plan that followed standardized cost sharing rules.[17] The primary way plans differed was on their networks of covered hospitals and doctors. 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 (Boston Medical Center plan (BMC), Neighborhood Health Plan (NHP), and Network Health) had comparably broad but differentiated networks, covering 75-85% of hospitals.[18] 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.

Conditional on participating in the market (and on their networks), insurers competed on prices (premiums) in a setup similar to our model. Insurers reset their premiums annually at the start of the year, which were then locked in until the end of the year. Premiums could vary only on specific factors like income group and region, not age or health status.[19] Consumers—who enrolled in the market throughout the year as needs for insurance arose—chose among competing plans at two times: (1) when they newly enrolled, and (2) at the start of each year ("open enrollment"), when premiums reset and they had an opportunity to switch plans.

CommCare's regulator oversaw the market using a strong form of the "managed competition" model envisioned by Enthoven (1993) — indeed, much more so than in ACA markets today. This strong regulation 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. At a high level, 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 model their counterfactual relevance in a setting (like the ACA) that does much less to soften and regulate price competition.

In what follows, we summarize the three main policy tools used by CommCare with relevance for adverse selection and price competition.

1) Risk Adjustment Like the ACA, CommCare applied *risk adjustment* to insurer revenue based on their enrollees' measured health risk, based on age and diagnoses. For an insurer that set a base premium of P_j and attracted enrollees with average risk score $\bar{\varphi}_j$, the insurer received revenues of $\bar{\varphi}_j P_j$. Risk adjustment can mitigate adverse selection by reducing differences in risk-adjusted costs across enrollees and thereby "flattening" the average cost curve (Mahoney and Weyl, 2017*a*). However, risk adjustment can be imperfect (Brown et al., 2014), and there is evidence of its imperfection in CommCare specifically (Shepard, 2022).

[17]Insurers were required to offer the same plan (with identical features) to all consumers in the state. They could, however, choose whether or not to participate in each of 38 "service areas," and there is significant variation in firm participation by area. Because this entry decision occurs at a lower geographic level than pricing, we have not yet explored it in this paper. For more on this type of "partial rating area" offering, see Fang and Ko (2018).

[18]See Appendix Figure A15 for a graph of network size over time.

[19]The degree of allowed variation narrowed over time. Premiums could vary: at the income group x region level (from 2007-09), at the regional level (in 2010), and statewide (2011-13).

2) Price Regulation Second, the exchange directly regulated prices using price ceilings and floors. Price ceilings were intended to limit price growth 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 from 2010 forward, 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.

3) Subsidies that Narrowed Premium Differences Third, premiums were *subsidized* to ensure affordability and encourage consumer take-up. CommCare's subsidies were more complex than in the ACA exchanges today, but we explain the details because they matter 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 price elasticities separately from any unobserved changes in plan quality. We use 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 3 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.

Figure 3. Variation in Enrollee Premiums in CommCare ($ per month)

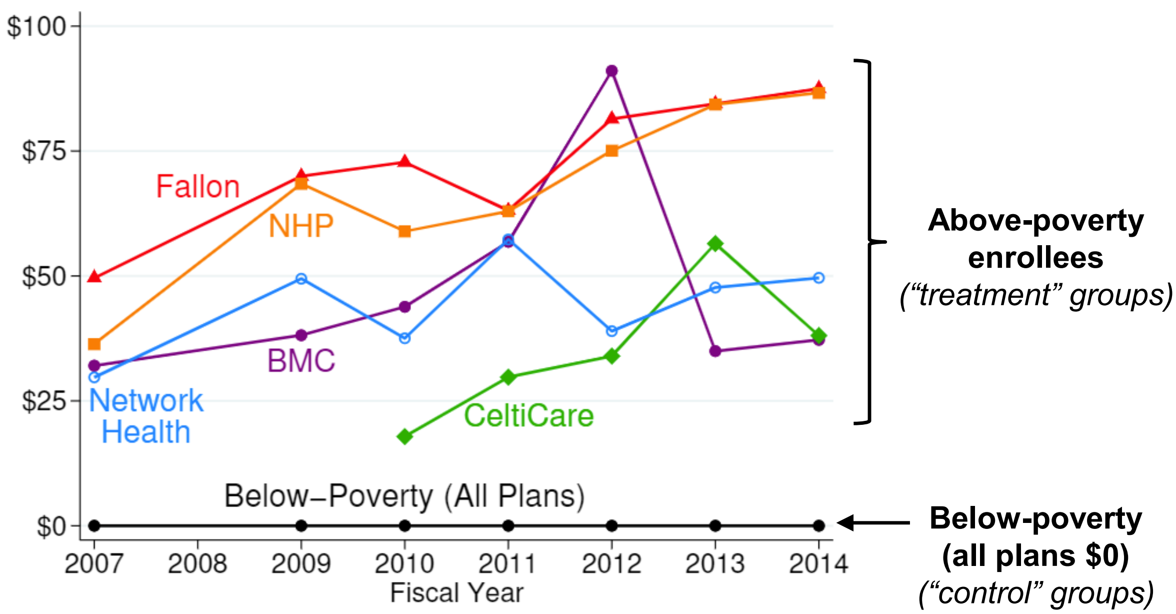

Note: The graph shows post-subsidy enrollee premiums for each insurer's plan in the CommCare market, by fiscal year and income group. Enrollment-weighted average premiums for above-poverty enrollees (100-300% of poverty) are shown in different colors by plan and labeled. As shown, these vary substantially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

3.2 Data: Administrative Enrollment and Insurance Claims

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 three '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 4. 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 4 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 simutaneously 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.

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 ($\eta_{j,P_j} = \frac{\partial \log D_j}{\partial P_j}$) and the average slope of the firm-specific average cost curve at the observed prices across all plan-years ($\frac{\partial AC_j}{\partial P_j}$). 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

Figure 4. Case Study: BMC and Network Health in 2012-2013

(a) Plan Bids and Relative Premiums

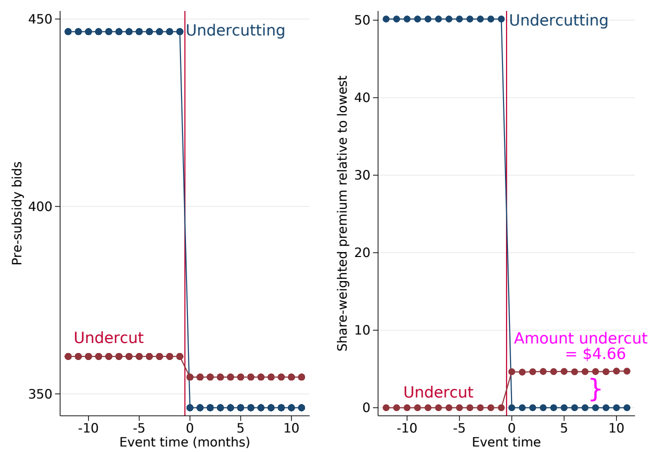

(b) Market Shares

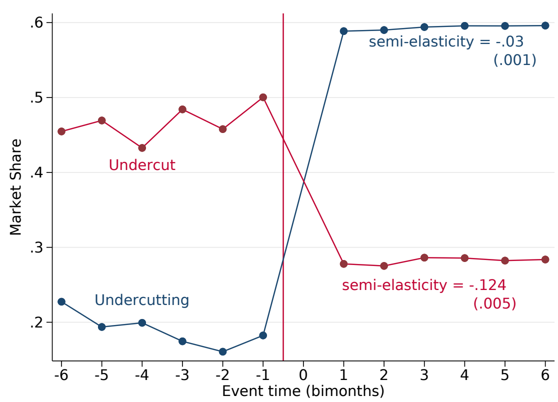

(c) Average 6-Month Cost

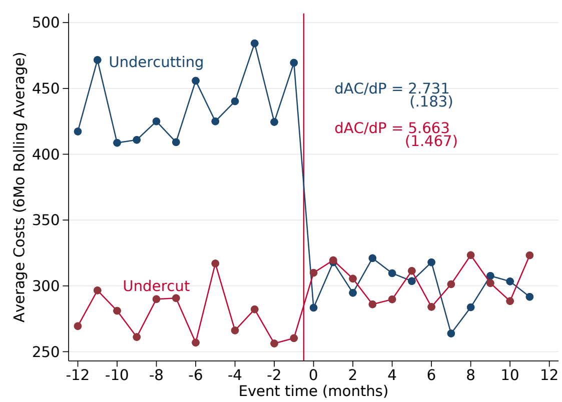

Note: Figure shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over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

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 at the start of each year, enrollment occurs monthly 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. 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, our estimator compares 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 $g \in \{0, 1, ...4\}$, 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:[20]

$$Y_{egt} = \tilde{\alpha}_{et} + \tilde{\gamma}_{eg} + \sum_{k \in \{-T,T\} \setminus \{-2\}} \tilde{\delta}_t \times 1\{g > 0, t = k\} + \tilde{\varepsilon}_{egt} \tag{13}$$

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

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

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 two months prior to the price change.[21] We estimate this regression for three key outcomes: relative premiums p, log market shares (log_sh), and average costs (AC).[22]

Figure 5 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 per month 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 increase in the monthly 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 rise, and when prices decrease, average costs decline. 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.[23]

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

[21] We 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 some evidence that these start to affect demand and costs in month -1.

[22] Costs for each enrollee-month are defined as the average monthly cost of the enrollee over the following 12 months, or until the enrollee leaves the dataset because they are no longer in the market.

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

Figure 5. Event Study Estimates for All Enrollees

(a) Relative Premiums ($/month)

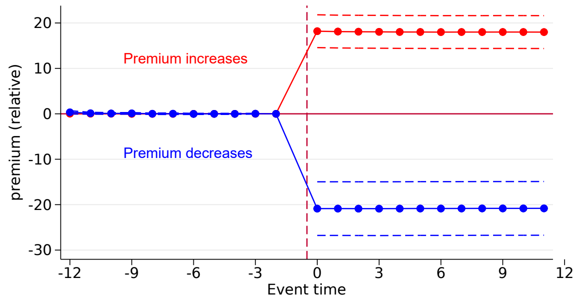

(b) Log Market Shares

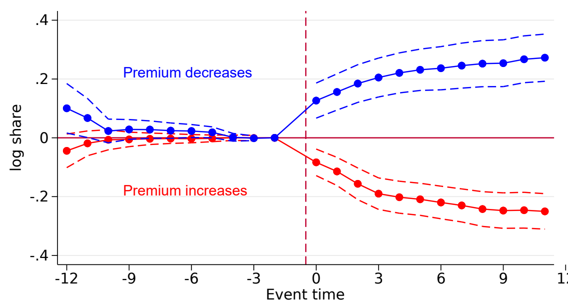

(c) Average Cost ($/month)

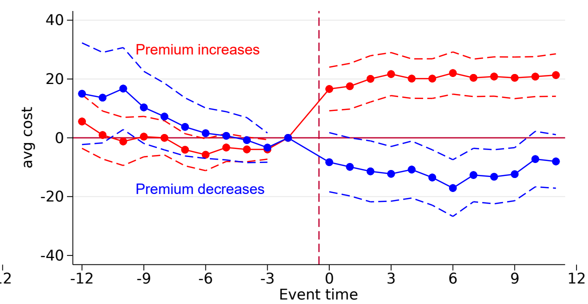

Note: Figure shows event study estimates of the impact of premium increases and decreases on relative premiums. Panel (a) shows results for relative premiums, Panel (b) shows results for log plan market shares, and Panel (c) shows results for average costs. The latter is defined as average costs per month (averaged over the subsequent year) of the set of enrollees who joined a plan in a given month.

equation (13). Figure 6 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.

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 new enrollees join the market (while prior enrollees drop out), leaving a larger share of enrollees who entered the market under the new prices. To illustrate this, in Appendix Figure A5 we replicate Figure 6, restricting only to new enrollees rather than including both incumbent and new enrollees. We find even larger price sensitivity (with 40% declines in market share for a $20 price increase) that are immediate and not dynamically changing.

Panel (c) of Figure 6 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

Figure 6. Pooled Event Study Estimates for All Enrollees

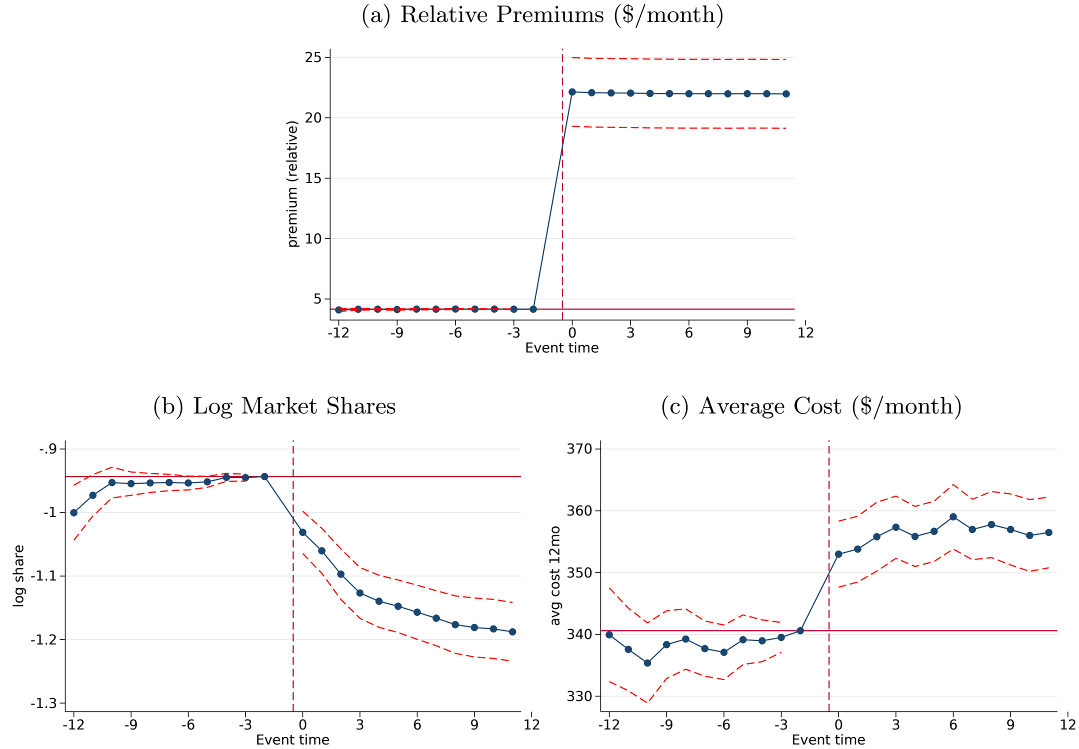

Note: Figure shows pooled event study estimates of the impact of all premium changes (increases and decreases), where outcomes are multiplied by -1 for premium decreases. Panel (a) shows results for relative premiums, Panel (b) shows results for log plan shares, and Panel (c) shows results for average costs. The latter is defined as average costs per month (averaged over the subsequent year) of the set of enrollees who joined a plan in a given month.

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 A5 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.

Table 1 summarizes overall difference-in-differences coefficients for a variety of outcomes and sub-samples. 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 6 and Column (2) corresponds to Figure A5. These estimates again suggest high levels of price sensitivity and strong adverse selection. Columns (3–5) return to using all enrollees but show results for selected subsets of price changes, including those where the identity of the cheapest plan changes and where price changes are relatively small (corresponding to "undercutting" style price changes). We continue to find large demand semi-elasticities and average cost curve slopes, suggesting our results are robust to these subsamples.

Table 1. Difference-in-Differences Results

	All Enrollees	New Enrollees	Cheapest Plan Changes	Price Changes < Median	Cheapest plan & Price Changes < Median
Panel (a): Price Sensitivity					
Premium	17.87***	17.90***	21.06***	7.135***	5.329***
	(0.287)	(0.318)	(0.611)	(0.113)	(0.829)
Market Share	-0.181***	-0.429***	-0.336***	-0.0459***	-0.141***
	(0.00522)	(0.0103)	(0.0100)	(0.00514)	(0.0311)
Demand Semi-Elasticity	**-0.0101**	**-0.0240**	**-0.0159**	**-0.00643**	**-0.0264**
Panel (b): Adverse Selection					
Premium	17.87***	17.90***	21.06***	7.135***	5.329***
	(0.287)	(0.318)	(0.611)	(0.113)	(0.829)
Average Cost	17.81***	32.48***	20.10***	9.689***	3.479
	(0.758)	(3.949)	(1.424)	(0.919)	(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 * $p < 0.10$.

Note: 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.

Ultimately, the results from the case studies and the diff-in-diff analysis 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 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.[24]

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.[25] 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:

$$u_{ijt} = -\alpha(Z_{it}) \cdot P_{ijt} + f(X_{jt}, Z_{it}; \beta) + \xi_j(W_{it}) + \varepsilon_{ijt}, \qquad j = 1, .., J \tag{14}$$

where P_{ijt} is plan j's post-subsidy premium for consumer i in year t (based on their income group

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

[25] The 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.

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 enrolle costs on medical observables (see Appendix Section D.1). 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 (6) 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

	Panel (a) All Enrollees		*Panel (b) New Enrollees*	
	Semi Elasticity	dAC/dP	Semi Elasticity	dAC/dP
By Plan				
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
By Income Group				
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 ($dACj/dPj$) 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 (8). 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:

$$C_{ijt}^{obs} = R_{it} \times \delta_{j,r(i)}. \tag{15}$$

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.[26] This 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.[27] We estimate the following Poisson regression specification:

$$E(C_{ijt}^{obs}|Z_{it}) = \exp\left(\alpha_i + \beta_t + Z_{it}\gamma + \lambda_{j,r}\right) \tag{16}$$

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 (15) 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.

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

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

Plan Cost Effect Estimates Table 3 shows estimates of cost heterogeneity.[28] 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 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.

[28]These 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.

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

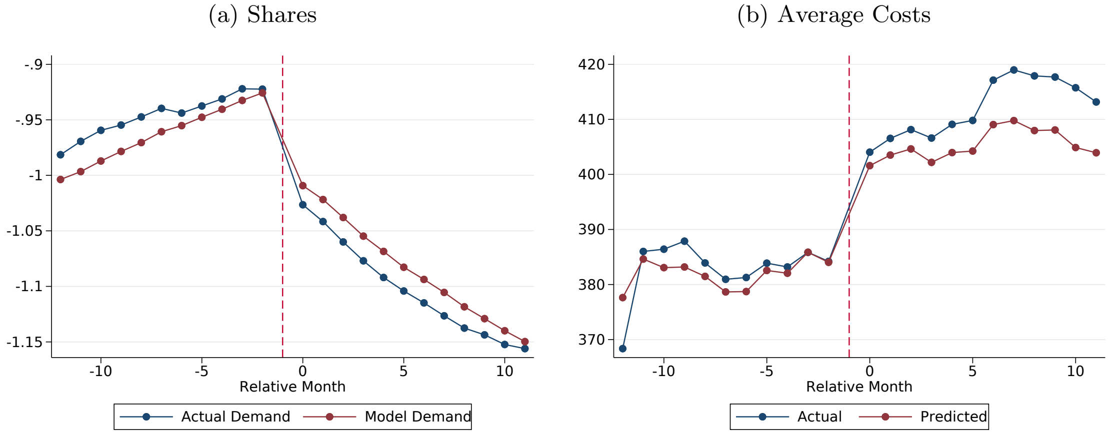

Notes: Figure shows how shares and average costs respond to increases and decreases in post-subsidy premiums, comparing actual shares and average costs for each spell with predictions using the demand estimates from Section 5. Panel (a) shows results for shares; Panel (b) shows results for average costs 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 (14) 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)

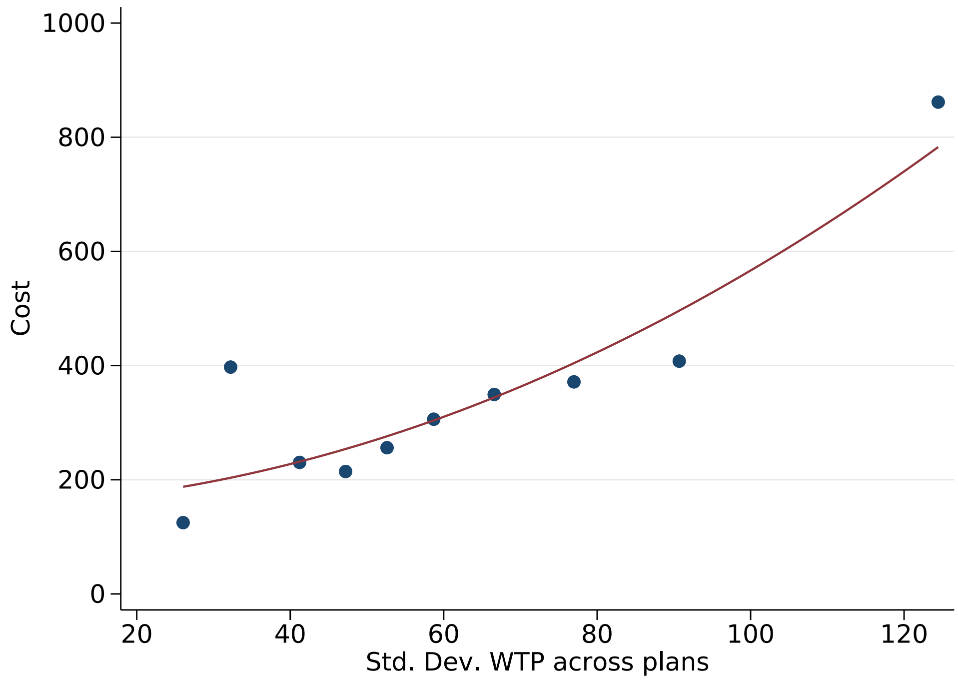

Notes: Figure shows a binned scatter plot of individual-level costs (for a given year t, and adjusted for plan effects) vs. deciles of the standard deviation of willingness to pay (V_{ijt}) *across plans* (j) in the individual's choice set. WTP is derived from the structural demand estimates of (14), 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.[29] 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

[29]To 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.[30]

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.[31]

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.[32] 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 E for additional details on the method of searching for equilibrium.

[30]Note 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).

[31]We 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.

[32]We 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:

$$\varphi_i = (C_{it}/\bar{C}_t)^{\lambda} \tag{17}$$

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 (5).

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 E 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 pure-strategy 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 ($\lambda = 1$) does allow the market to support all four firms. Among these three cases, consumer welfare is highest with moderate 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 (2017*a*) 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) Possible Entry Combinations	(2) Prices	(3) Shares	(4) Avg Price	(5) Welfare	(6) Remains with: $F=10$	(7) $F=30$
Panel (a) No Risk Adjustment						
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
Panel (b) Moderate Risk Adjustment ($\lambda = 0.50$)						
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	–
Panel (c) Perfect Risk Adjustment						
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 p1). 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

(a) No Risk Adjustment (b) Perfect Risk Adjustment

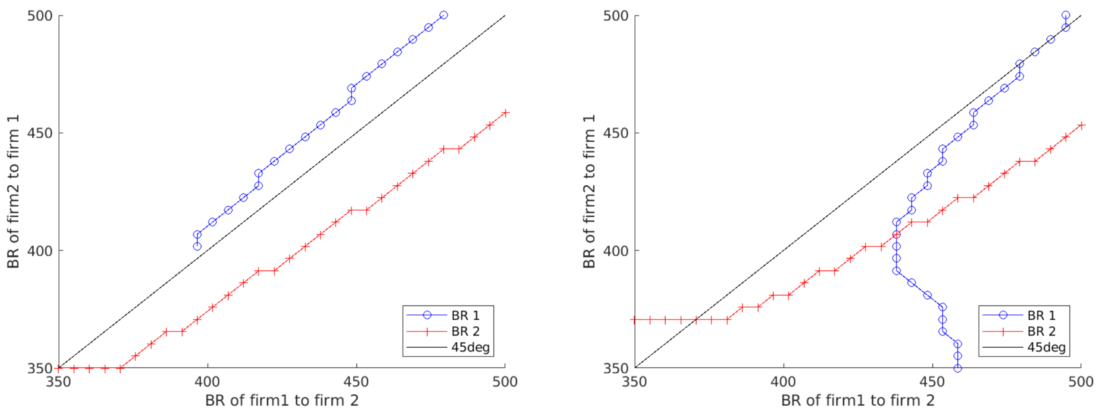

Notes: Panel (a) shows best response curves for BMC and Celticare in the case with no risk adjustment. Panel (b) shows the case with perfect risk adjustment. In both panels, the red curve labeled BR 2 shows the optimal pre-subsidy premium of Celticare on the Y-axis, given BMC's pre-subsidy premium on 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.[33]

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

[33]Extending 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

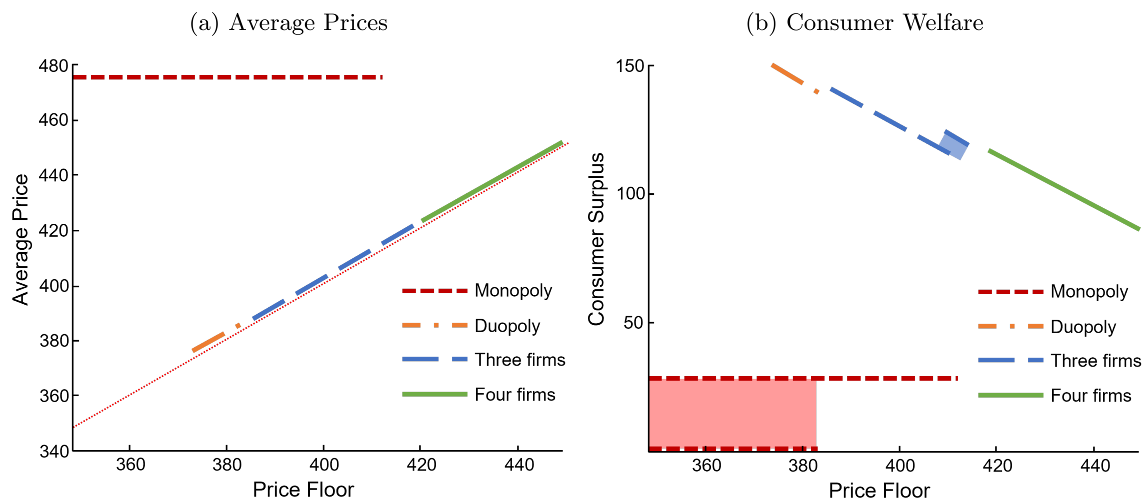

Notes: Figure shows equilibria as a function of the price floor, with no fixed costs and no risk adjustment for new enrollees. Panel (a) shows share-weighted average pre-subsidy premiums and Panel (b) shows consumer surplus per enrollee-month, normalized such that \$0 corresponds to the lowest-welfare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

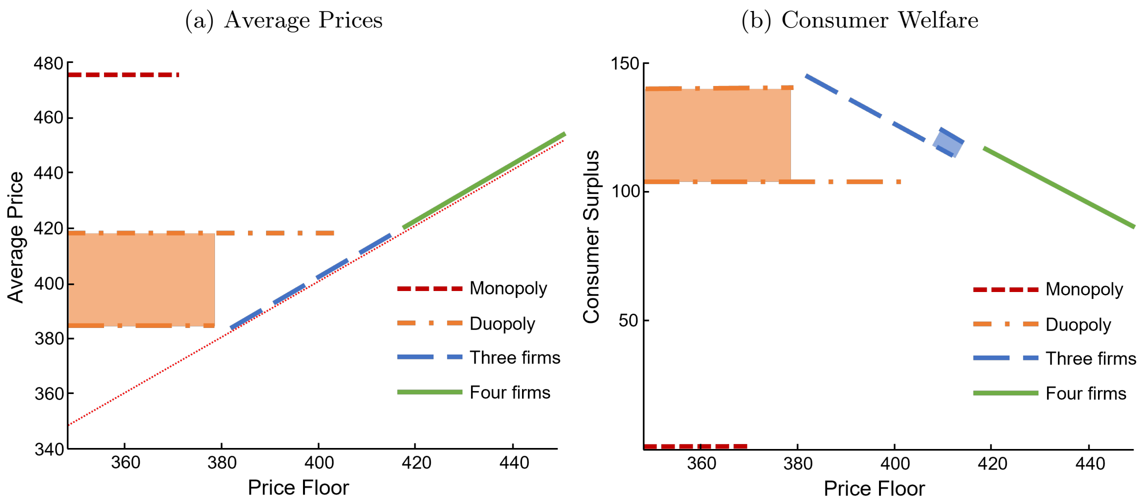

Notes: Figure shows equilibria from baseline simulations as a function of the price floor, with no fixed costs and moderate risk adjustment ($\lambda = 0.50$). Panel (a) shows share-weighted average pre-subsidy premiums and Panel (b) shows consumer surplus per enrollee-month, normalized such that \$0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

Table 5 shows how the consumer-welfare optimal price floor affects entry and welfare in a wide range of cases.[34] Panel (a) compares results between the sample of new enrollees and all enrollees for the case with moderate risk adjustment ($\lambda = 0.5$). 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

[34] We 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) Optimal Price Floor	(2) Firms	(3) Avg. Price	(4) Profits	(5) Welfare	(6) Firms Without Price Floor	(7) Welfare Gain from Price Floor
Panel (a) New Enrollees vs. All Enrollees							
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
Panel (b) Impact of Risk Adjustment							
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
Panel (c) Impact of Fixed Costs							
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
Panel (d) Impact of Cost Heterogeneity							
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.

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston. 2014. "How does risk selection respond to risk adjustment? New evidence from the Medicare Advantage Program." *American Economic Review*, 104(10): 3335–3364.

Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney. 2012. "Pricing and Welfare in Health Plan Choice." *American Economic Review*, 102(7): 3214–48.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight. 2014. "The impact of patient cost-sharing on low-income populations: evidence from Massachusetts." *Journal of health economics*, 33: 57–66.

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.

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

Einav, Liran, 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, 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.

Einav, Liran, and Amy Finkelstein. 2011. "Selection in Insurance Markets: Theory and Empirics in Pictures." *Journal of Economic Perspectives*, 25(1): 115–38.

Enthoven, Alain C. 1993. "The history and principles of managed competition." *Health affairs*, 12(suppl 1): 24–48.

Ericson, Keith, and Amanda Starc. 2015*a*. "Measuring Consumer Valuation of Limited Provider Networks." *American Economic Review Papers and Proceedings*, 105(5): 115–119.

Ericson, Keith M Marzilli, and Amanda Starc. 2015*b*. "Pricing regulation and imperfect competition on the massachusetts health insurance exchange." *Review of Economics and Statistics*, 97(3): 667–682.

Ericson, Keith M Marzilli, and Amanda Starc. 2016. "How product standardization affects choice: Evidence from the Massachusetts Health Insurance Exchange." *Journal of Health Economics*, 50: 71–85.

Fang, Hanming, and Ami Ko. 2018. "Partial Rating Area Offering in the ACA Marketplaces: Facts, Theory and Evidence." National Bureau of Economic Research Working Paper 25154.

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2019. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." *American Economic Review*, 109(1): 1530–67.

Frank, Richard G, and Thomas G McGuire. 2019. "Market concentration and potential competition in medicare advantage." *Issue brief (Commonwealth Fund)*, 2019: 1–8.

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.

Greenwald, Bruce C. 1986. "Adverse selection in the labour market." *The Review of Economic Studies*, 53(3): 325–347.

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.

Lustig, Josh. 2010. "Measuring welfare losses from adverse selection and imperfect competition in privatized Medicare." *Manuscript. Boston University Department of Economics.*

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.

McDermott, D, Anthony Damico, Cynthia Cox, Gary Claxton, and Larry Levitt. 2020. "Insurer participation on the ACA marketplaces, 2014–2021." *Data Note, Kaiser Family Foundation.*

McIntyre, Adrianna L., Mark Shepard, and Myles Wagner. 2021. "Can Automatic Retention Improve Health Insurance Market Outcomes?" *AEA Papers and Proceedings*, 111: 560–566.

Michaely, Roni, and Wayne H Shaw. 1994. "The pricing of initial public offerings: Tests of adverse-selection and signaling theories." *The Review of Financial Studies*, 7(2): 279–319.

Nash, John. 1951. "Non-Cooperative Games." *The Annals of Mathematics*, 54(2): 286–295.

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.

Saltzman, Evan. 2017. "The Welfare Implications of Risk Adjustment in Imperfectly Competitive Markets." University of Pennsylvania Working Paper.

Shepard, Mark. 2022. "Hospital network competition and adverse selection: evidence from the Massachusetts health insurance exchange." *American Economic Review*, 112(2): 578–615.

Shepard, Mark, and Ethan Forsgren. 2022. "Do insurers respond to active purchasing? Evidence from the Massachusetts health insurance exchange." *Journal of Risk and Insurance*.

Shepard, Mark, and Myles Wagner. 2021. "Reducing Ordeals through Automatic Enrollment: Evidence from a Health Insurance Exchange." Working paper.

Stantcheva, Stefanie. 2014. "Optimal income taxation with adverse selection in the labour market." *Review of Economic Studies*, 81(3): 1296–1329.

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

Tirole, Jean. 1988. *The theory of industrial organization.* MIT press.

A Undercutting Case Studies

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

(a) Plan Bids and Relative Premiums

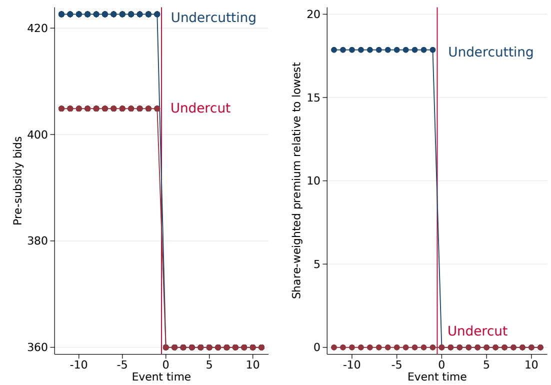

(b) Market Shares

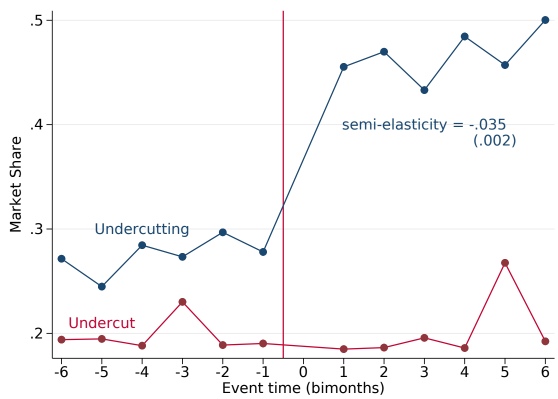

(c) Average 6-Month Cost

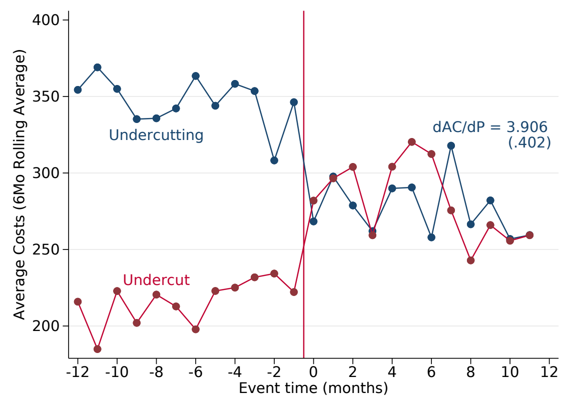

Note: Figure shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over 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

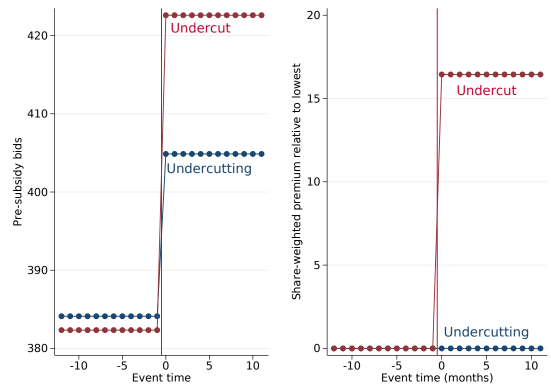

(b) Market Shares

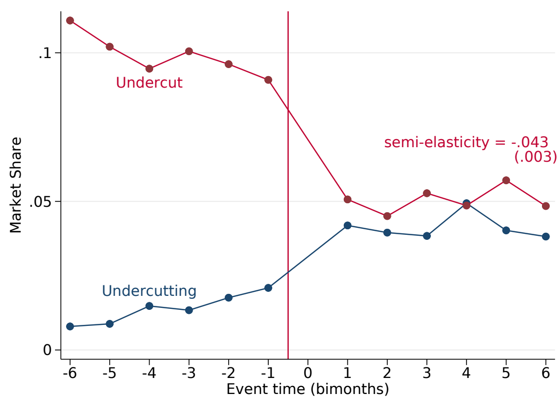

(c) Average 6-Month Cost

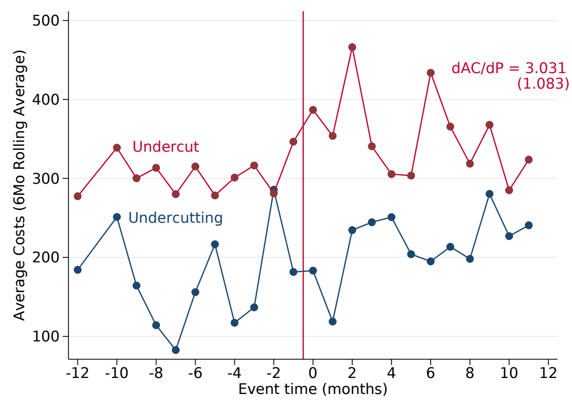

Note: Figure shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over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

B Regression Discontinuity Analysis

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 below 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 an imbalance in a key observable, 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.

Appendix Figure A3. RD's

(a) Total Number of Enrollees

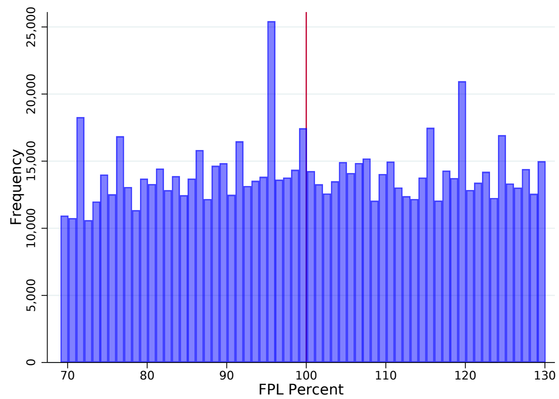

(b) Age

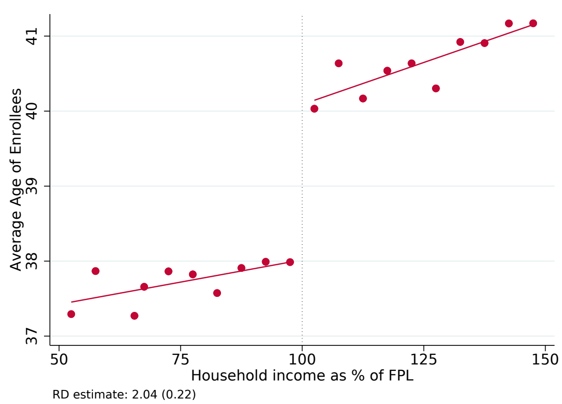

(c) Percent Female

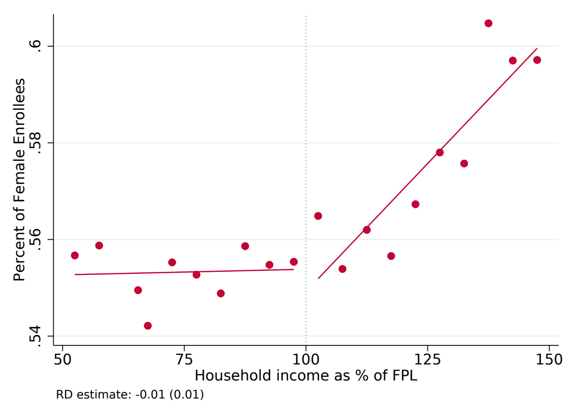

Because of these violations of the key identifying assumption, we interpret all RD results as descriptive — revealing patterns that are suggestive of strong price sensitivity but not cleanly identifying the key parameter of the firm-specific price semi-elasticity of demand. We also do not present RD 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:

$$Y_i = \beta_0 + \beta_1 Above100_i + \beta_2 FPL_i + \beta_3 FPL_i \times Above100_i + \epsilon_i \tag{18}$$

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

(a) All Enrollees (b) New Enrollees

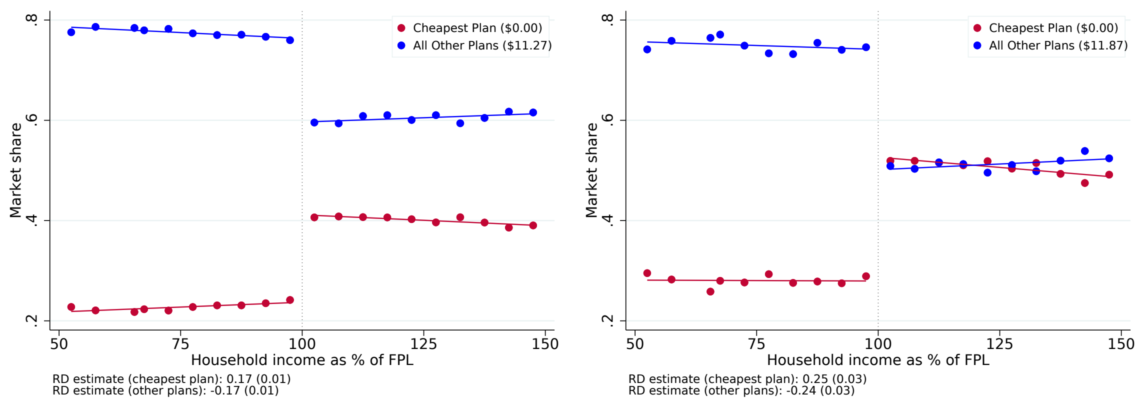

Note: Figure shows regression discontinuity plots for the market share of the lowest-priced plan (red) and the combination of all other plans (blue). Panel (a) shows results for all enrollees, Panel(b) shows results for new enrollees.

Appendix Table A1. RD Main Estimates

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

Appendix Figure A5. Pooled Event Study Estimates for New Enrollees Only

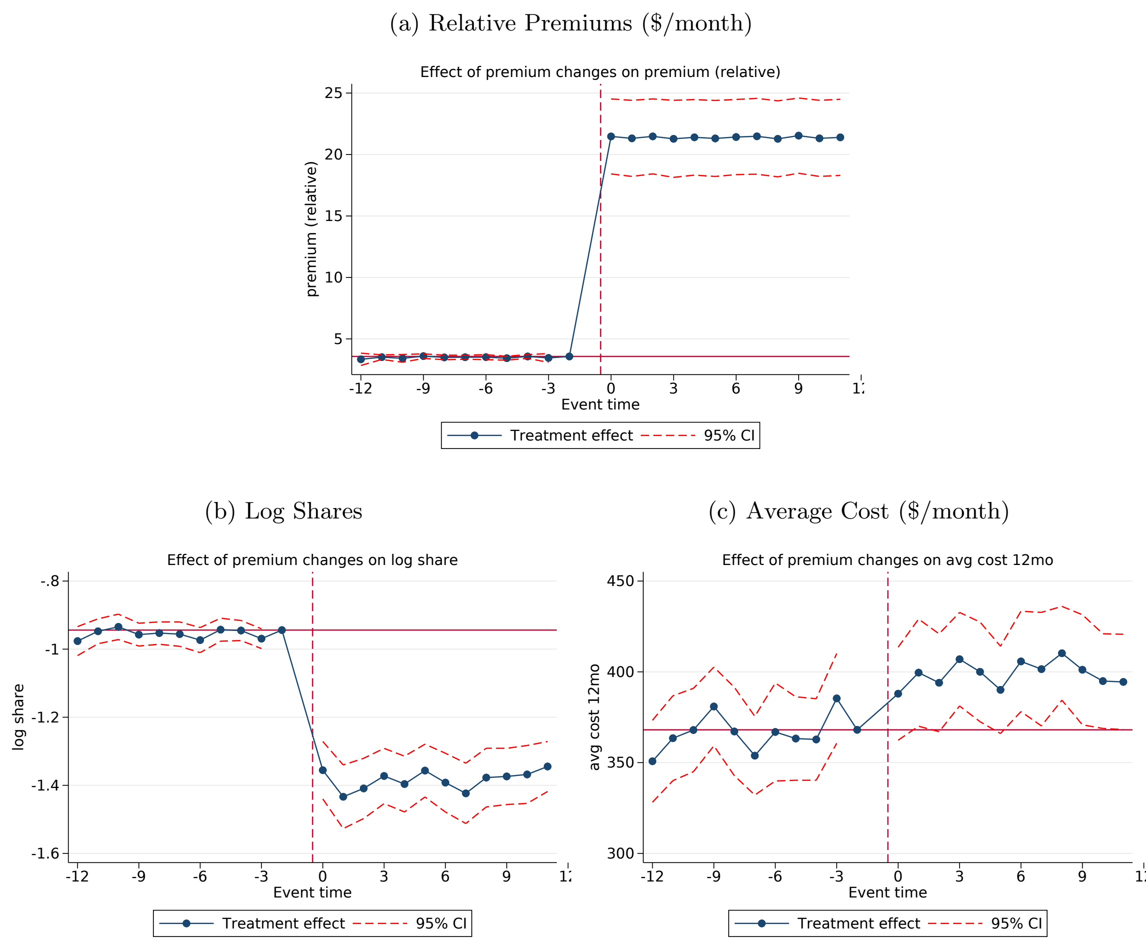

Note: Figure shows pooled event study estimates of the impact of premium increases and decreases, where all outcomes are multiplied by -1 for premium decreases. Sample is limited to new enrollees only. Panel (a) shows results for relative premiums, Panel (b) shows results for log plan shares, and Panel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

D Demand Estimation

D.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 16 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.

E Counterfactual Simulations: Detailed Methods

E.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. When the degree of adverse selection is high (e.g., in the no-risk-adjustment case), we generally find no pure-strategy Nash equilibrium. In such cases, we solve for mixed-strategy pricing equilibria, following the method in Subsection E.2 below. For all combinations of entrants, we are able to identify either a pure- or mixed-strategy pricing equilibrium, although the identified equilibrium does not always yield positive profits to all firms.

For a given combination of firms, we first attempt to find a pure-strategy pricing equilibrium. To do this, 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 tolerance[35] 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

[35] The 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).

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.[36]

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.

E.2 Mixed-strategy pricing equilibria

Non-existence of pure-strategy equilibria: In cases with a high degree of adverse selection and no corrective policies (e.g., risk adjustment or price floors), pure-strategy equilibria may not exist. In such cases, the best-response functions of one or more firms has a discontinuity, where a firm no longer finds it profitable to continue undercutting but would rather raise its price to a higher level (usually the price ceiling). As a result, there is no intersection of the best-response curves, and therefore no pure-strategy Nash equilibrium.

Figure A6 below shows an two-firm example where pure-strategy equilibria do not exist. Starting from the top right of the figure, both firms undercut each other, but the low-cost firm (firm 2, Celticare) undercuts to a greater degree. Eventually, firm 4 (NHP) no longer earns positive profits, and there is a discontinuity in its best response curve, where instead of undercutting, firm 4 would rather price at the price ceiling to minimize its losses. This demonstrates how non-existence of pure-strategy equilibrium corresponds to discontinuities in one or both best-response curves.

[36] Although 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 A6. Best Response Curves between Firms 2 and 4

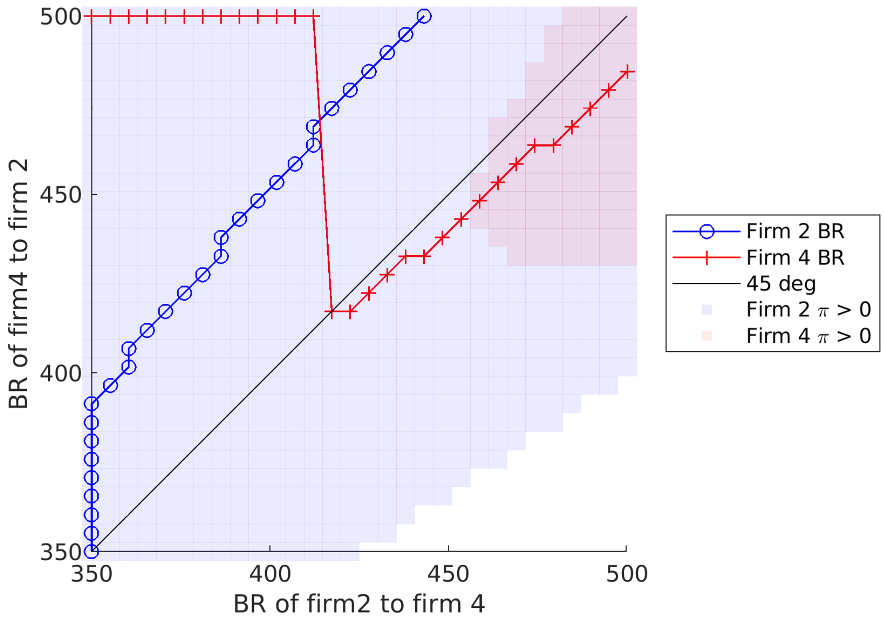

Notes: Figure shows best response curves for Celticare and NHP in the case with no risk adjustment. In both panels, the red curve labeled "Firm 4 BR" shows the optimal pre-subsidy premium of NHP on the Y-axis, given Celticare's pre-subsidy premium on the X-axis. The curve labeled "Firm 2 BR" shows the optimal pre-subsidy premium of Celticare on the X-axis, given NHP's pre-subsidy premium on the Y-axis. Red shaded regions correspond to prices where NHP earns positive profits, and blue shaded regions correspond to profitable regions for Celticare.

Although these cases do not permit pure-strategy equilibria, mixed-strategy equilibria do exist. In the case shown in Figure A6,

We therefore adopt an additional step-by-step approach for solving for mixed-strategy equilibria in these cases. For each set of entrants, we look for a *single* equilibrium, since the problem of enumerating *all* mixed-strategy equilibria is computationally infeasible.

Because our setting involves greater than two firms and a continuum of possible actions (prices), we adopt our own algorithm. We leverage the intuition that firms either want to undercut or "quasi-exit" by raising their prices to the price ceiling. For every candidate equilibrium, the last step of our algorithm checks all prices between the price floor and ceiling, in increments of $5, to ensure that there are no profitable global deviations. Below, we describe our process for finding mixed equilibria for each number of entrants.

For two firms, we allow for the following types of cases, where each firm can mix between up to 3 prices. We are able to find an equilibrium for all firm combinations using these cases.

1. Cases where only one firm (denoted firm i) mixes over two prices $p_{iL} and p_{iH}$, where p_{iH} is set to the price ceiling of $500. Firm j sets a single price p_j. In this case, p_{iL} and p_j found by solving the corresponding pricing first order conditions.

2. Both firms mix over up to two prices each. All four prices $p_{iL}, p_{iH}, p_{jL}, p_{jH}$ must satisfy the corresponding first order condition. Note: this allows one of the firms to play a pure strategy (e.g., $p_{iL} = p_{iH}$.
3. Case where both firms mix, but one firm mixes between three prices (with the highest price set to the price ceiling). This only applies in the case of BMC and Celticare with no risk adjustment.

In each case, prices that are not set to the price ceiling must satisfy the Nash first-order pricing conditions; we use the Matlab command *lsqnonlin* to find these prices.

For the three-firm case, we enumerate all possible combinations of which firms mix (we allow up to two firms may mix between up to 2 separate prices, yielding 7 cases). Within each of these cases, we allow up to 2 firms to set their highest price (either p, in the case of a single price, or p_H, in the case of two prices) to the price ceiling, yielding 6 cases. In total, this gives 42 cases for which we attempt to find an equilibrium. As before, we use *lsqnonlin* to solve for prices satisfying the pricing first order conditions for all prices that are freely set (i.e., not fixed at the price ceiling).

For the four-firm case, given the large number of possible combinations, we manually experiment with various cases until we arrive at a valid equilibrium. In practice, we begin by setting most prices to the price ceiling, allowing the lowest-cost firms to mix between two prices. This proves to be enough to locate an equilibrium.

F Testing our Model: Reinsurance and Participation in the ACA

F.1 Overview

How does adverse selection affect insurer participation in the ACA marketplaces? Our model predicts that, among horizontally differentiated issuers, areas with steeper average cost curves (larger dAC/dP) should have less firm participation, but that risk adjustment and reinsurance can correct for this.

We don't have firm-specific average cost data for the ACA marketplaces, so we use reinsurance payments from 2014-16 to get a proxy measure of dAC/dP. This proxy measure is d(reinsurance)/dP: the change in reinsurance payments received by a firm when it changes its price.

We combine our estimates of d(reinsurance)/dP (at the state level) with the removal of reinsurance after 2016 to estimate the effect of the slope of the average cost curve on firm participation. Firms in markets with larger d(reinsurance)/dP in the pre-period (2014-16) experience larger increases in adverse selection when reinsurance is removed. We expect that these areas will have larger decreases in firm participation in the post-period (2017-18).

We find that areas with larger increases in dAC/dP after the removal of reinsurance (i.e., areas where reinsurance was more important for flattening the average cost curve) have larger declines in participation, consistent with the predictions of our model. In the following subsections, we explain our identification argument in more detail (Appendix F.2), explain our estimation approach (Appendix F.3), show estimates of d(reinsurance)/dP (Appendix F.4) at the aggregate and state-specific levels, and then relate these estimates to changes in insurer participation after the expiration of reinsurance in 2016 (Appendix F.5).

F.2 Empirical Approach

F.2.1 Testing our model's predictions about firm entry: data and identification challenges

Testing the predictions of our model in the ACA marketplaces is challenging due to data limitations and identification concerns. Ideally, we would like to relate market-level insurer participation to exogenous changes in firm-level price elasticities $\frac{dD_{j,m}}{dP_{j,m}}$ and average cost curve slopes $\frac{dAC_{j,m}}{dP_{j,m}}$, which are the key empirical objects whose magnitudes reflect the degree of undercutting incentives and hence the degree to which firms do not participate in markets.

In terms of data availability, plan premiums are readily observed (QHP Landscape Files), but plan-specific enrollment is not released at the county level, which is the natural level at which to evaluate insurer participation, as insurers are allowed to selectively enter counties, despite only being able to set prices at the rating area level (Geddes 2023, Fang and Ko). Similarly, average costs are not observable at the plan level, only at the issuer level.

There are also identification challenges in this setting. Insurance premiums set by firms may affect their average costs via adverse selection, but the relationship between premiums and average costs may be confounded by reverse causality: firms may set higher premiums in response to changes in their costs. There also exists potential confounding due to straightforward omitted variable bias: unobserved plan quality differences will induce a positive correlation between premiums, costs, and demand. That is: plans offering broader provider networks or more generous coverage will tend to attract a larger number of costlier enrollees – these plans will need to set higher premiums as a result.

Lastly, we note that quasi-random price variation alone is not sufficient: while this would allow us to obtain unbiased estimates of demand elasticities and average cost curve slopes, cross-market variation in these quantities may still be correlated with unobservable market characteristics. For instance, an association between larger price elasticities and lower insurer entry could reflect confounding due to unobservably lower income.[37]

An ideal experiment would randomize *price elasticities* and *average-cost-curve slopes* across markets, but such an experiment is likely infeasible. Another potential source of variation in the average-cost-curve slope is variation in the degree of risk adjustment across markets, but this is also challenging to quantify.

With these data and identification challenges in mind, we develop a test of our model's predictions regarding the slope of the average cost curve that leverages the nationwide removal of reinsurance occurring after 2016 and pre-existing variation in the importance of reinsurance across states.[38] We describe the rationale of this approach in more detail in the following section.

F.2.2 Removal of reinsurance induces differential changes in the slope of the average cost curve

Reinsurance helps to reimburse firms for enrolling individuals with very high costs. Specifically, the ACA reinsurance program ran from 2014 to 2016 and reimbursed a fraction of insurers' costs for each individual with more than \$45K in annual spending (the "attachment point") up to a cap of \$250K.[39]

As we show below, the effect of reinsurance is to flatten the slope of the average cost curve. The degree of flattening can vary across markets. When reinsurance was removed, areas with more pre-period "flattening" experienced larger increases in the slopes of their average cost curves. It is this *heterogeneous increase in average cost curve slopes* that we use for identifying how the slope of the average cost curve affects participation.

To show this in mathematical terms, we can follow Section 2 to write firms' average total cost per enrollee, now including reinsurance, as:

[37] An additional nuance is that the price elasticity is an equilibrium object that depends on the number of competitors and applies only locally around the observed market equilibrium prices. An analysis of the effect of price sensitivity on entry would need to measure, market-by-market, the underlying price sensitivity primitives, which is one role of our structural model of demand.

[38] We do not provide a test for the effect of price sensitivity on entry; we leave this as a topic for future research.

[39] The timed phase-out of the reinsurance program after 2016 was determined in the original ACA legislation in 2010 and was common knowledge to all market participants.

$$ATC_j = AC_j(P) - RE_j(P)$$

Where $AC_j(P)$ denotes average claims cost per enrollee and and $RE_j(P)$ denotes the reinsurance payment given to firm j per enrollee, with fixed costs omitted for simplicity. We can further break down the reinsurance term as:

$$RE_j(P) = \frac{c}{D_j(P)} \times \sum_{i \in I(m)} y_{ij}(P) \times \mathbf{1}\{C_{ij} > \$45,000\} \times \min(\$250,000, C_{ij} - \$45,000)$$

Where c is the share of costs reimbursed each year,[40] $I(m)$ represents the set of enrollees in market m, $y_{ij}(P)$ denotes whether enrollee i chose plan j, and C_{ij} is the total annual cost of the enrollee to the insurer. Importantly, we can calculate $RE_j(P)$ at the state-by-insurer (i.e., "issuer") level, separately for 2014, 2015, and 2016.

The crucial empirical object — the slope of the average cost curve — can then be obtained by differentiating ATC_j with respect to P_j:

$$\frac{\partial ATC_j(P)}{\partial P_j} = \frac{\partial ATC_j(P)}{\partial P_j} - \frac{\partial RE_j}{\partial P_j}$$

As discussed in Section 2 of the main text, adverse selection in pricing occurs when $\frac{\partial ATC_j}{\partial P_j}$ is positive. Here, the first term $\frac{\partial AC(P)_j}{\partial P_j}$ represents adverse selection in terms of claims costs. The second term, $\frac{\partial RE_j(P)}{\partial P_j}$, encodes the ability of reinsurance to flatten the average cost curve (much like conventional risk adjustment). Taken together, the equation shows that an increase in P_j leads to an increase in claims costs (from enrolling disproportionately more high-cost individuals), but also an increase in reinsurance as the plan enrolls more individuals with costs above the attachment point.[41]

From the above equation, we can see that places where reinsurance payments respond strongly to price (large, positive values of $\frac{\partial RE_j(P)}{\partial P_j}$) are those where average cost curves will steepen more after the removal of reinsurance (after which we have $\frac{\partial ATC_j}{\partial P_j} = \frac{\partial AC(P)_j}{\partial P_j}$). Our model predicts that more insurer exit will occur in these markets. Specifically, our model predicts insurer exit where $\frac{\partial ATC_j}{\partial P_j}$ (including all risk adjustment, fixed costs, etc.) is less than 1 with reinsurance, but would be greater than 1 after the removal of reinsurance.

In the next section, we describe how we obtain state-specific estimates of $\frac{\partial RE_j(P)}{\partial P_j}$. We then show how we relate these estimates to various measures of insurer participation before vs. after the removal of reinsurance.

[40] This was 100% in 2014, 55% in 2015, and 53% in 2016 ()

[41] Under reasonable assumptions, $\frac{\partial RE_j(P)}{\partial P_j}$ must be weakly positive. For example, if price increases shift the C_{ij} distribution to the right, then $\frac{\partial RE_j(P)}{\partial P_j}$ is positive for distributions of any shape, as long as the support of the distribution includes the \$45K to \$250K region.

F.3 Estimation details

F.3.1 Estimating $\partial RE_j(P)/\partial P_j$ at the state level

To compute how reinsurance responds to premium changes, we first need to compile data on premiums and reinsurance at the issuer-state-year level for 2014-2016.[42]

Premium data: For premiums, we start with the ACA landscape files, which are at the plan-county level. We first limit to 1 observation per plan per rating area (only silver plans), since premiums do not vary within counties in the same rating area. There are 8 plans per issuer per rating area on average. We merge with plan-level enrollment data and keep the largest (participating) plan per issuer per rating area.[43] Starting with the individual premium for each plan, we subtract the cheapest plan's premium for each rating area. Finally, we collapse these relative premiums across rating areas within a state by weighting each rating area by its below-65 population shares (using 2010 Census population data). This gives us a measure of P_{jt}, or the premium for issuer j in year t.

Reinsurance data: We use reinsurance payment data at the insurer-state-year level from 2014-16, merged with the premium data above. We limit to 38 states with premium data and further to the 34 states with multiple years of data.[44] We divide by the number of enrollees at the issuer-year level, defined as the total "ever-enrolled" minus the total number who dis-enroll in a given year. This gives us a measure of RE_{jt}, the per-enrollee reinsurance payment for plan j in year t.

Note, both premiums and especially reinsurance are skewed variables with outliers, so we log transform both relative premiums P_{jt} and reinsurance RE_{jt}.[45]

Estimating the effect of premium changes on reinsurance: The baseline specification for estimating $\partial RE_j(P)/\partial P_j$ at the state level is as follows:

$$RE_{jst} = \delta_s \times P_{jst} + \eta_j + \eta_s + \eta_t + \eta_{jst} \tag{19}$$

Where $\delta_s \equiv \frac{dRE}{dP}_s$ is the coefficient of interest that describes the (state-specific) relationship between relative premium P_{jst} and per-enrollee reinsurance payment RE_{jst} for issuer j in state s and year t. The η_j, η_s, and η_t fixed effects control for fixed differences in reinsurance across

[42] The term "issuer" refers to a unique insurer-state combination.

[43] The enrollment data is at the issuer-state-year level, so we don't know if the largest plan is largest for all rating areas within a state. For each issuer-rating area combination, we compute the share of "ever-enrolled" individuals in each of that issuer's participating plans. There is 1 rating area with 2 equally large plans; we keep the cheapest plan in that rating area. In some cases, the largest plan for an issuer does not participate in all rating areas. In these cases, we keep the largest participating plan in each rating area, so different rating areas in a state may have different "largest plans" for the same issuer.

[44] About 11% of reinsurance payments are for issuers that do not exist in the ACA landscape files. About half of ACA-compliant plans are "off-exchange plans" sold by brokers. These tended to enroll wealthier individuals because they do not qualify for ACA subsidies.

[45] Because relative premiums are $0 for the lowest-cost plan, we use the transform $log(1+P_{jt})$. In practice, dropping observations with $0 relative premiums does not significantly change the results.

issuers and states,[46] and account for the overall decline in reinsurance payments over time from 2014-2016.[47] Identification comes from changes in the relative premium *within* an issuer over time.

We weight the above regression by the number of enrollees in each issuer-state-year observation, and use robust standard errors.[48]

We also estimate the following aggregate specification (replacing δ_s with δ), to obtain the average reinsurance slope over all states:

$$RE_{jst} = \delta \times P_{jst} + \tilde{\eta}_j + \tilde{\eta}_s + \tilde{\eta}_t + \tilde{\eta}_{jst} \tag{20}$$

Here, we cluster standard errors at the state level.[49] A positive estimate of δ, if causally identified, implies that insurers raising their premiums can expect to have larger reinsurance payments (because higher premiums attract more higher-cost enrollees, some of whom will qualify for reinsurance payments).

Applying Bayesian shrinkage to state-specific slopes: Finally, we perform Bayesian shrinkage to the (n=34) state-specific slopes using the estimator of Morris (1983) and code from Chandra, Finkelstein, Sacarny, and Syverson (2015).

F.3.2 Firm participation regressions

We start with data on issuer participation at the county-year level from 2014-2018. We focus on issuers operating silver plans in each county.[50] In the ACA marketplaces, all county-year observations have at least one issuer. We consider several measures of participation, including the number of issuers per county, log number of issuers per county, number of counties with 5 or more firms, and the number of counties that have no more than 1, 2, or 3 firms. This gives us six participation outcome variables in total.

We merge the county-level participation data with our state-level estimates of $\partial RE_j(P)/\partial P_j$ and run the following regression:[51]

$$Y_{cst} = \beta \times \text{post} \times \frac{\partial RE}{\partial P}_s + \alpha_s + \alpha_t + \varepsilon_{cst} \tag{21}$$

Where *post* is an indicator variable for years 2017-18, and $\frac{\partial RE}{\partial P}_s$ is the state-specific reinsurance slope after Bayesian shrinkage is applied. α_s and α_t are state and year fixed effects, respectively, and ε_{cst} is the error term.[52] The coefficient of interest β encodes the effect of reinsurance on

[46]In practice, state fixed effects are redundant with the issuer fixed effects, since each state includes multiple issuers, with each issuer mapping to a unique state.

[47]The decline in payments is a function of the coinsurance rate c.

[48]Here, the specific method of calculating standard errors in the above regression is important as it affects the Bayesian shrinkage procedure in the next step.

[49]State-level clustering is not possible in the specification where δ_s is state-specific.

[50]In practice, issuers participating in a given county are required to operate plans of all metal levels in that county.

[51]We run both unweighted and weighted versions. The weights are defined as the average number of enrollees in each state from 2014-16, using data from the Kaiser Family Foundation on state-specific marketplace enrollment.

[52]Including county-level fixed effects would not affect the estimation of β, since $\frac{\partial RE}{\partial P}_s$ only varies at the state level.

participation. We estimate the model via OLS and cluster standard errors at the county level.

F.4 Results - Estimating Reinsurance Slopes

F.4.1 Insurer participation in the ACA over time:

The figure below shows the number of issuers per county, separately by year. We can clearly see a large decline in participation in 2017-18, after reinsurance was removed after 2016. For example, there were more than 1,500 monopoly counties in 2018, compared to fewer than 250 in 2016. The timing of this wave of insurer exit is suggestive of a role for reinsurance for maintaining competition. However, further analysis is required to argue that reinsurance (1) played a causal role and (2) exerted its effect via the channel identified in our model (i.e., through its effect on adverse selection and the slope of the average cost curve).[53]

Appendix Figure A7. Issuer Participation Over Time

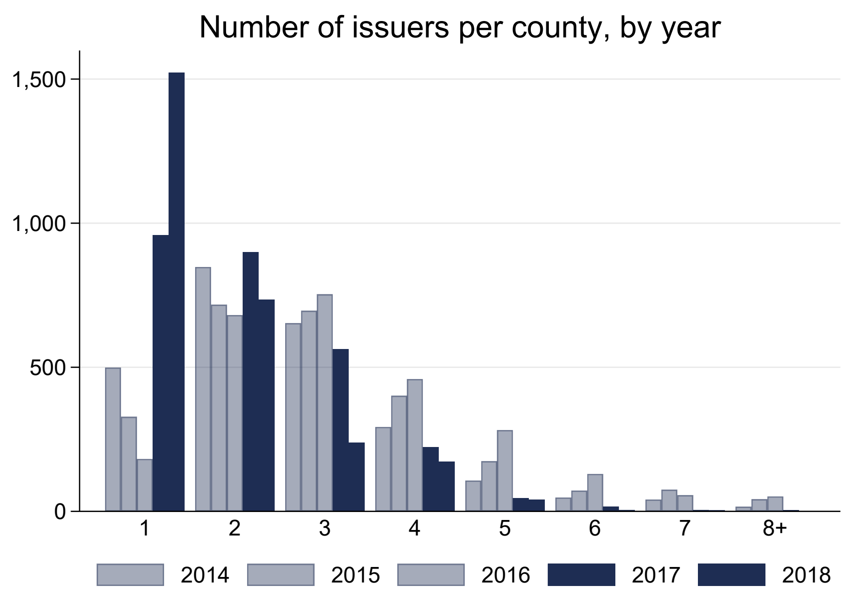

Notes: Figure shows the number of counties with 1, 2, ..., up to 8+ issuers. Within each group, bars (left to right) reflect the year (2014-2018).

F.4.2 Relationship between reinsurance and premiums:

Figure A8 below shows a positive relationship between (log) reinsurance per enrollee (log) relative premium. The left figure is a raw scatterplot, with no controls or fixed effects, where each circle

[53] In particular, one concern is that the reinsurance program in the ACA constituted a large *net* subsidy to marketplace plans. To address this, we focus on the *slope* of reinsurance with respect to price and use a difference-in-difference framework that not only compares pre- vs. post- 2016, but also compares across different states. Year and issuer fixed effects are used to control for fixed differences in the net subsidy of reinsurance across states.

represents an issuer-state-year observation, weighted by total enrollees. The right fiure plots the same data as a binned scatterplot, where both (log) reinsurance and (log) relative premium have been residualized on issuer, state, and year fixed effects.

Appendix Figure A8. Reinsurance per enrollee vs Premiums

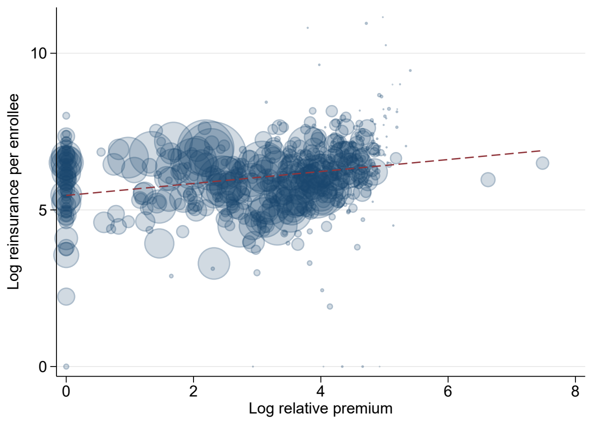

(a) Raw scatterplot, no controls

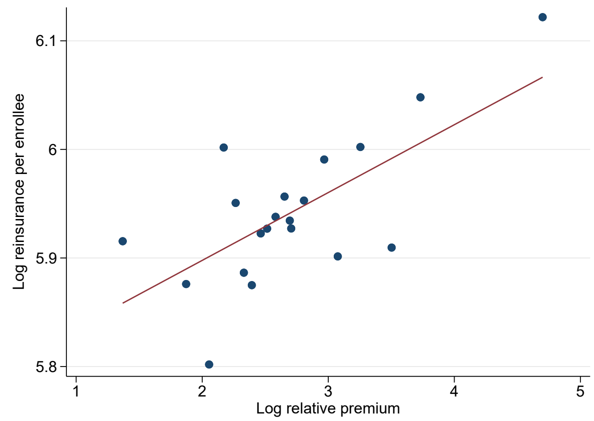

(b) Binned with controls

Notes: Figure shows a positive relationship between log(1 + reinsurance per enrollee) and log(1+ relative premium) at the issuer-state-year level, weighted by number of enrollees. Panel (a) shows a scatterplot of the raw data, with circle sizes reflecting the number of enrollees underlying each observation. Panel (b) shows a binned scatterplot where both variables are residualized on a set of year, state, and issuer fixed effects before plotting.

Table A2 below shows estimates for the specification in Equation 20, with different sets of fixed effects. Consistent with adverse selection, the sign of the reinsurance slope is positive across all specifications. In our preferred specification (column 5, with state, year, and issuer fixed effects), the elasticity of reinsurance with respect to premium is 0.062. Given an average relative premium of \$52 and average reinsurance per person of \$1,056, this elasticity corresponds to a \$1.28 increase in reinsurance for every \$1 increase in premiums. This suggests that reinsurance played a significant role in flattening the average cost curve in the ACA marketplaces prior to its expiration in 2016. Hence, removing reinsurance in 2016 resulted in a steepening of average cost curves; the magnitude of this effect varies from $\Delta dAC/dP = 1.02$ to 5.28, depending on the specification, with generally more modest effects when controlling for the full set of fixed effects.

Appendix Table A2. Estimates of the Aggregate Reinsurance Slope

(a) Weighted

	log Reinsurance per Enrollee				
log Relative Premium	0.050	0.137 **	0.106 **	0.216 **	0.062 **
	(0.031)	(0.043)	(0.032)	(0.028)	(0.021)
dReinsurance/dPremium	1.023	2.803	2.175	4.411	1.277
State FE		X		X	X
Year FE			X	X	X
Issuer FE					X

(b) Unweighted

	log Reinsurance per Enrollee				
log Relative Premium	0.191	0.258 **	0.182 **	0.251 **	0.109 *
	(0.035)	(0.038)	(0.033)	(0.035)	(0.042)
dReinsurance/dPremium	3.906	5.280	3.721	5.131	2.230
State FE		X		X	X
Year FE			X	X	X
Issuer FE					X

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

Notes: Table shows estimates of the aggregate reinsurance slope. The coefficient on log Relative Premium can be interpreted as the elasticity of reinsurance per enrollee with respect to relative premiums. Standard errors are shown in parentheses. The row titled "dReinsurance/dPremium" re-expresses the elasticity estimates as $ of reinsurance per enrollee for every $1 increase in relative premiums. Panel (a) reports regressions weighted by the number of enrollees in each issuer-state-year. Panel (b) reports unweighted regressions.

These results are consistent with the expiration of reinsurance playing a role in the decline in firm participation after 2016. of In the following section, we use state-specific reinsurance slope estimates to provide further evidence for a causal relationship between reinsurance and firm participation.

F.4.3 State-specific reinsurance slopes:

In this section, we show estimates of Equation 19, focusing on specifications that include issuer, state, and year fixed effects. We do this to ensure that our identifying variation comes from *within-issuer* variation over time.

Figure A9 below plots the state-specific coefficients for weighted and unweighted specifications. The median premium elasticity of reinsurance is 0.065 for the weighted specification and 0.095 for the unweighted specification. These values are consistent with the results shown previously in Table

A2. The unweighted specification gives a larger standard deviation for the estimates, potentially indicating that weighting by number of enrollees reduces sampling error. Reassuringly, we see that most of the state-specific slopes are positive.

Appendix Figure A9. State-specific slopes: reinsurance per enrollee vs premiums

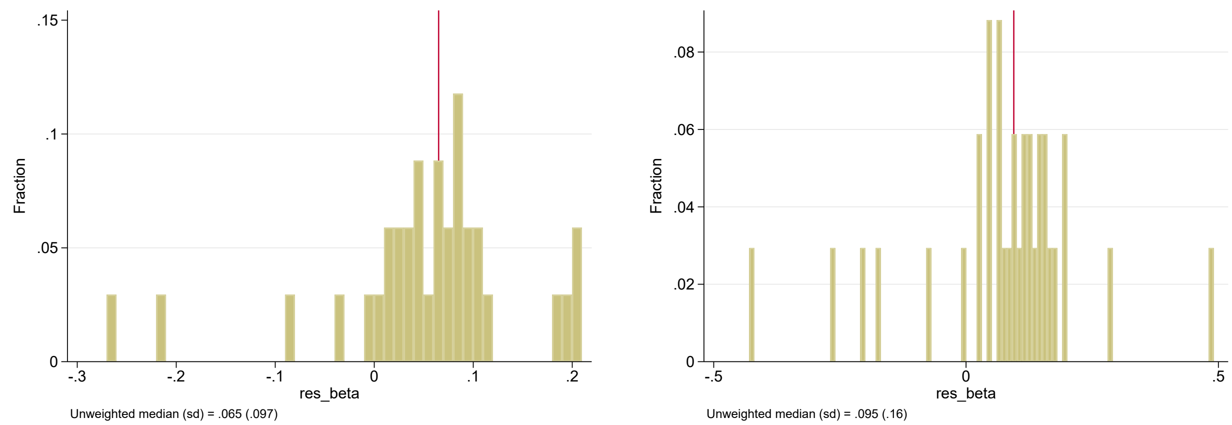

(a) Weighted (b) Unweighted

Notes: Figure shows the distribution of $\hat{\delta}_s$ from Equation 19. Plotted estimates are for $n = 34$ states using Healthcare.gov over multiple years. Regressions use log(1 + reinsurance per enrollee) and log(1+ relative premium) with state, year, and issuer fixed effects. Panel (a) shows results where the regression is weighted by the number of enrollees for each issuer-state-year observation. Panel (b) shows the unweighted regression. Bin widths are set to 0.01 for both figures.

F.4.4 Applying Bayesian shrinkage to the state-specific estimates:

Next, we apply Bayesian shrinkage to the estimates shown in Figure A9. While the original estimates are obtained using OLS, which is unbiased, certain small states may have extreme estimates as a result of sampling error, resulting in an overly dispersed distribution of reinsurance slope estimates. The shrinkage procedure selectively attenuates estimates towards the mean, based on the standard error of each estimate (estimates with larger standard errors are attenuated more).[54] Using the shrunk estimates in our participation regressions (Equation 21) also allows us to account for this differential sampling error.

Figure A10 below mirrors Figure A9, but shows the distribution of the estimates after shrinkage is applied. We see that dispersion is reduced for both weighted and unweighted distributions, with a much larger reduction in dispersion for the unweighted results (again, this likely reflects greater sampling error when weights are not used). The following Figure A11 shows the effects of shrinkage on each state-specific estimate, where shrinkage is represented as a rotation from the 45-degree line to the X-axis. Again, the unweighted regression in Panel (b) shows a much larger rotation, especially

[54] This attenuation reduces the mean-squared error of the estimates at the expense of introducing some bias.

for more extreme original estimates.

Appendix Figure A10. State-specific slopes: after shrinkage procedure

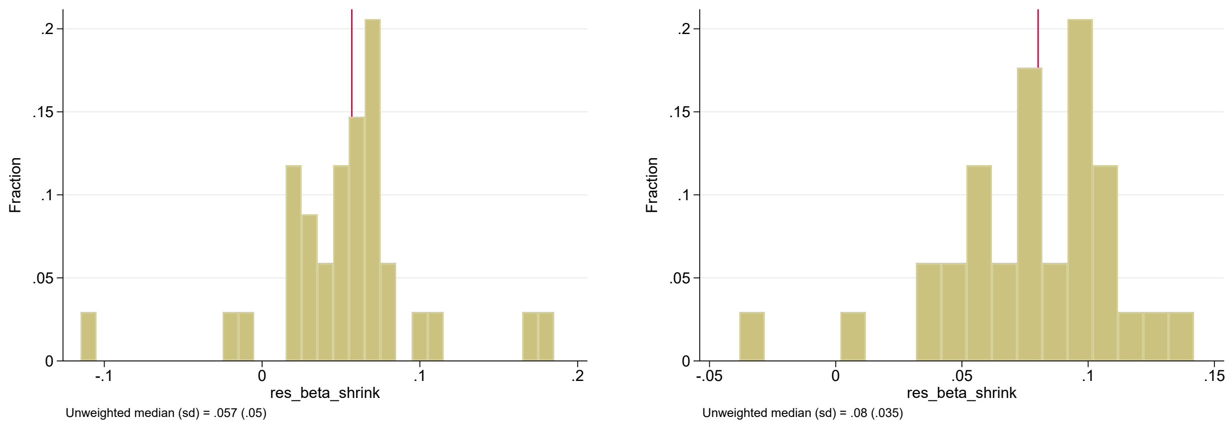

(a) Weighted (b) Unweighted

Notes: Figure shows the distribution of $\hat{\delta}_s$ from Equation 19 after the Bayesian shrinkage procedure has been applied. Plotted estimates are for $n = 34$ states using Healthcare.gov over multiple years. Regressions use log(1 + reinsurance per enrollee) and log(1+ relative premium) with state, year, and issuer fixed effects. Panel (a) shows results where the regression is weighted by the number of enrollees for each issuer-state-year observation. Panel (b) shows the unweighted regression. Bin widths are set to 0.01 for both figures.

Appendix Figure A11. Visualizing the effects of the shrinkage procedure

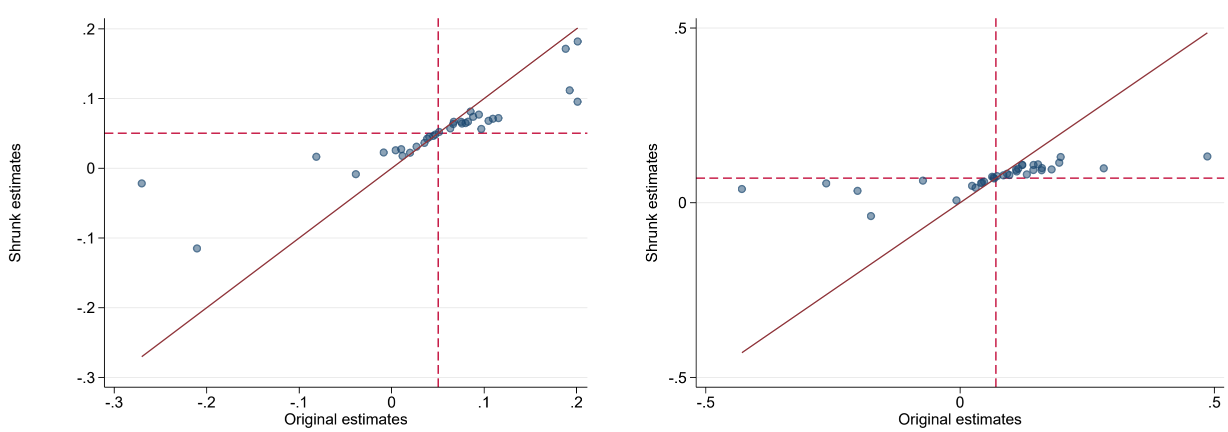

(a) Weighted (b) Unweighted

Notes: Figure plots shrunk estimates vs. original estimates of $\hat{\delta}_s$ from Equation 19. Panel (a) shows results where the underlying regression is weighted by the number of enrollees in each issuer-state-year. Panel (b) shows results for the unweighted specification.

F.5 Effects of Reinsurance on Firm Participation

We now return to the original test of our model: does increasing the slope of the average cost curve result in firm exit? Assuming that our state-specific estimates of dRE/dP are well identified, we can evaluate this question using the difference-in-differences specification given in Equation 21. While reinsurance expired nationwide after 2016, states with larger reinsurance slopes should experience more firm exit.

As detailed in the following sections, this is indeed what we find. Across all measures of participation, and in both weighted and unweighted regressions, counties in states with larger reinsurance slopes experienced larger declines in firm participation. This is consistent with our model, which ties firm participation to the slope of the average cost curve.

Firm participation outcome variables: We examine participation at the *county* level, because that is the relevant geographic unit where firm participation decisions are allowed to vary.[55] We consider 6 outcome variables: the number of issuers per county, the log number of issuers, and indicators for whether the county has ≥ 5 issuers, ≤ 3 issuers (triopoly), ≤ 2 issuers (duopoly), or only 1 issuer (monopoly).

Reinsurance slope estimates: In the specifications that follow, we use shrunk estimates of dRE/dP from the enrollee-weighted results (i.e., the estimates shown in Panel (a) of Figure A10). This is motivated by evidence (discussed above) that the weighted regression produces estimates with less sampling error. In addition, the enrollee-weighted estimates offer a more natural interpretation by weighting *enrollees* (rather than issuers) equally.

Weighting: Note that the choice of weights for the participation regression (Equation 21) is independent from the weighting scheme used for estimating the state-specific slopes (Equation 19). In the following results, we will consider two types of *state-level* weights when estimating equation 21. First, a weighted version weights states by their total pre-period ACA enrollment (Kaiser Family Foundation), averaged over 2014-16. This gives more weight to participation changes for counties in larger states. The second version is an unweighted specification: here, estimated effects treat all counties equally. We do not weight by county size, since this would upweight larger counties where insurer exit is generally less prevalent.

Clustering: We cluster standard errors at the county level. Clustering at the state level results in larger standard errors, but our results for N issuers and log N issuers remain statistically significant at the 1% and 5% levels respectively.

F.5.1 Firm participation results:

Table A3 below shows estimates from equation 21, for the outcomes and weighting schemes specified above. Across all specifications, we see that larger reinsurance slopes are associated with greater exit after the expiration of reinsurance. The reported coefficients can be interpreted as the effect

[55] In contrast, pricing decisions are only allowed to vary at the rating area level, where rating areas usually combine many counties.

of suddenly increasing the slope of the average cost curve dAC/dP by 1. In the unweighted specification, a increase in the average cost curve slope of 0.1 leads to -0.44 fewer insurers (-15% in the log specification), a 9.7pp decline in the probability of having ≥ 5 insurers, a +9.1pp chance of having ≤ 3 insurers, a +6.2pp chance of having ≤ 2 insurers, and a +9.3pp chance of monopoly. Put another way, the effect of reinsurance on the slope of the average cost curve can explain between 10% to 42% of the overall decrease in participation between 2014-16 and 2017-18.

Appendix Table A3. Effect of Reinsurance Expiration on Insurer Exit in the ACA

(a) Unweighted

	Firm Participation					
	N Issuers	log N Issuers	$\geq$ 5 Issuers	$\leq$ 3 Issuers	$\leq$ 2 Issuers	Monopoly
Post X Reinsurance Slope	-4.460 **	-1.613 **	-0.918 **	0.943 **	0.712 **	1.017 **
	(0.258)	(0.102)	(0.080)	(0.098)	(0.115)	(0.115)
Effect of 1SD Slope Incr.	-0.256	-0.093	-0.053	0.054	0.041	0.058
Share Pre-Post Explained	0.209	0.180	0.396	0.238	0.110	0.164

(b) Weighted

	Firm Participation					
	N Issuers	log N Issuers	$\geq$ 5 Issuers	$\leq$ 3 Issuers	$\leq$ 2 Issuers	Monopoly
Post X Reinsurance Slope	-4.850 **	-2.094 **	-0.768 **	0.655 **	0.912 **	1.800 **
	(0.462)	(0.171)	(0.137)	(0.168)	(0.186)	(0.182)
Effect of 1SD Slope Incr.	-0.279	-0.120	-0.044	0.038	0.052	0.103
Share Pre-Post Explained	0.237	0.253	0.312	0.159	0.153	0.338

** p<0.01, * p<0.05

Notes: Table shows estimates of Equation 21 for various participation outcomes. The coefficient on Post X Reinsurance Slope indicates the effect of a 1 unit increase in the slope of dRE/dP, or equivalently, a 1 unit increase in the slope of the average total cost curve $dATC/dP$. Standard errors are shown in parentheses. The row "Effect of 1SD Slope Incr." denotes the effect of a 1 standard deviation increase in the state-specific reinsurance slope dRE/dP. The row "Share Pre-Post Explained" computes the share of the overall decrease in the participation outcome variable that can be explained by the Post X Reinsurance Slope coefficient. Panel (a) reports the unweighted regression. Panel (b) reports the regression where states are weighted by their average ACA enrollment over the pre-period (2014-2016).

The following Figure A12 visualizes the data underlying Table A3. For each of the outcome variables, the figure plots a binned scatterplot of firm participation against estimated reinsurance slopes. The raw (binned) data are shown without weights or fixed effects. Because we omit year fixed effects, we can see that firm participation drops substantially in the post-period (2017-18) relative to the pre-period (2014-16) across the board, but the largest decreases in participation are for counties that had large reinsurance slopes in the pre-period. Because we omit state fixed effects, we can see that the states with large pre-period reinsurance slopes actually had *more* robust insurer participation. In some specifications, this pattern is not only reduced, but is reversed after the removal of reinsurance.

Appendix Figure A12. Caption for the figure

(a) N Issuers

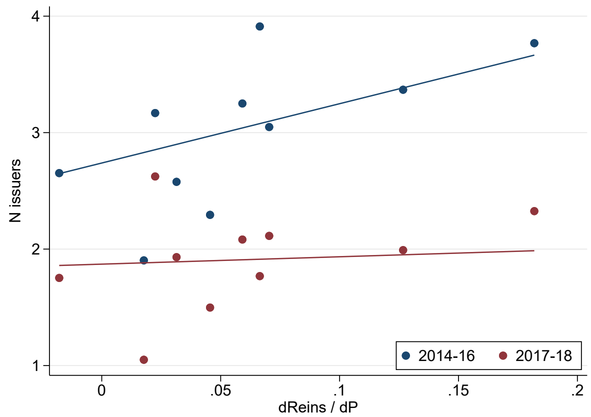

(b) Log N Issuers

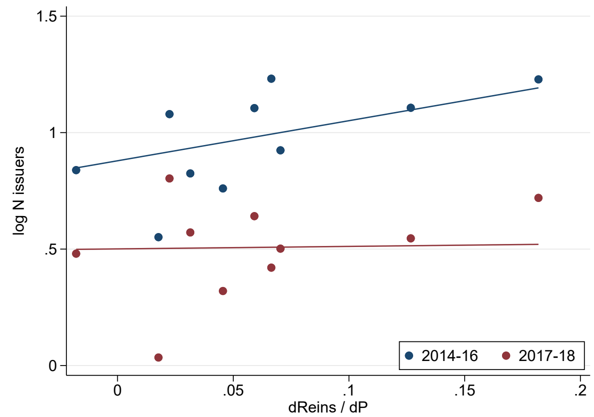

(c) 5+ Issuers

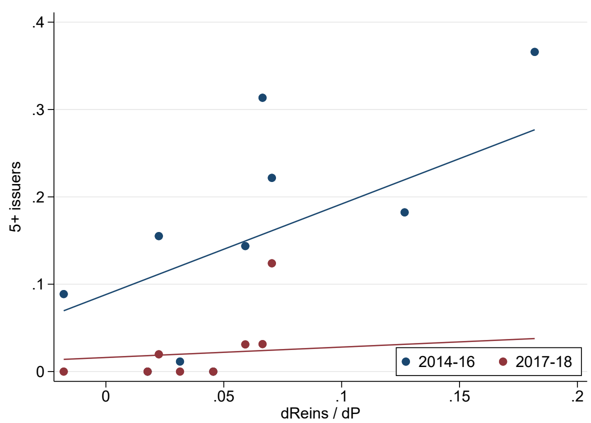

(d) ≤ 3 Issuers

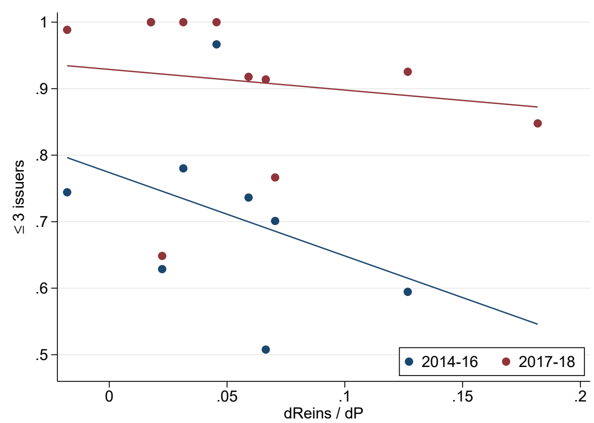

(e) ≤ 2 Issuers

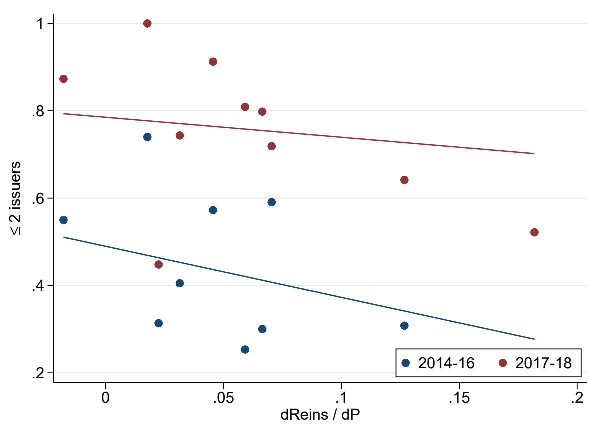

(d) Monopoly Issuers

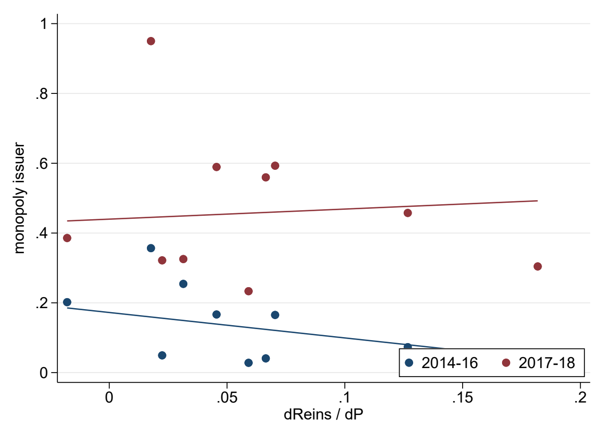

Notes: Figures show how county-level firm participation changes after reinsurance is removed from the ACA in 2016. The Y-axes plot one of 6 measures of firm participation, whereas the X-axes show the magnitude of the state-specific reinsurance slope for each of n =34 states (as shown in Figure A10). Each panel shows results for a different participation measure.

G Additional Exhibits

Appendix Figure A13. CommCare Plans Pre-Subsidy Prices

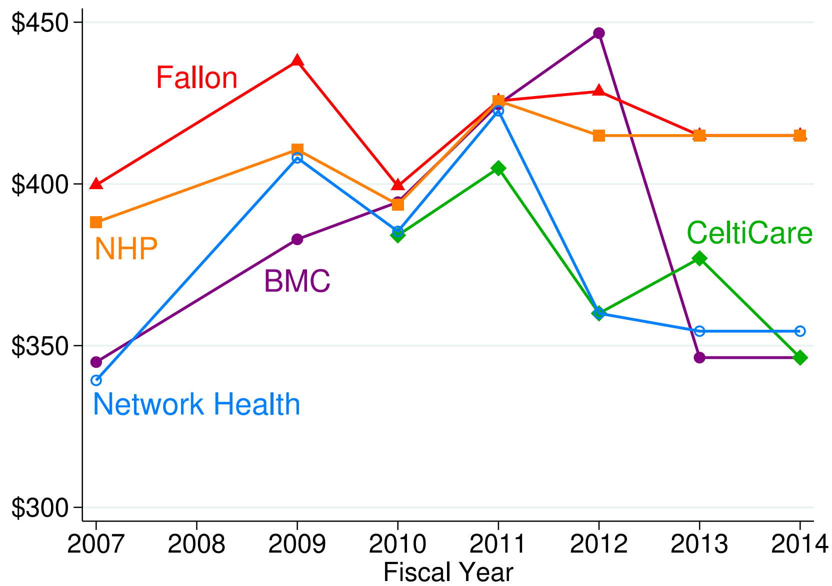

Note: The graphs show average pre-subsidy insurer prices for each insurer's plan in the CommCare market, by fiscal year. The five plans are shown in different colors and labeled. Values shown are averages for the plan's actual enrollees; underlying premiums and (in some years) prices vary by income group 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 A14. Premium Variation Example: Network Health (Boston region), 2010-13

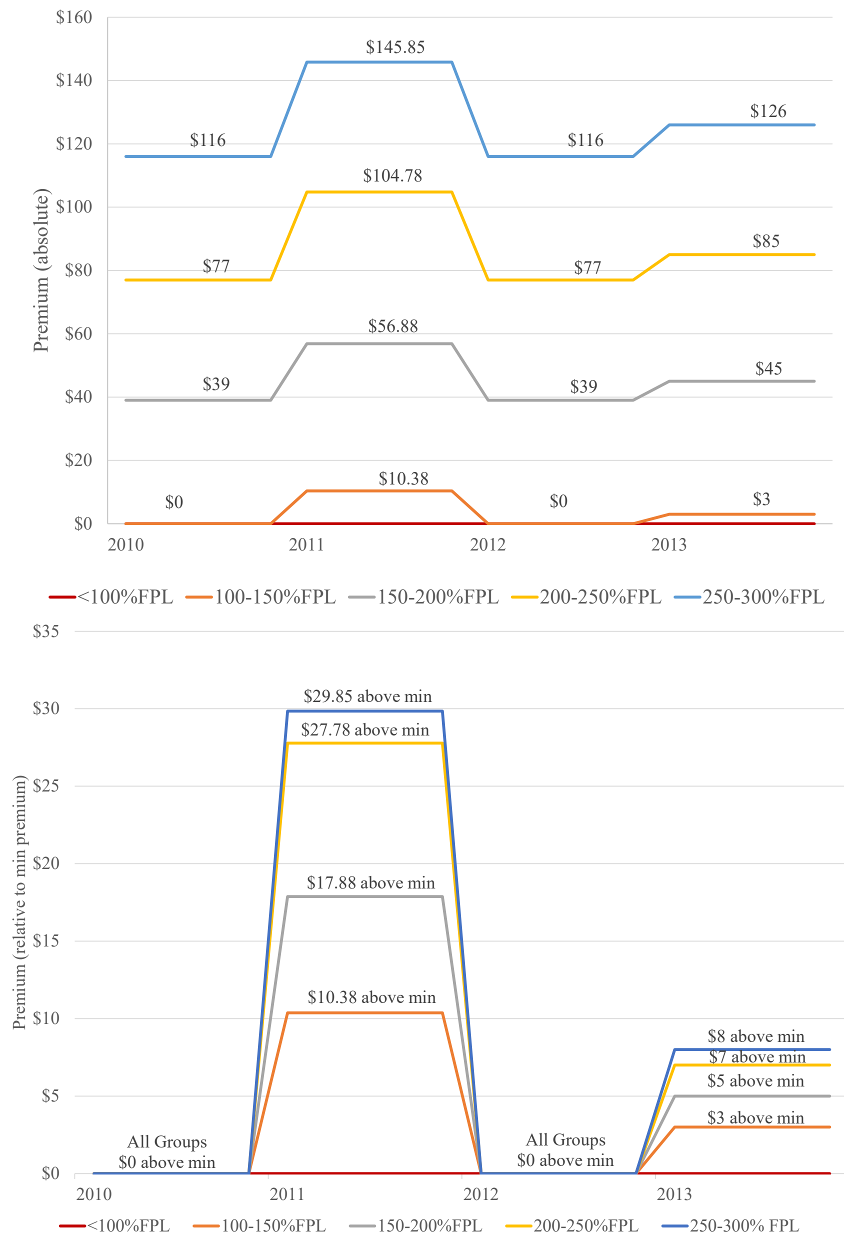

Note: The graphs shows the example of Network Health's (post-subsidy) enrollee premiums by income group over the 2010-2013 CommCare years. "FPL" refers to the federal poverty level. Pre-subsidy prices (and enrollee premiums) vary at the regional level in 2010, and the graph shows premiums specifically 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 A15. Hospital Coverage in Massachusetts Exchange Plans

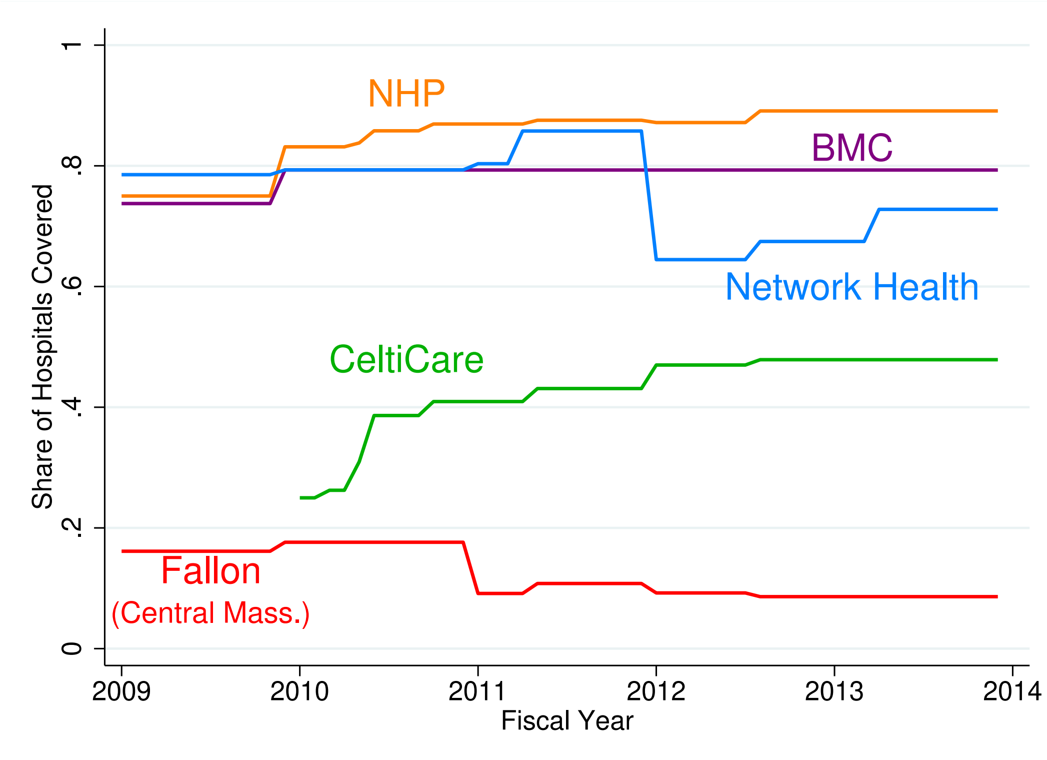

Note: The graph shows the shares of Massachusetts hospitals covered by each CommCare plan, where shares are weighted by hospital bed size in 2011. Fallon's hospital coverage share is much lower than other plans largely because it mainly operates in central Massachusetts and therefore does not have a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