Adverse Selection Pricing and Unraveling of Competition in Insurance Markets *

Edward Kong
Harvard University

Timothy J. Layton
Harvard and NBER

Mark Shepard
Harvard and NBER

August 11, 2022

Abstract

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

1 Introduction

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

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

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

In this paper, we study an additional feature of these markets that pushes against efficiency. The feature we study is a firm strategic pricing response to a type of consumer choice behavior that is related to but distinct from classic notions of adverse selection (Akerlof, 1970; Rothschild and Stiglitz, 1976). In the public economics literature, adverse selection is typically thought of as playing out in a vertical setting, where sick consumers have a stronger preference for more generous insurance than healthy consumers. Under a typical notion of perfect competition, this preference causes the prices of more generous insurance plans to be 'too high' and sometimes causes those plans to death spiral altogether (Einav et al., 2010*a*; Cutler and Reber, 1998). We refer to this type of consumer behavior as 'vertical selection.' We consider a different type of setting and consumer behavior, one where the preferences of sick and healthy consumers still differ in important ways, but where plans are more generally (horizontally) differentiated products that compete strategically on price. In this type of setting there may be a different type of horizontal selection, characterized by strong correlations between price sensitivity and risk, where healthy consumers will often be more price sensitive than sick ones. We show that in this type of setting it is important to consider firms' incentives to price strategically. Indeed, rather than taking plan demand and cost curves as market primitives, we bring insights from the industrial organization literature that consider these curves to be equilibrium objects, where demand and costs for a particular plan depend not only on that plan's price but also on the prices of the other plans in the market. We show that this is an important feature to incorporate into models of markets featuring horizontal selection, as strong correlations between price sensitivity and risk don't just influence plan prices or contract design but they cause plans to have strong price undercutting incentives that can destablize markets via downward price spirals. Ultimately, in contrast to the prior literature on vertical selection where it is well known the competitive equilibria are often characterized by distorted prices and distorted contracts, we show that strategic firm pricing in the presence of horizontal selection can cause equilibria where multiple plans participate in a market *to not exist*.

To fix ideas, consider a large urban market where there are two large hospital systems.

One system dominates the northern part of the city while the other system dominates the southern part of the city, though both systems have some presence throughout the city. While two *vertically* differentiated plans might consist of one plan that includes both hospital systems and one plan that includes only one, two *horizontally* differentiated plans would each cover one system. Consumers have heterogeneous preferences, with consumers living in the north on average preferring the plan that covers the northern hospital system and consumers living in the south on average preferring the plan that covers the southern hospital system. Selection may seem like a minor concern in such a market, as preferences are correlated with some dimension (geography) that is unlikely to be highly correlated with underlying health risk. However, geography is not the only dimension of preference heterogeneity. In such a market, it is likely that sicker consumers care deeply about which hospital system is in their plan's network – they interact with the healthcare system often, so they want to be sure that their preferred (i.e., nearby) providers are in the network. Healthier consumers, on the other hand, may care little about which hospital system is included in their plan's network – they rarely interact with the healthcare system, making them much more sensitive to price than to the network.

While this type of horizontal differentiation would typically result in low price sensitivity and an ability for plans to charge high mark-ups, in this setting the highly price sensitive healthy consumers prevent this from happening. If the northern plan charges a slightly higher price than the southern plan, the northern plan will lose many of its healthy price sensitive consumers, driving up its average cost and causing it to lose money. The northern plan thus has a strong incentive to *undercut* the southern plan and take the healthy consumers, thus disciplining the plan's desire to charge high mark-ups, a result that has been pointed out in the prior literature Starc (2014); Mahoney and Weyl (2017*a*); Einav et al. (2018). But, the key new insight here is that the southern plan has the same incentive—after the northern plan lowers its price, the southern plan still faces a strong incentive to drop its price even lower. If the correlation between price sensitivity and risk is strong enough, this incentive remains even after the plans are pricing below the market-wide average cost, as either plan can do better by pricing just below their competitor and taking all of the healthy enrollees in the market, driving their plan-specific average cost below the market-wide average. There is thus a kind of "race to the bottom" in price. Importantly, this race can often go "below the bottom" and result in no pure strategy Nash equilibrium existing where both plans are offered and earn non-negative profits. Thus, instead of disciplining the mark-up, selection leads to the elimination of competitive equilibria with multiple firms participating in the market. In these types of settings, the only valid equilibria may thus be ones where only one of the two plans is offered, resulting in losses to consumer surplus due to (1) the loss of one of the two plan options and (2) higher mark-ups that can now be charged by the remaining (monopolist) plan. Indeed, everyone in this market would

be better off if the two plans could coordinate on price and split the market, but such an equilibrium is not stable, as both plans have a strong incentive to undercut the other in order to attract the entire healthy group to themselves.

We formalize this idea by specifying a general model of an insurance market where multiple firms decide whether to offer fixed contracts (on which they then compete on prices) or to refrain from entering the market. We use the model to draw a novel connection between the economics of firm behavior under fixed costs and the economics of firm behavior under adverse selection. It is widely known that fixed costs lead to increasing returns to scale via a downward-sloping average cost curve. In such settings, firms have an incentive to undercut their competitors in order to increase their market share and spread the fixed costs across a larger set of consumers. This undercutting incentive may be difficult for the market to sustain many competitors due to the need to spread the firm-specific fixed costs of production across as many purchasers as possible. In the extreme, the market produces a natural monopoly. We show that markets with strong horizontal selection exhibit similar properties. It is widely known that markets with vertical selection also feature downward-sloping average cost curves at the market level(Einav et al., 2010*b*). What has not been previously understood, however, is that horizontal selection can also lead to *firm-level* downward-sloping average cost curves that, like in the case of fixed costs, can make it difficult for a market to sustain many competitors.

To assess the empirical importance of this conceptual possibility, we turn to data from Massachusetts Commonwealth Care (CommCare) insurance exchange, the state's pre-ACA health insurance marketplace for low-income people. The market consisted of 4-5 insurers that each offered a single plan with standardized cost sharing provisions but that differed in provider networks and drug formularies. Our model suggests that the strength of the undercutting incentive at a given set of prices is related to the degree of consumer price sensitivity and the slope of the firm-level average cost curve. We thus start by providing reduced form estimates of these parameters by leveraging two sources of exogenous, subsidy-driven variation in plan prices. First, we leverage a major change in subsidy structure for consumers with incomes above vs. below 100% of the federal poverty line (FPL). Consumers below 100% of FPL are fully subsidized, paying \$0 for all plans. Consumers just above 100% of FPL have the same plan choices but pay \$0 only for the *lowest-price* plan and must pay modestly higher premiums for higher-price plans. Using a regression discontinuity (RD) design, we show that even this small increase in premiums ($\approx$ \$12 per month) leads to large shifts in market shares toward the cheapest option. Specifically, an average premium difference of \$12 increases the cheapest plan's market share by an average of 20 percentage points (or 50-100% of the baseline market share). These results reveal that a subset of consumers in this market are extremely price sensitive, a key parameter for determining the effects of selection on firm participation.

Second, we implement a difference-in-differences (DD) design leveraging year-to-year premium changes faced by consumers with incomes above 100% of FPL (treatment group) but not by fully-subsidized consumers below 100% of FPL (control group), for whom all plans are free. We stack many market-year "experiments," leveraging only within-experiment (within-market-year treatment versus control) variation in premiums to avoid recent problems pointed out in the DD literature (Goodman-Bacon, 2019; de Chaisemartin and D'Haultfoeuille, 2020; Sun and Abraham, forthcoming). Our estimates again suggest strong price sensitivity, with each $10 increase in monthly premiums reducing market shares among (actively choosing) new enrollees by −24%, and among all enrollees (including inertial continuing enrollees) by 10%. A $10 premium change is quite small — just 2.6% of observed average medical costs of $390 per month — so this evidence again suggests small price changes lead to major shifts in market share.

Importantly, the DD design also allows us to estimate the average slope of the firm-specific average cost curves across firms in the market. We estimate the slope to be approximately equal to 1 on average. This estimate implies that when a firm lowers its price by $1, its average cost decreases by $1. When restricting to new enrollees making active choices, the slope is much higher, equal to 1.8. These slopes are strikingly high. These estimates suggest that, in the absence of any corrective policies, firms in this setting would have extremely strong undercutting incentives, as reductions in price lead to large reductions in average cost and large increases in market share.

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

We then use the model to show that in the absence of corrective policies, there are no pure strategy Nash pricing equilibria where all firms participate in the Massachusetts market. We then consider a two-stage game where in stage 1 insurers decide whether to enter the market and in stage 2 insurers set Nash-Bertrand prices to maximize profits. We posit an equilibrium concept where a combination of firms and prices is only an equilibrium if no non-entering firm can unilaterally enter and earn positive profits in the stage 2 Nash equilibrium that results when said firm enters. Under this equilibrium definition, we show that in the absence of corrective policies the undercutting incentives are so strong that the only surviving equilibria are ones where only one firm participates and prices at the monopoly price. The issues raised by our theoretical model thus appear to be empirically relevant in this setting.

We then use the model to perform counterfactual simulations exploring the effects of

various corrective policies on market equilibrium. The CommCare market featured two key corrective policies: risk adjustment and price floors.[1] Risk adjustment involves the regulator forcing transfers from firms that enroll healthier-than-average individuals to firms that enroll sicker-than-average individuals. Such transfers weaken the correlation between price sensitivity and risk and flatten the firm-specific average cost curve, weakening undercutting incentives. However, conditional on participation, risk adjustment also limits incentives for a firm to lower its price, potentially limiting gains in consumer surplus (Mahoney and Weyl, 2017*a*). The second policy is a price floor. During much of its existence, the Connector imposed both price ceilings and floors, which were technically applied via rate regulation to ensure actuarial soundness. If set correctly, a price floor could cause equilibria with multiple participating firms to exist, by effectively allowing firms to coordinate on a price at the floor and splitting the group of price-sensitive healthy consumers. Counterintuitively, if in the absence of the floor the only surviving equilibria are ones with fewer firms participating, a price floor could result in *lower* prices than without it, 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 make equilibria with multiple participating firms viable once again. In the absence of risk adjustment, there is no equilibrium where more than one firm participates. In other words, without risk adjustment this market is a natural monopoly. With moderate risk adjustment, however, equilibria with two participating firms survive, and perfect risk adjustment allows equilibria with three participating firms. Not surprisingly, consumer welfare always 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 also estimate welfare improvements but we also often find 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 pressure exerted by on the mark-up by selection. Ultimately, our results suggest that more risk adjustment is better when it comes to welfare. But when it comes to prices, while moderate levels of risk adjustment clearly reduce prices below monopoly levels, too much risk adjustment can lead to higher prices, while also increasing insurer participation.

Finally, we simulate various price floors. First, we show that the optimal price floor is almost never zero. Indeed, it typically exceeds the average cost across all consumers in the population. Price floors are more beneficial in the presence of fixed costs and in the absence of risk adjustment, but they are almost always helpful in terms of consumer welfare. Similar to risk adjustment, however, while sometimes the optimal price floor does actually generate

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

lower prices, other times the optimal price floor raises price above where they would be absent regulation.

Ultimately, we conclude that in the Massachusetts exchange, the undercutting incentives caused by horizontal selection have real and important effects for market stability and, ultimately, on the level of competition in the market. Our results indicate 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, however, allowing the market to support multiple firms who compete with each other on price, resulting in higher consumer welfare. More risk adjustment produces more entry and higher welfare but also often higher prices. Further, we find that non-zero price floors are optimal in this market, with the optimal level of the price floor often exceeding average cost.

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

Related Literature Our paper contributes to several literatures. First, we contribute to work on adverse selection in insurance markets. It has long been recognized that adverse selection can distort prices and contracts (Rothschild and Stiglitz, 1976; Einav et al., 2010*b*; 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 markups in settings with imperfect competition, implying that policies such as risk adjustment can reduce consumer welfare. Our analysis points out that selection may impact the *level* of competition, potentially outweighing the impacts of markups conditional on participation. When market structure is endogenous, corrective policies such as risk adjustment can sustain higher

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

participation and therefore lower markups.

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

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

2 Model

We start by developing a fairly general model of an insurance market where multiple firms decide whether to offer fixed contracts (on which they then compete on prices) or to refrain from competing in the market. Our goal is to derive an expression for insurer profits and show conditions under which adverse selection limits the ability of additional insurers to compete in the market.

Consider a set of potential firms $j \in \{1, ..., J\}$, each with the ability to offer a single contract (or "plan"). A contract consists of non-price attributes (ξ_j) and a single community-rated price that is constant across consumers (P_j). Firms compete in a stylized, two-stage entry game. In stage 1, each firm decides whether to enter the market. Entering involves incurring fixed costs $F \geq 0$ but gives it the opportunity to sell its product and earn profits.[3] Non-entering firms earn zero profits. We denoted entry choices by $E_j \in \{0, 1\}$ and the set of participating firms as $E = \{j : E_j = 1\}$. In stage 2, firms observe who entered, and participating firms set prices (P_j) to maximize profits in standard Nash-Bertrand competition.

There are $i \in \{1, ..., N\}$ consumers. Consumer i incurs expected medical costs C_{ij} in plan j, with costs varying because of heterogeneous medical risk and plan effects on costs. $D_{ij}(P)$ is a demand function that indicates i's probability of choosing plan j given the vector of plan prices, P, and (implicitly) the set of competitors E.[4] Define $D_j(P) \equiv \sum_i D_{ij}(P)$

[3] Fixed costs include items like designing contracts, negotiating provider networks, and running billing systems. Fixed costs are a standard feature of monopolistic competition models (e.g., Salop (1979)), since they help rationalize limited competition. This will be true in our model as well, but we will show that adverse selection can generate limited competition even without fixed costs.

[4] All demand and cost outcomes in our model depend on E; we suppress E from our notation for these

as total demand for plan j. The profits of a participating firm j equal:

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

where $AC_j(P) = \frac{1}{D_j(P)} \sum_i C_{ij} D_{ij}(P)$ is average medical costs of j's consumers. It is also useful to define average *total* costs, $ATC_j(P)$, which include fixed costs:

$$ATC_j(P) = AC_j(P) + \frac{F}{D_j(P)}, \tag{2}$$

which also lets us write profits as $\pi_j = [P_j - ATC_j(P)] \cdot D_j(P)$. The firm is profitable only if $P_j - ATC_j(P) > 0$, or equivalently if its variable profit margin exceeds its fixed costs per consumer, $P_j - AC_j(P) > \frac{F}{D_j(P)}$.

Adverse Selection Definition A key feature of insurance (and other selection markets) is that average costs are not fixed but vary with prices, since prices affect which types of (cost-varying) consumers choose each plan. We follow Einav et al. (2010*c*) in defining *adverse selection* as occurring if a firm's marginal consumers who leave when it raises prices tend to be relatively low-cost – that is, have costs below its average consumer. Formally, defining the expected cost of marginal consumers (or "marginal cost") as $MC_j(P) \equiv \frac{1}{\partial D_j / \partial P_j} \sum_i C_{ij} \frac{\partial D_{ij}}{\partial P_j}$, a firm faces adverse selection if:

$$AC_j(P) - MC_j(P) > 0 \tag{3}$$

As we note below, $AC_j - MC_j$ is a useful statistic that captures the degree of adverse selection. This definition of adverse selection is equivalent to having an average cost curve that is "downward sloping" with quantity – or equivalently, increasing with price ($\frac{\partial AC_j}{\partial P_j} > 0$).[5] Downward sloping average costs is also familiar from models of monopolistic competition, where the source of the downward slope is large fixed costs. One of our paper's key points is that adverse selection and fixed costs tend to generate similar implications of limited entry – and in the extreme, natural monopoly.

2.1 Adverse Selection and the Profitability of Entry

When will firms find it profitable to compete in the insurance market, and how does this relate to adverse selection? To understand this, we first solve for Nash-Bertrand pricing equilibrium *conditional* on a set of competitors. Firm j's first-order condition for profit

variables to improve parsimony.

[5] To see this, note that using the product rule to differentiate $AC_j(P) = \frac{1}{D_j(P)} \sum_i C_{ij} D_{ij}(P)$, one can show that $\frac{\partial AC_j}{\partial P_j} = \left(\frac{-1}{D_j} \frac{\partial D_j}{\partial P_j}\right)(AC_j(P) - MC_j(P))$. Because $\frac{-1}{D_j} \frac{\partial D_j}{\partial P_j}$ is positive by the law of demand, $\frac{\partial AC_j}{\partial P_j} > 0$ if and only if $AC_j - MC_j > 0$.

maximization is: $\frac{\partial \pi_j}{\partial P_j} = \left(1 - \frac{\partial ATC_j}{\partial P_j}\right) \cdot D_j(P) + [P_j - ATC_j(P)] \cdot \frac{\partial D_j}{\partial P_j} = 0$, which after rearranging can be written as:

$$\underbrace{P_j - ATC_j}_{\text{Profit margin}} = \underbrace{\frac{1}{\eta_{j,P_j}}}_{\text{Lerner Markup}} \times \underbrace{\left(1 - \frac{\partial ATC_j}{\partial P_j}\right)}_{\text{Effect of selection and fixed costs}} \tag{4}$$

where $\eta_{j,P_j} \equiv -\frac{1}{D_j}\frac{\partial D_j}{\partial P_j} > 0$ is the plan j's own-price semi-elasticity of demand at the equilibrium. The firm's profit margin equals the standard Lerner markup over marginal costs (the inverse of the semi-elasticity of demand) times a term to account for adverse selection and fixed costs (one minus the slope of the average total cost curve).

Equation (4) is a key expression we explore in this paper. It says that even at a firm's profit maximizing price, profits can be negative if $\frac{\partial ATC_j}{\partial P_j} > 1$. In this case, the current set of entrants will not be an equilibrium; one or more firms will need to exit the market for all participants to be profitable. In models without adverse selection or fixed costs, $\frac{\partial ATC_j}{\partial P_j} = 0$, and any differentiated firm will be able to profitably enter. But adverse selection and fixed costs push towards a larger slope of the average cost curve. Importantly, there is no reason $\frac{\partial ATC_j}{\partial P_j}$ need be less than 1. In perfect competition, an average cost curve with a slope exceeding 1 is typically referred to as "instability" and assumed away. In our model of imperfect competition, if $\frac{\partial ATC_j}{\partial P_j} > 1$ at the firm's optimal price, it indicates that *the firm would rather exit the market than continue to participate.*

When is this likely to occur? To gain insight, we differentiate ATC_j with respect to price and re-arrange terms to get:

$$\frac{\partial ATC_j}{\partial P_j} = \underbrace{\eta_{j,P_j}}_{\text{Price sensitivity}} \times \left[\underbrace{(AC_j - MC_j)}_{\text{Adverse selection}} + \underbrace{\frac{F}{D_j}}_{\text{Fixed costs}}\right] \tag{5}$$

This expression shows that the slope of average total costs with price relates to three features of demand and cost:

1. **Consumer price sensitivity**: Average cost is steeper when consumers are highly price-sensitive, as captured by a large semi-elasticity of demand (η_{j,P_j}).

2. **Adverse selection**: Average cost is steeper if a firm faces greater adverse selection, captured by the gap in costs between its average and marginal consumer.

3. **Fixed costs per consumer:** Average cost is steeper when fixed costs are large relative to demand.

When these forces are sufficiently large, we may have $\frac{\partial ATC_j}{\partial P_j} > 1$ even at j's optimal price, implying that firm j cannot profitability compete in the market. Notice that this may occur

even without fixed costs ($F = 0$) as long as there is strong enough adverse selection; adverse selection and fixed costs work together to limit profitable entry. This illustrates one of our paper's key points: that adverse selection reduces firm incentives to enter and compete in a similar way as fixed costs, and thus leads to fewer competitors in equilibrium.

Another way to view this result is to plug in equation (5) into (4) to yield:

$$\underbrace{P_j - ATC_j}_{\text{Profit margin}} = \underbrace{\frac{1}{\eta_{j,P_j}}}_{\text{Lerner Markup}} - \underbrace{(AC_j - MC_j)}_{\text{Adverse selection}} - \underbrace{\frac{F}{D_j}}_{\text{Fixed costs}} \tag{6}$$

Profit margins are negative if the standard Lerner markup that firms can earn (because they offer a differentiated product valued by many consumers) is exceeded by the degree of adverse selection and fixed costs. In insurance markets with a subset of highly price sensitive consumers who are also low cost (so drive adverse selection), markups are likely to be small and adverse selection severe, limiting the degree of profitable entry.

2.2 Adverse Selection with Horizontal Differentiation: A Simple Model

When is adverse selection likely to be important and therefore to limit entry? The standard understanding is that adverse selection applies primarily to generous contracts in settings of *vertical differentiation*. For instance, a high actuarial value plan is likely to face adverse selection when competing against a low actuarial value plan. This standard understanding has motivated quality regulation as a policy response. For instance, the ACA sets minimum standards for "essential health benefits" and requires insurers to offer coverage in four tiers of actuarial value (bronze/silver/gold/platinum).

Less appreciated, but equally relevant for insurer entry, is the way that adverse selection can be important when similar-quality but *horizontally differentiated* plans compete. This case appears to be quite relevant in ACA Marketplaces (Tebaldi, 2020), and it is more challenging to solve with direct quality regulation. We present a simple example here to show how adverse selection arises in this case and to develop intuition for the empirical results in our structural model, which allows for flexible horizontal and vertical differentiation.

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 $t_L < t_H$.[6] Consumers of type $i \in \{L, H\}$ have utility for firm j of

$$U_{ij} = V - P_j - t_i \cdot \|\ell_i - \ell_j\| \tag{7}$$

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

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

$$P^* = \underbrace{\frac{\sum_i(\theta_i\alpha_i)\cdot C_i}{\sum_i \theta_i\alpha_i}}_{=MC_j} + \underbrace{\frac{1/N}{\sum_i \theta_i\alpha_i}}_{\text{Lerner Markup}} \tag{8}$$

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.[7] To see this, note that by symmetry, $AC_j = \sum_i \theta_i C_i$,

[6] 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.

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

and

$$AC_j - MC_j = \sum_i \theta_i C_i - \sum_i s_i^{MC} C_i > 0 \tag{9}$$

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 (6) 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:

$$N < \frac{1}{(AC_j - MC_j) \cdot -\frac{\partial D_j}{\partial P_j}} \tag{10}$$

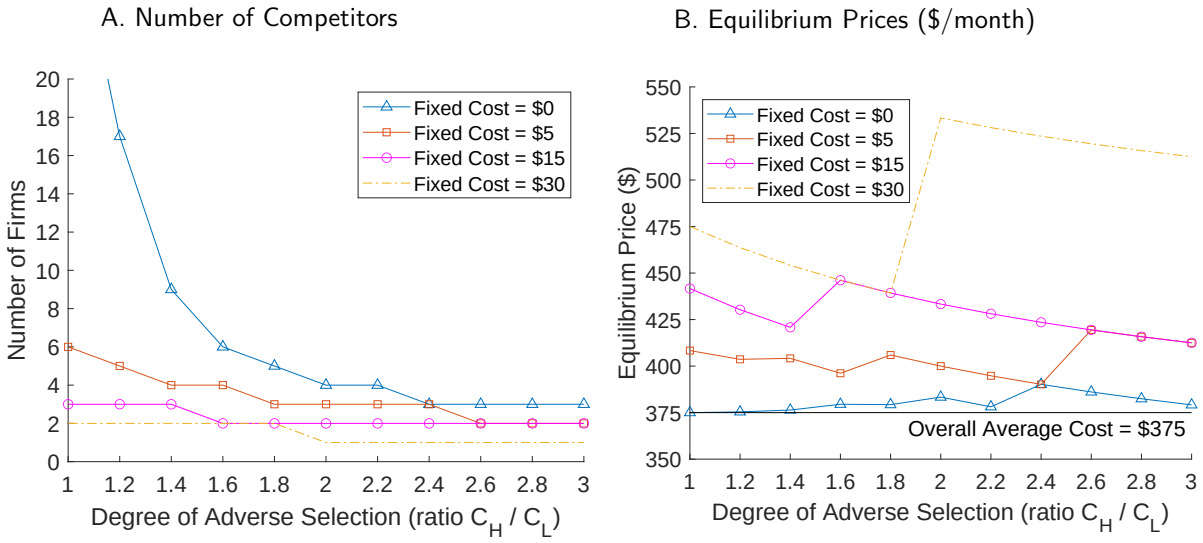
where note that AC_j, MC_j, and the demand slope $\frac{\partial D_j}{\partial P_j} = \sum_i \theta_i \alpha_i$ are all constants determined by primitives in this model. Thus, the maximum number of firms that can be sustained decreases with the degree of adverse selection and consumer price sensitivity.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

Calibrated Outcomes How does this outcome play out empirically? To understand this, Figure 1 presents results from a simple calibrated version of this model, with parameters calibrated based on estimates from our Massachusetts exchange data.[8] 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 1A shows the number of competing firms. Competition declines with stronger

[8] 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 1. Equilibrium Number of Competitors and Prices in Simple Model



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

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

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

3 Empirical Setting and Data

To investigate the empirical importance of adverse selection on insurer participation, we turn to data from the Massachusetts Health Connector, the state's precursor to the Affordable Care Act (ACA) Health Insurance Marketplaces. We start by providing 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)

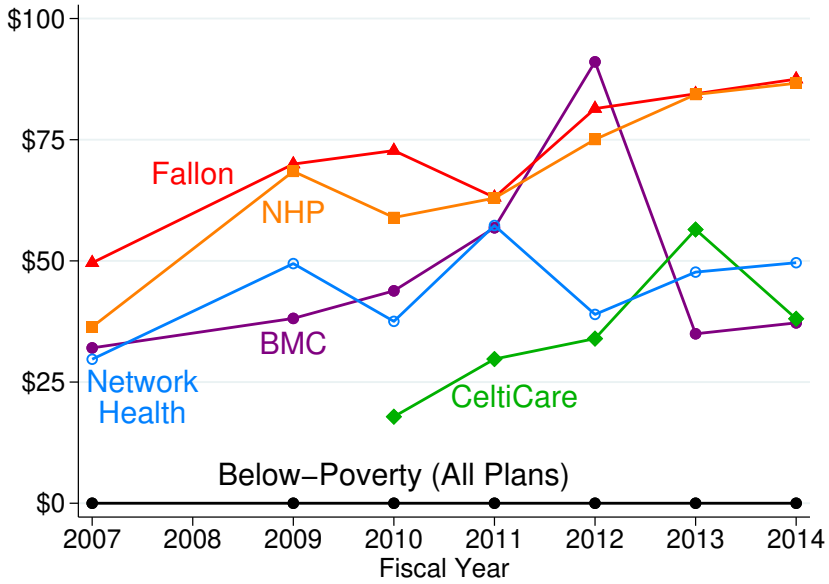
We study the pre-ACA subsidized Massachusetts health insurance exchange, a program called Commonwealth Care (or CommCare). CommCare provided subsidized health insurance for Massachusetts residents with incomes below 300% of FPL who did not qualify for Medicaid, Medicare, or job-based coverage. Consumers chose between four to five competing health insurers, with each insurer offering a single highly regulated plan that followed standardized cost sharing rules. The primary way plans differed was on their provider networks. Insurers were primarily Medicaid-based insurers offering limited networks similar to those of their Medicaid managed care plans. In the 2007-2011 period we study, three participating insurers (BMC, NHP, and Network Health) had comparably broad but differentiated networks, covering 75-85% of hospitals; see Appendix Figure A4 for a graph of network size over time. One plan (Fallon) was a regional carrier offering coverage only in central Massachusetts, and a final plan (CeltiCare) was a new entrant in 2010 offering a much narrower network.

Insurers set pre-subsidy premiums annually at the income group and region level (2007-09), regional level (2010), or statewide (2011+). Premiums could not vary with other factors including age or medical status. Premiums were then highly subsidized to ensure affordability. For the poorest enrollees (below 100% FPL), plans were fully subsidized — i.e. $0 for all plans. Higher-income enrollees paid premiums that were an increasing function of (pre-subsidy) plan prices. The cheapest plan was subsidized so that it cost an income-specific "affordable amount," which rose from $0 for the 100-150% of FPL group to $116 per month at 250-300% of FPL. Higher-price plans cost more, following a formula where price differences passed through at a higher rate to higher-income enrollees.

Figure 2 shows the evolution of the net-of-subsidy premiums over time for the Boston region. Each panel focuses on a different income group. Panel (a) shows that the lowest income group always faced zero dollar premiums for all plans. Panels (b)-(e) show that there

was significant variation in premiums within the market over time, including variation in which plan had the lowest bid. Similar figures for other rating areas can be found in the Appendix, also illustrating significant variation in premiums and the ordering of plans by premium over time. The set of firms offering plans in the market was consistent over the time period we study. As we describe in detail below, we use this quasi-experimental variation in prices over time and across income groups to estimate the key parameters of interest.

Figure 2. Variation in Enrollee Premiums in CommCare



Note: The graphs show average post-subsidy enrollee premiums for each insurer's plan in the CommCare market, by fiscal year and income group. 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.

The CommCare market had some unique features worth mentioning during the years we study. First, the exchange regulated prices by imposing price ceilings and floors. Price ceilings were used as a way of actively controlling costs and were gradually tightened so that they became binding for about half of plan prices set during 2011-13. Price floors were imposed under rules requiring that premiums be "actuarially sound," that is no lower than a minimum level defined by an independent actuary. These floors were often binding, especially for BMC, CeltiCare and Network Health. As a result it was quite common that 2+ plans tied for the lowest premium in a region. Although not explicitly intended to ensure participation, these floors may have had the effect of constraining a race-to-the-bottom in prices.

Second, as described above, CommCare's subsidies covered a portion of price differences between plans, rather than passing them fully onto consumers (as in the ACA). For enrollees below 100% of FPL, none of price differences were passed through (because all plans were free). Pass-through was positive for higher-income groups typically incomplete for lower-income groups (e.g., enrollees in 100-150% of FPL paid 39-59% of price differences during 2009-13). These "incremental" subsidies further softened price competition relative to what is likely to be true in the ACA.

Finally, CommCare automatically assigned plans for below-poverty enrollees prior to 2010 who qualified for coverage but failed to make an active plan choice (Shepard and Wagner, 2021). We exclude these enrollees from our sample because we want to capture demand among actively choosing individuals.

3.2 Enrollment and Claims Data

We acquired enrollment records and full medical and prescription drug claims data for the universe of Connector 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 Connector plan, which plan she enrolled in, and if she ever switched plans or left the market and returned again later. We also observe information about which income group and geographic market each enrollee is in, 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 Connector 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 Connector 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 Reduced Form Evidence

We use two quasi-experimental approaches to estimate our key parameters of interest, the firm-specific price semi-elasticity of demand and the slope of the firm-specific average cost curve. First, we use a difference-in-differences design leveraging variation in premiums over time, variation that affects consumers with incomes greater than 100% of FPL (treatment) but that does not affect consumers with incomes below 100% of FPL (control). Second, we use a regression discontinuity design, leveraging the fact that all premiums are fully subsidized for consumers with incomes below 100% of the Federal Poverty Line (FPL)

while consumers with incomes above 100% of FPL face different premiums for different plans.

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

We combine many price changes across many markets and several years in a "stacked" difference-in-differences design. Define an experiment e at the plan-region-consecutive year pair level. Basically, we consider each year-to-year price change for each plan in each market, which affects people in different income groups differently, as a single experiment. We restrict each experiment to the 12 months before and the 12 months after the price change. Within the experiment, we effectively compare changes in market share and average cost for the premium-paying groups to changes in market share and average cost for the non-premium-paying control group. Formally, each experiment contains multiple income groups $g \in \{0, 1, ...4\}$, 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:[9]

$$Y_{egt} = \tilde{\alpha}_{et} + \tilde{\gamma}_{eg} + \sum_{k \in \{-T,T\} \setminus \{-2\}} \tilde{\delta}_t \times 1\{g > 0, t = k\} + \tilde{\varepsilon}_{egt} \tag{11}$$

Each outcome Y_{egt} is measured at the experiment-income group-month level. The regression specification includes experiment-by-event time fixed effects $\tilde{\alpha}_{et}$ to ensure that we only use *within-experiment* variation in prices to identify the firm-specific demand response to the change in prices. We also include experiment-by-income group fixed effects to net out average differences in costs or preferences across income groups — if the higher-income treated groups have lower costs in general, these fixed effects ensure that the pre-period difference in costs is netted out, such that the treatment effects $\tilde{\delta}_t$ only represent post-period differences in the outcome variable (i.e., changes due to the premium change). The coefficients of interest are the $\tilde{\delta}_t$s. These are the event study coefficients, and they reflect how the gap in Y_{egt} between the control group and the other groups changes over time, relative to the gap 2 months prior to the price change.[10] We estimate this regression for three key outcomes: relative premiums p, (log) market shares (log_sh), and average costs (AC).[11]

Figure 3 presents the event study plots for the positive and negative price change experiments. Panel (a) shows the average price change for each group, around \$20 for both price increases and price decreases. Panel (b) shows the changes in log market share. The event study plot shows that differential market share trends between the treatment and control groups are steady throughout the pre-period for both the positive and negative price change experiments, suggesting that the parallel trends assumption is likely to be satisfied here. At time $t = 0$, however, market shares diverge between the treatment and control groups. For the price increases, market shares decrease by around 20%. For the price decreases, market shares increase by slightly more than 20%. The effects of price changes on market shares thus appear to be symmetric. This provides evidence of the credibility of our empirical strategy, as spurious trends in market share are unlikely to be positively correlated with price decreases and negatively correlated with price increases. The effects are also quite strong, with a \$20 (115%) increase in the premium causing a 20% shift in market share.

Panel (c) shows the changes in the average cost of plan enrollees that correspond to the shifts in market share documented in Panel (b). Here, estimates are noisier, and there is some evidence of a pre-trend for the premium decreases. However, the plot suggests

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

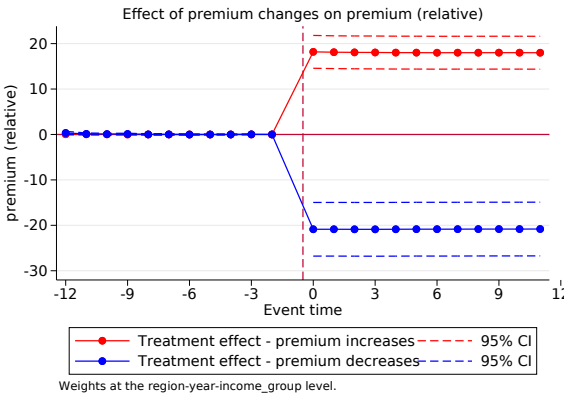
[10] We normalize to month -2 and exclude month -1 because prices for the following year were publicized one month prior to the open enrollment period. We some evidence that these start to affect demand and costs in month -1.

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

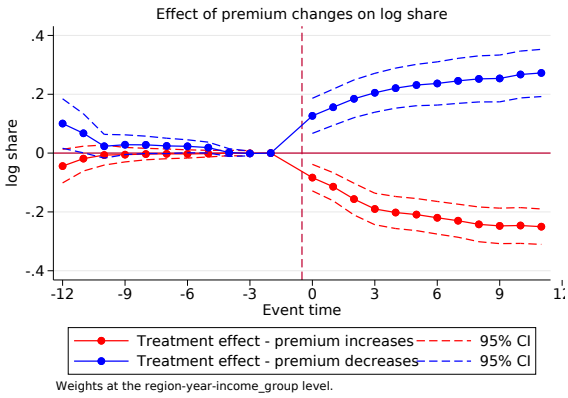
that when prices increase, average costs increase and when prices decrease, average costs decrease. These results are consistent with adverse selection. The magnitudes are also large: A $20 price increase leads to an increase in the average cost of a plan's enrollees of around $20, suggesting that the slope of the firm-specific average cost curve is close to one. As discussed in Section 2, this raises concerns about the ability of this market to support multiple competing plans in the absence of policies used to combat selection, such as price floors and risk adjustment.[12]

Figure 3. Event Study Estimates for All Enrollees

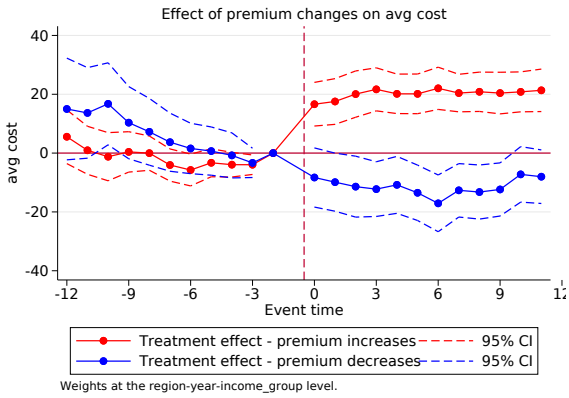
(a) Relative Premiums



(b) Log Shares



(c) Average 12-Month Cost



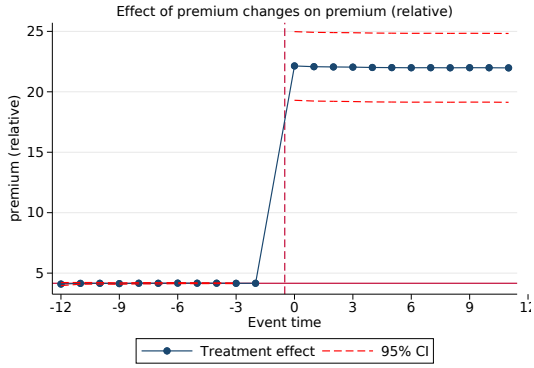
Note: Figure shows event study estimates of the impact of premium increases and decreases on relative premiums. Panel (a) shows results for relative premiums, Panel (b) shows results for log plan shares, and Panel (c) shows results for average costs.

[12]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.

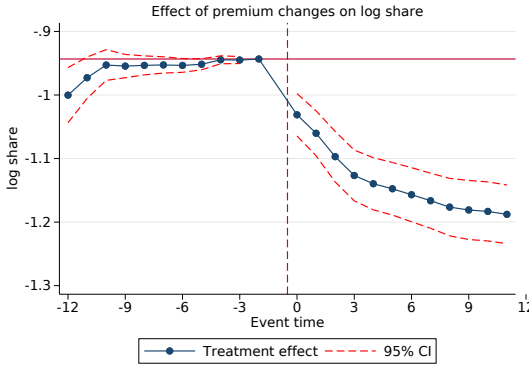
Next, we leverage all experiments in a single unified regression in order to maximize power. Specifically, we multiply all outcomes from negative price change experiments by -1, stack *all* experiments (both positive and negative price changes), and re-run the specification described in Equation 11. Figure 4 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.

Figure 4. Pooled Event Study Estimates for All Enrollees

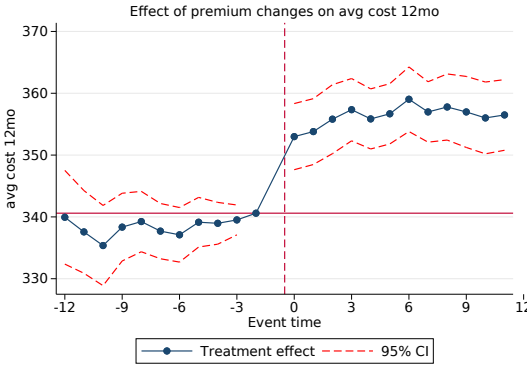
(a) Relative Premiums



(b) Log Shares



(c) Average 12-Month Cost



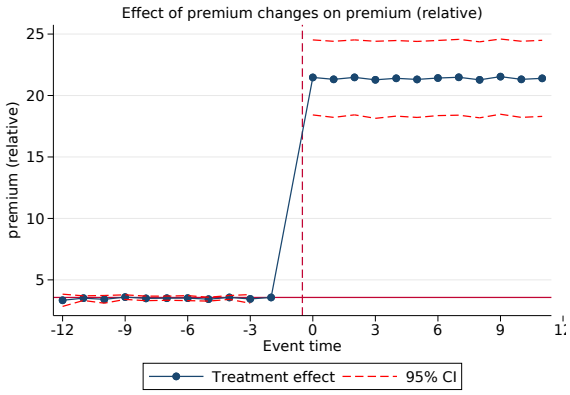
Note: Figure shows pooled event study estimates of the impact of premium increases and decreases, where all outcomes are multiplied by -1 for premium decreases. Panel (a) shows results for relative premiums, Panel (b) shows results for log plan shares, and Panel (c) shows results for average costs.

Here, a dynamic effect is clear, with the effect growing over time. Such a dynamic may

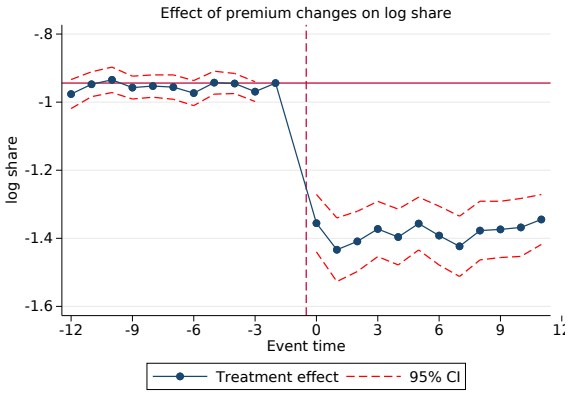
seem odd, but it is due to enrollment churn. In the first months of the new year, a large share of the enrollees are incumbent enrollees subject to inertia in their plan choices. As the new year goes on, however, more and more of the inertial enrollees drop out of the market, leaving a larger share of enrollees who entered the market under the new prices. To illustrate this, in Figure 5 we replicate Figure 4, restricting only to new enrollees rather than estimating the effects of prices on both incumbent and new enrollees combined. Again, the change in prices is a little over $20. The effects on market share, however, are much larger: The same $20 increase in prices decreases a plan's market share by around 40% among new enrollees compared to the 20% decrease across all enrollees, revealing that inertia is strong in this setting. Further, Panel (b) of Figure 5 shows that among new enrollees, the treatment effect is not dynamic. Instead, as soon as the price change, the market share among new enrollees drops by 40% and continues at that level in perpetuity. These results combine to show that we would expect the effect of a price change on market share to increase over time as fewer of the consumers in the market are inertial.

Figure 5. Pooled Event Study Estimates for New Enrollees Only

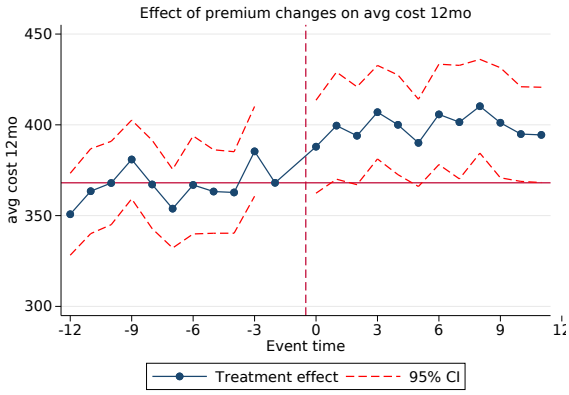
(a) Relative Premiums



(b) Log Shares



(c) Average 12-Month Cost



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

Panel (c) of Figure 4 shows the effect of the price increase on the average cost of plan's enrollees. Here, after pooling across the price increases and decreases, results are less noisy than before. Now we see that treatment and control groups have similar average cost trends prior to the price changes. And, as before, we find that the average cost of a plan's enrollees increases markedly following the price increase, though now the estimates are much cleaner. Specifically, we again estimate that a \$20 increase in the premium results in an increase in the average cost of a plan's enrollees of around \$20. Panel (c) of Figure 5 shows that selection is also stronger among new enrollees, with a \$20 price increase producing an increase of around \$40 in the average cost of a plan's enrollees. Thus, the overall slope of the firm-specific average cost curve appears to be around 1, while the slope of the firm-

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

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

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 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 * $p < 0.10$.

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

Ultimately, these results 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 will find it difficult to enter and earn non-negative profits. 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? In the next section, we leverage this quasi-experimental variation in prices to estimate a structural model of consumer demand and plan costs in this market. We then use that model to perform counterfactual simulations and assess the roles of various corrective policies in achieving an equilibrium with multiple competing plans.

Regression Discontinuity. As discussed above, 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 A1 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 A1, we find imbalance in a key , age, on either side of the discontinuity. Panel (c) shows that we do not find such a discontinuity for gender, but there is a clear shift in slopes for this characteristic at 100% of FPL. These results combine to suggest that the key assumption for a valid RD design — that individuals are as good as randomly assigned to one side of the discontinuity versus the other — is violated in this setting.

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

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

$$Y_i = \beta_0 + \beta_1 Above100_i + \beta_2 FPL_i + \beta_3 FPL_i \times Above100_i + \epsilon_i \tag{12}$$

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 6, 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

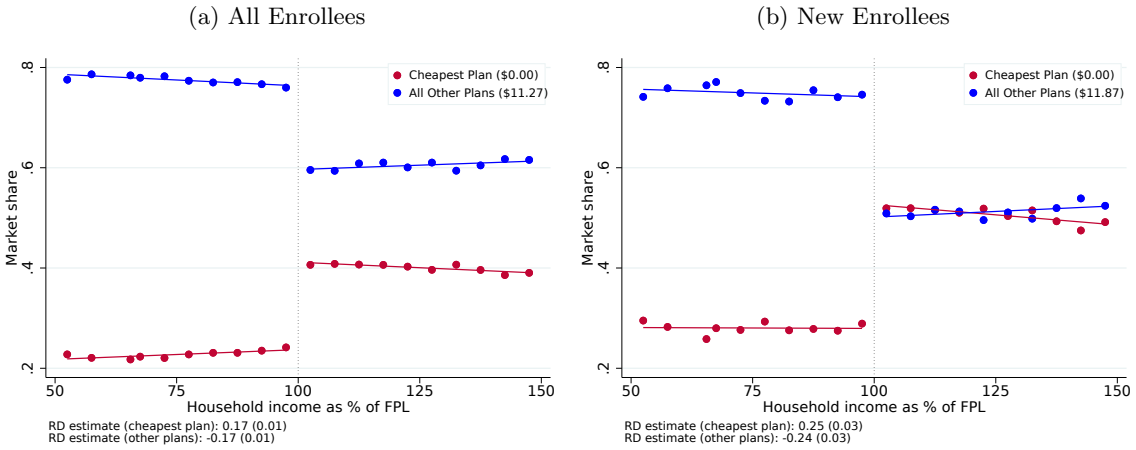
stronger preferences for the more generous, higher-priced plans), not more, suggesting that our estimate of the shift in market share may be an *under*-estimate of the true shift caused by the change in price.

In Panel (b) of Figure 6, 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 2, we present the RD coefficient estimates. Columns 1 and 2 present estimates corresponding to Panels (a) and (b) of Figure 6. 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.

Figure 6. RD Estimates - Market Shares



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

Table 2. 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.

5 Demand Estimation

While the reduced form estimates of price sensitivity and adverse selection suggest that the market will have major difficulties supporting multiple competing plans, this market did in fact achieve modest levels of insurer participation, with 4-5 insurers competing at any

given time. How could this be so? In this section, we seek to uncover the reasons why this market was able to overcome the selection-related issues we highlight and document above. We do this by leveraging the quasi-experimental variation in prices illustrated in Section 4 to estimate a structural model of consumer demand and plan costs in this market. We then combine that model with a model of insurer entry to perform counterfactual simulations and assess the roles of various corrective policies in achieving an equilibrium with multiple competing plans. We also use demand model estimates to study the effects of these policies on plan prices and consumer welfare. We start in this section by presenting the model, estimating it using the CommCare data, and assessing how well the estimated model reproduces the degree of price sensitivity and adverse selection shown in the reduced form results as illustrated by actual versus model-predicted event studies.

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 the month when an enrollee newly enrolls in the market, when an enrollee re-enrolls after a break of at least one month, and when continuing enrollees face open enrollment periods when they are able to switch plans.[13] We limit the sample to plan years 2008–2011. We exclude enrollees 65 years of age or older, as they are eligible for Medicare.

We estimate a multinomial logit choice model over insurance plans, where enrollees choose between one of five CommCare health insurance plans. We specify the utility enrollee i receives from enrolling in plan j as:

$$u_{ij} = \alpha(Z_i) \cdot P_{ij} + \mathbf{X_{ij}}\beta + \xi_j(Z_i) + \varepsilon_{ij}, \qquad j = 1, .., J \tag{13}$$

where P_{ij} is plan j's post-subsidy premium for consumer i. Recall that post-subsidy premiums depend on the pre-subsidy premium and each enrollee's income group. Price sensitivity is allowed to vary by income bins, demographics (age-sex bins), enrollment status (new, re-enrollee, or current enrollee), and HCC risk scores. The variables X_{ij} capture enrollee observables that vary by plan, including dummy variables for whether re-enrollees and continuing enrollees had previously chosen the plan as well as the interactions of these variables with income, age-sex bins, and HCC risk scores. Plan quality $\xi_j(Z_i)$ is captured by plan-region-year, plan-region-income, and plan-risk score fixed effects. We 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. Note that we model individuals' utility conditional on each choice instance, which implicitly controls for any fixed characteristics of individuals in each choice instance that do not also vary across

[13] 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.

alternative plans.[14]

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

Table 3 below shows the implied own-price semi-elasticities and own-price average cost derivatives dAC/dP for each plan, for all enrollees. The model implies that a \$1 increase in a plan's monthly premium lowers its demand by an average of 1.9 percent. These estimates reproduce semi-elasticities from **?** 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 \$1.07. These estimates imply margins of 18% above marginal cost before risk adjustment. As a result of adverse selection, the implied margins above *average* cost are much smaller, at -1%.

For new enrollees, we find that a \$1 increase in monthly premium lowers demand by 3 percent. Without risk adjustment, a \$1 *decrease* in monthly premium would lower average cost by \$1.76. These estimates imply margins of 11% 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%.

[14] We also include fixed effects within plan-region-income for plan year 2010 and plan months 10, 11, and 12. This controls for potential auto-assignment of enrollees that occurred in these months, for which we do not observe auto-assignment status.

Table 3. Implied own-price semi-elasticities and dAC/dP

Panel (a) All Enrollees		
	Semi Elasticity	dAC/dP
By Plan		
BMC	-0.014	0.753
Celticare	-0.043	1.155
NHP	-0.024	1.296
Network	-0.017	1.704
By Income Group		
IIA (100.1–150% FPL)	-0.028	0.949
IIB (150.1–200% FPL)	-0.012	0.913
IIIA (200.1–250% FPL)	-0.013	1.261
IIIB (250.1–300% FPL)	-0.011	1.149
Overall	-0.019	1.070

Panel (b) New Enrollees		
	Semi Elasticity	dAC/dP
By Plan		
BMC	-0.025	1.193
Celticare	-0.041	1.259
NHP	-0.036	2.401
Network	-0.028	3.216
By Income Group		
IIA (100.1–150% FPL)	-0.048	1.658
IIB (150.1–200% FPL)	-0.019	1.513
IIIA (200.1–250% FPL)	-0.015	1.627
IIIB (250.1–300% FPL)	-0.013	1.478
Overall	-0.030	1.762

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 5. 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

Plan Cost Heterogeneity. 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. We could assume that a given consumer would cost the same under each plan, but this might lead to mis-leading

conclusions, especially for the plan widely believed to be the lowest-cost plan in the market, Celticare. Instead, we estimate plan effects on costs. Since costs for the same enrollee may differ across plans, we assume that each plan has a constant proportional effect on costs and estimate these proportional plan spending effects separately by plan and region.

We start with actual realized costs (instead of modeling costs as a function of observables) to retain both observed and unobserved relationships between predicted demand and costs. This also allows us to more closely fit patterns of adverse selection in the data (which relies on preserving correlations between costs and demand) and avoid issues regarding lack of predictive power when modeling costs. We identify the effect of plans on costs by leveraging cases where the same individual switches plans for different enrollment spells. We could instead leverage variation in the plan an individual is enrolled in generated by changes in prices from one year to the next, as in Section 4, but inertia makes this type of analysis (1) underpowered and (2) non-representative of the inertial consumers. Thus, we leverage this alternative strategy that involves observing changes in costs within a person across different enrollment spells, all of which involve an active choice and thus avoid the inertia problem. Our estimation sample has observations at the enrollee x enrollment spell level, and we limit to individuals observed in at least 2 spells, separated by a gap in CommCare enrollment.[15]

We estimate the following Poisson regression specification:

$$E(C_{ijt}^{obs}|X_{it}, Z_{it}) = \exp(\alpha_i + \beta_{jt} + X_{it}\gamma + \eta_j(Z_{it})) \tag{14}$$

This specification controls for individual fixed effects (α_i) and plan-year fixed effects (β_{jt}). The additional controls X_{it} capture observed risk factors that may vary across spells, including age-sex bins, splines in HCC risk score deciles (splitting the top decile into two bins), FPL group, and service area. The $\eta_j(Z_{it})$ coefficients represent the plan-specific cost multipliers, which we allow to vary by region ($Z_{it} =$). We estimate this specification on all enrollees, since by definition an enrollee can only be a "new enrollee" during their first enrollment spell.

Table 4 below shows estimates of cost heterogeneity.[16] 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 substantial heterogeneity in costs across plans, consistent with prior work focusing on cost heterogeneity across Medicaid plans in New York City (Geruso et al., 2020).

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

[16] 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.

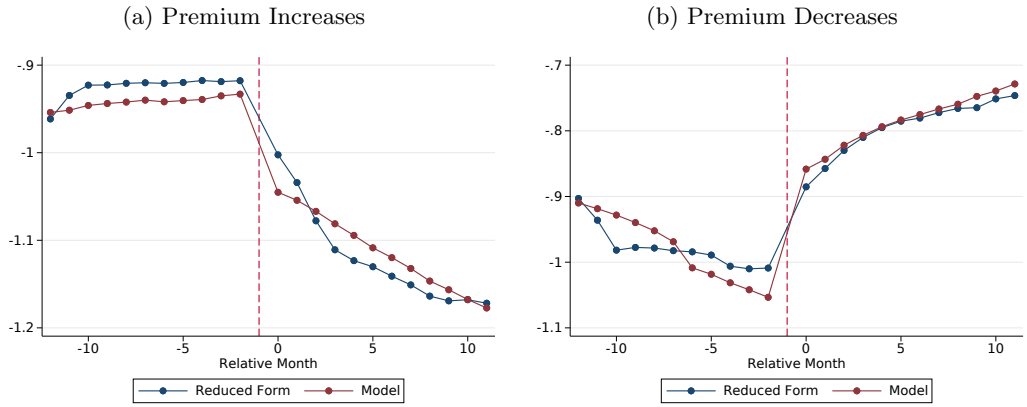
Table 4. Plan Cost Heterogeneity Estimates

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

Comparing Demand Estimates 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. To conduct this comparison, we apply the predicted choice probabilities from the demand model to all months in each corresponding enrollment spell. Then, we collapse the resulting data set and run event studies following the reduced form analysis, described previously in Section 4.

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

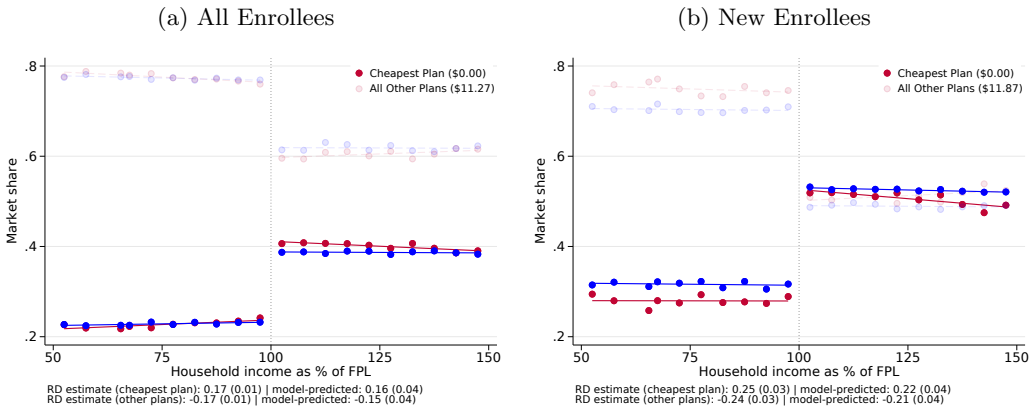


Notes: Figure shows how shares respond to increases and decreases in post-subsidy premiums, comparing actual shares from the reduced-form results with shares predicted using the demand estimates from Section 5. Panel (a) shows results for premium increases; Panel (b) shows results for premium decreases.

We further validate our model's ability to match the extent of price sensitivity in the data by comparing model predictions to the regression discontinuity results presented above

in Section 4. Figure 8 shows that our model reproduces the regression discontinuity results to within one standard error, in both the sample of all enrollees (Panel (a)) and new enrollees (Panel (b)). This provides further evidence that our demand model accurately captures the effect of price changes.

Figure 8. Comparing Model-Predicted Shares with RD Results

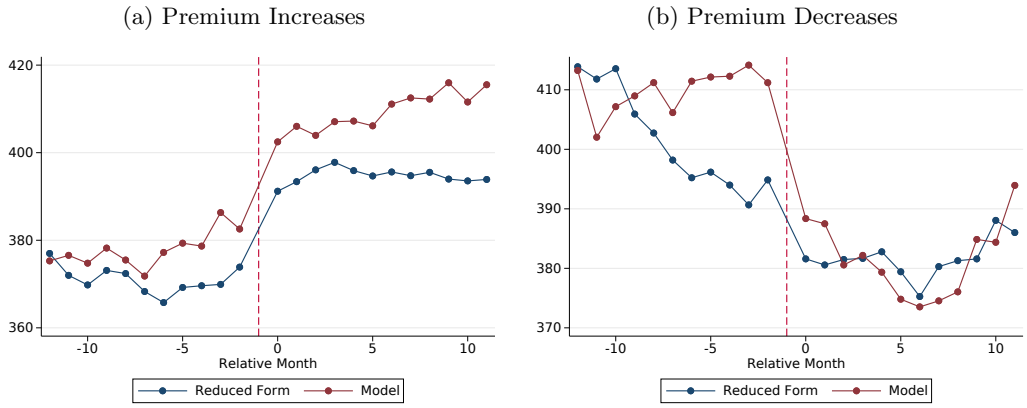


Notes: Figure shows regression discontinuity plots for the market share of the lowest priced plan (solid) and the combination of all other plans (transparent) for observed shares (red) compared to model-predicted shares (blue).

Next, we combine our predicted choice probabilities with enrollee costs to assess whether our model can reproduce the patterns of adverse selection observed in the reduced form results. Each observation at the plan-region-income group-month level is a weighted average of enrollee costs, where the weights are each enrollee's predicted shares from the demand model.

Figure 9 below demonstrates that our demand model reproduces similar increases in average costs when premiums increase, and similar decreases in average costs when premiums decrease. The model does not fully reproduce the trajectory of costs over time for premium decreases. This is not surprising as the reduced form results allow average costs to vary month by month, whereas the model-predicted costs are assumed to be constant within each enrollment spell. More importantly, we are assured by the model's ability to match the jump in costs that occurs when prices increase or decrease, suggesting that the model can generate similar magnitudes of adverse selection.

Figure 9. Comparing Model-Predicted Average Costs with Reduced Form Results



Notes: Figure shows how a plans' average costs respond to increases and decreases in its post-subsidy premiums, comparing actual costs from the reduced-form results with predictions using the demand estimates from Section 5. Panel (a) shows results for premium increases; Panel (b) shows results for premium decreases.

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 three stages. In the first stage, single-plan insurers choose whether to enter the market. In the second stage, each insurer sets Nash-Bertrand prices. In the last stage, consumers choose plans and incur health care costs, which determines insurer profits.

Our baseline simulations are conducted on all enrollees in the market and do not impose risk adjustment. We also conduct simulations under perfect risk adjustment, various degrees of partial risk adjustment, and with different levels of fixed costs. For each level of risk adjustment, we assess the impact of imposing price floors on average insurer premiums, total insurer profits, and consumer welfare. We assume that all firms are potential entrants in all geographic areas, with the exception of one insurer (Fallon) that only operates in a handful of geographic markets. We exclude Fallon from all simulations. Lastly, as in Sections 2 and 5, we assume that all enrollees must choose a plan, i.e., that there is no outside option. As we discussed previously, this is a realistic assumption for insurance markets with generous subsidies or strong mandates, and is useful for studying market outcomes under systems with universal health insurance coverage.

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. Given the set of four possible insurer entrants and

a set of bounds for pre-subsidy premiums, 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, insurers set Nash-Bertrand equilibrium prices (subject to any policies restricting prices). 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.

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 $525 and assume that all monopoly firms price at the ceiling, due to the lack of an extensive margin response to price from consumers.

In general, there may be a multiplicity of valid equilibria corresponding to different combinations of firms that could profitably enter. In these cases, we report all valid equilibria.[17] Thus, if the combination of all four firms constitutes a valid pricing equilibrium, then it is the unique equilibrium. However, if there is no valid four-firm equilibrium, then there may be more than one valid equilibrium with three firms or fewer. See Appendix Section B for details.

As discussed above, insurers may have direct effects on enrollee costs. In some simulations, we allow costs to vary across insurers for a given individual according to our cost model estimates described in Section 5. In other simulations, we assume away all heterogeneity in costs across insurers.

6.2 Solving for Equilibria

We solve the model backwards, starting with enrollee demand. For enrollee demand, we use the demand model and estimated coefficients from Section 5. We exclude below-poverty individuals from our simulation sample to better match the population of the ACA exchanges. We assume that consumers receive fixed subsidies of $350. While subsidies were linked to prices in this market, the fixed subsidies allow changes in pre-subsidy price levels to affect consumer welfare and shed light on the broader conceptual issues we hope to highlight in the simulations.[18]

[17] Multiplicity only arises in terms of different combinations of entrants. Conditional on a set of entrants, we do not find cases where multiple price vectors satisfy the equilibrium conditions.

[18] Without an outside option, fixed subsidies yield equivalent equilibria (in terms of participating plans and pre-subsidy prices) to price-linked subsidies that hold the price of the cheapest plan fixed. Only consumer

Given a set of insurers entering the market, 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 (4). Due to adverse selection, not all solutions to the FOCs will solve insurers' maximization problem. Indeed, some solutions to the FOCs are local minima for certain firms. To surmount this issue, we use 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 satisfy the equilibrium conditions outlined above. We then solve the FOCs within a local region of each candidate price vector to obtain exact equilibrium price vectors. For each equilibrium, we compute consumer welfare as consumer surplus (which accounts for the value of plan choice to consumers and consumer premiums) less the (fixed) premium subsidy (which accounts for the government's costs in providing premium subsidies).

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 5 shows the set of possible equilibria that exist in the absence of risk adjustment, in the population of new enrollees, with no fixed costs. Our baseline simulations focus on new enrollees to abstract away from consumer inertia and dynamic considerations. Due to high price sensitivity and adverse selection, the market can only support one firm in equilibrium without risk adjustment. This result holds with or without fixed costs. The single monopoly firm prices at the ceiling, \$525. 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

welfare differs.

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 goes even further, allowing the market to support three firms instead of two. All three firm equilibria survive low fixed costs, but only one survives high fixed costs. Consumer welfare is often, but not always, higher with perfect risk adjustment. This reflects the fact that risk adjustment leads to more firm participation (welfare-increasing) but can sometimes also generate higher prices (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, which Mahoney and Weyl (2017*a*) assumed to be exogenously determined.

Table 5. Baseline Simulated Equilibria

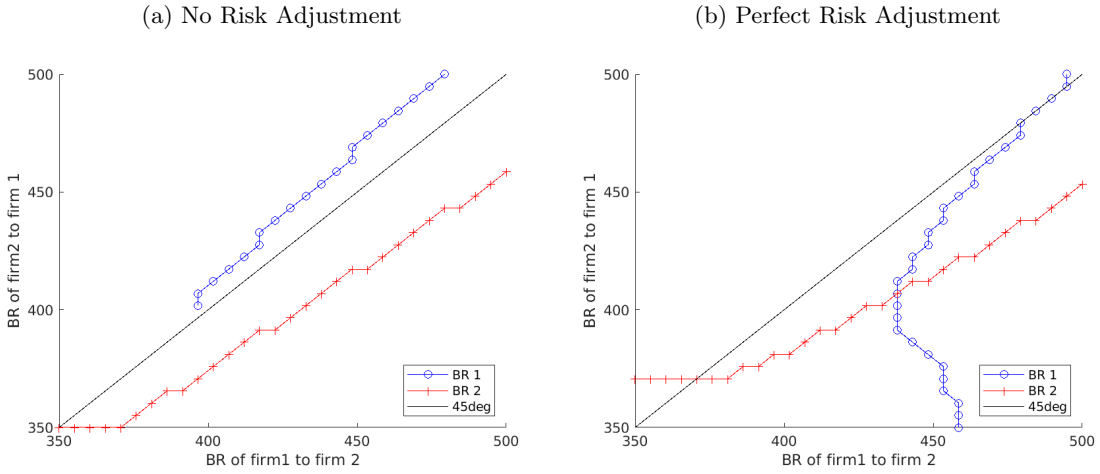
(1) Possible Entry Combinations	(2) Prices	(3) Shares	(4) Avg Price	(5) Welfare	(6) Remains with: $F = 10$	(7) Remains with: $F = 30$
Panel (a) No Risk Adjustment						
Monopoly [BMC]	[475]	[1]	475	29	Yes	Yes
Monopoly [Celticare]	[475]	[1]	475	0	Yes	Yes
Monopoly [NHP]	[475]	[1]	475	31	Yes	Yes
Monopoly [Network]	[475]	[1]	475	31	Yes	Yes
Panel (b) Moderate Risk Adjustment ($\lambda = 0.50$)						
Two firms [BMC,NHP]	[449, 471]	[0.67, 0.33]	456	73	–	–
Two firms [BMC,Network]	[415, 418]	[0.52, 0.48]	416	110	Yes	–
Two firms [Celticare,NHP]	[452, 503]	[0.72, 0.28]	466	42	–	–
Two firms [NHP,Network]	[455, 438]	[0.35, 0.65]	444	84	Yes	–
Panel (c) Perfect Risk Adjustment						
Three firms [BMC,Celticare,Network]	[426, 395, 464]	[0.34, 0.54, 0.11]	414	110	Yes	Yes
Three firms [BMC,NHP,Network]	[440, 500, 443]	[0.49, 0.07, 0.44]	446	92	Yes	–
Three firms [Celticare,NHP,Network]	[404, 506, 437]	[0.60, 0.06, 0.34]	422	98	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 entrants. 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 10 plots the best response curve of BMC to Celticare 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.

Figure 10. Best Response Curves for BMC and Celticare



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

In Figure 10 above, we see that BMC and Celticare have a (unique) valid equilibrium under perfect risk adjustment but not in the case with no risk adjustment. 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. 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.[19].

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. Price floors can also be represented on this graph by preventing firms' best response curves from dropping lower than some threshold.

The best response curves can also be used to understand why models of adverse selection that assume perfect competition, in which firms set prices equal to average cost, may not accurately reflect incentives in real-world insurance markets with a limited number of horizontally differentiated insurers. If BMC set its price equal to market average costs (about $400 on the x-axis), then Celticare's best response (with adverse selection) would still be to undercut BMC's price and enjoy higher profits by enrolling individuals with below-average costs. However, with perfect risk adjustment, there is no longer a cost advantage to undercutting. In this case, given a low price from BMC (e.g., $400), Celticare's best response is to price much higher (around $450) to capture a larger markup.

6.5 Price Floors

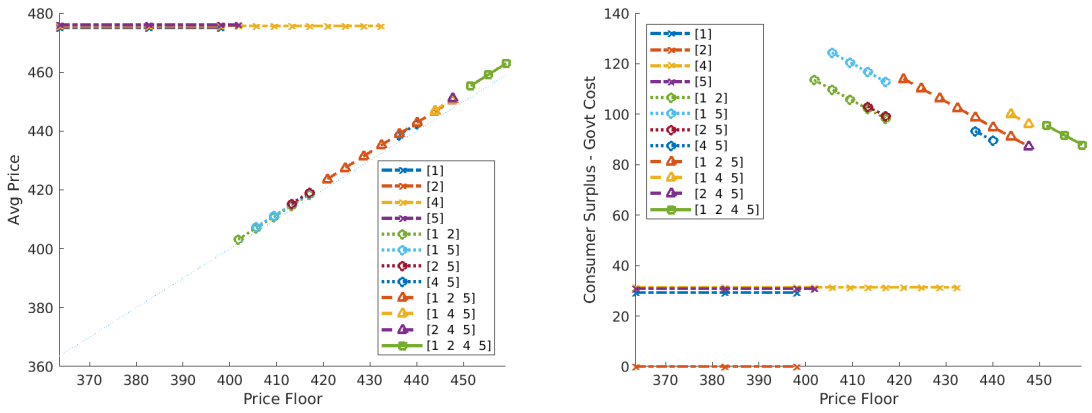
Our results and our model highlight that while adverse selection can sharpen price competition and reduce markups, it also can also generate pricing externalities (a firm that lowers its price increases costs for other firms while reducing its own). We have shown that in the absence of strong risk adjustment 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

[19]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.

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 is not $0 given most levels of risk adjustment.

Figure 11. Impact of Price Floors in Simulations with No Risk Adjustment

(a) Average Prices (b) Consumer Welfare



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 zero corresponds to the case where Celticare operates as a monopoly in the market.

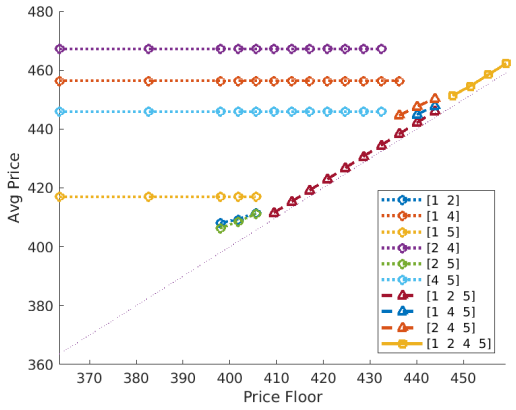
Figure 12 plots all possible market outcomes against increasing levels of price floors in the case with no fixed costs and no risk adjustment ($\lambda = 0.5$). The figure enumerates all possible equilibria that satisfy the conditions introduced above. Without a price floor, or with very low price floors, only equilibria where one firm participates survive. In these equilibria, prices are high (at the ceiling) and consumer welfare is low. But when price floors begin to bind (just below market average costs), they can recover equilibria where more than one firm participates. These equilibria actually involve *lower* prices than what would be charged in the absence of a binding price floor. Consumer welfare is also significantly higher, both due to the lower prices and due to the additional plans available to consumers.

Figure 12 mimics Figure 12 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, the only equilibria that survive involve two firms participating, with average prices ranging from $418 to $470. As the price floor increases and becomes binding for certain firms, the equilibria [1 2] (BMC, Celticare) and [2 5] (Celticare, Network) become available, with average prices that are actually *lower* than the lowest-priced equilibrium that previously existed ([1 5], BMC & Network), despite the higher price floor. Moreover, a price floor of $409 (about 4% above market average costs) allows three firms to enter (BMC, Celticare, and Network). Due to the increase in entrants, the average price decreases a bit further. This illustrates that

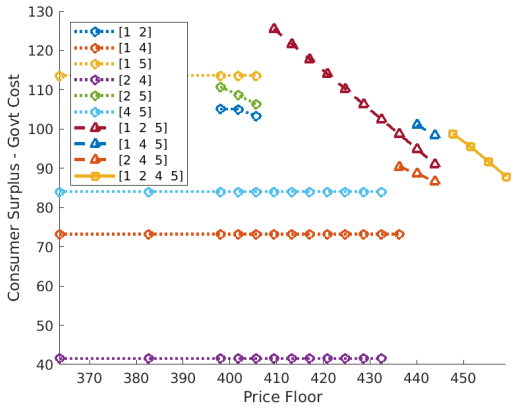
while undercutting incentives reduce markups, they can also selectively eliminate the most competitive equilibria. Panel (d) of the figure shows that price floors increase the number and quality of entrants, which can increase consumer welfare by expanding plan choice. On the other hand, price floors also increase premiums; this is reflected in declining consumer welfare with large price floors. Consumer welfare is maximized using a price floor just above average cost.

Figure 12. Impact of Price Floors in Baseline Simulations

(a) Average Prices



(b) Consumer Welfare



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 zero corresponds to the case where Celticare operates as a monopoly in the market.

Table 6 below shows how the optimal price floor affects entry and welfare in a wide range of cases.[20] Panel (a) compares results between the sample of new enrollees and all enrollees for the case with moderate risk adjustment. 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 \$383 among new enrollees, we find that the optimal price floor is about 7% above average cost.

The remainder of Table 6 shows that the optimal price floor is rarely \$0. Indeed, only in the cases of strong, perfect, and HCC risk adjustment 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 instability/reduced entry is due to fixed costs rather than adverse selection. Panel (d) further shows that price floors improve welfare when firms' costs are assumed to

[20] 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.

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 6. 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)	409	Three firms [1,2,4]	409	30	126	Two firms [1,4]	12
All Enrollees	383	Three firms [1,2,4]	397	36	148	Three firms [1,3,4]	8
Panel (b) Impact of Risk Adjustment							
No Risk Adj	406	Two firms [1,4]	406	19	124	Monopoly [3]	93
Lambda = 0.25	406	Two firms [1,4]	406	19	124	Monopoly [3]	93
Lambda = 0.50 (baseline)	409	Three firms [1,2,4]	409	30	126	Two firms [1,4]	12
Lambda = 0.75	None	Three firms [1,2,4]	411	53	115	Three firms [1,2,4]	0
Perfect Risk Adj	None	Three firms [1,2,4]	414	58	110	Three firms [1,2,4]	0
HCC Risk Adj	None	Four firms [1,2,3,4]	416	46	123	Four firms [1,2,3,4]	0
Panel (c) Impact of Fixed Costs							
F = 0 (baseline)	409	Three firms [1,2,4]	409	30	126	Two firms [1,4]	12
F = 10	417	Three firms [1,2,4]	417	37	118	Two firms [1,4]	4
F = 30	429	Three firms [1,2,4]	429	49	106	Monopoly [3]	75
Panel (d) Impact of Cost Heterogeneity							
With cost het (baseline)	409	Three firms [1,2,4]	409	30	126	Two firms [1,4]	12
Without cost het	411	Three firms [1,3,4]	411	17	133	Two firms [3,4]	21

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.

6.6 Sensitivity to Model Assumptions

We test the sensitivity of our results to alternative modelling assumptions. First, we consider a version of the model in Section 5 without insurer cost heterogeneity (i.e., the same enrollee has identical costs whether they choose Celticare or NHP). This removes one source of insurer differentiation that brings our counterfactuals closer to our theoretical results from Section 2, where firms are assumed to be symmetric. As shown in Panel (d) of Table 6, removing insurer cost heterogeneity does not change our qualitative conclusions. If anything, undercutting is somewhat more intense without cost heterogeneity (total insurer profits are lower) and price floors are somewhat *more* effective at recovering entry and increasing welfare (as evidenced by the larger welfare gain from price floors in the case without cost heterogeneity). We also estimate and simulate a version of the model with an outside option, where enrollees are allowed to exit the market.

7 Conclusion

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

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

Our findings also generalize to other markets with downward-sloping average cost curves, such as markets with significant fixed costs (e.g., pharmaceuticals). We show how market competition is stymied by undercutting incentives, which endogenously determine the number of firms that can exist in a market. In both insurance markets and pharmaceuticals, firms have an additional incentive to undercut their rivals: not only does the undercutting firm acquire larger market shares, but they also move down their own average cost

curve. As others have observed (**?**), undercutting can accentuate competition by reducing markups. However, we show that undercutting can lead to lower welfare overall by limiting the number of entrants and can even lead to higher prices. We show that two types of regulation can limit these problems. First, in some markets regulators are able to directly manipulate the average cost curve (e.g., with risk adjustment). However, this is not possible in cases where average cost curves slope downwards because of fixed costs. For these cases, regulators could ensure sufficient entry by setting price floors high enough to cover fixed costs for a desired number of firms.

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

References

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

Azevedo, Eduardo M., and Daniel Gottlieb. 2017. "Perfect Competition in Markets With Adverse Selection." *Econometrica*, 85(1): 67–105.

Berry, Steven, and Peter Reiss. 2007. "Empirical models of entry and market structure." *Handbook of industrial organization*, 3: 1845–1886.

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

Chiappori, Pierre-André, and Bernard Salanié. 2000. "Testing for asymmetric information in insurance markets." *Journal of political Economy*, 108(1): 56–78.

Cutler, David M, and Sarah J Reber. 1998. "Paying for health insurance: the trade-off between competition and adverse selection." *The Quarterly Journal of Economics*, 113(2): 433–466.

de Chaisemartin, Clement, and Xavier D'Haultfoeuille. 2020. "Two-way Fixed Effects Estimators with Heterogeneous Treatment Effects." *American Economic Review*, 110(9): 2964–96.

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

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

Einav, Liran, Amy Finkelstein, and Mark R. Cullen. 2010*c*. "Estimating welfare in insurance markets using variation in prices." *Quarterly Journal of Economics*, 125(3): 877–921.

Einav, Liran, Amy Finkelstein, and Pietro Tebaldi. 2018. "Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies."

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 Layton, Grace McCormack, and Mark Shepard. 2021. "The Two Margin Problem in Insurance Markets." *Review of Economics and Statistics*, Forthcoming.

Goodman-Bacon, Andrew. 2019. "Difference-in-Differences with Variation in Treatment Timing."

Gruber, Jonathan. 2017. "Delivering Public Health Insurance Through Private Plan Choice in the United States." *Journal of Economic Perspectives*, 31(4): 3–22.

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

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.

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

Salop, Steven C. 1979. "Monopolistic competition with outside goods." *The Bell Journal of Economics*, 141–156.

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

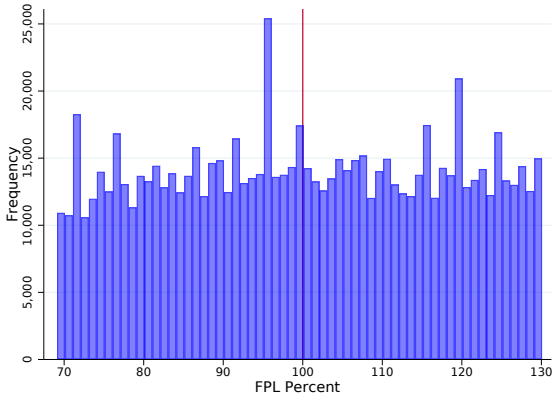
Starc, Amanda. 2014. "Insurer pricing and consumer welfare: Evidence from Medigap." *RAND Journal of Economics*, 45(1): 198–220.

Sun, Liyang, and Sarah Abraham. forthcoming. "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects." *Journal of Econometrics*.

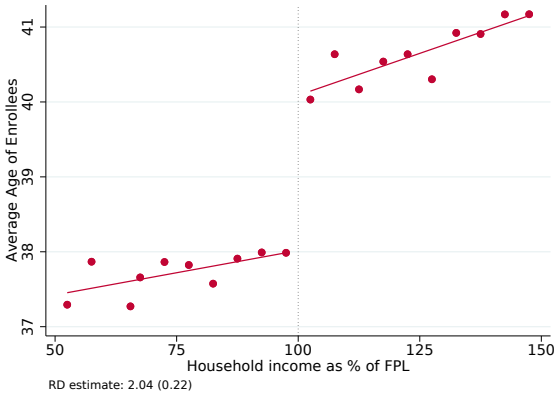
A Regression Discontinuity: Further Analysis

Appendix Figure A1. RD's

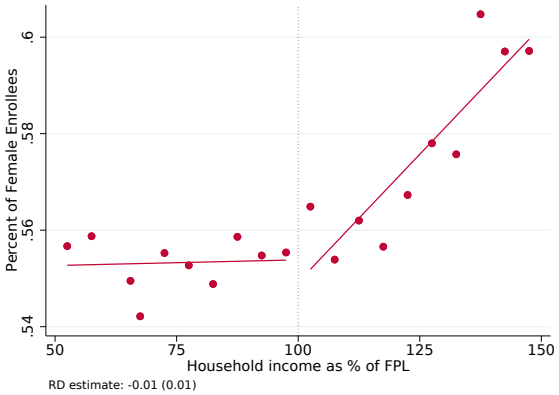
(a) Total Number of Enrollees



(b) Age



(c) Percent Female



B Counterfactual Simulations: Detailed Methods

B.1 Solving for Equilibria

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

To solve for the insurer combinations that deliver valid equilibria (as defined above in

Section 6.1), we search for valid price equilibria for all possible combinations of entrants. We adopt the following grid search approach.

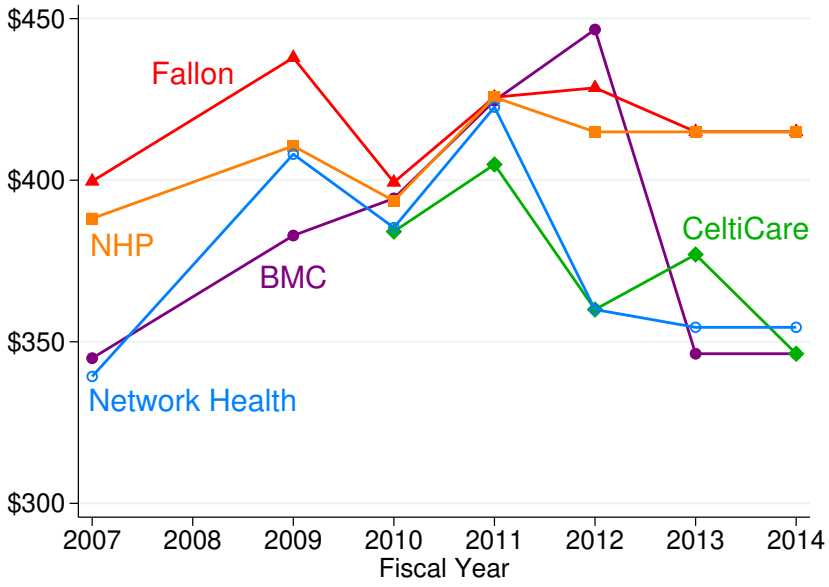
For each possible combination of plans, we evaluate first order conditions and profits for a grid of all possible price vectors, where each price takes one of 30 evenly spaced values from \$350 to \$500. For 4 firms, there are $30^4 = 810{,}000$ possible price vectors. We then identify candidate price vectors that satisfy the following four conditions. First, all firms' prices fall between the price floor and \$500. Second, all firms make positive profits. Third, the FOCs are satisfied within a pre-defined tolerance[21] Fourth, no firm can deviate to a higher or lower price (within the bounds of the price floor and \$500) and make higher profits. To account for imprecision from the grid, we allow price vectors in a +/- 1 grid point box around each of these candidate vectors. For each of the resulting candidate price vectors, we solve for the exact equilibrium prices using the fmincon function in Matlab to solve the system of FOCs within a box of +/- 2 grid points around the candidate price vector. Finally, we check whether this exact price vector continues to satisfy the four conditions above. This procedure delivers up to one candidate equilibrium price vector for each possible combination of firms.[22]

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.

[21] 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).

[22] Although we do not formally prove uniqueness of the equilibria we identify using this procedure, in practice we do not observe any cases where multiple different price vectors satisfy all of the conditions.

Appendix Figure A2. CommCare Plans Pre-Subsidy Prices



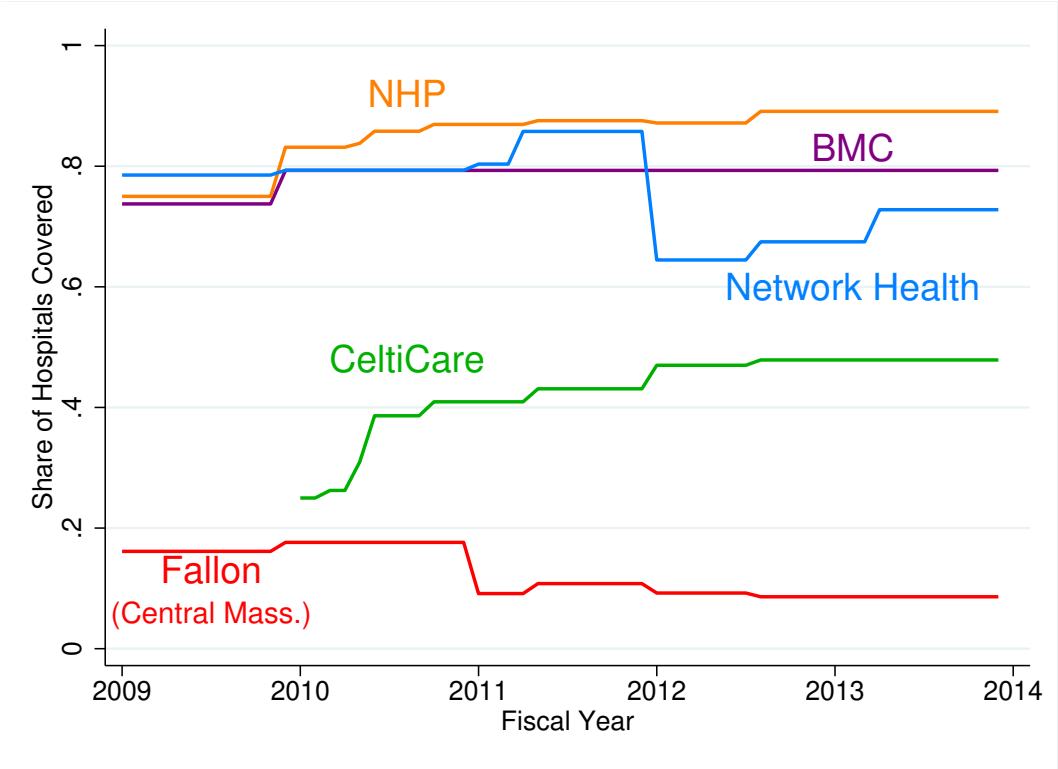
Note: The graphs show average pre-subsidy insurer prices for each insurer's plan in the CommCare market, by fiscal year. The five plans are shown in different colors and labeled. Values shown are averages for the plan's actual enrollees; underlying premiums and (in some years) prices vary by income group 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 A3. 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 A4. 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.