NBER WORKING PAPER SERIES

ADVERSE SELECTION AND (UN)NATURAL MONOPOLY IN INSURANCE MARKETS

Edward Kong
Timothy Layton
Mark Shepard

Working Paper 33187
http://www.nber.org/papers/w33187

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
November 2024

We are grateful to David Anderson, Chris Avery, Peter Blair, Andrew Ching, Shooshan Danagoulian, Coleman Drake, Liran Einav, Amy Finkelstein, Nathan Hendren, Kate Ho, Ami Ko, Lee Lockwood, Neale Mahoney, Joe Newhouse, Jim Poterba, Amanda Starc, Pietro Tebaldi, Nick Tillipman, Richard Zeckhauser for helpful comments and discussions. We especially thank Richard for suggesting the term "un-natural monopoly" used in our title. We thank David Anderson and Coleman Drake for sharing data on the ACA Marketplaces. We also thank seminar participants at MIT, Georgetown University, the Congressional Budget Office, Duke University, Northwestern University, the University of Chicago, the University of Virginia, Johns Hopkins University, the University of Illinois-Chicago, the University of Wisconsin-Madison, Brigham Young University, Erasmus University Rotterdam, the University of Pennsylvania, Stanford University, Tufts University, the Federal Trade Commission, the American Economic Association meetings, the Annual Meeting of the American Society of Health Economists, and the NBER Insurance, NBER Industrial Organization, and NBER Economics of Health meetings for comments. Ilana Salant, Bradley Katcher, and Kai Shen Lim provided superb research assistance. We gratefully acknowledge funding from Arnold Ventures and the Agency for Healthcare Research and Quality (K01-HS25786-01). Edward Kong acknowledges support from award Number T32AG51108 from the National Institute on Aging. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute on Aging or the National Institutes of Health. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Adverse Selection and (un)Natural Monopoly in Insurance Markets
Edward Kong, Timothy Layton, and Mark Shepard
NBER Working Paper No. 33187
November 2024
JEL No. D4, I11, I13, L1, L40

ABSTRACT

Adverse selection is a classic market failure known to limit or "unravel"' trade in high-quality insurance and many other economic settings. While the standard theory emphasizes quality distortions, we argue that selection has another big-picture implication: it unravels competition among differentiated firms, leading to fewer surviving competitors—and in the extreme, what we call "un-natural" monopoly. Adverse selection pushes firms toward aggressive price cutting to attract price-sensitive, low-risk consumers. This creates a wedge between average and marginal costs that (like fixed costs in standard models) limits how may firms can profitably survive. We demonstrate this insight in a simple model of insurer entry and price competition, estimated using administrative data from Massachusetts' health insurance exchange. We find a large "selection wedge" of 20-30% of average costs, which (without corrective policies) unravels the market to monopoly. Our analysis suggests a surprising policy implication: interventions that limit price-cutting can improve welfare by supporting more entry, and ultimately lower prices.

Edward Kong
Harvard University
Department of Economics
1805 Cambridge St
Cambridge, MA 02138
edkong@nber.org

Timothy Layton
Frank Batten School of Leadership
and Public Policy
University of Virginia
Charlottesville, VA 22904-4893
and NBER
timothyjlayton@virginia.edu

Mark Shepard
Harvard Kennedy School
Mailbox 114
79 JFK Street
Cambridge, MA 02138
and NBER
mark_shepard@hks.harvard.edu

1 Introduction

Adverse selection is a classic market failure in economics. Dating back to seminal insights by Akerlof (1970) and Rothschild and Stiglitz (1976), a growing body of theory and empirics shows how welfare-improving trade can break down when either buyers or sellers have private information relevant to the other party's payoffs. The insight that adverse selection limits or "unravels" trade (Hendren, 2013) is relevant in a wide variety of settings, from used cars (Akerlof, 1970) to labor markets (Greenwald, 1986; Stantcheva, 2014) to corporate finance (Michaely and Shaw, 1994). But perhaps the most prominent application of adverse selection is insurance, and especially health insurance. While this literature has provided motivation for public health insurance systems, in practice essentially all high-income nations have meaningful health insurance markets in one form or another, whether for comprehensive coverage (as in the Netherlands, Switzerland, Germany, Chile, Israel, and the U.S.) or for supplementary insurance (as in Canada, France, and the U.K.). Addressing adverse selection — and more generally, "making insurance markets work" — is therefore a critical policy objective (Einav et al., 2022).

While policymakers have developed many tools to address selection, most policies focus on ensuring *trade* and adequate *quality* of insurance benefits, since these are the key problems in standard models of adverse selection.[1] For instance, boosting trade (e.g., reducing uninsurance) and ensuring quality (e.g., covering essential benefits without exclusions for pre-existing conditions) were the central focus of the Affordable Care Act (ACA) and its suite of subsidies, mandates, and regulations.

However, as trade has increased (e.g., as uninsurance declines in the U.S.), there is growing recognition of another concern: *limited competition*. By conventional antitrust standards, health insurance markets are often highly concentrated.[2] Nowhere is this more evident than in the health insurance exchanges created by the ACA where as of 2021, one-fifth of enrollees lived in areas with just 1-2 participating insurers and 24 whole states had three or fewer competitors (McDermott et al., 2020). Given these facts, there is growing interest in understanding the industrial organization of "selection markets" and the implications for insurance market policy (Einav et al., 2021).

In this paper, we ask whether there may be a deep conceptual link between the economics of adverse selection and limited competition (or "natural monopoly") in insurance markets. We show that adverse selection, through its impact on strategic pricing incentives, creates a barrier to robust firm participation – operating much like (and in tandem with) fixed costs in standard models of natural monopoly. In short, adverse selection may unravel not just trade and quality (as in standard models) but also *competition* among differentiated firms, even when quality is fixed or regulated. Central to our analysis is the idea that price competition creates unraveling problems in selection markets, incentivizing aggressive price competition in a way that limits markets' ability to support a robust

[1]Unraveling of trade is the key problem in the line of work following Akerlof (1970), including Einav et al. (2010) and the many papers following it. Unraveling of quality is the key problem in Rothschild and Stiglitz (1976) and the (fewer) papers that model quality determination, either with perfect or imperfect competition (e.g., Handel et al. (2015); Veiga and Weyl (2016)). A few papers, such as Azevedo and Gottlieb (2017), seek to bridge these two approaches.

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

number of participating firms. Addressing adverse selection, therefore, may involve policy that limits or softens price-cutting incentives, which (counter-intuitively) may lead to *lower* equilibrium prices and higher consumer welfare by stimulating firm entry.

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

We then show that adverse selection — which occurs when low-risk (healthy) consumers are more price sensitive than high-risk (sick) consumers — intensifies price competition, a result pointed out by Starc (2014) and Mahoney and Weyl (2017). The intuition mimics the standard competition logic, but with a twist. With adverse selection, when a firm raises its price, it doesn't just lose marginal consumers, it loses its *lowest-cost* consumers, driving up its average costs and reducing profit margins. With a fixed set of competitors, therefore, adverse selection *constrains* the use of market power to raise prices and thus augments price competition. Thus, holding the number of participants fixed, more selection implies lower prices and better outcomes for consumers.

Our model's central point is that this conclusion is incomplete: Adverse selection affects not just price competition among a given set of firms, but *how many* firms are willing to enter a market. Precisely because it leads to more extreme price competition and lower prices, adverse selection makes it difficult for a robust number of firms to profitably compete in a market while covering fixed costs and the high cost of (relatively sick) inframarginal consumers. In the extreme, a market may be unable to support more than a single firm, resulting in limited consumer choice and higher (monopoly) prices than if more firms entered.[3]

Consider the following example to illustrate our logic. Suppose there is an insurance market with two insurers, each of which offers a single plan that appeals to an (equal-risk) subset of consumers in the market. For instance, one plan may have a better network of covered doctors on the east side of the market, while the other is better on the west side. Both plans have equal cost structures, and neither is superior for all consumers. Optimally, both plans would participate in the market and set equal prices (e.g., at the market-wide average cost, AC^{Mkt}), with consumers able to sort into their preferred option. Indeed, this would be the result in a model of perfect competition, in which firms are assumed to mechanically set prices at average costs.

With *strategic* price competition, however, adverse selection can undermine this desirable outcome. At equal prices, each firm has an incentive to undercut the other to attract price-sensitive healthy

[3] A natural question is whether the *threat* of entry can constrain prices, even if few firms actually enter. In our model, this does not occur because the threat is non-credible: monopoly is an equilibrium only when all non-entrants know they will lose money if they enter. Of course, the degree to which entry threat (or "contestable markets") constrain prices in reality is an empirical question with a long history in economics (see e.g., Spence (1983)).

consumers — a type of "cream-skimming" via price cuts. With strong enough adverse selection, a firm can *increase* its profit margin by *cutting* prices because a price decrease of $1 leads to a decrease in the firm's average cost of more than $1. While privately attractive, this cream-skimming is ultimately zero-sum and can make it impossible for the other firm to survive unless it also follows suit. We show that when the cream-skimming incentive exceeds the standard incentive to extract markups from inframarginal consumers (the Lerner markup), the market can no longer support a Nash equilibrium in prices where both firms participate and break-even.[4] Instead, the only equilibrium is one with a single participating insurer, which deprives consumers of variety and also results in higher prices.

We formalize this cream-skimming incentive in what we call the "selection wedge," which is the gap between average and marginal costs when a firm j competes on price ($= AC_j - MC_j$). We show that this selection wedge enters the conditions for equilibrium entry in a similar way as fixed costs, the standard entry-limiting force. Formally, for a firm to profitably compete, its Lerner markup (its inverse semi-elasticity of demand) must exceed the sum of its selection wedge plus its fixed costs per consumer.[5] Stronger adverse selection, therefore, limits entry just like large fixed costs in the standard theory of natural monopoly.

After laying out the theory, we then assess empirically whether adverse selection could actually lead to limited firm participation in the real world. To do so, we turn to data from Massachusetts' Commonwealth Care (CommCare) insurance exchange, the state's pre-ACA health insurance marketplace for low-income individuals. The market consisted of 4-5 competing insurers, each offering a single plan with standardized cost sharing provisions but differing hospital and physician networks. Consistent with our model, there were not clear "vertical" quality rankings among plans; instead, they differed *horizontally* on networks, with different plans appealing to different consumers.

We start by providing transparent descriptive evidence of the strong undercutting incentives induced by adverse selection. To do so, we estimate summary measures of price sensitivity and the effects of price changes on average plan costs in a quasi-experimental difference-in-differences design based on year-to-year plan price changes. Our analysis leverages a control group of individuals below 100% of the federal poverty line (FPL), who have access to the same plan menu but at fully subsidized ($0) premiums. These quasi-experimental estimates indicate that consumers in this market are very price sensitive and that price sensitivity is much stronger among healthy versus sick consumers, the key feature of adverse selection. At observed prices, a price decrease of $1 generates a shift in the composition of consumers choosing the firm's plan that lowers the plan's average medical cost by around $1 — a critical condition that our model suggests leads to unraveling of firm participation in the absence of corrective policies, even without fixed costs.

[4] Instead, both firms have incentives to undercut each other's prices to below overall market average costs. This undercutting incentive results in either a pure strategy Nash pricing equilibrium where one or both firms loses money or no pure strategy pricing equilibrium for a given set of firms. In cases with no pure strategy equilibrium, there is always a mixed strategy pricing equilibrium (Nash, 1951), and we find that at least one firm loses money in expectation.

[5] The Lerner markup is a function of product differentiation, so in principle adverse selection might result in greater product differentiation. In practice, this may be limited in health insurance markets where product attributes are highly regulated. Further, differentiation on quality attributes can also increase the adverse selection wedge, offsetting the larger Lerner markup.

We then use the same quasi-experimental price variation to estimate a full structural model of demand and costs for insurance plans in the Massachusetts market. We use the model to simulate equilibria with and without various corrective policies that were used by Massachusetts, including risk adjustment, incremental subsidies for higher-price plans, and price floors (imposed under rules requiring "actuarial soundness"). We show that these corrective policies were critical for achieving the levels of insurer participation observed in this market. Without them, the undercutting incentives are so strong that the only surviving equilibrium is one with a single monopoly firm. In the absence of these policies, choice would thus be limited and prices would be high. The issues raised by our theoretical model thus appear to be empirically relevant in this setting.

We use the structural model to perform counterfactual simulations exploring the effects of two key corrective policies featured in the Massachusetts market: risk adjustment and price floors. Risk adjustment reduces adverse selection by enforcing transfers from firms that attract (observably) healthier consumers to those that attract sicker people. Price floors directly limit price undercutting, much like benefit regulation (e.g., essential health benefits rules) prevents quality unraveling.

Our simulations show that consistent with the theory, these corrective policies can indeed increase firm participation. Strong risk adjustment that offsets 60% of cost selection (which is far stronger than was the market's actual risk adjustment) allows the market to support duopoly, while near-perfect risk adjustment (offsetting $> 80\%$ of selection) can induce full entry among all the firms we simulate. Price floors just above market-wide average costs can induce three firms to participate, while floors about 16% above average costs can induce full participation. Interestingly, across a range of simulations, we find that the optimal policy almost always involves some (binding) price floors, unless risk adjustment is near-perfect (which is likely infeasible). Indeed, even ignoring risk adjustment feasibility constraints, the global optimal policy is *not* perfect risk adjustment (i.e., the complete removal of selection) but instead involves *moderate* risk adjustment (offsetting half of cost variation) and moderate price floors (just above marketwide average costs). The intuition for this result is that a little selection can be a good thing, as it encourages firms to keep prices low (Starc, 2014; Mahoney and Weyl, 2017), as long is it is not so strong as to cause firms to leave the market entirely.

Ultimately, our results suggest that undercutting incentives caused by adverse selection have important implications for market stability in health insurance markets like the Massachusetts exchange. Our simulations suggest that without policies like risk adjustment and price floors, this market would have low firm participation, high premiums, and limited choice. Moderate risk adjustment goes a long way toward improving the situation, and price floors are typically welfare improving.

These results suggest a novel explanation for why so many insurance markets may have trouble sustaining robust participation, alongside the traditional explanations associated with fixed and sunk entry costs. Indeed, our results may be particularly helpful for understanding limited insurer participation in the ACA health insurance exchanges. In the last section of the paper, we consider trends in ACA market participation in light of our theory. We start by showing that broad patterns of insurer participation in this market are consistent with our model. Specifically, we show that participation dropped significantly just after the end of a federal reinsurance program, which reimbursed insurers for high-cost enrollees and thus cushioned adverse selection. We also show that in more recent years participation has rebounded, and, importantly, that this rebound coincides with the proliferation of "zero-dollar" plans, for which a consumer's subsidy exceeds the plan's gross premium. In this market, there is an effective price floor at zero, and in recent years this price floor has become increasingly binding, as subsidies (which are linked to "silver" tier plan prices) have grown to exceed prices of "bronze" tier plans for most consumers. We argue that the increasing relevance of this effective zero-dollar price floor may help explain the rebound in firm participation since 2019. Finally, we offer a more well-identified test of the theory, using variation in exposure to the end of the reinsurance program across states to implement a difference-in-differences design. Consistent with our theory, we find that the end of the reinsurance program led to lower insurer participation, accounting for about 20% of the overall decline in the number of insurers per county from 2014-16 to 2017-18.

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

Related Literature Our paper contributes to several literatures. First, we contribute to work on adverse selection in insurance markets. It has long been recognized that adverse selection can distort prices and contracts (Rothschild and Stiglitz, 1976; Einav et al., 2010; Bundorf et al., 2012; Azevedo and Gottlieb, 2017; Chade et al., 2022) and can make offering choice less desirable (Ericson and Sydnor, 2017; Marone and Sabety, 2022; Hendren et al., 2021). Previous work has also shown that selection can fully unravel trade in insurance (Akerlof, 1970; Hendren, 2013). Our paper shows that even when subsidies or mandates ensure that trade occurs, adverse selection can still limit the market's ability to support multiple competing firms, with important implications for consumer welfare.[6]

Our work also builds on previous work studying the interaction of imperfect competition and selection (Starc, 2014; Mahoney and Weyl, 2017). 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 reduce the number of competing firms, potentially outweighing the impacts of markups conditional on participation. When market structure is endogenous, corrective policies such as risk adjustment can sustain higher participation and therefore *lower* markups.

Our paper also contributes to a growing literature studying policies used to combat selection, such as risk adjustment, subsidies, and contract and price regulation (see Geruso and Layton (2017); Einav et al. (2018)). This literature has shown that these policies can sometimes be beneficial. Some work has also established a variety of unintended consequences of these policies (Geruso and Layton, 2020;

[6] As far as we know, this is a new insight. There is prior work showing that adverse selection can be a barrier to entry in banking (Dell'Ariccia et al., 1999), but the mechanism (information advantages for incumbent firms) is quite different.

Geruso et al., 2023*b*). Our paper introduces an additional benefit of these policies: They can improve consumer welfare by allowing the market to support more competitors. We also propose a new policy to combat selection problems — price floors.

Finally, our paper contributes to the literature on firm entry in industrial organization (see Berry and Reiss (2007) for a review). This work has focused on fixed and sunk costs and the nature of competition as explanations for limited entry. Our work shows the role adverse selection can play in shaping entry outcomes, including leading to limited entry in settings without fixed costs. Importantly, while our results may help explain a variety of facts about firm participation in insurance markets, our model is (intentionally) not a perfect model of entry and exit in insurance markets, and we do not incorporate many features like dynamics, tacit collusion, etc. Our primary goal is to extend the canonical simple model of firm participation to incorporate a key issue for understanding insurance markets (adverse selection) and to show that it can explain a variety of previously difficult-to-understand facts. Future work may wish to incorporate these additional features to better understand the interactions between adverse selection, price competition, and entry.

Outline of Paper The paper proceeds as follows. Section 2 presents a model that conveys the paper's main conceptual points. Section 3 introduces the empirical setting and data, and Section 4 shows reduced form evidence on key parameters in this market. Sections 5-6 estimate a structural model and use it to simulate policy counterfactuals in the Massachusetts market. Finally, Section 7 discusses and applies our ideas to the ACA exchanges and other settings, and Section 8 concludes.

2 Theory: Adverse Selection and Limited Competition

In this section, we present a model of firm entry and competition in a market with adverse selection. Our goal is to show how adverse selection works to limit equilibrium entry — very much like fixed costs in the classic theory of natural monopoly and monopolistic competition. The key insight is that adverse selection *intensifies price competition*, since it pushes firms to aggressively cut prices to attract (or "cream-skim") price-sensitive low-cost consumers. Just as adverse selection pushes firms to cut *quality* in order to cream-skim in classic selection models (Rothschild and Stiglitz, 1976), it likewise pushes firms to cut *prices* to cream-skim in markets with flexible price competition.

While low prices directly benefit consumers *conditional* on the set of competing firms, they also limit how many firms can profitably compete and survive, thus reducing desirable product variety. When strong enough, selection unravels all multi-firm competition, leading to an (undesirable) equilibrium we call "un-natural" monopoly. A key insight of our model is that price competition can become *too* intense in selection markets, yielding un-competitive equilibria with low variety and (counterintuitively) high prices. We show that policies that soften or limit downward price competition – notably, price floors – can, if used carefully, generate strict improvements in consumers welfare.

We begin in Sections 2.1-2.2 with a simple but flexible model that, in a concrete way, illustrates all of our key insights. Next, Section 2.3 shows how these ideas apply in a general model, which is the basis of our empirical work in the remainder of the paper.

2.1 Simple Model: Salop Circle with Adverse Selection

We start by using a stylized model to provide intuition for our core ideas. The model starts with the classic Salop (1979) model of monopolistic competition and simply incorporates cost heterogeneity and risk selection. In the Salop model, a mass-1.0 population of consumers (indexed by i) reside uniformly around a small city, modeled as a unit-circumference circle. A set of N competing firms locate equidistantly around the circle, where N is an equilibrium object. Each firm ($j = 1, ..., N$) sells a homogeneous product (or "plan") of value V_i to each consumer i (i.e., there are no "vertical" quality differences between products), but consumers dislike travel so prefer nearby firms. Firm location, therefore, captures pure "horizontal" differentiation. In health insurance, this includes features like local provider network differences, varying brand preferences, and consumer loyalty based on past experience.

Figure 1 visualizes this setup in the simplest case with $N = 2$ firms located on the west ($\ell_W = 0$) and east ($\ell_E = \frac{1}{2}$) ends of the city. One way to think about this particular configuration is as a city with two hospitals on opposite sides, each exclusively covered by one of the two insurers. At equal premiums, consumers living on the west side prefer plan $j = W$, while east-side consumers prefer plan $j = E$. Consumers near the middle are marginal and can be swayed by a lower price to the opposite-side plan.

Firms compete in a simple two-stage entry and pricing game. In stage 1, firms decide whether to enter, which involves incurring fixed costs $F \geq 0$. In stage 2, the N entrants compete on prices, P_j, to maximize profits, $\pi_j(P)$, in standard Nash-Bertrand equilibrium, where $P = \{P_1, P_2, ..., P_N\}$ is the full vector of prices. In this classic model, prices and markups above *marginal* costs are determined based on consumers' price elasticity of demand (the inverse of the travel cost). But (symmetric) firms must cover not just their marginal/variable costs but also their *fixed* costs to profitably participate. Thus, the consumer price elasticity of demand and the markup implied by this elasticity implies a cap on the number of firms that can profitably participate in the market (derived below).

Extension to include adverse selection: The classic Salop model includes a single consumer type with homogeneous cost, $C > 0$, and travel disutility, $t > 0$. We enrich this setup by allowing for two risk types, $r \in \{L, H\}$: (1) healthy (low-risk) type L with costs C_L who comprise share θ_L of consumers, and (2) sick (high-risk) type H with costs $C_H > C_L$ who comprise share $\theta_H = 1 - \theta_L$. Each type is uniformly distributed around the city (preserving firm symmetry), but types may differ in their travel disutility, $t_L \neq t_H$ (and thus their price elasticity). Firms cannot price discriminate across types, implying that there is cost-relevant asymmetric information (or "community rating" regulation banning price discrimination). We denote the overall market average cost as $\overline{C} \equiv E[C_r] = \sum_r \theta_r C_r$, and likewise for average travel disutility, $\bar{t} \equiv E[t_r]$.

Demand is determined by consumer preferences. The utility of consumer i of risk type r for firm j is: $U_{irj} = (V_i - t_r \cdot d_{ij}) - P_j$, where d_{ij} is travel distance from consumer i to firm j. Assuming

Figure 1. Salop Circle Model with Adverse Selection ($N = 2$ firm case)

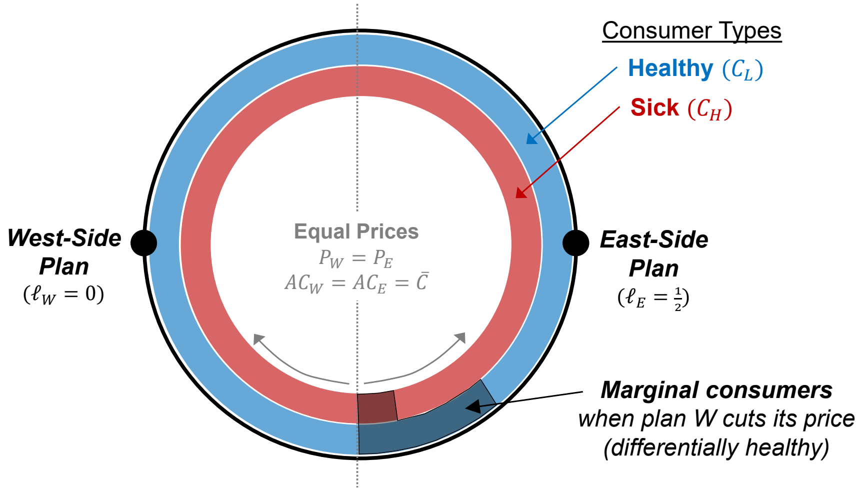

Note: The figure visualizes our simple model, which extends the Salop (1979) circular city model to include adverse selection. Consumers locate uniformly around a city, and N firms ($N = 2$ is shown) locate equidistantly around the city and compete on prices to sell a product. We extend the model by including two consumer types: (1) low-risk (healthy) consumers with costs C_L and price-sensitivity α_L, and (2) high-risk (sick) consumers with costs $C_H > C_L$ and price-sensitivity α_H. The key driver of adverse selection is that *healthy consumers are more price-sensitive* ($\alpha_L > \alpha_H$). When one firm cuts its price (e.g., as shown for plan W), their *marginal* consumers are differentially price-sensitive healthy types with lower costs. This gives firms a "cherry picking" incentive to compete aggressively on prices.

for simplicity that all consumers buy exactly one good (i.e., no outside option), the uniform location distribution yields a simple linear demand (market share) function for each type r:[7]

$$D_{rj}(P) \;=\; \tfrac{1}{N} \;-\; \alpha_r \cdot \left(P_j - \overline{P}_{j'}\right) \tag{1}$$

where $\overline{P}_{j'}$ is the average price of j's adjacent firms, and $\alpha_r \equiv 1/t_r > 0$ is the "*price sensitivity*" (or demand slope) of type r, equal to their inverse travel disutility. t_r captures how much a consumer values *product differences* (i.e., firm location in the model), so α_r is increasing in *product similarity*, with more similar products leading to more elastic demand. In our example of the city with two hospital systems, consumers with lower α_r (higher t_r) are more "attached" or "loyal" to a particular hospital and thus less likely to leave the plan that covers this hospital as the price increases.

Summing across types, firm j's total market share is $D_j(P) = \sum_r \theta_r D_{rj}(P) = \frac{1}{N} - \overline{\alpha} \cdot \left(P_j - \overline{P}_{j'}\right)$, which is again linear in (relative) price with slope $-\overline{\alpha} = -\sum_r \theta_r \alpha_r$, market-level average price-sensitivity. Total *variable* firm profits equal type-specific profit margins times type-specific demand, summed over types,

$$\pi_j(P) \;=\; \sum_r \left[P_j - C_r\right] \cdot \theta_r D_{rj}(P) \;=\; \left[P_j - AC_j(P)\right] \cdot D_j(P) \tag{2}$$

[7] See Appendix A.1 for derivations for the Salop model. Note that this and subsequent demand formulas assume non-zero market shares for all products (i.e., $D_{rj}(P) \in (0,1)\ \forall j, r$), which is the relevant case around a symmetric pricing equilibrium, and in general whenever $\left|P_j - \overline{P}_{j'}\right| < \frac{1}{N \cdot \alpha_r}$ for all j and r.

where $AC_j(P)$ is the firm's average (variable) costs at prices P. Importantly, average costs can vary with (all firms') prices because prices affect which types of consumers select into each plan. This is a defining feature of insurance and other "selection markets" (Einav et al., 2021).

Pricing Incentives under Adverse Selection: The key driver of adverse selection in the model is that high-cost H types place greater value on product differences or, equivalently, have higher travel costs ($t_H > t_L$) implying lower price sensitivity ($\alpha_L > \alpha_H$). Stated in general terms:

$$\textbf{Adverse Selection:} \qquad Cov(t_r, \alpha_r) > 0 \iff Cov(C_r, \alpha_r) < 0 \tag{3}$$

In our two-type model, this simplifies to $t_H > t_L$ or, equivalently, $\alpha_L > \alpha_H$ (given that $C_L < C_H$), but (3) captures the general condition relevant for richer multi-type models. This condition, that low-cost consumers tend to be more price-sensitive, is natural (and testable) in health insurance markets, and there is substantial evidence that it holds in practice, both in our data and in prior work.[8] It generalizes the classic notion of adverse selection in vertical models – that high-risk types have greater demand for *more generous* insurance – to settings with more general (horizontal) product differentiation.

Before deriving equilibrium firm participation and premiums under the stylized model, we explore general firm pricing incentives under adverse selection, as these incentives are critical to understanding the equilibrium results we present later. Consider what happens if starting from equal prices $P_j = P$ $\forall j$, one firm j unilaterally undercuts its competitors, as shown for plan W in Figure 1. Because $\alpha_L > \alpha_H$, its marginal consumers from this price cut are differentially low-cost L types. In the example, the healthy east-side consumers move to W because they value lower premiums and don't care too much about travel time while the sick ones stay in E due to their high weight on travel time (likely because they expect to go to the hospital much more than the healthy). It is straightforward to show that the cost of the marginal consumers, or the firm's "*marginal cost*" at prices P, equals $MC_j(P) = \frac{1}{\overline{\alpha}} \sum_r (\theta_r \alpha_r) \cdot C_r$.[9] If the adverse selection condition holds, this marginal cost is less than the firm's prior average cost, $AC_j(P) = \overline{C} = \sum_r \theta_r C_r$, because there are more healthy, low-cost consumers than sick, high-cost consumers in the marginal group. Importantly, this implies that a price cut reduces the firm's average costs, or $\partial AC_j / \partial P_j > 0$. This is the defining feature of what we call the "*adverse selection undercutting*" incentive.

A simple measure of the strength of the undercutting incentive is the gap between average and marginal cost, which we call the "*selection wedge*":

$$\textbf{Selection Wedge} \equiv AC_j(P) - MC_j(P) \overset{\text{Salop}}{=} \sum_r \left[\theta_r - \left(\tfrac{\theta_r \alpha_r}{\overline{\alpha}}\right)\right] \cdot C_r = \tfrac{-Cov(C_r, \alpha_r)}{\overline{\alpha}} \tag{4}$$

Under the adverse selection condition in (3), the selection wedge is positive and grows one-for-one with the covariance $-Cov(C_r, \alpha_r)$. A related measure of the strength of the undercutting incentive is

[8] See e.g., Finkelstein et al. (2019), Saltzman (2021), and Tebaldi (2024).

[9] Because of linear demand in the Salop model, marginal cost is a price-invariant constant as long as each plan gets positive market share for all risk types. More generally, $MC_j(P)$ may vary with prices (see general theory in the next subsection).

the price-slope of the firm-specific average cost curve, which shows how responsive firm average costs (i.e. the selection of healthy versus sick consumers) are to changes in price. It is straightforward to show that this slope is equal the following:

$$\frac{\partial AC_j}{\partial P_j} = \underbrace{\eta_{j,P_j}(P)}_{\text{Demand semi-elasticity}} \times \underbrace{[AC_j(P) - MC_j(P)]}_{\text{Selection wedge}} \overset{\text{Salop}}{=} -Cov(C_r, \alpha_r) \times N \quad (5)$$

where $\eta_{j,P_j}(P) \equiv -\frac{\partial D_j/\partial P_j}{D_j(P)} > 0$ is the firm's own-price semi-elasticity, which is positive by the law of demand. The first equality in (5) is general and applies regardless of market structure, prices, and demand/cost functional forms (i.e., it is not specific to this stylized model). It says that the "steepness" of the firm's average cost curve is proportional to the demand semi-elasticity (how *many* consumers are swayed by price cuts) times the selection wedge (how much *cheaper* those marginal consumers are than the average consumer). The second equality in (5) shows how this slope simplifies in the Salop model. The average cost curve slope is proportional to the selection covariance, but it also grows with the number of firms, N. This illustrates an insight, first noted in work by Lustig (2010), that (cross-firm) adverse selection undercutting incentives tend to grow stronger in markets with more competitors. As we show below, this strengthening selection with more competitors plays a key role in the logic behind how selection limits the number of firms a selection market can support.

Parallel: Adverse Selection and Fixed Costs Before proceeding further, we note a key parallel between the economics of adverse selection pricing and that of fixed costs, which holds in both the Salop model and more generally. Adverse selection implies that a firm's average costs increase with prices ($\partial AC_j/\partial P_j > 0$), or in quantities, a *downward-sloping* average cost curve ($\partial AC_j/\partial D_j < 0$). It also implies a positive "wedge" between average and marginal costs. Both of these features are familiar from the classic textbook model of adverse selection in Einav and Finkelstein (2011).

We observe that the same two features are also true of *fixed costs* in the classic theory of natural monopoly. To see this, define a firm's average *total* costs (including fixed costs) as $ATC_j(P) = AC_j(P) + \frac{F}{D_j(P)}$, which lets us write net profits as $\pi_j^{Net}(P) = [P_j - ATC_j(P)] \cdot D_j(P)$. In a non-selection market (with constant consumer marginal costs C), fixed costs create a wedge between average total costs and marginal costs, and they also imply that $\partial ATC_j/\partial P_j > 0$, or a downward-sloping ATC_j curve in quantity. In general, using (5), we can write the price-slope of average total costs as:

$$\frac{\partial ATC_j}{\partial P_j} = \underbrace{\eta_{j,P_j}(P)}_{\text{Demand semi-elasticity}} \times \left[\underbrace{AC_j(P) - MC_j(P)}_{\text{Selection wedge}} + \underbrace{\frac{F}{D_j(P)}}_{\text{Fixed costs}} \right] \quad (6)$$

which is the product of the demand semi-elasticity times the sum of the selection wedge and fixed costs per consumer. Appendix Figure A1 depicts these parallels graphically, with the left panel showing a textbook graph of natural monopoly due to fixed costs, and the right panel showing a textbook graph of adverse selection. This close parallel between the two concepts has (to our knowledge) not been previously highlighted. It lies at the core of our paper's argument about how adverse selection tends to push markets towards limited competition.

2.2 Implications of adverse selection for firm entry and other market outcomes

We now solve for equilibrium firm participation and premiums in the Salop model. As is standard, we proceed backwards, starting with pricing (given N competitors) and then equilibrium entry (solving for N). We derive four implications for the "comparative statics" of stronger adverse selection, captured by a larger (more negative) selection covariance, $-Cov\left(C_r, \alpha_r\right)$, while holding fixed other model parameters (including $\overline{C}$ and $\overline{\alpha}$).[10]

Start with optimal pricing. Assuming a symmetric equilibrium (with $P_j^* = P^*$ and $D_j^* = \frac{1}{N}$ and $AC_j\left(P^*\right) = \overline{C}$ for all j), the FOC for profit maximization ($\frac{\partial \pi_j}{\partial P_j} = 0$) implies that:

$$P^* = \underbrace{\sum_r \left[\left(\frac{\theta_r \alpha_r}{\overline{\alpha}}\right) \cdot C_r\right]}_{\text{Marginal costs } = MC_j} + \underbrace{\left(\frac{1/N}{\overline{\alpha}}\right)}_{\text{Lerner Markup}} = \overline{C} + \underbrace{\left[\left(\frac{1/N}{\overline{\alpha}}\right) - \left(\overline{C} - MC_j\right)\right]}_{\text{Profit Margin}} \quad (7)$$

This equation shows two ways of understanding the implications of adverse selection for pricing. The first equality shows that — just as in standard non-selection markets — firms set prices equal to marginal costs (MC_j) plus the standard Lerner Markup, equal to one over the price semi-elasticity of demand, $\eta^{-1}_{j,P_j} \equiv \frac{D_j}{-\partial D_j / \partial P_j} = \frac{1/N}{\overline{\alpha}}$. However, with adverse selection, marginal costs are less than average costs, with the gap being the selection wedge, $\overline{C} - MC_j > 0$. Importantly, the second equality shows that the equilibrium price is actually a function of this selection wedge, with a larger selection wedge implying lower prices and lower profits margins, reflecting firms' stronger incentive to cut prices to attract price-sensitive low-cost types. This is our first implication of adverse selection:

Implication #1: *Conditional on the set of competing firms, adverse selection intensifies price competition and therefore reduces prices and profit margins.*

Implication #1 aligns closely with prior work on the interaction between adverse selection and imperfect competition (Starc, 2014; Mahoney and Weyl, 2017). This prior work shows that adverse selection disciplines market power, creating an interesting "theory of the second best"-style trade-off between mitigating selection incentives (e.g., with policies like risk adjustment) versus reducing market power. The key take-away from that prior work is that in settings with market power, policymakers should not seek to eliminate adverse selection, but to balance selection and competition.

Importantly, however, this prior work *holds fixed* the set of competitors in the market. With endogenous entry, the second implication of our analysis is that adverse selection limits equilibrium entry and competition:

[10] In the model, this corresponds to a larger cost variance or "spread" between C_L and C_H, while holding mean costs ($\overline{C}$) fixed. In the real world, this could correspond to a weakening of selection-mitigating policies like risk adjustment transfers. We formally simulate the impact of stronger/weaker risk adjustment in our structural analysis in Section 6. Our comparative statics approach follows that of Mahoney and Weyl (2017), who study a similar mean-preserving spread for costs that "rotates" the average cost curve.

Implication #2: *Adverse selection limits how many competing firms (N) a market can support in equilibrium.*

To see this, consider the conditions for equilibrium entry. In equilibrium, all firms must earn sufficient variable profits to cover their fixed costs, or $\pi_j(P^*) = [P^* - AC_j(P^*)]\, D_j(P^*) \geq F$. Rearranging, this implies:

$$\underbrace{P_j^* - AC_j(P^*)}_{\text{Profit margin}} = \underbrace{\frac{1}{\eta_{j,P_j}}}_{\text{Lerner Markup}} - \underbrace{[AC_j(P^*) - MC_j(P^*)]}_{\text{Selection wedge}} \geq \underbrace{\tfrac{F}{D_j(P^*)}}_{\text{Fixed costs per consumer}} \tag{8}$$

We call (8) the "profitable participation" condition. We write it in general notation, since its logic is general, though in the Salop model certain terms simplify as noted above. This condition implies a maximum number of firms that can profitably participate in the market. With $F = 0$, the limit on entry can be derived explicitly for the Salop model as:

$$N_{max}^{AS} \leq \frac{1}{\overline{\alpha}} \times \frac{1}{(AC_j - MC_j)} = \frac{1}{-Cov(C_r, \alpha_r)} \tag{9}$$

Even with $F = 0$ the number of entrants cannot grow larger than N_{max}^{AS}, and this limit is smaller as adverse selection grows stronger (a larger selection wedge and/or covariance). Further — and consistent with the parallel noted above — adverse selection and fixed costs work together to limit how much competition a market can sustain. This can be observed by deriving the (implicit) limit with both fixed costs and adverse selection under the Salop model:

$$N \leq \frac{1}{-Cov(C_r, \alpha_r) + \overline{\alpha} \cdot F \cdot N} \tag{10}$$

These characterizations on the maximum number of firms that can profitably participate in a market show that adverse selection and fixed costs enter in an additive way in this formula to constrain entry.

The logic here is the same as in the classic Salop model: Lower profit margins lead to less entry. The key difference is that now margins can be lowered not just by stronger price sensitivity among consumers but also by stronger adverse selection. We visualize one way this might play out in Figure 2, which is based on our stylized model. In each panel, the x-axis is the number of competitors, N, and the graphs show two sets of curves: (1) profit margins in blue, and (2) fixed costs per consumer (at the firm level) in orange. Panel A shows the standard case of *no selection*, in which $AC_j = MC_j$. As a result, profit margins equal the Lerner markup, which is always positive. As more firms enter, markups decline as price competition intensifies (larger η_{j,P_j}), while fixed costs per consumer rise as more firms split the fixed market size. Equilibrium occurs at the largest integer, N^0, to the left of these two curves' intersection point. Notably, as $F \to 0$ (corresponding to lower fixed costs or a growing market size) N^0 grows in an unlimited way, as it does in the Salop model.

Panel B shows the impact of adverse selection. Because of the selection wedge, profit margins $(P_j - AC_j)$ are less than the Lerner markup $(P_j - MC_j)$. This downward shift means both that fewer

Figure 2. Adverse Selection and Equilibrium Firm Participation

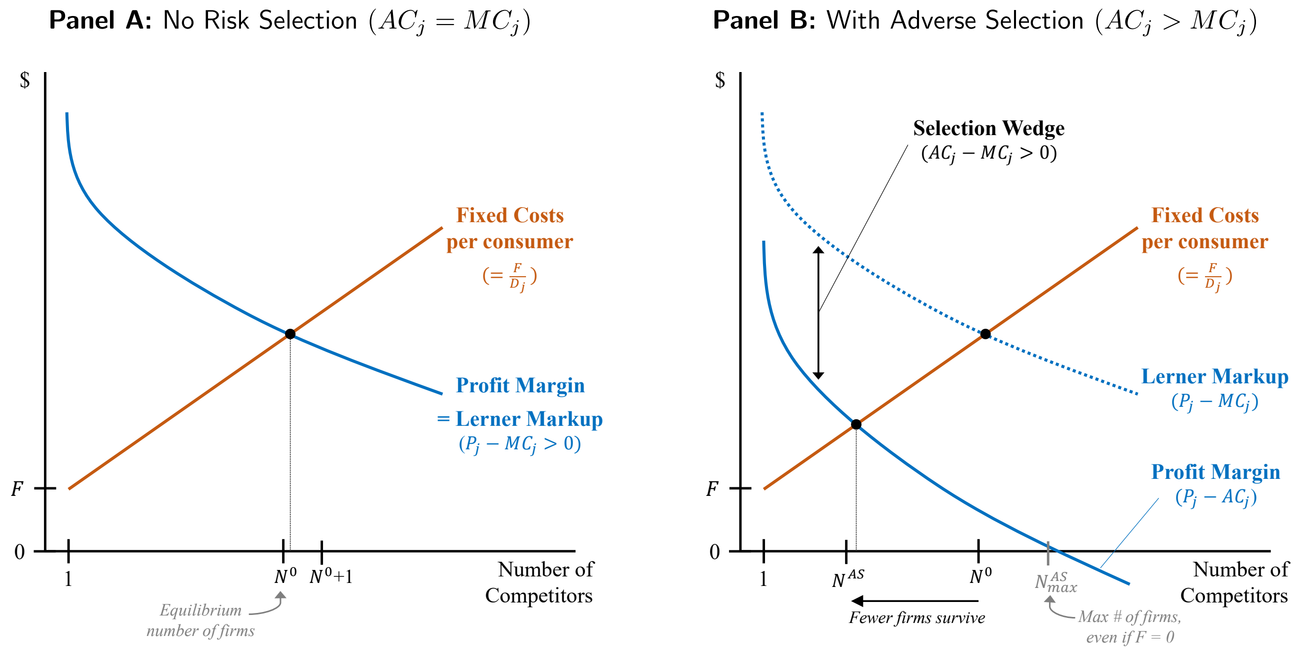

Note: The figure shows the implication of adverse selection for equilibrium firm entry into a market, based on the condition in equation (8). See the body text for a detailed discussion. An equilibrium occurs at the maximum integer number of firms, N, at which the profit margin ($P_j - AC_j$) exceeds fixed costs per consumer (F/D_j). Panel A shows this for a non-selection market, where average and marginal costs are equal, and the profit margin equals the (positive) Lerner markup over marginal costs. Panel B shows that adverse selection (which implies that $AC_j > MC_j$) drives a wedge between profit margins and the Lerner markup, reducing the number of firms that can survive. Moreover, there may be a maximum number of firms that can enter, even with $F = 0$, which occurs where the profit margin curve becomes negative.

competitors can profitably enter ($N^{AS} < N^0$) and that the profit margin curve need not always be positive. The integer just below where the profit margin curve goes negative, N^{AS}_{max}, reflects the largest number of firms the selection market can support as $F \to 0$.

The entry limit can become quite strong given realistic parameters. To see this, Figure 3 presents results from a simple calibrated version of this model, based on parameters from our empirical setting.[11] In each panel, the x-axis is the degree of adverse selection, captured by the ratio C_H/C_L, which varies from 1.0 (equal cost; no selection) up to 5.0 (sick H enrollees have costs 5x that of the healthy L types), while holding fixed average costs at $\overline{C} = \$400$. Panel A shows equilibrium firm entry according to the limit on the number of firms derived in equation (10). Consistent with implication #2, entry declines steeply with stronger adverse selection, declining from $N = 6$ firms with no selection to $N = 2$ (duopoly) at $C_H/C_L = 2.0$ and to $N = 1$ (monopoly) for $C_H/C_L > 3.6$. This amount of cost variation

[11] Specifically, we set market average costs at $\overline{C} = \$400$ per month, near the average in our data. We set $\theta_H = \theta_L = 0.5$ and set the price semi-elasticity of demand as $\eta_{j,P_j} = 2.4\%$ per \$1 of price increase (see Table 1), which with four firms (as in our data), implies $\overline{\alpha} = -\frac{\partial D_j}{\partial P_j} = \eta_{j,P_j} * D_j = 0.024 * \frac{1}{4} = 0.006$. We assume $\alpha_L = 2\alpha_H$ (based roughly on estimates in Figure A10), which then implies $\alpha_L = 0.008$ and $\alpha_H = 0.004$. Our main estimates in Figure 3 use modest positive fixed costs of 1% of variable costs ($F = 0.01 * \overline{C} = \$4$), which equals half of administrative costs reported by insurers to the regulator. Appendix Figure A3 shows results for other estimates of F.

is realistic in health insurance. For instance in our data, the top half of consumers ranked by medical risk score have costs about 4x that of the bottom 50% risks (\$595 vs. \$150 per month). Thus, even realistic selection – if not offset by corrective policy – can severely constrain firm entry. We find similarly tight constraints in our more flexible structural model simulations in Section 6.

Prices, Welfare, and Un-Natural Monopoly If stronger adverse selection lowers prices conditional on entry (implication #1) but also reduces entry (implication #2), it suggests a more nuanced relationship between selection, prices, and consumer welfare than previously recognized. This is our third implication:

Implication #3: *With endogenous entry, there is a complex (non-monotonic) relationship between the strength of adverse selection, prices, and consumer welfare.*

To visualize this relationship, we continue with results from the calibrated Salop model in Figure 3. Panel B shows that as selection strengthens, equilibrium prices follow a non-monotonic "saw-tooth" pattern. Within each entry band (with a given N), prices decline with stronger selection (consistent with implication #1). As the graph shows, this occurs because stronger selection reduces MC and widens the selection wedge ($AC-MC$), while the Lerner markup ($= \frac{1/N}{\overline{\alpha}}$) remains constant conditional on N. But at each threshold where a firm exits, prices jump up as the Lerner markup increases (due to less competition). The result is a saw-tooth pattern with prices varying over a narrow range from \$408 to \$440 (a 2-10% margin) as long as $N \geq 2$. Panel C shows that consumer surplus follows a similar (but inverted) saw-tooth pattern — rising within each entry segment as prices fall, but jumping downward when each firm exits, reflecting higher markups and less variety.[12]

However, the logic changes when the market collapses to monopoly. Without multiple competitors, the price-disciplining effect of cross-firm adverse selection goes away, as indicated by the MC curve jumping up and the selection wedge disappearing. At this point, the monopolist (in theory) has unlimited pricing power, so we assume that the regulator caps prices — a common policy in other natural monopoly settings (e.g., utilities).[13] Beyond this point, stronger selection no longer constrains prices and profits; instead, the market becomes stuck in a high-price, low-entry equilibrium we call "un-natural" monopoly.

It is worth unpacking the source of this monopoly equilibrium — and how it differs from the standard concept of *natural* monopoly due to high fixed costs. With fixed costs, there is a *real cost* ($F > 0$) of having each additional firm operate and compete in the market. When F is large enough, monopoly is the socially efficient outcome — and with appropriate price regulation, it can be the best

[12] In the model, average consumer surplus is $E(CS) = V - P - \overline{t} \cdot \overline{d}(N)$, where $\overline{d}(N) = \frac{1}{4N}$ is average travel distance given N firms. We set $V = \$600$ arbitrarily to set the level of the plots.

[13] The price cap is somewhat arbitrary, but we choose a cap of \$475 (or a markup of \$75, or 19%) to be roughly consistent with the two-firm Lerner markup of \$83. While in principle we could model monopolist pricing based on competition with an outside option (e.g. uninsurance), in practice, this channel is often shut down in health insurance markets where maximal participation is a goal. For instance, both Massachusetts and the ACA use "price-linked" subsidies that adjust to make post-subsidy premiums affordable and therefore shut down competition with the outside option (Jaffe and Shepard, 2020).

Figure 3. Implications of Adverse Selection in Calibrated Salop Model

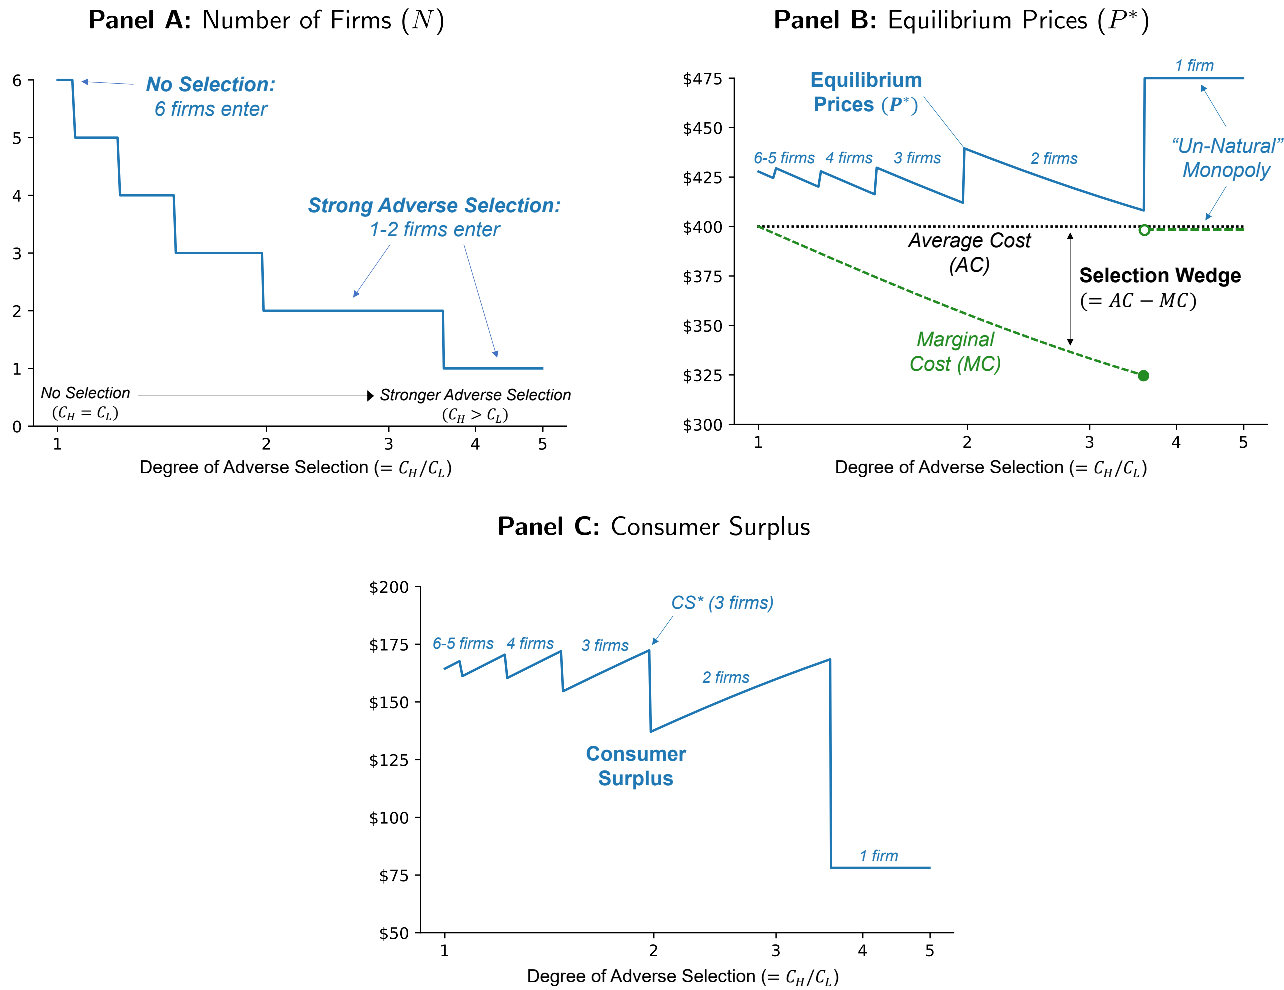

Note: The graphs show results from a calibration of our simple model, based on parameters from our empirical health insurance setting (see footnote #11 for details). In each panel, the x-axis is the degree of adverse selection, captured by the cost ratio (C_H/C_L) between high- and low-risk types, while holding fixed average costs at $\overline{C}$ = \$400 (i.e., a mean-preserving cost spread).

outcome for consumer welfare. By contrast, our "un-natural" monopoly due to adverse selection does not involve a real cost of more firms but a *coordination failure* in price competition. To see this, notice that if firms could coordinate on a single fixed price $\overline{P}$ (which can be well below the monopoly price), it would be technically feasible to support up to $\overline{N} = \frac{\overline{P} - \overline{C}}{F}$ firms in the market. For instance in our calibration (with $F = \$4$), a fixed price of $\overline{P} = \$420$ allows up to $\overline{N} = 5$ firms to enter. This would be a win-win for consumers, with more variety *and* lower prices. However, without pricing coordination and when adverse selection is strong, a price of \$420 with 5 competing firms is not a *strategic* equilibrium. Each firm has an incentive to undercut the others to cherry pick low-risk types, as shown in Figure 1.[14]

[14] We visualize this undercutting incentive in Appendix Figure A4, which shows the pricing best-response functions with two plans. Even though there are many prices at which both firms can earn profits (the green shaded area), each firm has an incentive to undercut at these prices. The only Nash equilibrium is at a price below what is needed to cover fixed costs. Indeed, if selection is strong enough, selection can lead to *non-existence* of pure-strategy Nash equilibrium,

This zero-sum "cream-skimming pricing" unravels a desirable pooling equilibrium, just as free entry of low-quality plans unravels a desirable pooling equilibrium in the classic Rothschild and Stiglitz (1976) model. Just as adverse selection can lead to a "race-to-the-bottom" in quality, in our model it can lead to a race-to-the-bottom in price competition — with the result being a market with few surviving competitors.

It is also worth discussing, briefly, whether the monopoly price can be sustained in the monopoly equilibrium. One might suggest that the threat of entry would cause the monopolist to keep its price lower than the monopoly price (the price ceiling). However, this threat is not credible here (at least in our model). The potential entrant knows that if they enter, they will lose money at the equilibrium prices that would occur, and thus the monopolist firm knows they will not enter. Their profitable entry would require some sort of coordination between the firms. Such coordination is not more feasible here than it would be in other settings. In fact, it may be more difficult here due to the much larger rewards from deviating from any agreed upon price (due to adverse selection). Further, even tacit collusion is likely to be quite difficult here due to the fluctuating nature of healthcare spending from one year to the next that makes it difficult to identify some kind of implicitly agreed-upon coordination point.

Corrective Policy Responses The problems with price competition under adverse selection suggest a role for corrective policies. The key way policy can help is by rectifying the coordination failure in which (inefficient) price undercutting to cream-skim unravels desirable pricing and entry outcomes. In other words, policy needs to limit or soften downward price competition. This is our fourth implication:

Implication #4: *Strong adverse selection motivates policies to soften or limit (downward) price competition, including price floors. If set carefully, a price floor can increase entry, raise consumer welfare, and even lower prices.*

Limiting price competition — and especially price floors — are not the typical way policymakers think about addressing adverse selection, so this bears discussion. First, we note that many standard selection-mitigating policies already involve softening price competition. For instance, *risk adjustment* transfers — by compensating plans more when they attract higher-risk consumers — reduce selection and thereby softens price competition (see implication #1). Similarly, *incremental* subsidies that differentially subsidize higher-price plans also soften competition by reducing the effective price-elasticity of demand.[15] While often seen as a downside or trade-off of these policies (Cutler and Reber, 1998; Mahoney and Weyl, 2017), our analysis suggests softer price competition may be desirable to support more entry.

Our analysis also suggests that *price floors* may be a useful policy for addressing the coordination failure involved with price competition. While price floors may seem strange to economists, they are in fact common in subsidized insurance programs like Medicare Advantage and the ACA that limit

a familiar result from Rothschild and Stiglitz (1976). When this is true, the only equilibria are in mixed strategies.

[15]Incremental subsidies are used in our setting (see Section 3) and also commonly in employer-sponsored insurance. This competition-softening feature has been frequently critiqued by Alain Enthoven (see e.g., Enthoven (1993)).

post-subsidy prices to be non-negative – a limit that is often binding. More indirectly, insurance regulators (including in our Massachusetts exchange setting) can and do impose rules about premium ranges that are "actuarially sound," effectively disallowing price cuts below a floor.

We provide a full analysis of various corrective policies via counterfactual simulations in Section 6, but here we provide a brief discussion of how policy responses fit in our basic model. To do so, we show how price floors affect our calibrated model in Appendix Figure A2. We compare market outcomes without floors (dashed blue lines, which replicates results from Figure 3) versus a medium floor of $\overline{P} = \$420$ (solid orange) and a high floor of $\overline{P} = \$450$ (solid green). Panel A shows that, when binding, price floors promote more entry. Additionally, they *stabilize* the number of entrants even as adverse selection becomes stronger. As noted in our discussion above, the price floor effectively allows firms to coordinate on a fixed price and avoid the (sometimes inefficient) price undercutting spiral that would otherwise unravel entry.

Price floors can have an important trade-off, however: Both the medium and high floors stabilize prices as selection strengthens. The stable price level (at the floor) may be higher or lower than the free pricing equilibrium, but even the high floor is well below the monopoly price.[16] With adverse selection, therefore, *price floors can lower prices,* because they allow for coordination and support more entry. Because this entry is itself valuable to consumers (via variety), price floors can also improve consumer welfare, even without lowering prices. Indeed, the medium price floor at \$420 results in CS near the upper-envelope of what is feasible with free price competition. The high price floor goes too far in promoting entry (at high prices) when selection is modest, but even the high floor benefits consumers relative to the monopoly outcome when selection is strong.

This discussion suggests that price floors — if used carefully — may be a useful (and novel) addition to the policy toolkit for addressing market failures due to adverse selection. However, one limitation of our simple model is that it assumes *symmetric* firms with equal cost structures (i.e., purely horizontal differences). If some firms are also vertically differentiated with lower cost and quality (e.g., via lower actuarial value), price floors may limit their ability to efficiently compete with lower prices. If cost/quality differences are observable, the regulator can in principle condition the floor on quality or generosity – e.g., in the ACA markets, set different floors for gold vs. silver vs. bronze plans — but this may be infeasible if quality differences are non-contractible. In our structural model, we allow for (and find) this type of mixed horizontal and vertical differentiation and evaluate whether price floors still improve consumer welfare.

2.3 General Theory

We now show how our basic ideas developed in the simple model extend to a much more general setup. Because the logic follows closely to the Salop model analysis, we present results quite briefly. The

[16]Indeed, there is an interesting paradox for the monopoly case. We assumed that the monopolist limited its prices because of a price *ceiling* (at \$475 in the calibration). If the regulator also imposed a price *floor* at this same value, many more firms could enter (up to 18 firms at $\overline{P} = \$475$) and earn positive net profits, while consumers would benefit from more variety at the same price. This would be a win-win for both consumers and 17 of the 18 firms (all except the monopolist).

general theory is useful for showing the key (measurable) statistics on which the selection-pricing-entry relationship depends. This creates a bridge to the empirical work in the remainder of the paper.

Model Setup We consider a general (insurance) market where a set of potential firms $j \in \{1, ..., J\}$ decide whether to enter by offering a fixed product (or "plan") with differentiated attributes X_j. As in the simple model, firms compete in a two-stage entry game, following a workhorse approach in the IO literature (Berry and Reiss, 2007).[17] In stage 1, firms decide whether to enter, which involves costs $F_j \geq 0$. In stage 2, entrants $j \in E$ compete on prices, P_j, to maximize profits, $\pi_j(P)$, in Bertrand-Nash equilibrium. An equilibrium involves a set of entrants E^* and Bertrand-Nash prices P^* where: (1) all entrants are profitable net of fixed costs, and (2) no non-entrant $j' \notin E^*$ can unilaterally enter and earn net profits at the prices that result among firms $E^* \cup j'$.

Consumers (i) are characterized by their WTP for each firm's plan (V_{ij}) – and therefore demand functions, $D_{ij}(P)$, at prices P — and their risk type, r_i, which in turn affects a firm's expected costs, $C_{ij} = C(X_j; r_i)$. Firms cannot price discriminate on risk, something often disallowed in practice, so the market features cost-relevant asymmetric information.[18]

Firms compete on prices to maximize (variable) profits, $\pi_j(P) = [P_j - AC_j(P)] \cdot D_j(P)$, where $D_j(P)$ is firm j's total market share (demand) and $AC_j(P) = \frac{1}{D_j(P)} \sum_i [C_{ij} \cdot D_{ij}(P)]$ is its average costs at these prices.[19] As in the simple model, average costs vary with (all firms') prices because prices affect which types of consumers select into each plan. At a given price vector P, we say that a firm faces *adverse selection undercutting incentives* if $\frac{\partial AC_j(P)}{\partial P_j} > 0$. Equivalently, a firm's marginal costs when it adjusts price, $MC_j(P) \equiv \frac{1}{\partial D_j / \partial P_j} \sum_i \left[C_{ij} \cdot \frac{\partial D_{ij}}{\partial P_{ij}}\right]$, are less than its average costs. We again call $AC_j(P) - MC_j(P)$ the "selection wedge." The following equation (replicating (5)) connects the two definitions:

$$\textbf{Adverse Selection:} \quad \frac{\partial AC_j(P)}{\partial P_j} \;=\; \eta_{j,P_j}(P) \;\times\; [AC_j(P) - MC_j(P)] \;>\; 0 \tag{11}$$

where $\eta_{j,P_j}(P) \equiv -\frac{\partial D_j / \partial P_j}{D_j} > 0$ is the firm's (own price) semi-elasticity of demand, which is positive by the law of demand. Therefore, $\frac{\partial AC_j(P)}{\partial P_j} > 0$ if and only if the selection wedge is positive, and the two are proportional to each other.

Implications for Market Outcomes We now show how the implications of selection for market outcomes generalize. Starting with pricing (conditional on entry), firms optimally price at marginal

[17] We believe, however, that our model's key ideas extend much more generally to other entry games, including games with dynamics and with endogenous product design (see e.g., Wollmann 2018). This is because selection influences the *FOC for profit-maximizing pricing*, which is a common element of nearly all models with flexible price competition. However, extension of these ideas to more complicated games is beyond the scope of this paper, and we leave it for future research.

[18] We abstract from policies like partial price discrimination (e.g., by age), risk adjustment transfers, and reinsurance that can mitigate – but rarely eliminate – risk selection. We return to these in the structural model. All of our math below carries through if (unadjusted) average costs are replaced with "risk-adjusted" average costs and demand with risk-scaled demand, as shown by Curto et al. (2021).

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

costs plus the Lerner markup, $P_j^* = MC_j(P^*) + \frac{1}{\eta_{j,P_j}}$. This in turn implies profit margins of:

$$\begin{aligned} \underbrace{P_j^* - AC_j(P^*)}_{\text{Profit Margin}} &= \frac{1}{\eta_{j,P_j}} \times \left(1 - \frac{\partial AC_j}{\partial P_j}\right) \\ &= \underbrace{\frac{1}{\eta_{j,P_j}}}_{\text{Lerner Markup}} - \underbrace{\left[AC_j(P^*) - MC_j(P^*)\right]}_{\text{Selection wedge}} \end{aligned} \tag{12}$$

where the second equality follows from plugging in $\partial AC_j/\partial P_j$ from equation (11). Stronger adverse selection — a larger average cost curve slope and selection wedge — disciplines market power by reducing equilibrium prices and profit margins. This corresponds to our implication #1 in the simple model.

With endogenous entry, however, lower profit margins mean that there are fewer firms willing to compete (implication #2). The "profitable participation" condition, which we stated in general terms in (8) above, continues to apply and to imply a limit on firm entry, as discussed in our analysis of Figure 2. We can summarize this limit on entry in the following general proposition:

Proposition 1. *Suppose that there is adverse selection in pricing, and prices are set in Bertrand-Nash equilibrium. Then the number of competing firms (N) in a market is limited by:*

$$N \leq \frac{1}{\mathbb{E}_j\left[\left(-\frac{\partial D_j}{\partial P_j}\right) \times \left((AC_j - MC_j) + \frac{F}{D_j}\right)\right]} \tag{13}$$

where $\mathbb{E}_j[.]$ is an (unweighted) average across competing firms $j \in E$.

The proof is in Appendix A.2. Its logic follows from rearranging the profitable participation condition in (8), which must be true for each firm so must also be true on average. Basically, there needs to be sufficient revenues in the market for all costs in the market to be covered. N in equation (13) is the maximum number of firms where total market revenues are greater than or equal to total market costs. Interestingly, this condition is equivalent to $\frac{\partial ATC_j}{\partial P_j} \leq 1$, as can be seen by plugging in equation (6). This suggests an empirical test for when adverse selection has become "too strong" to sustain a given market structure: the average price-slope of average costs cannot exceed one.

Proposition 1 brings together in one expression the three determinants of market structure we have highlighted. Entry is limited by the price-sensitivity of demand ($-\frac{\partial D_j}{\partial P_j}$, capturing differentiation) times the sum of the adverse selection wedge and fixed costs per consumer (which again enter together). Stronger adverse selection and larger fixed costs limit how many firms can profitably participate.

The remaining two implications of adverse selection are harder to show mathematically, but they follow naturally from the analysis so far. If adverse selection (1) lowers prices conditional on entry but (2) also reduces entry, which generally implies larger markups, then it is natural that the pricing-selection relationship will be complex and non-monotonic. At the simplest level, if selection results in the market collapsing to monopoly, prices are likely to be quite high, and policies that stimulate entry by limiting or softening price competition are likely to improve consumer welfare. We evaluate this further using our structural model simulations in Section 6.

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, in the next two sections, by providing descriptive evidence of strong adverse selection in pricing, using quasi-experimental variation in prices. In Sections 5-6, we leverage this variation to estimate a full structural model of consumer demand and plan costs and use that model to run counterfactual simulations illustrating the effects of policies like risk adjustment and price floors.

3.1 Setting: Subsidized Massachusetts Exchange (CommCare)

We study the pre-ACA subsidized Massachusetts health insurance exchange, a program called Commonwealth Care (or "CommCare"). From 2007-2014, CommCare provided subsidized insurance for state residents with incomes below 300% of FPL without access to Medicaid, Medicare, or job-based coverage. Because of its ACA-like structure, rich policy variation, and detailed administrative data, the Massachusetts exchange has been a fruitful setting for research on health insurance markets.[20]

The market featured four to five competing health insurers, with each insurer offering a single highly regulated plan that followed standardized cost sharing rules.[21] Plans were differentiated primarily in their networks of covered hospitals and doctors. Insurers were primarily Medicaid-based insurers offering limited networks similar to those of their Medicaid managed care plans. Three participating insurers — Boston Medical Center plan (BMC), Neighborhood Health Plan (NHP), and Network Health — had comparably broad but differentiated networks, covering 75-85% of hospitals.[22] 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 narrower network.

Insurers in CommCare chose whether to participate in the market annually, as part of a Request for Proposals (RFP) process run by the market regulator. Conditional on participating (and on their networks), insurers competed on prices (premiums) in a setup similar to our model. Insurers reset their premiums annually at the start of the year, which were then locked in until the end of the year.

[20] Prior work on CommCare includes Chandra et al. (2010, 2014); Finkelstein et al. (2019); Jaffe and Shepard (2020); Shepard (2022); McIntyre et al. (2021); Shepard and Forsgren (2022); Geruso et al. (2023*b*); Shepard and Wagner (2022). Other work has studied the pre-ACA *un*subsidized Massachusetts health insurance exchange, a program known as "CommChoice" (Ericson and Starc, 2015*a*,*b*, 2016).

[21] Insurers were required to offer the same plan (with identical features) to all consumers in the state. They could, however, choose whether or not to participate in each of 38 "service areas," and there is significant variation in firm participation by area. Because this entry decision occurs at a lower geographic level than pricing, we have not explored it in this paper. For more on this type of "partial rating area" offering, see Ko and Fang (2023); Geddes (2024).

[22] See Appendix Figure A21 for a graph of network size over time.

Premiums could vary only on specific factors like income group and region, not age or health status.[23] Consumers—who enrolled in the market throughout the year as needs for insurance arose (due to job loss, changes in Medicaid eligibility, etc.)—chose among competing plans at two times: (1) when they newly enrolled, and (2) at the start of each year ("open enrollment"), when premiums reset and they had an opportunity to switch plans.

CommCare's regulator oversaw the market using a strong form of the "managed competition" model envisioned by Enthoven (1993) — indeed, much more so than in ACA markets today. This strong regulation may help explain the market's ability to sustain a robust and stable set of competitors, despite the forces we highlight in our model. At a high level, we use CommCare (and its rich premium variation and detailed data) as a "laboratory" to infer key demand and cost elasticities relevant for the theory. This lets us assess the strength of these forces and model their counterfactual relevance in a setting (like the ACA) that does much less to soften and regulate price competition. In Section 8 we briefly study the ACA market (where only aggregate data is available and premium variation is less clean) directly.

While CommCare had many regulatory features, we focus on those relevant for adverse selection and price competition. Specifically, we summarize the following three main policies:

1) Risk Adjustment Like the ACA, CommCare enforced *risk adjustment* transfers between insurers based on their enrollees' health risk, as measured by age and diagnoses. For an insurer that set a base premium of P_j and attracted enrollees with average risk scores $\bar{\varphi}_j$, the insurer received post-transfer revenues of $\bar{\varphi}_j P_j$. Risk adjustment can mitigate adverse selection by making post-transfer profits less dependent an insurer's enrollee risk mix than its pre-transfer profits. By doing so, risk adjustment "flattens" the average cost curve in price (Mahoney and Weyl, 2017) or decreases the selection wedge from Section 2. However, risk adjustment is usually imperfect (Brown et al., 2014), and there is evidence of its imperfection in CommCare specifically (Shepard, 2022; Geruso et al., 2023*b*).

2) Price Regulation Second, the exchange directly regulated prices using price ceilings and floors. Price ceilings were intended to limit price growth and were gradually tightened to become binding for about half of plan prices during 2011-13. Price floors were imposed under rules requiring that premiums be "actuarially sound," or no lower than a minimum level deemed adequate by an independent actuary. These floors were often binding from 2010 forward, especially for BMC, CeltiCare, and Network Health. As a result it was common that 2+ plans tied for the lowest premium in a region. Although not explicitly intended to ensure participation, these floors may have helped halt a race-to-the-bottom in prices. While price regulation is not explicitly used in ACA markets today, there is a $0 floor on *post-subsidy* premiums that has become increasingly binding over time.

3) Incremental Premium Subsidies Third, premiums were *subsidized* to ensure affordability and encourage consumer take-up. CommCare's subsidies were more complex than subsidies in the ACA

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

Marketplaces, but we explain the details because they matter for the economics of the market and provide our key source of identifying variation. Unlike the ACA exchanges, CommCare's subsidies were not a flat amount across plans but instead followed an income-varying formula that affected both premium levels and *differences* across plans. Specifically:

- "Below-poverty" enrollees (0-100% of FPL) were *fully subsidized*; all available plans were $0.
- "Above-poverty" enrollees (100-300% of FPL) were *partly subsidized.* The cheapest plan cost an income-varying "affordable amount," which varied from $0 for 100-150% FPL to $116 per month at 250-300% FPL. Higher-price plans cost more, following a progressive formula that smoothed price differences more for those with lower incomes.

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

This subsidy structure, while complex, has two important implications for our analysis. First, the income-based subsidies create useful identifying variation in the premiums different consumers pay *for the same plan choice set.* In particular, below-poverty enrollees, who are insulated from prices, serve as a sort of "control group" for estimating the impact of premium changes on demand. By comparing demand responses to plan price changes for above- and below-poverty groups, we can infer price elasticities separately from any unobserved changes in plan desirability. We use this identification strategy in both our descriptive analysis and our structural demand model.

Second, this subsidy structure softened insurer price competition, a key force in our model. Because *all* plans were fully subsidized for below-poverty enrollees, who comprised about half of the market, these individuals were completely inelastic to firm prices. This substantially lowers firms' effective price elasticity of demand, thus reducing the undercutting incentives highlighted in our model.

Figure 4 plots average net-of-subsidy premiums for each plan as paid by above-poverty enrollees (see Appendix Figure A19 for the underlying pre-subsidy premiums). The black line at $0 indicates that all plans were free for below-poverty enrollees, even as premiums varied across plans for above-poverty enrollees. There is substantial variation across plans and over time, including in the identity of the cheapest plan, which we make use of in our empirical analysis.

3.2 Data: Administrative Enrollment and Insurance Claims

We acquired enrollment records and full medical and prescription drug claims data for the universe of CommCare enrollees (Massachusetts Health Connector, 2014). The enrollment records provide demographic and geographic information for each enrollee, with monthly enrollment information so we can observe when the individual first enrolled, which plan she enrolled in, and if she ever switched plans or left the market and returned later. We also observe each enrollee's income and geographic market, letting us identify net-of-subsidy prices of each plan in the enrollee's choice set.

Figure 4. Variation in Enrollee Premiums in CommCare ($ per month)

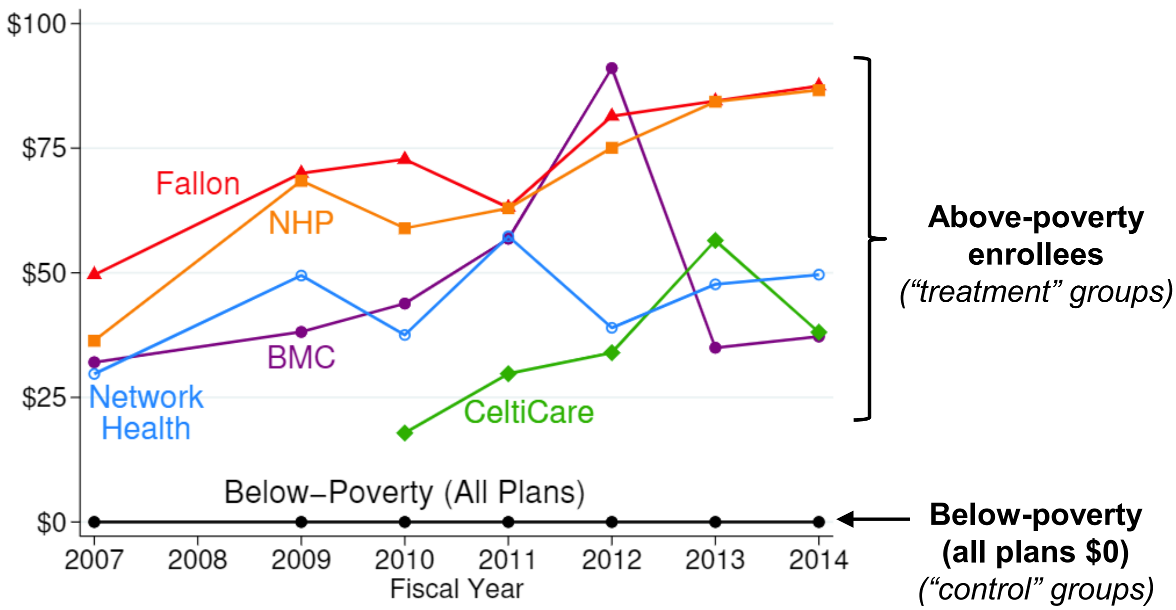

Note: The graph shows post-subsidy enrollee premiums for each insurer's plan in the CommCare market, by fiscal year and income group. Enrollment-weighted average premiums for above-poverty enrollees (100-300% of poverty) are shown in different colors by plan and labeled. As shown, these vary substantially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

We also have full claims data for all CommCare enrollees and plans. This lets us construct measures of healthcare utilization and spending for each person, including total insurer claims costs. These data also let us construct diagnosis-based risk scores to analyze the effect of risk adjustment. We use two risk scores in our analysis. First, we use CommCare's *actual* risk scores (as used in its risk adjustment), which is a retrospective model based on an enrollee's prior-year diagnoses (and for new enrollees, based on age-sex factors). Second, we construct *concurrent* risk scores using the HHS-HCC model used by the ACA for its risk adjustment (Kautter et al., 2014). The HCC model is more predictive — e.g., its R^2 for costs in our data is 26.8%, compared to 5.9% for CommCare's actual risk scores. In some analyses, we use the gap between actual and HCC risk scores as a measure of "residual" risk missed by CommCare's risk adjustment.

We use the raw administrative data to construct two main datasets for further analysis, a reduced form dataset and a structural model datasets. The datasets are largely similar in terms of restrictions applied. The primary difference is the format of the datasets, which matches the methods used. See Appendix B for details.

4 Reduced Form Evidence

We now show descriptive evidence of the type of adverse selection undercutting incentives implied by our model in Section 2. Specifically, our model suggests that adverse selection is most likely to affect insurer entry if: (1) consumers are highly price sensitive (high price-elasticity of demand) and (2) firms' average costs are strongly increasing in their price (large slope $\frac{\partial AC_j}{\partial P_j}$), with a large "selection wedge" between average and marginal costs. To estimate these statistics in CommCare, we use the premium variation shown in Figure 4 to implement a difference-in-differences analysis that identifies the average effect of plan premium changes on a plan's market share and average costs. We use the same variation in our structural estimation in Section 5, so our results here both validate and provide intuition for what identifies the structural model.

Our difference-in-differences (DD) design leverages two key sources of variation in premiums: (1) changes in plan (pre-subsidy) prices over time and (2) differences in how those pre-subsidy prices translate into net-of-subsidy prices across income groups. Specifically, changes in plan (pre-subsidy) prices over time affect net-of-subsidy premiums faced by *above-poverty* enrollees, but not *below-poverty* enrollees who have access to the same plans but are fully subsidized (all plans are \$0). While prices shift at the start of each year, enrollment occurs monthly throughout the year, suggesting an event study design analyzing how monthly enrollment patterns change at the start of each new year for consumers facing premium changes (above-poverty "treament" group) versus those not facing premiums (below-poverty "control" group). The inclusion of a control group lets us difference out any shifts due to time-varying factors other than premium changes (e.g., changes in plan networks or other unobserved plan differences). We note, however, that plan designs and enrollee composition were fairly consistent over time, resulting in essentially flat trends for control group outcomes.

We combine many price changes across plans, regions, and years in a "stacked" DD setup. Define an "experiment" e as a plan (j) x region (r) x year-pair (yr, $yr+1$) cell. Within each experiment $e \equiv \{j, r, yr\}$, premiums vary across the five income groups g in CommCare: (0) $g = 0$ for below-poverty (control) enrollees (for whom $P_{e,g,t} = 0$ in all cases), and $g \in \{1, 2, 3, 4\}$ for above-poverty (treatment) enrollees with incomes (1) 100-150%, (2) 150-200%, (3) 200-250%, and (4) 250-300% of FPL. We restrict each experiment to the 12 months before and after each price change.

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}.[24] We stack all experiments and estimate the following event-study regression:

$$Y_{e,g,t} \;=\; \alpha_{e,t} \;+\; \gamma_{e,g} \;+\; \sum_{k=-T\setminus\{-2\}}^{T} \delta_t \times Treat_g \times 1_{t=k} \;+\; \varepsilon_{e,g,t} \qquad (14)$$

where $Treat_g$ is a dummy for above-poverty treatment groups, $g \in \{1, 2, 3, 4\}$. Each outcome $Y_{e,g,t}$ (market shares or average costs) is measured at the experiment x income group x month level. The experiment-by-event time fixed effects (α_{et}) ensure that we only use *within-experiment* price variation (i.e., across income groups for a given plan-region-month) for identification. The experiment-by-income group fixed effects ($\gamma_{e,g}$) net out any fixed differences in costs or preferences across income groups within a plan — e.g., if higher-income groups have lower costs in general. The coefficients

[24]Price changes are defined as *relative* price changes, or the change in the *difference* between the plan's net-of-subsidy price and the net-of-subsidy price of the lowest price plan in the market.

of interest are the δ_t's, which capture how treatment vs. control group outcomes change, relative to the gap two months prior to the price change.[25] We estimate this regression for two key outcomes: log market shares for plan j ($\log(Share_{j,g,t})$, which captures price semi-elasticites of demand), and average costs of a plan's enrollees ($AC_{j,g,t}$, which captures average cost curve slopes, $\frac{\partial AC_j}{\partial P_j}$).[26]

Figure 5 presents the event study plots. For both price increases and decreases, the average enrollee premium change is about \$18-21 per month at time $t = 0$, which is about 5% of average costs (\$380 per month) and 40% of average post-subsidy premiums for above-poverty enrollees (\$49). Panel (a) shows results for log market share. After essentially zero pre-period trends, market shares jump sharply up (for price cuts) and down (for price increases) at $t = 0$ when premiums change. The symmetric pattern adds credence to our strategy, since any spurious trends are unlikely to generate similar results in opposite directions.[27] The implied price-sensitivity of demand is large, with the typical \$20 premium change shifting market share by about 20%.

Panel (b) shows the changes in the average cost of plan enrollees that correspond to the market share shifts in Panel (a). While estimates are noisier and show some pre-trends for premium decreases, the plot suggests that when prices increase, average costs rise, and when prices decrease, average costs decline. These results are consistent with adverse selection in pricing, or $\frac{\partial AC_j}{\partial P_j} > 0$ in our model. The magnitudes are also large: An average \$20 monthly net-of-subsidy premium increase results in an increase in the plan's average cost of around \$20, suggesting an average cost curve slope close to 1.0. As discussed in the model, this raises concerns about the market's ability to support many competing plans in the absence of policies used to combat selection, such as price floors and risk adjustment.

We next leverage all experiments in a single "pooled" regression to maximize power and report headline DD estimates. Specifically, we stack all price change experiments, with treatment status multiplied by -1 for price decreases.[28] We then run summary DD regressions of the form:

$$Y_{e,g,t} \;=\; \alpha_{e,t} \;+\; \gamma_{e,g} \;+\; \delta \times Treat_g \times Post_t \;+\; \varepsilon_{e,g,t} \tag{15}$$

where $Post_t$ is a dummy for periods after the price change ($t > 0$) and all other variables and samples are analogous to those in the event studies above. The main coefficient of interest is δ, which captures the differential change in outcomes for above-poverty treatment groups (who experience a premium change) vs. below-poverty control groups (who do not).

These pooled event studies appear as panels (c) and (d) of Figure 5. These results mimic the results in panels (a) and (b) but are less noisy and send an even clearer message: consumers are highly price

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

[26] Costs for each enrollee-month are defined as the average monthly cost of the enrollee over the following 12 months, or until the enrollee leaves the dataset because they are no longer in the market. To reduce noise, we winsorize 12-month average costs at the 99.9th percentile and add back a constant to all observations to preserve the original sample mean for costs. Results without winsorization have similar point estimates but larger standard errors.

[27] The effect on shares grows somewhat during the post period. This reflects the combination of choices for a stock of incumbent enrollees, who can switch plans in month $t = 1$ (but tend to be inertial) and new enrollees joining the market each month (who are highly price-sensitive). We split out results new enrollees separately in analyses below.

[28] In other words, $Treat_g = 1\{g \in \{1, 2, 3, 4\}\}$ for premium increases and $Treat_g = -1\{g \in \{1, 2, 3, 4\}\}$ for decreases.

Figure 5. Price Change Event Study Estimates

(a) Log Market Shares

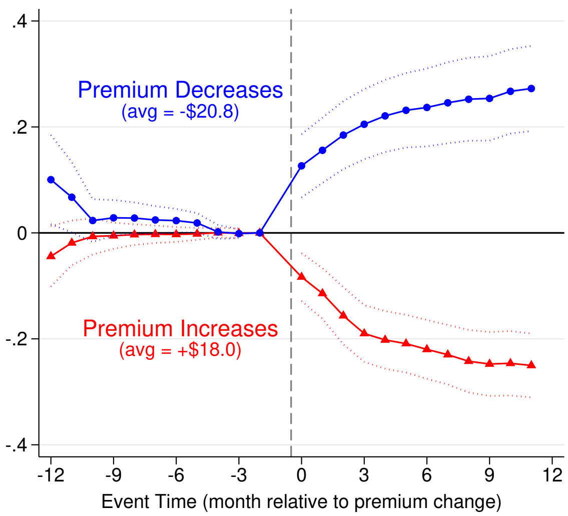

(b) Average Cost ($/month)

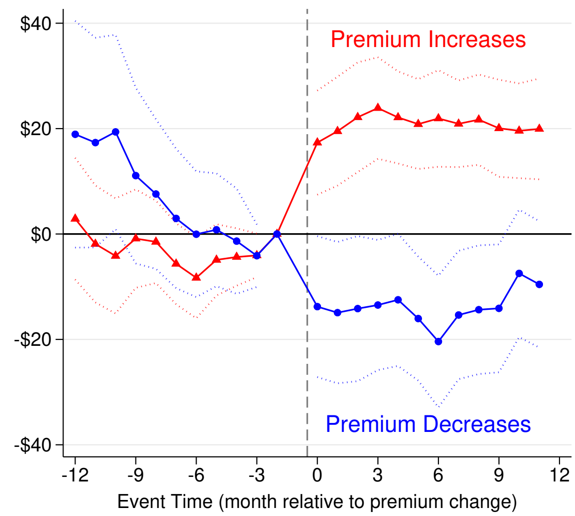

(c) Log Market Shares (pooled increases + decreases)

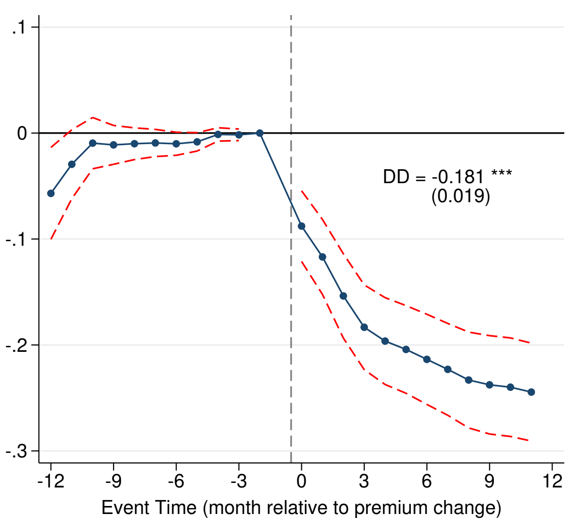

(d) Average Cost (pooled increases + decreases)

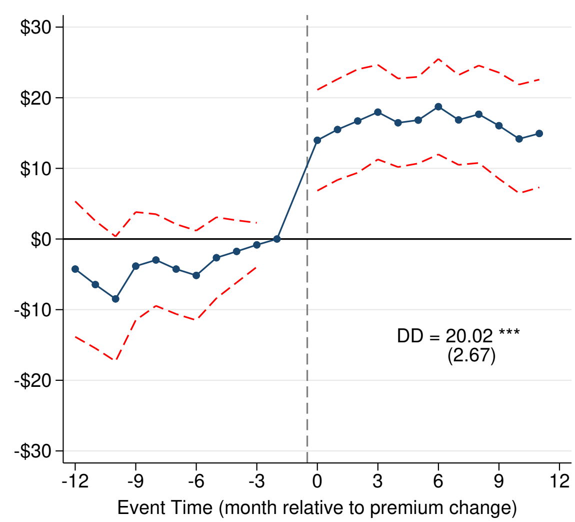

Note: Figure shows event study estimates (regression (14)) of the impact of plan premium changes on demand (log market share) and risk selection (average costs of plan enrollees). Panels (a)-(b) show estimates separately for premium increases vs. decreases, while panels (c)-(d) show our main estimates pooling both types of changes (with decreases multiplied by -1) and reports our main difference-in-difference estimates. Average costs are defined as average costs per month (averaged over the subsequent year) of the set of enrollees who joined a plan in a given month. Regressions are weighted by enrollment at the year-region-income group level, and 95% confidence intervals are shown using standard errors clustered at the plan x year-region-income group level.

Table 1. Difference-in-Differences Results

	Baseline	**By Enrollee Type**		**By Enrollee Risk**		
	All Enrollees	New Enrollees	Current Enrollees	Low Risk *(0-25%)*	Mid Risk *(25-75%)*	High Risk *(75-100%)*
	(1)	(2)	(3)	(4)	(5)	(6)
Panel (a): Regression results						
Premium	17.87***	17.90***	17.08***	17.25***	18.09***	17.79***
	(1.45)	(1.56)	(1.43)	(1.47)	(1.44)	(1.51)
Log Market Share	-0.181***	-0.429***	-0.080***	-0.257***	-0.178***	-0.130***
	(0.019)	(0.040)	(0.015)	(0.026)	(0.018)	(0.016)
Average Cost	20.02***	37.15***	16.04***	---	---	---
	(2.67)	(5.08)	(3.44)			
Panel (b): Theory-Relevant Statistics						
Demand Semi-Elasticity	-0.0101	-0.0240	-0.0047	-0.0149	-0.0098	-0.0073
Slope of Avg Costs (=dAC/dP)	1.12	2.08	0.94	---	---	---
Adverse Selection Wedge	**$110.9**	**$86.5**	**$201.5**	---	---	---
[% of Avg Cost]	**[30%]**	**[21%]**	**[56%]**			
Num. Observations	5,888	4,922	5,750	5,359	5,819	5,612
Average Cost ($/month)	$374	$411	$360	$134	$239	$865

*Standard errors reported in parentheses. *** p < 0.01, ** p < 0.05, and * p < 0.10.*

Note: Table shows estimates of the pooled DD specification (15) for premium change experiments, as described in the text. Panel (a) presents estimates, with each row corresponding to a separate regression with a different outcome variable (enrollee premiums, log market shares, and average costs). Columns correspond to different enrollee samples, including new vs. current enrollees (cols. 2-3) and enrollee risk groups (cols. 4-6), defined by indicated quartiles of the enrollee's HCC risk score. Panel (b) shows key empirical statistics relevant to our model, including the semi-elasticity of demand ($= \partial \log D_j / \partial P_j$), the slope of average costs ($= \partial AC_j / \partial P_j$), and the adverse selection wedge ($= AC_j - MC_j$), all of which are computed from estimates in panel (a). All regressions are weighted by the number of enrollees in the relevant cell from which averages are calculated, and standard errors are clustered at the plan-region-income group level. The number of observations differs across columns because of small $\{e, g, t\}$ cells dropped because of zero observations.

sensitive and selection is strong, with a $20 increase in premiums leading to around a $20 increase in average cost.

Table 1 reports results from the pooled DD regression specification (15), with columns showing results for different enrollee subsamples. Column (1) shows estimates for our main sample of all enrollees (new and incumbent), corresponding to Figure 5. These estimates again suggest high levels of price sensitivity and strong adverse selection. The typical enrollee premium increase (of $17.87 per month) results in a plan's market share falling by 18.1 log points (about 20%) and $20.02 higher average costs among all enrollees. Columns (2)-(3) shows that price-sensitivity and average cost (selection) responses are even larger for new enrollees (who must make active plan choices) and somewhat smaller for current enrollees (who are often inertial), implying that the relative sizes of these groups matters greatly for overall levels of price-sensitivity and selection. Columns (4)-(6) show the underlying source of the selection responses: low-risk (healthy) consumers are much more price-sensitive than high-risk (sick) types. The lowest-risk quartile of enrollees has an implied semi-elasticity of demand (-0.0149) about twice as large as that of the highest-risk quartile (-0.0073). This provides evidence that the key adverse selection condition highlighted in our theory (see condition (3)) holds here.

Panel (b) reports the key theory-relevant statistics implied by these estimates, including the price semi-elasticity of demand ($\eta_{j,P_j} = -\frac{\partial \log D_j}{\partial P_j}$), average cost curve slope ($\frac{\partial AC_j}{\partial P_j}$), and selection wedge ($AC_j(P) - MC_j(P)$), all of which play a prominent role in our theory in Section 2.[29] Notably, the average cost curve slope of 1.12 exceeds the key threshold of 1.0, and the selection wedge is \$110.9 per month, or about 30% of average medical spending. This suggests that without corrective policies, adverse selection exerts a strong entry-deterring influence equal to fixed costs of 30% of variable costs.

The CommCare exchange, like many other programs, used risk adjustment to try to mitigate adverse selection. However, there appears to be meaningful adverse selection in pricing even after risk adjustment. To show this, we replicate Table 1 using *risk-adjusted* costs (equal to an observation's average cost divided by its average risk score) in Appendix Table A1. Because of data limitations, we cannot use CommCare's actual risk scores (which we observe only starting in 2011, the last year of the DD analysis). We instead use HHS-HCC risk adjustment scores, the method used by the ACA (which we replicate in our data). Because the HCC method is much more predictive (e.g., its R^2 for costs in our data is 27%, versus 6% for CommCare's scores), this is likely a lower bound on the amount of residual selection in CommCare.

Appendix Table A1 shows that even after risk adjustment, price increases result in statistically higher average costs, with a coefficient of \$6.19 (s.e. = 2.28). The risk-adjusted AC curve slope is 0.35 for all enrollees (and 0.76 for new enrollees), and the selection wedge is \$45.1 per month, or 12% of average costs (and \$36.2, or 9% of AC, for new enrollees). These values are about one-third as large as without risk adjustment, suggesting that it "flattens" the AC curve by about two-thirds. However, even after risk adjustment, there is still meaningful residual selection on pricing.

Ultimately, these reduced form estimates provide strong evidence of high price sensitivity and strong adverse selection among consumers in the CommCare market. Without the robust policies CommCare used to mitigate selection — like risk adjustment, incremental subsidies, and price regulation (see Section 3) — the demand and cost elasticities suggest the market would have difficulty sustaining robust entry even among 4-5 competing insurers. To better understand the role of these corrective policies for sustaining entry, we next estimate a full structural model of the market, which we use to simulate outcomes with various counterfactual policies.

5 Structural Model and Estimation

In this section, we describe and estimate our structural model of insurance plan choice (demand) and enrollee-level insurer costs. In Section 6, we combine these estimates with a model of equilibrium entry and pricing to study the implications of adverse selection (and corrective policies) for insurer participation, prices, and consumer welfare. Our demand and cost models follow closely the approaches

[29] These statistics are simple functions of the DD estimates in panel (a) of the table. The demand semi-elasticity is the ratio of the log market share and premium coefficients, and the average cost slope is the ratio of the average cost and premium estimates. The selection wedge is the ratio of these two, $AC_j - MC_j = \frac{\partial AC_j}{\partial P_j} / \eta_{j,P_j}$ (see equation (4)).

of Shepard (2022) and Jaffe and Shepard (2020), who also study the CommCare market. We therefore discuss the model somewhat briefly and refer readers to Appendix E for further details.

5.1 Insurance Demand Model

We model enrollees' plan choices as a function of premiums and prior plan choices, all interacted with enrollee characteristics. To construct our demand estimation sample, we restrict the data to "choice instances," defined as one of two times when consumers can choose/switch plans: (1) when an enrollee newly enrolls in the market, or re-enrolls after a break in coverage, and (2) when continuing enrollees have the option to switch plans during the annual open enrollment period.[30] A single enrollee may have multiple choice instances; we index unique enrollee-choice instance pairs by (i,t).

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

$$u_{ijt} = \underbrace{-\alpha(Z_{it}) \cdot P^{cons}_{ijt}}_{\text{Subsidized Premium}} + \underbrace{f(X_{jt}, Z_{it}; \beta)}_{\text{Plan Attributes}} + \underbrace{\xi_j(W_{it})}_{\text{Unobs. Differences (FE)}} + \varepsilon_{ijt}, \qquad j = 1, .., J \qquad (16)$$

where P^{cons}_{ijt} is plan j's post-subsidy premium for consumer i in year t (based on their income group and region). Following Shepard (2022), price sensitivity $\alpha(Z_{it}) = Z_{it}\alpha$ is allowed to vary by income bins, medical diagnoses (dummies for chronic disease and cancer), demographics (5-year age-sex bins), medical (HCC) risk scores, and immigrant status. Relative to Shepard (2022), we add one key covariate to Z_{it} on which price sensitivity can vary: an estimate of enrollee's unobserved health risk, which we separately identify from plan effects on costs using the methods described in the plan cost model section below (see Appendix E.2 for details). We bin these residuals into deciles and include them in Z_{it}. We find that this additional covariate helps match the empirical patterns of adverse selection in response to price changes.

The function $f(X_{jt}, Z_{it}; \beta)$ includes interactions of observed plan and consumer factors that affect demand. These include terms capturing the desirability of a plan's hospital network for a given enrollee, derived from a hospital demand system and from existing relationships with physicians and hospitals. We also include dummy variables for the immediate prior plan choice of continuing enrollees (capturing switching costs and other drivers of state dependence) as well as the interactions of these variables with income, age-sex bins, and risk scores. Finally, we capture unobserved plan demand differences with $\xi_j(W_{it})$, which are plan-region-year, plan-region-income, plan-risk score, plan-age-sex, and plan-immigrant-status fixed effects. These allow us to flexibly capture any unobserved demand shocks across plans, areas, years, and demographic and health groups.

Given this choice model specification, the price-sensitivity parameters are identified by *within-plan* premium variation across income groups created by subsidy rules, as discussed in Section 4 above.[31]

[30] 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.

[31] The plan-region-year and plan-region-income dummies in $\xi_j(W_{it})$ in (16) align closely with the dummies in our DD

Below-poverty enrollees pay no premiums for any plan, whereas higher-income groups pay positive premiums and face higher premiums when a plan's pre-subsidy price goes up. 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 (see Figure 5) lend support to our identification strategy.

Demand Estimates. We estimate the plan choice model using maximum likelihood. We report the full set of demand estimates in Appendix E.1. Table 2 reports the implied own-price semi-elasticities of demand and average cost curve slopes, both overall and for each plan and income group (averaged across plans). These are calculated using the levels of demand evaluated at observed prices in the data and are based on a \$1 change in a plan's *post-subsidy* enrollee premium. The analysis excludes below-poverty enrollees, for whom we cannot estimate premium coefficients because they never pay premiums.

The model implies that each \$10 increase in a plan's monthly premium (a 2.5% increase as a share of average pre-subsidy prices) lowers its demand among all enrollees by an average of 16 percent, or a semi-elasticity of -0.016. For new enrollees, who are less subject to inertia, the demand semi-elasticity is even larger at -0.029. Both estimates are similar to the comparable numbers from our difference-in-difference regressions in Section 4 (see Table 1), consistent with the similar identification strategies used in both approaches. The remaining rows in the table show that demand semi-elasticities vary somewhat across plans and income groups. Notably, semi-elasticities are larger for lower-income consumers and for CeltiCare, which is the low-premium, narrow-network plan. These demand elasticities are large, though within the range of prior estimates for health insurance markets.[32] The -0.016 semi-elasticity for all enrollees implies a Lerner markup of $1/0.016 = \$62.50$ above marginal costs.

Panel (b) of the table shows average cost curve slopes, which are based on a combination of the demand and cost model (see below). They suggest adverse selection that is quite strong, with estimates of $\partial AC_j/\partial P_j$ of 0.899 on average for all enrollees, and 1.685 for new enrollees. These estimate suggest strong undercutting incentives that could make it difficult for the market to support a robust number of competing firms. Average cost slopes are positive and high for *all* plans, consistent with our model where all plans compete for price-sensitive low-risk types. Notably, the slope is somewhat larger for NHP (generally seen as the broadest-network plan) and a bit lower for CeltiCare (the narrow-network plan), but they are uniformly positive and large.

specification in (14) for experiment-time periods ($\alpha_{e,t}$) and experiment-income groups ($\gamma_{e,g}$). We use slightly less rich fixed effects in our logit model to avoid incidental parameters problems, but the fixed effects are still sufficient to absorb all premium variation except for within-plan-region-year variation across income groups.

[32] Most prior work reports premium *elasticities* of demand; for instance, Ho (2006) (reviewing the literature) reports a typical range of -1.0 to -5.0, and Tebaldi (2024) finds elasticities of -1.4 to -2.0 for silver plans in ACA markets. There are two relevant elasticities in our subsidized market. For *consumer-side* elasticities, multiplying semi-elasticities times average enrollee premiums (\$48 per month) yields -0.77 for all enrollees and -1.39 for new enrollees, on the low end of this range. For *firm-side* elasticities relevant for markups, multiplying times average total premiums (\$380 per month) yields -6.1 for all enrollees and -11.0 for new enrollees, somewhat above the usual range.

Table 2. Implied Demand Semi-elasticities and Average Cost Slopes

	(a) Demand Semi-Elasticity		*(b) Avg. Cost Slope (dAC/dP)*	
	All Enrollees	New Enrollees	All Enrollees	New Enrollees
Overall	**-0.016**	**-0.029**	**0.899**	**1.685**
By Plan				
BMC	-0.013	-0.024	0.647	1.107
CeltiCare	-0.037	-0.041	0.891	1.057
NHP	-0.021	-0.037	1.376	3.060
Network	-0.015	-0.028	0.881	1.682
By Income Group				
100-150% Poverty	-0.022	-0.043	0.843	1.763
150-200% Poverty	-0.013	-0.022	0.778	1.335
200-250% Povery	-0.012	-0.019	0.960	1.365
250-300% Poverty	-0.010	-0.016	0.843	1.242

Notes: Table reports own-price semi-elasticities of demand and average cost curve slope ($\partial AC_j/\partial P_j$) by plan, by income group, and for the market as a whole. Own-price semi-elasticities are computed for each plan using the formula $\eta_j = -\sum_i(\partial s_{ij}/\partial p_j)/(\sum_i s_{ij})$. The average cost price derivatives are computed without risk adjustment using the formula $\eta_j * (AC_j - MC_j)$ introduced in Equation (6). The "Overall" row shows averages of the plan-specific values, weighted by plan market shares, and we report results separately for all enrollees vs. new enrollees. Results for Fallon are omitted because they enroll few enrollees and are omitted in our simulations in Section 6.

5.2 Cost Model

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

$$C_{ijt}^{obs} = \Gamma_{it} \times \delta_{j,r(i)}. \tag{17}$$

We then proceed in two steps. First, we estimate $\delta_{j,r}$. To do so, we leverage cases where the same individual enrolls in the market in two separate spells in which they choose different plans.[33] This lets us estimate a model of plan effects on costs after controlling for both time-varying enrollee observables and also individual fixed effects. Our estimation sample has observations at the enrollee-by-enrollment spell level, and we limit to individuals observed in at least two spells, separated by a gap in CommCare

[33] We also explored using individuals who switch plans within a given spell (i.e., at open enrollment). However, we found that this sample was small and non-representative, likely due to the large role of inertia.

enrollment.[34] We estimate the following Poisson regression specification:

$$E(C_{ijt}^{obs}|Z_{it}) = \exp\left(\alpha_i + \beta_t + Z_{it}\gamma + \lambda_{j,r}\right) \tag{18}$$

This specification controls for individual fixed effects (α_i), year fixed effects (β_t), and time-varying enrollee observables Z_{it} (age-sex bins, a spline in risk score, income group, and enrollee location). The $\lambda_{j,r}$ coefficients represent the plan-specific cost effects, which we allow to vary across regions r to account for differential cost structures based on a plan's regional provider network. The estimated multiplicative plan cost effect of interest is $\hat{\delta}_{j,r} = \exp(\hat{\lambda}_{j,r})$. We normalize the scale of these fixed effects so that $\hat{\delta}_{j,r}$ has an (enrollment-weighted) mean of 1.0 across all plans. The model assumes that plan cost effects are constant over time, though they can vary by region. This is reasonable only if the determinants of costs — in our setting, primarily networks — are stable, which is roughly true over our 2007-2011 sample period. Additionally, we have explored allowing plan effects to vary across years in the data and found similar results.

Having estimated $\hat{\delta}_{j,r}$, our second step is to predict enrollees' costs in counterfactual plans. To do so, we simply follow the specification in (17) to estimate enrollee risk as $\hat{\Gamma}_{it} = C_{ijt}^{obs}/\hat{\delta}_{j,r}$. The cost model's prediction for enrollee i's cost in a counterfactual plan k, therefore, simply equals their observed costs times the ratio of the two plan effects, $\hat{\delta}_{k,r}/\hat{\delta}_{j,r}$.

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

Plan Cost Effect Estimates Table 3 shows estimates of the cost heterogeneity parameters, $\hat{\delta}_{j,r}$, for the four main statewide plans used in our simulations.[35] We normalize the plan effects such that the average plan effect is 1.0. As expected, CeltiCare has the lowest cost effect, with costs that are 27% lower than average (i.e., average $\hat{\delta}_{j,r} = 0.73$). On the other hand, NHP has costs that are 11% higher than average. The estimated cost effects for each plan are broadly similar across regions. These estimates imply meaningful heterogeneity in costs across plans, consistent with prior work focusing on cost heterogeneity across Medicaid plans (Geruso et al., 2023*a*). On the other hand, BMC and Network—which are the two largest plans by market share—have relatively similar cost structures, as in the "horizontal" differentiation case we highlighted in the theory in Section 2.

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

[35] We exclude Fallon, which is a small plan that is only available in a subset of regions.

Table 3. Plan Cost Effects 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**

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

5.3 Model Validation and Analysis.

Comparison to Reduced Form In order to test the validity of our demand and cost model estimates, we compare our model predictions to the actual data, following the method in Appendix Section E.3. Figure 6 below shows that our estimated demand model is capable of reproducing the extreme price sensitivity evident in our reduced form results, as well as the large response of average costs to premium changes.

Figure 6. Comparing Model-Predicted Shares and Costs with Reduced Form Results

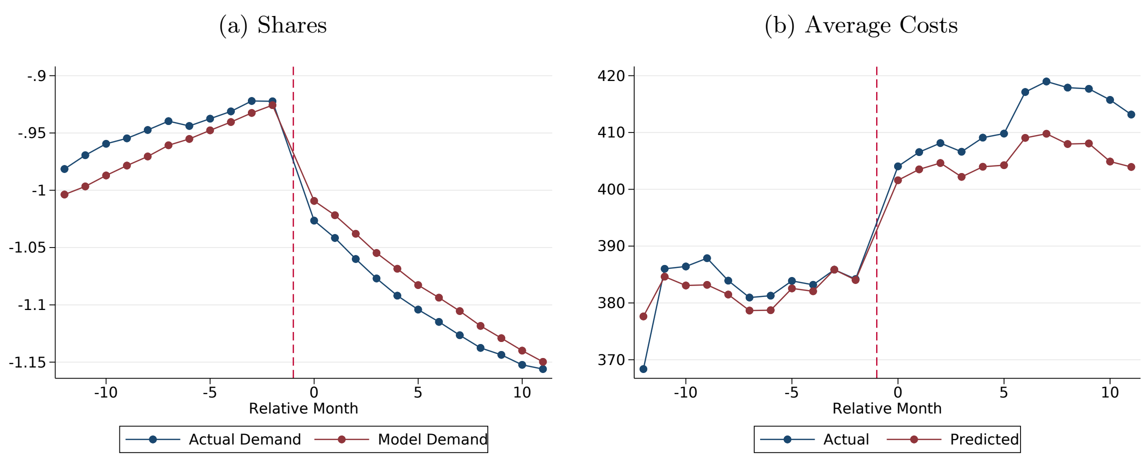

Notes: Figure shows how shares and average costs respond to increases and decreases in post-subsidy premiums, comparing actual shares and average costs for each spell with predictions using the demand estimates from Section 5. Panel (a) shows results for shares; Panel (b) shows results for average costs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

Analysis of heterogeneity generating adverse selection In Section 2, we showed that the key factor generating adverse selection under general differentiation is the (negative) correlation between enrollee risk or cost and their price-sensitivity of demand, or $-\alpha(Z_{it})$ in the demand model. Appendix Figure A10 shows a binned scatter plot of this relationship. Individual costs are strongly negatively correlated with price-sensitivity coefficients, confirming that adverse selection is strong in this market.

6 Counterfactual Simulations

We use our model estimates from Section 5 to simulate equilibrium insurer entry and price competition under various policies (or lack thereof) to address adverse selection. Our overall goal is to understand how adverse selection shapes entry, pricing, and welfare outcomes in a health insurance market using realistic insurance demand and cost primitives estimated from the Massachusetts data. Our exercise seeks to balance realism (capturing the key market institutions and policies) with simplicity and tractability (which leads us to remain close to the simple model of Section 2).[36]

6.1 Model and Simulation Assumptions

As in our theoretical model (Section 2), we model entry and price competition in a two-stage game. In stage one, a set of potential insurers ($j \in J$) simultaneously choose whether to participate in the market by offering their single (differentiated) plan with fixed non-price attributes. In stage two, participating insurers ($j \in E \subset J$) compete on prices in Nash-Bertrand equilibrium. Conditional on a set of entrants and prices, consumers choose plans and incur health care costs, which determines insurer profits. We now elaborate on key model and simulation details.

Potential Entrants A perennial challenge in applied work on entry is specifying the set of *potential* entrants and their product attributes, especially when considering counterfactual policies. We address this challenge by treating the four *actual* CommCare insurers that operated statewide as our potential entrants in all simulations.[37] We then simulate counterfactual policy changes that tend to *reduce* participation relative to the actual CommCare experience (e.g., removing price floors, weakening risk adjustment). This lets us evaluate the effect of these selection-relevant policies on entry (our main goal) while side-stepping the challenging issue of modeling entry and product design among out-of-sample firms. We view modeling the implications of selection policies for the *combination* of entry and product design as an important avenue for future work.[38]

[36]The exercise is intended as a theoretical analysis to help understand the economic impact of corrective policies that address adverse selection incentives. It does not necessarily reflect the views or strategies of the Massachusetts Health Connector or the carriers that participated in the CommCare program.

[37]This excludes one small insurer (Fallon) that operates in less than one-fifth of the state's service areas. Excluding this insurer substantially reduces simulation complexity and is unlikely to qualitatively affect our conclusions.

[38]For recent work modeling entry and product design in Medicare Advantage, see Zahn (2024). See also Miller et al. (2023), who models endogenous plan design responses subsidy policies, holding entry fixed, also in Medicare Advantage. For work studying (within-rating area) insurer entry in ACA markets, see Ko and Fang (2023) and Geddes (2024).

Baseline Specification and Policies We simulate equilibrium for the set of consumers and firms in the market in 2011, the last year covered by our structural estimates and the only year for which we observe CommCare's actual risk adjustment scores.[39] We restrict to new enrollees to avoid dynamic pricing considerations with inertial continuing enrollees, though we test robustness of key results to using a sample of all new and returning enrollees (with inertia included, but without dynamic pricing). We limit to enrollees with incomes 100-300% of FPL, which is the population for whom we can estimate demand responses to premiums.[40] The 100-300% of poverty group also better matches the relevant population for ACA Marketplaces, where subsidies are available to 100-400% of poverty enrollees. For our baseline simulations, we use CommCare's actual risk scores from 2011.

Following ACA policy, we assume a flat subsidy across all plans, which preserves pre-subsidy price differences. We do not include CommCare's incremental subsidies, which narrow price differences and tend to soften price competition.[41] Our simulations, therefore, tend to capture the stronger pricing incentives likely to prevail in the ACA setting.

In all simulations, we assume that all enrollees must choose a plan — i.e., there is no extensive margin. This is consistent with the price-linked subsidies in both the ACA and CommCare, which fix consumer premiums for the cheapest (or second-cheapest) plan regardless of insurers' prices, thereby shutting down the extensive margin as a competitor (Jaffe and Shepard, 2020; Tebaldi, 2024). However, this creates a challenge for the monopoly case, since a monopolist faces no constraint on its pricing power, other than what is imposed by regulators. Therefore, we impose a price ceiling of $475 per month, about 25% above overall average medical costs (of $379).[42] We find the ceiling to be binding only on monopolists; even with two firms, prices are competed down well below the ceiling.

Finally, in our baseline we (conservatively) assume zero fixed cost of participation ($F = 0$). This is clearly a lower bound but is not implausible given that nearly all CommCare insurers offer other health insurance plans in the state (e.g., in Medicaid, or in small group markets). As a robustness check, we consider fixed costs of $10 and $30 per enrollee-month, based on values observed in insurer financial reports.[43]

Defining and Solving for Equilibrium We solve for standard full-information (subgame perfect) Nash equilibrium in the two-stage entry and pricing game specified above. In Stage 2, conditional on a set of entrants E, insurers set Nash-Bertrand equilibrium prices (subject to any price floor or

[39]For computational tractability, simulations use a random 10% sample of N = 5,053 enrollees. Our method is slower than usual because of the need to do a careful grid search for pricing equilibria, given the undercutting incentives created by selection (see below).

[40]Individuals with incomes below 100% of poverty do not face prices (all plans are free for this group). Individuals with incomes above 300% of poverty are not eligible for the CommCare program.

[41]Because we require all enrollees to choose a plan (see next), the subsidy amount is arbitrary; we set it to $350 per month, based on average subsidies in CommCare.

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

[43]CommCare insurers report average administrative expenses of about $30 per enrollee-month. The upper value assumes this is all fixed costs; the lower value assumes one-third of this is fixed costs.

ceiling policies). In Stage 1, insurers decide whether to enter, fully anticipating outcomes in Stage 2. Although straightforward in principle, there are several important details involved.

We define a "valid pricing equilibrium" as a combination of entrants (E) and prices (P_E^*) such that three criteria hold:

1. **Nash Pricing:** P_E^* is a pure- or mixed-strategy Nash equilibrium among entrants E.
2. **Profitable Participation:** All entrants make non-negative expected profits net of fixed costs.
3. **Rational Non-entry:** No non-entering firm can unilaterally enter and earn positive profits in the Stage 2 Nash equilibrium that results when said firm enters.[44]

We solve the model backwards, starting with price competition (Stage 2). We search for Stage 2 pricing equilibria using a multi-step approach starting with a grid search, described further in Appendix Section F.[45] We then refine this set of pure-strategy candidate equilibria using Criterion #3 (Rational Non-entry), by testing whether any additional firm can enter and earn positive profits in the resulting Stage 2 pricing equilibrium. These "breaking equilibria" can be pure- or mixed-strategy equilibria.[46] In rare cases, no pure-strategy equilibrium satisfies all criteria. When this occurs, we look for a mixed strategy equilibrium, following Appendix Section F.2. In principle, there could be multiple valid equilibria that satisfy the above three conditions, but in practice, we always find a unique equilibrium in our setting.

Key Policies: Price Floors and Risk Adjustment We simulate variation in two selection-relevant policies: price floors and risk adjustment. Price floors limit how low firms can cut prices, thereby potentially stopping a price-undercutting death spiral. We consider floors ranging from $350 to $460 per month (about 5% below to 20% above market average costs of $379). In practice, we find that floors below average costs are non-binding (and therefore do not induce entry), while floors 3-4% above average costs induce higher entry with all participating firms pricing at the floor.

Risk adjustment mitigates adverse selection by adjusting firm revenues to better match each enrollee's predicted costs. For enrollee i with risk score φ_i (a measure of predicted cost relative to the average), firm j receives $\varphi_i * P_j$ for covering person i.[47] We consider two specifications for risk scores.

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

[45] Starting with a grid search helps ensure that we do not miss candidate solutions to the FOCs. Due to adverse selection, not all solutions to the FOCs will be global optima for the insurers' maximization problem. Indeed, we find that some solutions to the FOCs are local minima for certain firms.

[46] For example, in order to establish that a Celticare monopoly is the unique equilibrium, we test whether any other firm, say BMC, can enter and make profits in the (potentially mixed) [BMC, Celticare] Stage 2 equilibrium. A common scenario has BMC mixing between the price ceiling of $475 and a lower price. If BMC's expected net profit is negative, then the Celticare monopoly survives. If the Celticare monopoly also survives hypothetical entry by NHP and Network, then it satisfies Criterion 3 and is an equilibrium.

[47] Curto et al. (2021) show that the pricing equilibrium with this style of risk adjustment is defined by standard Nash-Bertrand conditions, but replacing raw enrollee costs, C_{ij}, with risk-adjusted costs, $C_{ij}^{RA} = C_{ij}/\varphi_i$, and raw demand D_{ij} with risk-scaled demand, $D_{ij}^{RA} = \varphi_i D_{ij}$.

In our baseline, we use CommCare's actual risk scores, which we observe in our data. Additionally, we consider a broader range of hypothetical risk adjustment strengths (some of which may not be technologically feasible), from none to perfect. To define this, we set φ_i as a scaled function of enrollee costs relative to the mean:

$$\varphi_i = (C_{it}/\bar{C}_t)^{\lambda} \tag{19}$$

where $\bar{C}_t$ is the market-wide average cost at time t and $\lambda \in [0, 1]$ is a factor that scales the "level" of risk adjustment from none ($\lambda = 0$, implying $\varphi_i = 1$ for all i) up to perfect ($\lambda = 1$, which implies that φ_i aligns perfectly with costs).[48]

6.2 Policy Counterfactuals Results

We start by simulating equilibrium in our baseline case with limited corrective policies (actual CommCare risk adjustment and no price floors) and test robustness to a variety of assumptions. We then illustrate how outcomes change with increasing levels of risk adjustment and price floors. Finally, we identify the optimal combination of risk adjustment and price floor policies.

Baseline case: In our baseline case, shown in Row 1 of Appendix Table A6, we find that only one insurer, Celticare — the lowest-cost and narrowest-network plan — participates in equilibrium and prices at the price ceiling of $475.[49] We calculate "consumer welfare" as enrollee surplus (which accounts for both consumer premiums and plan utility, using the standard inclusive value formula)[50] minus the government's subsidy spending. Because subsidies are fixed, consumer welfare moves one-for-one with enrollee surplus. We normalize consumer welfare in this baseline case to $0.

Robustness: In panel (b) of Appendix Table A6, we explore the robustness of this monopoly equilibrium result to key assumptions. First, we eliminate Celticare, the low-cost, narrower-network option. Given its low-cost status, Celticare can often enter and "break" candidate equilibria with multiple firms. Indeed, we see that when Celticare is eliminated, two firms (BMC and Network Health) participate in equilibrium and price significantly below the ceiling, resulting in higher consumer welfare ($135 per enrollee-month). This highlights the fact that vertical differentiation can worsen the effects of selection on participation. While selection *can* lead to monopoly in the presence of purely horizontal differentiation (as shown in Section 2), the amount of selection required to achieve this outcome is higher than the amount required when there is also some degree of vertical differentiation (as when Celticare is present). In addition, we show that duopoly cannot be achieved by simply removing

[48] The parameter λ can be interpreted as the coefficient in a regression of log risk scores on log costs. As a benchmark, for CommCare's actual risk scores in our data in 2011, this regression coefficient is 0.10. For the more predictive HHS-HCC risk scores (as used by the ACA), it is 0.24.

[49] One might hypothesize that the threat of entry should constrain the monopolist's price, a concept known as "contestable markets." But in our two-stage model it does not because entry is not a credible threat, as any entrant would lose money at the second-stage equilibrium price vector that would occur.

[50] The formula for enrollee surplus for enrollee choice instance (i, t) is $\log(\sum_j (u_{ijt} - \varepsilon_{ijt}))$, where u_{ijt} is defined as in Equation (16).

Celticare's cost advantage (the row labeled "No Cost Heterogeneity"); complete removal of Celticare from the set of potential entrants is necessary.

Finally, we consider whether the baseline monopoly equilibrium result is robust to including inertial consumers who face switching costs when moving away from their current plans. We find that inertia is indeed protective in that it allows for a duopoly equilibrium with lower prices and higher consumer welfare ($133 per enrollee-month). Recall that the model in Section 2 showed that the participation constraint depended critically on the price elasticity of demand. This price elasticity is substantially lower with consumer inertia, producing this result. Similar to Handel (2013), this result shows that consumer inertia can sometimes lead to better market outcomes.

Risk Adjustment: Figure 7 shows market outcomes as we vary the level of risk adjustment from none ($\lambda = 0$) to perfect ($\lambda = 1$) (see also panel (a) of Appendix Table A6). The monopoly equilibrium persists at low and moderate levels of risk adjustment. Only when risk adjustment compensates firms for $\lambda = 0.60$ of the differences in average cost across firms do we get a duopoly equilibrium, with Network Health joining alongside Celticare. When this occurs, prices fall markedly from the ceiling ($475) to around $390 (panel (a)) and consumer surplus rises by over $100 (panel (b)). Increasing risk adjustment transfers to $\lambda = 0.70$ results in three firms participating. To get all four firms to participate, transfers compensating for $\lambda = 0.80$ of average cost differences are necessary, and this is the consumer-welfare optimal level of risk adjustment without price floors. Such a risk adjustment policy is likely infeasible with modern risk adjustment technologies that generate risk scores based on the relationship between diagnosis codes and current spending (see Geruso and Layton (2017)), suggesting that achieving full participation may not be feasible using risk adjustment alone.

Price floors: In Figure 8 we show the effects of price floors on equilibrium prices and consumer surplus, holding risk adjustment constant at the actual 2011 level.[51] With relatively low price floors (which end up being non-binding), the CeltiCare-only monopoly equilibrium persists. But when the price floor reaches $393 (about 4% above the average cost in the market), two additional firms (BMC and Network Health) enter and compete prices down to the floor. Consumer surplus is maximized at this $393 price floor: while higher floors achieve more entry, the associated welfare gains are exceeded by the higher prices involved.

Interactions and Optimal Policy Finally, we consider how risk adjustment and price floors interact to affect participation, prices, and consumer surplus by simulating a full grid of policy combinations, presented as heatmaps in Figure 9. In each heatmap, the price floor varies on the x-axis, the level of risk adjustment varies on the y-axis, and the z-axis (the color value) shows the specified outcome.

Figure 9(a) shows how firm participation depends on the interaction of risk adjustment and price floors. At most levels of risk adjustment, price floors are important to achieve robust firm participation, though they become unnecessary with very strong risk adjustment ($\lambda \geq 0.80$).

[51] Results are similar if we instead use no risk adjustment, or any level of $\lambda \in [0, 0.4]$ (see Figure 9).

Figure 7. Counterfactual Simulations: Effect of Risk Adjustment

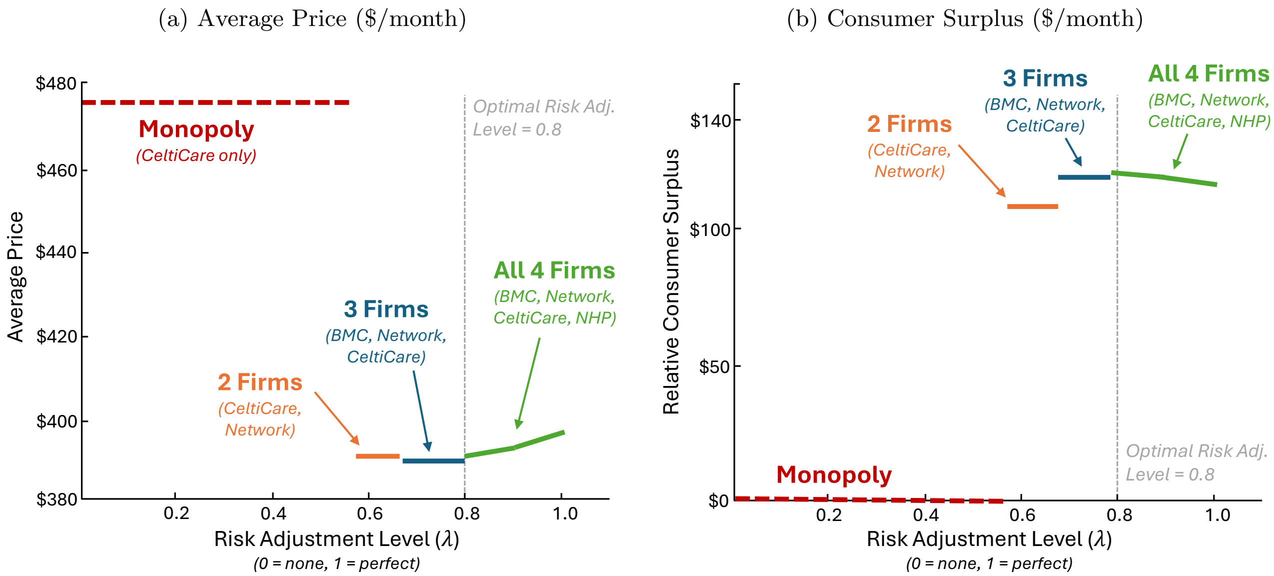

Notes: Figure shows how the market equilibrium changes when strengthening risk adjustment (x-axis) from zero risk adjustment ($\lambda = 0$) to perfect risk adjustment ($\lambda = 1$). The simulations do not include price floors. Panels (a) and (b) respectively show average prices and "relative" consumer surplus, defined as surplus relative to the worst-case outcome of monopoly with the least-preferred insurer (CeltiCare). Without risk adjustment, the market unravels to monopoly with very high prices. Stronger risk adjustment allows more firms to enter but can also push up markups. The optimum is at a risk adjustment level of 0.8. Note that we conduct simulations for a discrete grid of risk adjustment levels, so this value is approximate.

Figure 9(b) shows consumer surplus. Optimal policy is shown by the red line, which indicates the optimal price floor for each level of risk adjustment. With most levels of risk adjustment, the optimal price floor does not vary greatly across simulations, typically falling just above the average cost in the market (= $379). With very strong risk adjustment ($\lambda \geq 0.80$), price floors lower consumer surplus because they aren't necessary to achieve full participation and only result in higher average prices.[52] Overall, we conclude (perhaps surprisingly) that price floors usually benefit consumers, as long as there is meaningful adverse selection remaining after risk adjustment.

Looking across the full 2x2 policy grid, it is interesting to note that the optimal policy for consumers is *not* perfect risk adjustment with no price floor (i.e., completely eliminating risk selection). This is because some adverse selection is beneficial by inducing firms to constrain markups, consistent with the ideas of Starc (2014) and Mahoney and Weyl (2017). While perfect risk adjustment achieves full entry, it does so at the cost of higher prices. Instead, the optimal policy is *modest* risk adjustment ($\lambda = 0.5$) combined with a *modest* price floor of $380 (almost exactly equal to average costs), which is sufficient to induce three of four firms to enter. The lone non-entrant (NHP) is the highest-cost plan, and the additional surplus from its availability is outweighed by the price increase needed to sustain

[52]This can be seen more clearly in Appendix Figure A22, which shows heatmaps for additional outcomes, including prices and profits. Panels (a) and (b) of this figure present the average and minimum equilibrium prices respectively.

Figure 8. Counterfactual Simulations: Effect of Price Floors

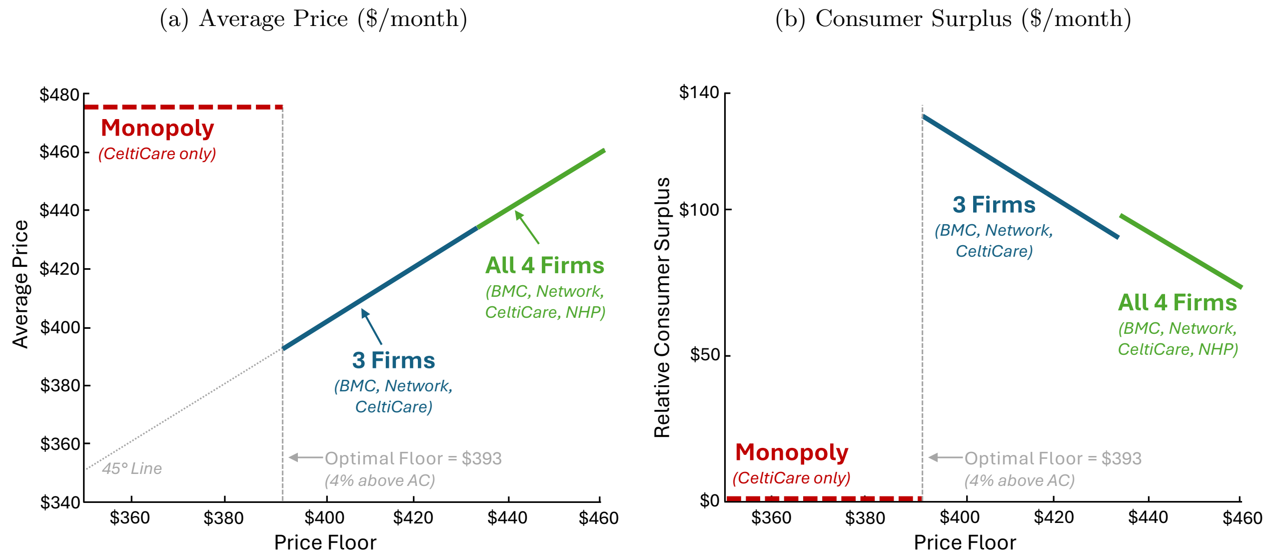

Notes: Figure shows how the market equilibrium changes with an increasing price floor, from a low (and non-binding) level of $350 (about 8% below the overall market average cost of $379) to a high level of $460 (about 20% above average costs). Risk adjustment is held at actual 2011 levels, as in our baseline case in Table A6. Panels (a) and (b) respectively show average prices and "relative" consumer surplus, defined as surplus relative to the outcome of monopoly with the least-preferred insurer (CeltiCare). The consumer-surplus optimal floor is $393, about 4% above average costs. Below this, the market unravels to monopoly at a very high price. At $393, two additional firms are willing to enter, and they compete prices down to the (now binding) floor. Above this, a higher floor can increase entry but also raises prices, reducing consumer surplus on net. Note that we conduct simulations for a discrete grid of price floors, so the value for the optimal floor is approximate.

its participation. This result highlights the insight that more entrants may not always be better, if enticing the marginal firm to enter requires too much softening of price competition.

7 Discussion and Application to the ACA Experience

Thus far, we have presented a new theory suggesting that adverse selection, by stimulating aggressive price competition, can limit the number of firms that can profitably participate in a market (Section 2). We have also estimated elasticities in a real-world and policy-relevant setting (the Massachusetts exchange) that suggest that adverse selection is strong enough to sharply limit participation without corrective policies (Sections 3-4). Our structural model simulations confirm this "un-natural monopoly" finding and evaluate the role of two corrective policies, risk adjustment and price regulation, for sustaining competition and improving consumer welfare (Sections 5-6). Together, these analyses all point towards a new theoretical implication of adverse selection: that it can limit firm participation in insurance markets.

In this section, we take this largely theory-based argument and show how it can help explain some key empirical facts and institutions in health insurance markets. We start by examining limited

Figure 9. Counterfactuals: Market Outcomes across Grid of Risk Adjustment and Price Floors

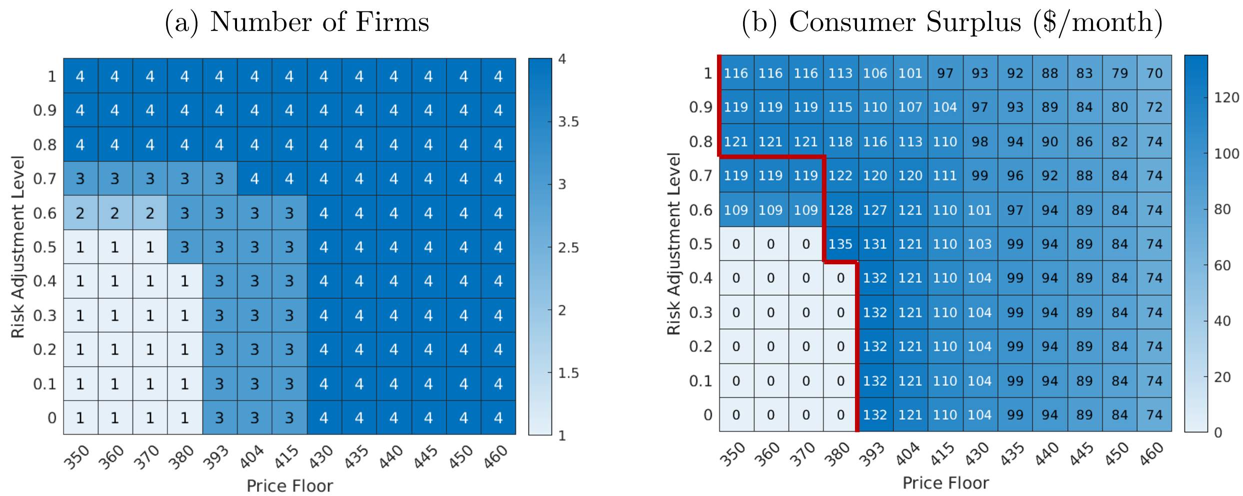

Notes: Figures show how equilibrium outcomes vary as a function of the level of risk adjustment $\lambda \in (0, 1)$ (shown on the Y-axis) and price floors (X-axis). For each panel, the Z-axis (i.e., heatmap color) plots a different outcome of interest. Each cell corresponds to a unique equilibrium. Panel (a) shows the number of firms that enter; Panel (b) shows the resulting consumer surplus (relative to the worst-case outcome of monopoly with the least-preferred firm). The red-shaded line indicates the optimal price floor for each level of risk adjustment. Notably, the optimal floor is positive (and binding) except at strong risk adjustment ($\lambda \geq 0.8$) that is itself sufficient to induce full entry by all four firms.

competition in the ACA Marketplaces — the setting that motivated our paper — and then discuss several other examples more briefly. While we intend this section fundamentally as a discussion (not a formal empirical test), we hope it can stimulate further thinking and empirical work on the determinants of insurance market competition. We cite a few examples of work on the ACA where this is starting to occur.

7.1 Insurer Competition in ACA Marketplaces

A key motivating fact for our paper is the low levels of competition in the state-level ACA health insurance Marketplaces (or "exchanges"). These markets, like our Massachusetts setting, were established as regulated markets where private insurers can compete to offer plans to individuals without access to affordable job-based coverage. Like Massachusetts, many aspects of plan design (including covered benefits and financial generosity) are regulated, but insurers can compete on premiums and less-regulated features like networks. The ACA Marketplaces have been described and studied extensively in other work (see Handel and Kolstad (2022) for a review). Here, we focus on trends in insurer competition and how they relate to policies relevant for adverse selection and price competition.

Figure 10 shows the main trends in insurer participation in the ACA markets.[53] There are two

[53]We focus on the HealthCare.gov states, for which insurer participation and plan premium data are most readily available. We limit the sample to the 30 states that are stably observed in this dataset from 2014-24. See Appendix G.1 for a description of our dataset.

Figure 10. ACA Markets: Insurer Participation and Zero-Premium Floors

(a) Insurer Participation over time (b) Share with Plans at Zero-Premium Floor

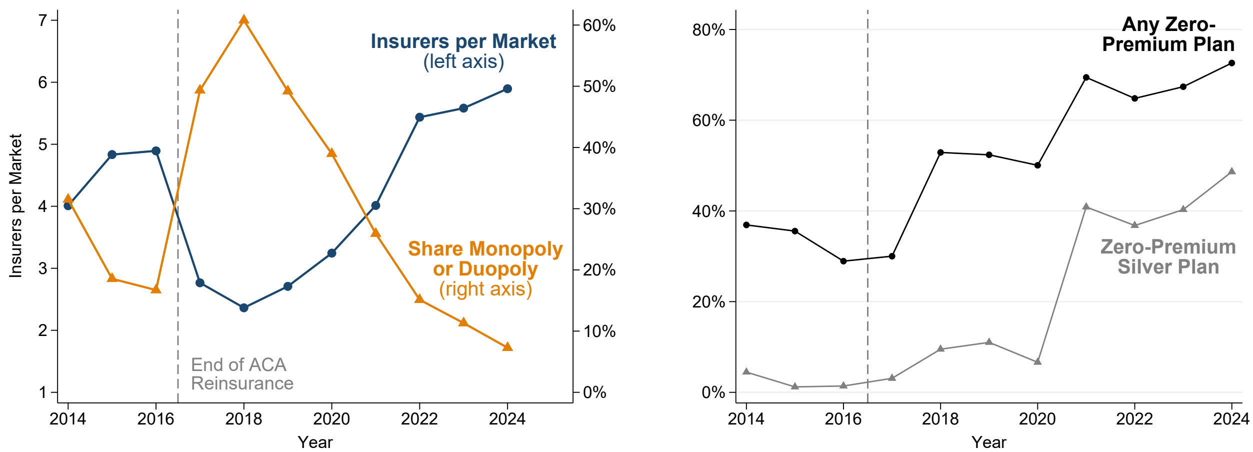

Notes: Figure shows trends for ACA insurance Marketplaces in insurer participation (panel a) and the share of enrollees with plans that hit the zero-premium floor after subsidies (panel b). All estimates are for the 30 states covered by Healthcare.gov over the 2014-24 period, and averages are weighted by county population. Panel (b) shows estimates of the share of enrollees with at least one zero-premium gold/silver/bronze plan, and one zero-premium silver plan. Estimates are based on premiums and subsidies for a 40-year-old individual (approximately the median age) and the actual distribution of enrollees across income groups. See Appendix G.2 for details.

salient facts. First, although limited in all years, insurer participation falls sharply in 2017-2018. The average number of insurers per county market (the level at which entry occurs) falls from 4.9 in 2016 to 2.4 in 2018, and the (population-weighted) share of markets that were monopoly or duopoly triples from 17% to 61%. Second, participation gradually recovers from 2018-2022 after which it stabilizes around 5.5-6.0 firms per market. Although widely reported, these trends are poorly understood, partly because there is not a standard theoretical framework for the determinants of insurer entry.[54]

Our ideas can help account for these facts about ACA insurer participation. First, we note that relative to Massachusetts and other market-based insurance programs, the ACA features relatively strong and flexible price competition. The ACA subsidies are designed to be a flat amount across all plans (for a given individual) that therefore preserves pre-subsidy price differences and avoids softening price competition. This differs from the *incremental* subsidies used for many of the poorest consumers in Massachusetts and other settings (e.g., the Low-Income Subsidy enrollees in Medicare Part D). Additionally, unlike Massachusetts, the ACA does not feature explicit pre-subsidy premium floors — though, as we argue below, an *indirect* floor at a $0 consumer premium has become increasingly relevant. Given the ACA's price-sensitive low-income consumer population and strong evidence of adverse selection (see e.g., Tebaldi (2024), Saltzman (2021), Panhans (2019)), it is sensible to argue that our theory that adverse selection limits insurer participation applies here.

[54] An exception to this is work by Ko and Fang (2023) and Geddes (2024), who study a particular component of insurer entry: *partial* insurer entry among counties within larger pricing (rating) areas.

Second, our ideas can help explain the timing of participation changes. The fall in insurer competition has long been associated in press coverage and industry reports with the 2017 expiration of two temporary policies intended to stabilize premiums and mitigate adverse selection: reinsurance and risk corridors.[55] Consistent with these popular accounts, Holmes (2024) provides causal empirical evidence that these policies' expiration contributed to higher premiums and lower insurer competition in ACA marketplaces relative to small-group insurance markets in the same states. However, previously there was no theoretical link between these policies and insurer participation. Our model helps provide a theoretical foundation for these empirical findings based on the idea that they mitigate adverse selection and selection limits competition.[56]

In Appendix G, we provide additional suggestive evidence that reinsurance's impact on *adverse selection* specifically contributed to this fall in competition. Although reinsurance was a national policy, we show that it had differential impacts across states on the price-slope of (net-of-reinsurance) average costs, $\frac{dAC^{Net}}{dP} = \frac{dAC}{dP} - \frac{dReins}{dP}$, our key model statistic for adverse selection, via heterogeneity in the elasticity of reinsurance to plan premiums. We find that pre-2017 participation was higher in states where reinsurance had a larger selection-mitigating impact, and it fell more during 2017-18 after reinsurance expired. This finding complements those of Holmes (2024) and suggests selection as a key mechanism for the decrease in firm participation. However, we note, as a caveat, that research on (smaller) state-specific reinsurance programs introduced since 2019 finds that these policies reduced premiums but did not spur entry (Schwab et al., 2018; Oyeka and Wehby, 2023). More work is needed to sort out this mixed evidence.

Finally, our ideas may also help explain the rebound in ACA insurer competition since 2019. An under-appreciated feature of ACA subsidies is that they are capped at a plan's full premium: post-subsidy consumer premiums may not go below \$0. This creates an indirect price floor that can bind for a subset of consumers and plans — especially for lower-income consumers (who get larger subsidies) and for lower-premium Bronze plans. In the ACA's early years, this price floor was rarely binding for silver plans and binding for only about one-third of enrollees in any plan (usually a Bronze plan). But starting in 2018, the benchmark second-lowest premium silver plan (to which subsidies are linked) rose sharply due to a variety of factors.[57] Finally, policy changes starting in 2021 increased subsidies across the board and ensured that all enrollees with incomes below 150% of poverty qualified for two zero-premium silver plans. All together, this sharply increased the share of consumers and plans for whom the zero-premium floor was binding.

[55] Reinsurance is a type of "insurance for insurers" that mitigated selection by covering part of the actual expenses of high-cost enrollees (above an attachment point of \$45,000 to \$90,000 per year). Risk corridors compensated insurers whose claim costs fell short of premium revenues. In practice, reinsurance was much larger than risk corridors (which were never fully funded), so we focus on it in the discussion below.

[56] Insurer exits have also been associated with market turmoil around the Trump administration's 2017 (unsuccessful) push to "repeal and replace" the ACA. However, Figure 10 shows that most of the fall in competition occurred in *2017*, for which insurer decisions were made during the spring/summer of 2016, well before the November 2016 election. The ACA repeal debate no doubt contributed to the further decline in competition in 2018.

[57] These include the end of reinsurance, insurer exits, and "silver-loading" policies (beginning in 2018) in which state regulators pushed insurers to raise silver premiums to compensate for lost cost-sharing reduction subsidies that were cancelled by the Trump administration.

We illustrate this pattern in Figure 10(b), which estimates the share of Marketplace enrollees (across all income groups) eligible for zero-premium plans. By 2024, this share had risen to 73% for any plan, and to 49% for silver plans. The rebound of insurer entry coincides with a sharp rise in the role of price floors, the key policy that emerged from our analysis of the selection and price competition problem. Although more work is needed to test for a causal connection, this evidence is at least consistent with the main ideas of our paper and suggests that our model can help explain previously not-well-understood swings in firm participation in these markets.

7.2 Relevance For Other Insurance Markets

In addition to the ACA experience, there are several other stylized facts from other insurance markets that are consistent with our theory.

Medicare Advantage: The Medicare Advantage (MA) program is another large individual insurance market in the United States that allows Medicare beneficiaries to take a (fixed-value) voucher and enroll in a private health plan instead of the traditional fee-for-service Medicare plan. As of 2023, more than 50% of Medicare beneficiaries are enrolled in a MA plan (Kaiser Family Foundation, 2023). Unlike the ACA Marketplaces we study in this paper, MA has relatively high levels of insurer participation.[58]

There are many potential reasons why insurer participation is less of an issue in MA, but we highlight three here that are consistent with our theory. First, in contrast to the ACA setting, only a very small share of MA beneficiaries are new enrollees making active choices, with a much larger role of inertia in MA. This decreases overall levels of price sensitivity and thus the size of the Lerner markup, making it easier for that markup to overcome the AC-MC wedge we highlight in Section 2. Second, MA has a soft premium floor at $0 (above the standard Part B amount), and Stockley et al. (2014) finds that roughly 50% of MA consumers are in plans whose premium binds at this floor. This price floor halts any premium undercutting and may increase firm participation as we show in Section 2 and Section 6. Finally, MA enrollees also always have the option to enroll in the traditional fee-for-service Medicare plan. As its price is set exogenously (rather than strategically) and its network is notably broad, traditional Medicare's presence may also interrupt the undercutting spiral.

Massachusetts Exchange: Our model is also consistent with experience in the (post-ACA) Massachusetts Connector exchange, the successor to the CommCare program we study. According to market regulators, two insurers reported concerns about large swings caused by small premium differences (of just $1-15 per month), with the insurer charging the slightly higher premium getting much lower market share and higher average-cost set of enrollees (consistent with our findings of high price sensitivity and average cost slopes in Section 4). The insurers expressed concerns about their ability

[58] In 2023, the average beneficiary had access to plans offered by 9 insurers. Almost 90% of beneficiaries had access to at least six insurers; more than 65% of beneficiaries had access to at least eight insurers; and a full 40% of consumers had access to 10 or more insurers (Kaiser Family Foundation, 2022). That said, MA markets are still highly concentrated, with two firms dominating nationally and locally.

to remain in the market given this instability. In response, starting in 2019 the regulator instituted a "premium smoothing" policy that narrowed premium differences for plans with (pre-subsidy) prices within a small threshold of the lowest. In practice, this meant that consumers faced *equal* premiums for the 2-3 lowest-price plans in most regions of the state (Connector, 2018). This mitigated selection undercutting incentives among low-price plans, and ultimately helped the concerned insurers remain in the market.

"Actuarial Soundness" Regulations: Our model also offers an explanation for insurance regulators' widespread use of "actuarial soundness" rules in which they reject premiums deemed too low, effectively acting as a price floor. Often, the stated reason is to ensure that premium income can cover costs, so that insurers remain financially solvent. However, another potential rationale for actuarial soundness rules may be their ability to avoid inefficient selection-induced price undercutting. By imposing a de-facto price floor, these rules may help to ensure greater firm participation and higher consumer welfare, as we argue in Section 6.

8 Conclusion

Adverse selection has been shown to cause many problems in insurance markets (Rothschild and Stiglitz, 1976; Cutler and Reber, 1998; Einav et al., 2010; Geruso et al., 2023*b*, 2019). The prior literature has focused on two key sets of problems: unraveling of trade and contract quality distortions.[59] In this paper, we show that selection can cause a third problem that may be equally important: it can limit the number of firms that the market can support, even without vertical differentiation (i.e., "high-coverage" vs. "low-coverage" plans). Indeed, in the extreme case, it can cause a market to become an (un)natural monopoly, so called because the lack of competition does not stem from any "real" fixed costs but rather a sort of coordination failure between undercutting firms. We show this via a general model of an insurance market that highlights the effects of selection on entry and pricing. We also show that the un-natural monopoly result is not just theoretical—it is actually the outcome predicted by our empirical model 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 using risk adjustment, a common policy in health insurance markets, or price floors.

These findings have important implications for health insurance markets. Individual health insurance markets are now highly prevalent in health insurance programs in the U.S. and globally. Our results show how fragile these markets can be and just how much they rely on corrective policies such as risk adjustment to succeed. Our results also suggest that an additional policy, price floors, could improve outcomes in settings where they are not already in place (directly or indirectly).

Our findings may also generalize to other markets with downward-sloping average cost curves, such as markets with significant fixed costs and high price-sensitivity (e.g., generic pharmaceuticals). We show how market competition is stymied by undercutting incentives, which endogenously determine

[59] Contract distortions include the well-known results showing that more generous insurance plans can often unravel (Rothschild and Stiglitz, 1976).

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 (Mahoney and Weyl, 2017). 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 can directly affect 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, the best policies to improve competition in these markets may be those that target adverse selection.

References

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

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

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

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

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

Chade, Hector, Victoria R Marone, Amanda Starc, and Jeroen Swinkels. 2022. "Multidimensional Screening and Menu Design in Health Insurance Markets." National Bureau of Economic Research Working Paper 29886.

Chandra, Amitabh, Jonathan Gruber, and Robin McKnight. 2010. "Patient Cost-Sharing and Hospitalization Offsets in the Elderly." *American Economic Review*, 100(1): 193–213.

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

Connector, Massachusetts. 2018. "Final Award of 2019 Seal of Approval." https://www.mahealthconnector.org/wp-content/uploads/board_meetings/2018/09-13-18/Final-Award-of-2019-Seal-of-Approval-VOTE-091318.pdf.

Curto, Vilsa, Liran Einav, Jonathan Levin, and Jay Bhattacharya. 2021. "Can health insurance competition work? evidence from medicare advantage." *Journal of Political Economy*, 129(2): 570–606.

Cutler, David M, 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.

Dell'Ariccia, Giovanni, Ezra Friedman, and Robert Marquez. 1999. "Adverse selection as a barrier to entry in the banking industry." *The RAND Journal of Economics*, 515–534.

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

Einav, Liran, Amy Finkelstein, and Neale Mahoney. 2021. "The IO of selection markets." In *Handbook of Industrial Organization*. Vol. 5, 389–426. Elsevier.

Einav, Liran, Amy Finkelstein, and Pietro Tebaldi. 2018. "Market Design in Regulated Health Insurance Markets: Risk Adjustment vs. Subsidies." *NBER Working Paper 32586*.

Einav, Liran, Amy Finkelstein, and Ray Fisman. 2022. *Risky Business: Why Insurance Markets Fail And What To Do About It.* New Haven:Yale University Press.

Einav, Liran, Amy Finkelstein, Stephen P Ryan, Paul Schrimpf, and Mark R Cullen. 2013. "Selection on moral hazard in health insurance." *The American economic review*, 103(1): 178–219.

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

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

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

Ericson, Keith Marzilli, and Justin Sydnor. 2017. "The questionable value of having a choice of levels of health insurance coverage." *Journal of Economic Perspectives*, 31(4): 51–72.

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

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

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

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

Geddes, Eilidh. 2024. "The Effects of Price Regulation in Markets with Strategic Entry: Evidence from Health Insurance Markets." Working Paper.

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. 2020. "Upcoding: evidence from Medicare on squishy risk adjustment." *Journal of Political Economy*, 128(3): 984–1026.

Geruso, Michael, Timothy J Layton, and Jacob Wallace. 2023*a*. "What difference does a health plan make? Evidence from random plan assignment in Medicaid." *American Economic Journal: Applied Economics*, 15(3): 341–379.

Geruso, Michael, Timothy J Layton, Grace McCormack, and Mark Shepard. 2023*b*. "The two-margin problem in insurance markets." *Review of Economics and Statistics*, 105(2): 237–257.

Geruso, Michael, Timothy Layton, and Daniel Prinz. 2019. "Screening in contract design: Evidence from the ACA health insurance exchanges." *American Economic Journal: Economic Policy*, 11(2): 64–107.

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

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

Handel, Benjamin, and Jonathan Kolstad. 2022. "The affordable care act after a decade: industrial organization of the insurance exchanges." *Annual Review of Economics*, 14(1): 287–312.

Handel, Benjamin R. 2013. "Adverse selection and inertia in health insurance markets: When nudging hurts." *American Economic Review*, 103(7): 2643–82.

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

Hendren, Nathaniel, Camille Landais, and Johannes Spinnewijn. 2021. "Choice in Insurance Markets: A Pigouvian Approach to Social Insurance Design." *Annual Review of Economics*, 13: 457–486.

Ho, Katherine. 2006. "The welfare effects of restricted hospital choice in the US medical care market." *Journal of Applied Econometrics*, 21(7): 1039–1079.

Holmes, Jonathan. 2024. "How Insurer Subsidies Affect Market Equilibria."

Jaffe, Sonia, and Mark Shepard. 2020. "Price-linked subsidies and imperfect competition in health insurance." *American Economic Journal: Economic Policy*, 12(3): 279–311.

Kaiser Family Foundation. 2022. "Medicare Advantage 2023 Spotlight: First Look." *https://www.kff.org/medicare/issue-brief/medicare-advantage-2023-spotlight-first-look/*.

Kaiser Family Foundation. 2023. "Medicare Advantage in 2023: Enrollment Update and Key Trends." *https://www.kff.org/medicare/issue-brief/*

medicare-advantage-in-2023-enrollment-update-and-key-trends/.

Kautter, John, Gregory C Pope, Melvin Ingber, Sara Freeman, Lindsey Patterson, Michael Cohen, and Patricia Keenan. 2014. "The HHS-HCC risk adjustment model for individual and small group markets under the Affordable Care Act." *Medicare & Medicaid research review*, 4(3).

Ko, Ami, and Hanming Fang. 2023. "Pricing Regulation and Insurer Entry on the ACA Health Insurance Marketplaces." Working Paper.

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

Mahoney, Neale, and E Glen Weyl. 2017. "Imperfect competition in selection markets." *Review of Economics and Statistics*, 99(4): 637–651.

Marone, Victoria, and Adrienne Sabety. 2022. "When Should There Be Vertical Choice in Health Insurance Markets." *American Economic Review*, 112(1): 304–42.

Massachusetts Health Connector. 2014. "Commonwealth Care Program: Administrative Enrollment and Health Insurance Claims Records." *Accessed via DUA with the Massachusetts Health Connector.*

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

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

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

Miller, Keaton, Amil Petrin, Town Robert, and Chernew Michael. 2023. "The Optimal Geographic Distribution of Managed Competition Subsidies." Technical Report.

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

Oyeka, Onyinye, and George L Wehby. 2023. "Effects of state reinsurance programs on health insurance exchange premiums and insurer participation." *Health services research*, 58(5): 1077–1088.

Pakes, Ariel, Jack R Porter, Mark Shepard, and Sophie Calder-Wang. 2021. "Unobserved heterogeneity, state dependence, and health plan choices." National Bureau of Economic Research.

Panhans, Matthew. 2019. "Adverse Selection in ACA Exchange Markets: Evidence from Colorado." *American Economic Journal: Economic Policy*, 11(2): 1–36.

Riley, John G. 1979. "Informational Equilibrium." *Econometrica*, 47(2): 331–359.

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

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

Saltzman, Evan. 2021. "Managing adverse selection: underinsurance versus underenrollment." *The RAND Journal of Economics*, 52(2): 359–381.

Schwab, Rachel, Emily Curran, and Sabrina Corlette. 2018. "Assessing the effectiveness of state-based reinsurance: case studies of three states' efforts to bolster their individual markets." *Georegtown University Health Policy Institute.*

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

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

Shepard, Mark, and Myles Wagner. 2022. "Do ordeals work for selection markets? Evidence from health insurance auto-enrollment." National Bureau of Economic Research working paper.

Spence, Michael. 1983. "Contestable markets and the theory of industry structure: A review article."

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

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

Stockley, Karen, Thomas McGuire, Christopher Afendulis, and Michael Chernew. 2014. "Premium Transparency in the Medicare Advantage Market: Implications for Premiums, Benefits, and Efficiency." National Bureau of Economic Research Working Paper 20208.

Tebaldi, Pietro. 2024. "Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the ACA." *Review of Economic Studies*.

Veiga, André, and E. Glen Weyl. 2016. "Product Design in Selection Markets." *Quarterly Journal of Economics*, 131(2): 1007–1056.

Zahn, Matthew V. 2024. "Entry and Competition in Insurance Markets: Evidence from Medicare Advantage." *Working Paper*.

Appendix

Table of Contents

A Theory Appendix

This appendix provides supplementary material for Section 2 presenting our model. We include derivations for the Salop Model with adverse selection, proofs, and additional figures.

A.1 Derivations for Salop Model with Adverse Selection

This appendix presents derivations for the Salop model with adverse selection in Section 2. For completeness, we recap the model setup and show how to derive key equations and functions mentioned in the body text. Each element of the derivation is called out in a separate bold-heading paragraph.

In the model, a mass-1.0 population of consumers (indexed by i) reside uniformly around a small city, modeled as a unit-circumference circle, with consumer locations indexed by $\ell_i \sim U[0,1]$. A set of N competing firms (indexed by j) locate equidistantly around the circle at locations $\ell_j \in \{0, \frac{1}{N}, ... \frac{N-1}{N}\}$. Firms sell a homogeneous product (or "plan") of value V_i to each consumer i (i.e., there are no "vertical" quality differences between products), but consumers dislike travel so prefer nearby firms.

We enrich the classic Salop model setup by allowing for two risk types, $r \in \{L, H\}$: (1) healthy (low-risk) type L with costs C_L who comprise share θ_L of consumers, and (2) sick (high-risk) type H with costs $C_H > C_L$ who comprise share $\theta_H = 1 - \theta_L$. Each type is uniformly distributed around the city (preserving firm symmetry), but types may differ in their travel disutility, $t_L \neq t_H$. Firms cannot price discriminate across types, implying that there is cost-relevant asymmetric information. We denote overall market average cost as $\overline{C} \equiv E[C_r] = \sum_r \theta_r C_r$, and likewise for average travel disutility, $\bar{t} \equiv E[t_r]$.

The utility function for consumer i of risk type r for firm j is:

$$U_{irj} = [V_i - t_r \cdot d_{ij}] - \ P_j \tag{20}$$

where $d_{ij} = |\ell_i - \ell_j|$ is travel distance from consumer i to firm j's location. We assume that each consumer buys exactly one good, which occurs if V_i is sufficiently large and it is a discrete-good setting like insurance (where each consumer needs only one plan).

Derivation of Type-Specific Demand We solve for demand by finding the location ℓ_r^* at which type-r consumers are indifferent between their two adjacent firms given prices P, and then calculating implied market shares given the uniform distribution of consumers. Since all firms are symmetric, we can do this for consumers living between a single pair of firms: firm 1 (at location $\ell_1 = 0$) and firm 2 (at location $\ell_2 = \frac{1}{N}$). A consumer of type-r is indifferent between these two firms if

$$U_{ir1} = [V_i - t_r \cdot d_{i1}^*] - \ P_1 = U_{ir2} = [V_i - t_r \cdot d_{i2}^*] - \ P_2$$

rearranging terms yields:

$$d_{i1}^* - d_{i2}^* = -\frac{(P_1 - P_2)}{t_r} = -\alpha_r \cdot (P_1 - P_2)$$

where we use the definition $\alpha_r = \frac{1}{t_r}$. Plugging in $d_{i1}^* = \ell_r^* - 0$ and $d_{i2}^* = \frac{1}{N} - \ell_r^*$ yields:

$$\ell_r^* = \frac{1}{2} \cdot \left[\frac{1}{N} - \alpha_r \cdot (P_1 - P_2)\right] \tag{21}$$

which is the valid equation for the indifferent type as long as $\ell_r^* \in [0, \frac{1}{N}]$.[60] Assuming this holds, the set of type-r consumers living in $\ell_i \sim [0, l_r^*]$ all choose firm 1 (instead of firm 2), and firm 1's demand on this margin is ℓ_r^*. By a symmetric argument, its indifferent type-r consumer in competing with its neighboring firm on the other side ($j = N$) lives at location

$$1 - \ell_r^{**} = \frac{1}{2} \cdot \left[\frac{1}{N} - \alpha_r \cdot (P_1 - P_N)\right]$$

and this is also firm 1's demand on that margin. The total demand among type-r consumers for firm $j = 1$, therefore, is:

$$D_{rj}(P) \;=\; \ell_r^* + (1 - \ell_r^{**}) \;=\; \frac{1}{N} - \alpha_r \cdot \left(P_j - \overline{P}_{j'}\right) \tag{22}$$

where $\overline{P}_{j'} \equiv \frac{1}{2}(P_2 + P_N)$ is the average price of firm $j = 1$'s neighboring competitors. This replicates body-text equation (1) for type-r demand for firm $j = 1$, and by symmetry, the same function holds for all other firms j.

We note, as a technicality, that equation (22) is only valid for demand if price differences are not too large – which is true locally around all the symmetric pricing equilibria we consider. More completely, type-r demand for firm j equals:

$$D_{rj}(P) \;=\; \max\left\{0,\; \min\left\{1,\; \frac{1}{N} - \alpha_r \cdot \left(P_j - \overline{P}_{j'}\right)\right\}\right\} \tag{23}$$

Moving forward, we focus on the relevant case in which price differences are not too large and equation (22) holds for demand

Total Demand, Costs, and Profits Summing across types, firm j's total market share is

$$D_j(P) = \sum_r \theta_r D_{rj}(P) = \frac{1}{N_f} - \overline{\alpha} \cdot \left(P_j - \overline{P}_{j'}\right)$$

where $\overline{\alpha} = \sum_r \theta_r \alpha_r$ is the population-average price-sensitivity. Total firm (variable) profits equal

$$\pi_j(P) \;=\; \sum_r [P_j - C_r] \cdot \theta_r D_{rj}(P) \;\; = [P_j - AC_j(P)] \cdot D_j(P) \tag{24}$$

where

$$AC_j(P) = \frac{1}{D_j(P)} \cdot \sum_r [C_r \cdot \theta_r D_{rj}(P)] \tag{25}$$

is the firm's average (variable) costs at prices P.

[60] This will hold as long as price differences are not too large; specifically, as long as $|P_1 - P_2| < \frac{1}{N \cdot \alpha_r}$.

Marginal Costs Marginal costs are defined as the cost of marginal consumers who switch to plan j if it lowers its price marginally. To derive this, note that if firm j lowers its price by dP, it attracts $\theta_r \cdot dD_{rj} = \theta_r \alpha_r dP$ additional consumers of type r, and $dD_j = \overline{\alpha} \cdot dP$ total additional consumers across all types. The average cost across all marginal consumers is:

$$MC_j\left(P\right) \;=\; \frac{1}{\overline{\alpha}\cdot dP} \times \sum_r \left(\theta_r \alpha_r dP\right) \cdot C_r \;=\; \frac{1}{\overline{\alpha}} \sum_r \left(\theta_r \alpha_r\right) \cdot C_r \tag{26}$$

which replicates the formula for marginal cost in the text. Notice that marginal cost is a constant in the Salop model (across all firms and prices), as long as demand for all firms is positive for all types, as is true in any symmetric equilibrium.

Selection Wedge (around a symmetric equilibrium) Note that at the special case of a symmetric pricing equilibrium ($P_1 = ... = P_N = P^*$), firms split demand ($D_j\left(P^*\right) = D_{rj}\left(P^*\right) = \frac{1}{N}$), and average costs for each firm simply equals the overall population average, $AC_j\left(P^*\right) = \overline{C} = \sum_r \theta_r C_r$. The "selection wedge" between average and marginal cost is

$$AC_j\left(P\right) - MC_j\left(P\right) \;=\; \sum_r \left[\theta_r - \left(\tfrac{\theta_r \alpha_r}{\overline{\alpha}}\right)\right] \cdot C_r \;=\; -\tfrac{1}{\overline{\alpha}} \cdot Cov\left(C_r,\, \alpha_r\right) \tag{27}$$

The first equality follows from the equations for AC_j and MC_j around a symmetric equilibrium, derived above, and the second equality follows from the covariance formula (weighted by type shares θ_r):

$$Cov\left(C_r,\, \alpha_r\right) \;= E\left[\left(\alpha_r - \overline{\alpha}\right)\left(C_r - \overline{C}\right)\right] \;=\; \sum_r \theta_r \cdot \left(\alpha_r - \overline{\alpha}\right) \cdot \left(C_r - \overline{C}\right) \;= \sum_r \theta_r \cdot \left(\alpha_r - \overline{\alpha}\right) \cdot C_r$$

where the final equality follows from the fact that $\sum_r \theta_r \left(\alpha_r - \overline{\alpha}\right) \overline{C} = 0$.

Selection Wedge-Average Cost Slope Relationship We now derive the relationship between $\frac{\partial AC_j}{\partial P_j}$ and the selection wedge, $AC_j\left(P\right) - MC_j\left(P\right)$, per equations (5) and (11) in the body text. This is easiest to derive generally. Define total firm costs as $TC_j\left(P\right) = AC_j\left(P\right) \cdot D_j\left(P\right)$, implying that

$$AC_j\left(P\right) = \frac{TC_j\left(P\right)}{D_j\left(P\right)}$$

. Also note that since marginal cost is the cost of marginal consumers when demand increases, $MC_j(P) = \frac{\partial TC_j/\partial P_j}{\partial D_j/\partial P_j}$ by definition. Now, differentiating $AC_j(P)$ using the quotient rule, we get:

$$\begin{aligned}\frac{\partial AC_j}{\partial P_j} &= \frac{(\partial TC_j/\partial P_j)\cdot D_j - (\partial D_j/\partial P_j)\cdot TC_j}{D_j^2} \\ &= \tfrac{1}{D_j}\tfrac{\partial D_j}{\partial P_j} \times \left[\frac{\partial TC_j/\partial P_j}{\partial D_j/\partial P_j} - AC_j\right] \\ &= -\eta_{j,P_j}(P_j) \times [MC_j - AC_j] \\ &= \eta_{j,P_j}(P_j) \times [AC_j - MC_j]\end{aligned}$$

where $\eta_{j,P_j}(P_j) \equiv -\frac{1}{D_j}\frac{\partial D_j}{\partial P_j}$ is the semi-elasticity of demand. This, therefore, derives equations (5) and (11) in the text.

Equilibrium Pricing The first-order condition for profit-maximization prices is $\frac{\partial \pi_j}{\partial P_j} = 0$. Using the formula for profits in (24), we have:

$$\frac{\partial \pi_j}{\partial P_j} = \sum_r \theta_r D_{rj}(P) + \sum_r [P_j - C_r]\cdot\theta_r \tfrac{\partial D_{rj}}{\partial P_j} = 0$$

Using the fact that $\frac{\partial D_{rj}}{\partial P_j} = -\alpha_r$ (linear demand) and that $\sum_r \theta_r D_{rj}(P) = D_j(P) = \frac{1}{N}$ in symmetric equilibrium, and that $\sum_r \theta_r\alpha_r = \overline{\alpha}$, this simplifies to:

$$\tfrac{1}{N} - P_j\overline{\alpha} + \sum_r \theta_r\alpha_r C_r = 0$$

or

$$P_j^* = \underbrace{\frac{1}{\overline{\alpha}}\sum_r \theta_r\alpha_r C_r}_{\text{Marginal Cost}=MC_j} + \underbrace{\frac{1/N}{\overline{\alpha}}}_{\text{Lerner markup}} \tag{28}$$

which replicates the first part of (7) in the body text. The second half, then, follows from adding and subtracting $\overline{C}$:

$$P_j^* = \overline{C} + \left[\tfrac{1/N}{\overline{\alpha}} - \left(\overline{C} - MC_j\right)\right] \tag{29}$$

Firm Entry Limit In equilibrium, all firms must earn sufficient variable profits to cover their fixed costs, or $\pi_j(P^*) = [P^* - AC_j(P^*)]\,D_j(P^*) \geq F$. Rearranging, this implies the "entry profitability" equation (8) in the body text:

$$\underbrace{P_j^* - AC_j(P^*)}_{\text{Profit margin}} = \underbrace{\frac{1}{\eta_{j,P_j}}}_{\text{Lerner Markup}} - \underbrace{[AC_j(P^*) - MC_j(P^*)]}_{\text{Selection wedge}} \geq \underbrace{\frac{F}{D_j(P^*)}}_{\text{Fixed costs per consumer}} \tag{30}$$

We then plug in the equations for the selection wedge, Lerner markups, and symmetric demand for the Salop model to yield:

$$\underbrace{\frac{1/N}{\overline{\alpha}}}_{\text{Lerner markup}} \quad - \underbrace{\left[-\frac{1}{\overline{\alpha}} \cdot Cov\left(C_r,\, \alpha_r\right)\right]}_{\text{Selection wedge}} \geq \underbrace{F \cdot N}_{\text{Fixed costs per consumer}}$$

When $F = 0$, this can be explicitly solved to yield the following upper bound for N:

$$N \;\leq\; \frac{1}{\overline{\alpha}} \cdot \frac{1}{-Cov\left(C_r,\, \alpha_r\right)}$$

which replicates equation (9) in the body text. Note that the direction of this inequality requires that $Cov\left(C_r,\, \alpha_r\right) < 0$, i.e. adverse selection. More generally, with $F > 0$, we can derive the following (implicit) inequality for participation:

$$N \leq \frac{1}{-Cov\left(C_r, \alpha_r\right) \;+\; \overline{\alpha} \cdot F \cdot N}$$

which replicates equation (10) in the text.

A.2 Proof of Proposition 1

To prove Proposition 1, we start from the entry profitability condition (8) in the body text:

$$\underbrace{P_j^* - AC_j\left(P^*\right)}_{\text{Profit margin}} \;=\; \frac{1}{\underbrace{\eta_{j,P_j}}_{\text{Lerner Markup}}} \;-\; \underbrace{\left[AC_j\left(P^*\right) - MC_j\left(P^*\right)\right]}_{\text{Selection wedge}} \;\geq\; \underbrace{\frac{F}{D_j(P^*)}}_{\text{Fixed costs per consumer}} \tag{31}$$

which must hold for all entrants $j \in E$. With minor rearranging, this condition implies that for all $j \in E$

$$\eta_{j,P_j} \;\times\; \left[\left(AC_j - MC_j\right) \;+\; \tfrac{F}{D_j}\right] \;\leq\; 1$$

Using the fact that $\eta_{j,P_j} = -\frac{\partial D_j / D_j}{\partial P_j}$ and $D_j > 0$, we can equivalently write for all j:

$$-\tfrac{\partial D_j}{\partial P_j} \;\times\; \left[\left(AC_j - MC_j\right) \;+\; \tfrac{F}{D_j}\right] \;\leq\; D_j$$

Summing across firms $j \in E$ and dividing each side by the number of firms (N) yields:

$$\tfrac{1}{N}\sum_{j\in E}\left[\left(-\tfrac{\partial D_j}{\partial P_j}\right) \;\times\; \left\{\left(AC_j - MC_j\right) \;+\; \tfrac{F}{D_j}\right\}\right] \;\leq\; \tfrac{1}{N}\sum_{j\in E} D_j \;=\; \tfrac{1}{N}$$

where the equality uses the fact that $\sum_{j\in E} D_j = 1$ because the D_j's are market shares. Using the definition of a simple average $\mathbb{E}_j\left[X_j\right] = \frac{1}{N}\sum_{j\in E} X_j$, we have:

$$\mathbb{E}_j\left[\left(-\tfrac{\partial D_j}{\partial P_j}\right) \times \left\{\left(AC_j - MC_j\right) \;+\; \tfrac{F}{D_j}\right\}\right] \;\leq\; \tfrac{1}{N}$$

Rearranging this, assuming that the LHS is positive – as it is if there is adverse selection (or not too strong advantageous selection) on average – yields the desired expression in (13):

$$N \;\leq\; \frac{1}{\mathbb{E}_j\left[\left(-\frac{\partial D_j}{\partial P_j}\right)\times\left\{(AC_j - MC_j)\;+\;\frac{F}{D_j}\right\}\right]} \qquad \blacksquare$$

A.3 Additional Figures for Theory Section

Appendix Figure A1. Conceptual Parallel: Fixed Costs and Adverse Selection

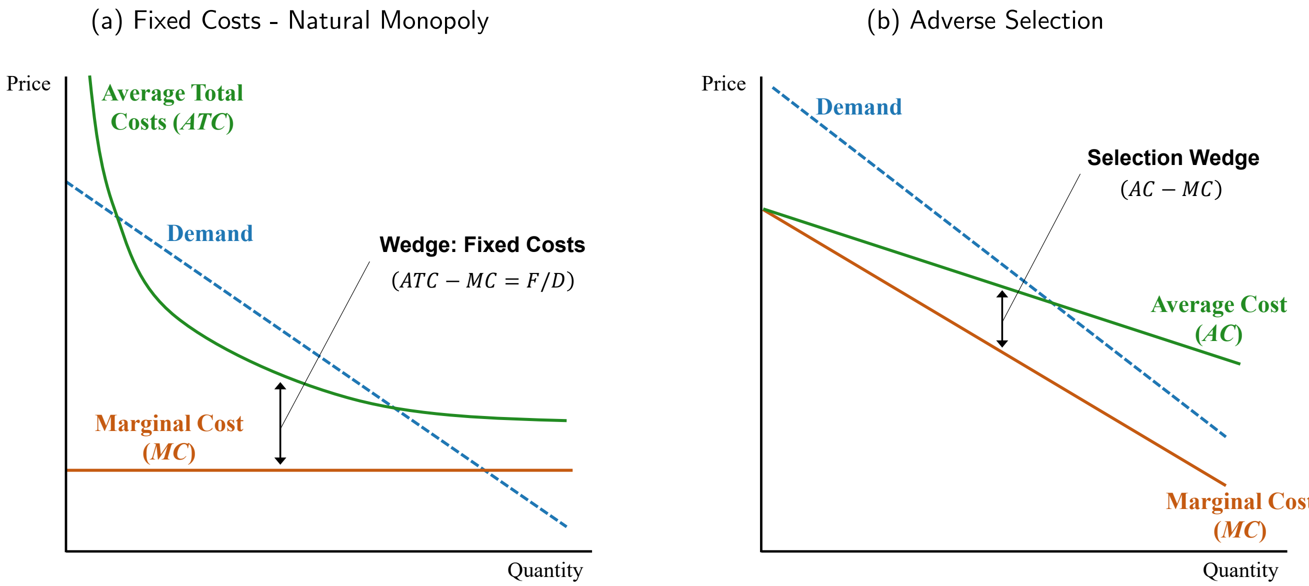

Note: The graphs show the parallel between the economics of adverse selection and fixed costs that lead to natural monopoly. Panel (a) shows a textbook graph of natural monopoly due to high fixed costs. Panel (b) shows a textbook adverse selection market, following Einav and Finkelstein (2011). Both graphs share two key features: (1) the average (total) cost curve is downward-sloping in quantity, and (2) there is a positive "wedge" between average (total) costs and marginal costs.

Appendix Figure A2. Policy Responses to Selection – Impact of Price Floors

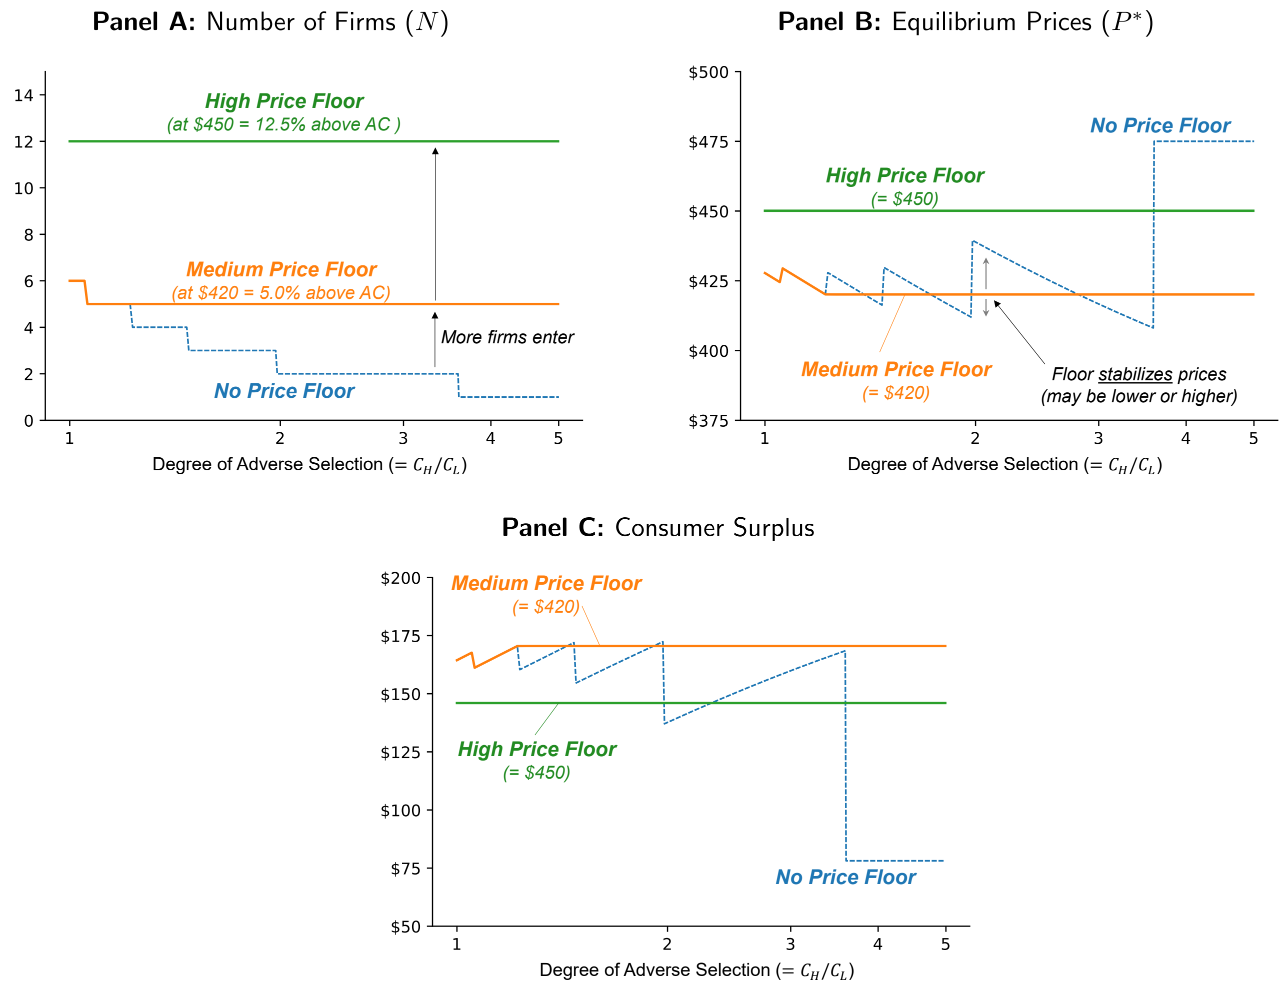

Note: The graphs show results from the calibrated Salop model with adverse selection when we add price floors at $\overline{P}$ = \$420 (medium floor, in orange) and \$450 (high floor, in green), or 5% and 12.5% above average costs ($\overline{C}$ = \$400). Outcomes without floors (which replicate Figure 3) are shown in blue dashed lines. See footnote #11 for calibration details. Price floors increase entry (panel A), stabilize prices (panel B), and (when not set too high) can raise consumer welfare (Panel C).

Appendix Figure A3. Calibrated Salop Model: Varying Fixed Costs

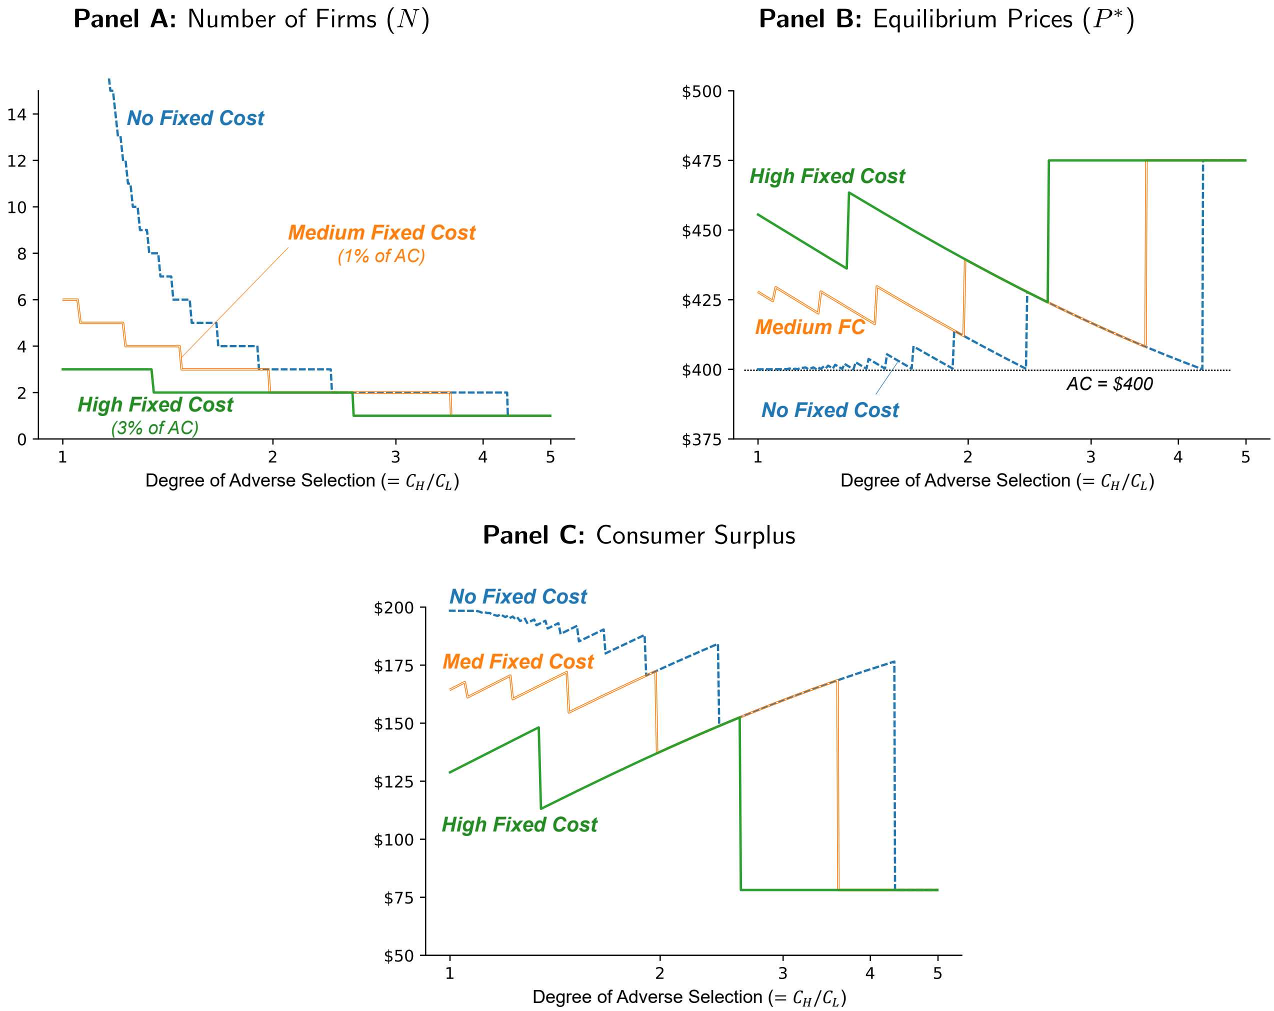

Note: The graphs show results from the calibrated Salop model with adverse selection when we vary fixed costs. See footnote #11 for calibration details. We show results for three levels of fixed costs: (1) No fixed costs ($F = 0$), (2) Medium fixed costs of $F = 0.01 \times AC$, and (3) High fixed costs of $F = 0.03 \times AC$. In our calibration, $AC = \$400$, so the latter two cases are $F = \$4$ and $F = \$12$ per enrollee-month in the market.

Appendix Figure A4. Pricing Best Response Functions

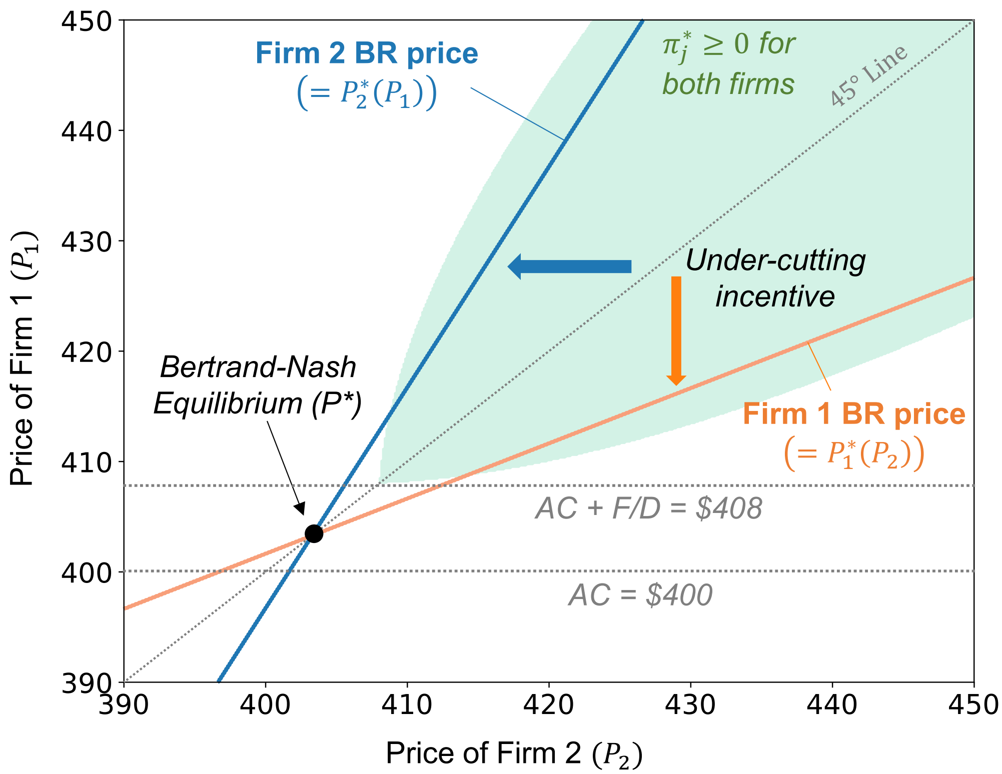

Note: The figure shows pricing best-response functions for a two-firm case of our calibrated Salop model. The orange line is firm 1's best-response price given firm 2's price value (on the x-axis), and the blue line is firm 2's best-response given firm 1's price (on the y-axis). A Nash equilibrium occurs where the two best-response lines intersect. The shaded area is the region where net profits are non-negative for both firms, which is required for firms to be willing to participate. As shown, the only Nash equilibrium occurs at point where prices are too low to cover fixed costs. Therefore, the market cannot support two firms, and the only equilibrium in our model is monopoly.

B Data Construction

In this appendix, we briefly describe the construction of our two primary datasets. We construct separate datasets for our reduced form and structural analyses. The primary differences between the datasets are in the structure of the data, which is tailored to the estimation methods used.

1) Reduced Form Dataset For the reduced form analysis (Section 4), we use the enrollment data to calculate market shares, $MktShare_{j,r,g,t}$, for each plan (j) in each region (r), income group (g), and monthly time period (t) cell. We then use the claims data to calculate average costs per month over the subsequent 12 months for the individuals (in the relevant region-income cell) enrolled in plan j at time t, or $AvgCost_{j,r,g,t} = E[Cost_i \,|\, i \in \{j, r, g, t\}]$.[61]

For some analyses, we also construct risk-adjusted average costs by dividing each enrollee's average costs by their HCC risk score. Our main analyses use all enrollees in the market (new and incumbent), but for some specifications we calculate these variables just for new enrollees (who enter the exchange in month t) or for incumbent ("current") enrollees who joined in a prior year. For our main reduced form analysis, we limit the sample to 2008-2011, a period over which market rules and plan provider networks were relatively stable.[62]

2) Structural Model Dataset For the structural model (Sections 5-6), we construct two datasets for demand and cost estimation. For demand, the data are at the level of consumers' plan choice instances, which occur at new enrollment and annual open enrollment. For each choice instance, we code the available plans, their characteristics (networks and premiums), and the enrollee's chosen option, which we then use to estimate our multinomial logit choice model. For costs, we calculate (insurer-paid) costs per month at the consumer-plan-year level ($C^{obs}_{i,j,t}$), which we use in a panel regression model described in Section 5.2.

[61] We use insurer-paid costs (not including patient cost sharing), since this is the relevant measure for insurer profits. However, total costs are quite similar, since actuarial values exceed 95% in this market. For individuals who exit the market before being enrolled 12 months, we use average monthly cost over the (shorter) period enrolled.

[62] Beginning in 2012, the exchange began limiting choice sets for below-poverty new enrollees, and (partly in response) a large plan (Network Health) narrowed its network by excluding a top-ranked "star" hospital system, a change studied by Shepard (2022). By limiting to 2008-2011, we avoid needing to model these changes. Additionally, we exclude from the sample individuals who joined the market via passive auto-enrollment, a policy studied by Shepard and Wagner (2022).

C Undercutting Case Studies

To illustrate the consequences of undercutting for a plan's market share and average cost, we identify three cases in CommCare where one plan undercuts another on price. The primary case focuses on Network Health and BMC in plan-years 2012 and 2013. In 2013, BMC dropped its bid from just under \$450 to just under \$350 to undercut Network Health, previously the cheapest plan in the market. This can be seen in Panel (a) of Figure A5. In 2013, the after-subsidy price gap between these two plans was just under \$5 per month on average, ranging from \$3 for the second-lowest-income group (the lowest-income group always paid \$0 for both plans) to \$8 for the highest-income group. Despite this small difference in 2013 prices between Network Health and BMC, Network Health's (the under-*cut* plan) market share plummeted in 2013, dropping from around 50% to around 30% (Panel (b)). BMC (the under-*cutting* plan), on the other hand, saw its market share spike from around 20% to around 60%. These shifts in market share correspond with price semi-elasticities of -0.124 and -0.03, respectively.

Panel (c) of Figure A5 shows how the average cost of the BMC and Network Health enrollees shifted around the time of the price change. BMC saw an enormous drop in average cost from around \$450 to around \$300 at the time of the price change. Given that the drop in BMC's price was around \$100, this shift in average cost makes it clear why BMC would want to undercut Network Health in this way — With a \$100 drop in price, BMC simultaneously increased its market share by 300% and *increased* its average profit margin. This is exactly the type of adverse selection undercutting incentive implied by our model in Section 2. Network Health's average cost likewise increased, though only by around \$50. While this change in average cost is much smaller than BMC's change, Network Health's (relative) price only moved slightly between 2012 and 2013, increasing by only around \$5. This shift in average cost thus also implies a steep own-firm average cost curve for Network Health. It also implies that *both* plans were adversely selected on price.

The additional two case studies show similar findings. Appendix Figure A6 shows a case where Network Health was priced around \$18 per month above Celticare in 2011 but cut its price to tie Celticare in 2012. In response, Network Health saw a large increase in market share (around 20 percentage points, or more than 60%) and an enormous decrease in average cost (around \$100, or around 30%).

Appendix Figure A7 shows a case where Celticare undercut Network Health in the prior year, going from effectively being tied with Network Health in 2010 to being priced around \$16 below Network Health in 2011. Relative shifts in market share were again substantial, while shifts in average cost were noisy but with point estimates still suggesting very steep own-firm average cost curves.

Overall, these case studies illustrate the strong under-cutting incentives we discuss in Section 2. In all cases, the undercutting plan simultaneously increased its market share *and* increased its profit margin per enrollee by dropping its price below the price of its competitor. The under-cut plans simultaneously saw decreases in market share and decreases in their profit margins.

Appendix Figure A5. Case Study: BMC and Network Health in 2012-2013

(a) Plan Bids and Relative Premiums

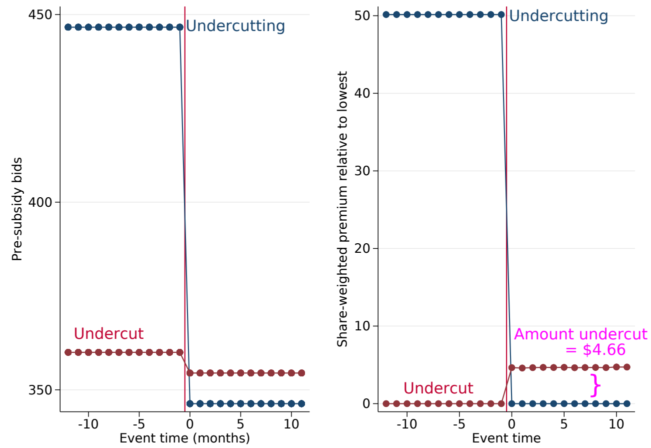

(b) Market Shares

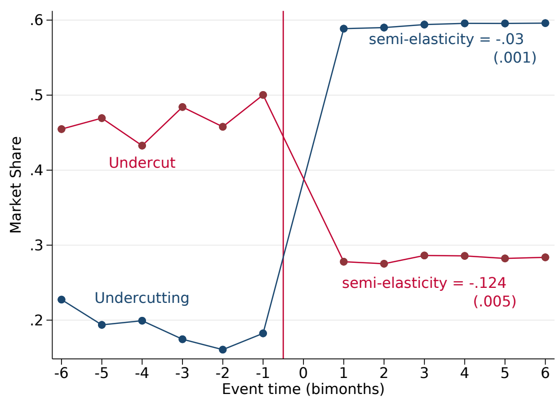

(c) Average 6-Month Cost

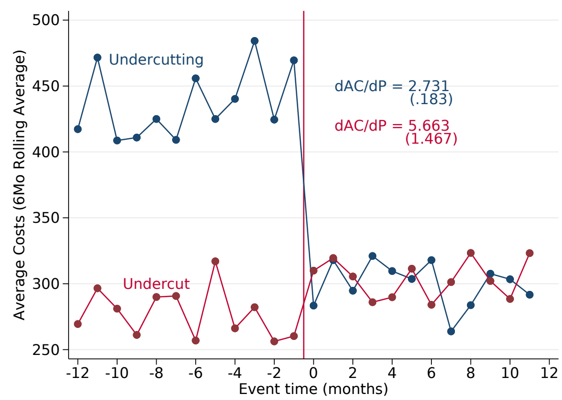

Note: Figure shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over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

Appendix Figure A6. Case Study: CeltiCare and Network Health in 2011-2012

(a) Plan Bids and Relative Premiums

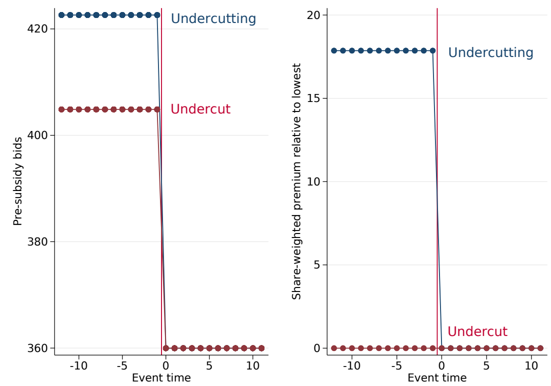

(b) Market Shares

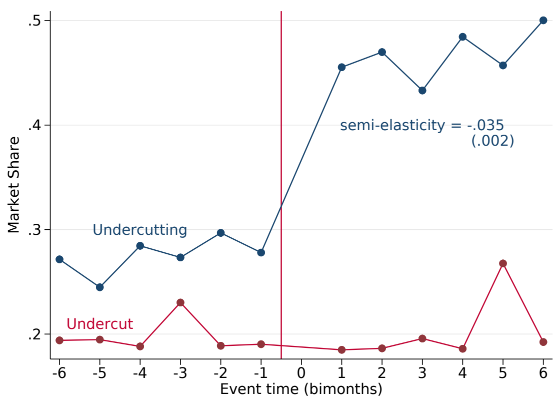

(c) Average 6-Month Cost

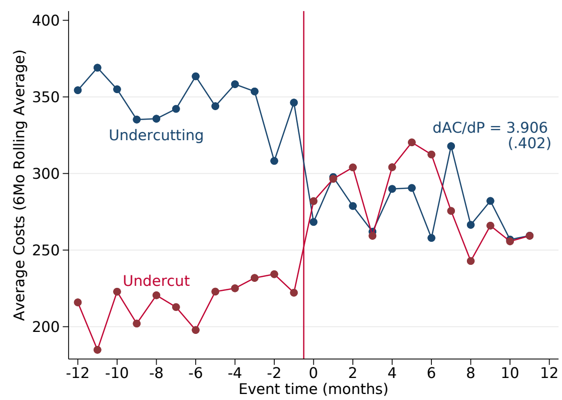

Note: Figure shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over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

Appendix Figure A7. Case Study: CeltiCare and Network Health in 2010-2011

(a) Plan Bids and Relative Premiums

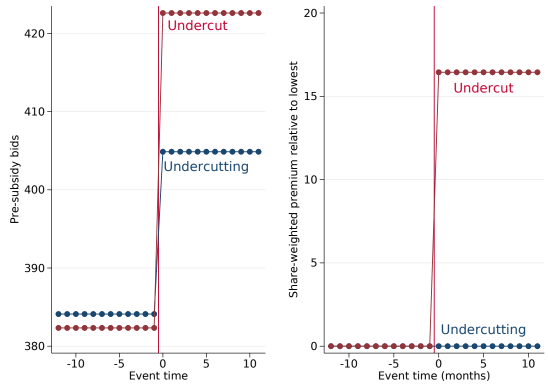

(b) Market Shares

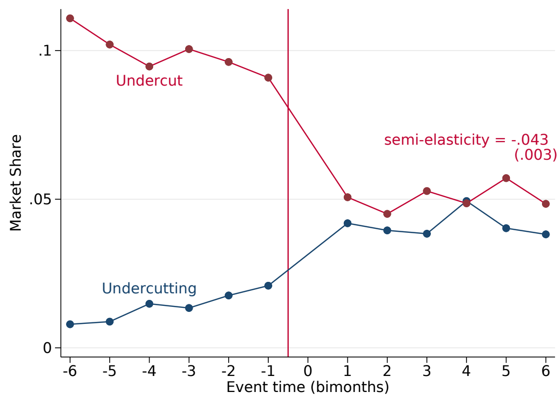

(c) Average 6-Month Cost

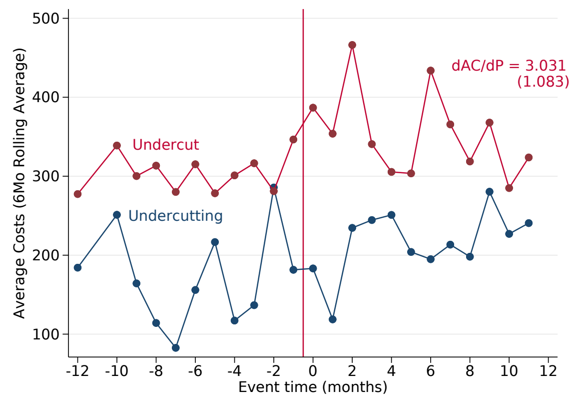

Note: shows how plans strategically respond to each other's bids/premiums and how market shares and average costs evolve as a result. Panel (a) shows the change in plan bids and relative premiums, Panel (b) shows the evolution of shares, and Panel (c) shows the evolution of average costs over time. In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

D Additional Reduced Form Results

In this appendix, we provide additional results for the reduced form section of the paper, Section 4. That analysis examines the causal effect of plan premium changes on demand and average cost (selection) outcomes. Appendix Table A1 replicates the main difference-in-difference estimates (in body-text Table 1) using risk-adjusted outcomes. Figure A8 shows pooled event study plots (pooling premium increases and decreases) for new enrollees only (whereas results in Figure 5 were for all enrollees).

Appendix Table A1. Difference-in-Differences Estimates, Risk Adjusted

	Baseline	**By Enrollee Type**		**By Enrollee Risk**		
	All Enrollees	New Enrollees	Current Enrollees	Low Risk *(0-25%)*	Mid Risk *(25-75%)*	High Risk *(75-100%)*
	(1)	(2)	(3)	(4)	(5)	(6)
Panel (a): Regression results						
Premium	17.87***	17.90***	17.08***	17.25***	18.09***	17.79***
	(1.45)	(1.56)	(1.43)	(1.47)	(1.44)	(1.51)
Log Market Share (risk wgt.)	-0.137***	-0.376***	-0.049***	-0.248***	-0.169***	-0.118***
	(0.017)	(0.038)	(0.014)	(0.025)	(0.018)	(0.017)
Average Costs (risk-adjusted)	6.185**	13.60***	8.345*	---	---	---
	(2.28)	(3.49)	(3.33)			
Panel (b): Theory-Relevant Statistics						
Demand Semi-Elasticity	-0.0077	-0.0210	-0.0029	-0.0144	-0.0093	-0.0066
Slope of Avg Costs (risk adj.)	0.35	0.76	0.49	---	---	---
Adverse Selection Wedge	**$45.1**	**$36.2**	**$168.9**	---	---	---
[% of Avg Cost]	**[12%]**	**[9%]**	**[44%]**			
Num. Observations	5,888	4,922	5,750	5,359	5,819	5,612
Risk Adj Average Cost ($/month)	$383	$394	$385	$131	$239	$897

*Standard errors reported in parentheses. *** p < 0.01, ** p < 0.05, and * p < 0.10.*

Note: Table shows a risk-adjusted version of our difference-in-difference analysis shown in Table 1 in the body text. The 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 regression estimates (with each row corresponding to a different outcome variable). Panel (b) presents implied theory-relevant statistics. 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.

Appendix Figure A8. Pooled Event Study Estimates, New Enrollees Only

(a) Log Market Shares (b) Average Cost ($/month)

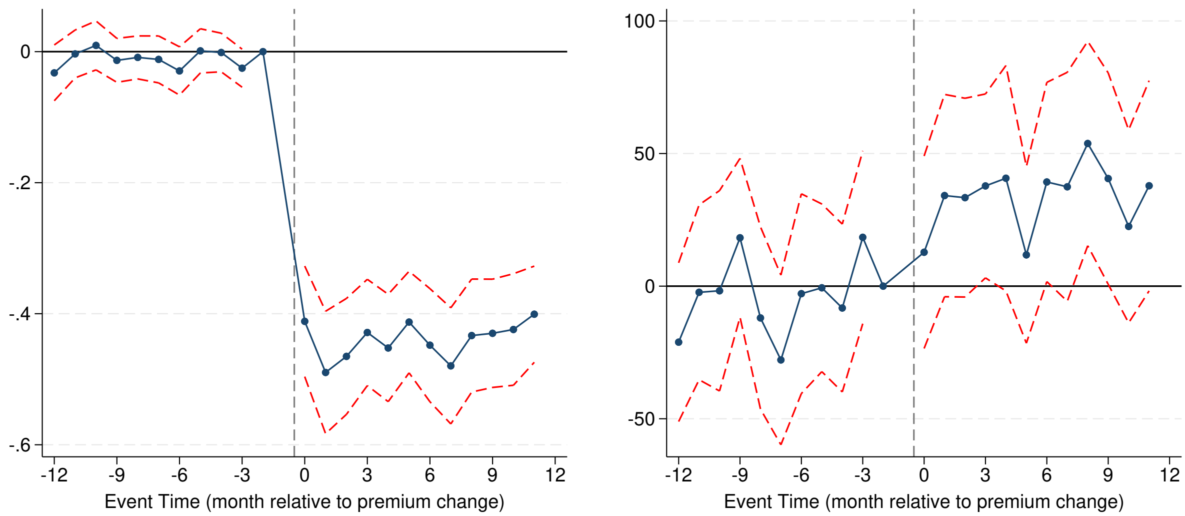

Note: Figure shows pooled event study estimates of the impact of premium increases and decreases. These correspond to Figure 5(c)-(d) in the body text, but with the sample limited to new enrollees only. Panel (a) shows results for log plan market shares, and Panel (b) shows results for average costs (in $ per month).

E Structural Model Details

In this appendix, we provide details of the structural model. The demand model is quite similar to the model estimated in Shepard (2022), with some modifications to incorporate heterogeneity in price sensitivity that is correlated with observed and unobserved consumer risk. The plan cost effects model is quite similar to that of Jaffe and Shepard (2020). Parts of the text in this appendix is taken directly from those papers to allow readers to more conveniently find model details.

E.1 Plan Choice Model Details and Estimates

The plan choice model is described in Section 5.1. This appendix describes additional model details. Table A2 below shows plan choice model estimates.

The model is a standard multinomial logit choice model that allows for preference heterogeneity across consumers based on observables. The choice utility specification, extending the specification reported in equation (16) in the body text, is:

$$u_{ijt} = \underbrace{-\alpha(Z_{it}) \cdot P^{cons}_{ijt}}_{\text{Subsidized Premium}} + \underbrace{V\left(N_{j,t}; Z_{i,t}, \beta\right)}_{\text{Provider Network}} + \underbrace{\delta\left(Z_{i,t}\right) \cdot 1\{CurrPlan_{i,j,t}\}}_{\text{Inertia (current enrollees)}} + \underbrace{\xi_j(W_{it})}_{\text{Plan FE}} + \varepsilon_{ijt} \tag{32}$$

where $P^{cons}_{i,j,t}$ is the enrollee's subsidized premium, $V\left(N_{j,t}; Z_{i,t}, \beta\right)$ is plan provider network attributes, $\delta\left(Z_{i,t}\right)$ captures the impact of inertia in one's current plan (for "current" enrollees who have the opportunity to switch plans at annual open enrollment), $\xi_j(W_{it})$ are plan fixed effects capturing unobserved differences, and $\varepsilon_{i,j,t}$ is the type 1 extreme value error that gives shares their logit form. Coefficients on these plan characteristics are allowed to vary with consumer observables, Z_{it}. Here are some additional details about each of these terms and the consumer observables their coefficients can vary with:

1. Subsidized Premiums Post-subsidy consumer premiums are observed and included directly in the choice model. As in Shepard (2022), premium coefficients, $\alpha\left(Z_{it}\right)$, are allowed to vary with: (1) income groups (100-150%, 150-200%, 200-250%, and 250-300% of poverty), (2) quantile of the HCC risk score (quintiles, plus an extra group for the highest 5% risk enrollees), (3) dummies for having any chronic illness and for cancer, (4) age-sex groups, and (5) immigrant status. Additionally, we allow price coefficients to vary with one other factor: (6) deciles of consumer-level "unobserved" risk, captured by a consumer-level residual from our plan cost model (after accounting for demographics, observed risk, and plan effects). See Appendix Section E.2 below for the cost model and definition of this residual.

As discussed in Section 3, premiums vary not just across plans and years (j, t) but also *across consumers* for a given plan-year. Insurers (who each offer a single plan) are limited to setting pre-subsidy premiums at either the plan-year-region level (from 2007-2010) or at the plan-year level (from 2011-2013). Thus, pre-subsidy premiums vary only at the plan-region-year level. The exchange applies a subsidy schedule that varies across income groups and that also affects prices differences across plans. Subsidies are set so that the lowest-price plan always costs a targeted "affordable" amount by income

— e.g., in 2009-2012 this amount is $0 per month for enrollees with incomes below 150% of poverty, $39 for 150-200% of poverty, $77 for 200-250% of poverty, and $116 for 250-300% of poverty.

Subsidies for higher-price plans follow a schedule that also varies across income groups and leads to variation in *premium differences* for the same plans across incomes. For enrollees in the 0-100% of poverty group, *all plans* are subsidized to be $0 — i.e., there are no premium differences. For enrollees in the 100-300% of poverty groups, higher-price plans cost more than the cheapest plan, but the gap between plans is adjusted in a "progressive" way so that premium gaps are smaller for lower-income groups and larger for higher-income groups. This subsidy schedule implies that the same (pre-subsidy) price changes generate *differential changes* across income groups in post-subsidy premiums. An example of these differential premium changes, for Network Health in the Boston region, is given in Appendix Figure A20.

2. Network Valuation Networks are observed and modeled using two sets of variables, which are identical to those used in Shepard (2022). The first variable is the "network utility" measure from a hospital demand model, which is a multinomial logit choice model estimated using CommCare enrollee's hospital choices observed in the insurance claims data (see Shepard's paper for a full specification). The second are variables for whether the plan covers the hospitals with which the consumer has past outpatient relationships (or the share covered if there are multiple). These variables are all observed and vary across consumers and years, so identification comes from the relationship between this panel variation and consumer plan choices.

Coefficients on network utility are allowed to vary by: (1) income groups, (2) HCC risk score quantiles, and (3) dummies for having any chronic illness and for cancer. We do not vary coefficients with age-sex groups because the illness probabilities used to define network utility already vary by age-sex groups. Coefficients on coverage of hospitals with which a consumer has relationships are allowed to vary with these same three sets of characteristics, and we also further interact these coefficients with whether the hospital is a Partners hospital to allow for special loyalty to the star hospitals (a key issue in the analysis of Shepard (2022)).

3. Inertia (current enrollees) To capture inertia, which is well known to affect health insurance choices, we include a dummy for current enrollees' current plan. Coefficients, $\delta(Z_{it})$, are allowed to vary with the same observables as premium coefficients, with the exception of unobserved risk: (1) income groups, (2) HCC risk score quantiles, (3) chronic illness and cancer dummies, (4) age-sex groups, and (5) immigrant status. Inertia is separately identified from persistent preference heterogeneity by comparing the choices of new enrollees (for whom inertia is zero) and current enrollees, as in prior work (e.g., Handel (2013)).[63] Shepard (2022) shows that price/network valuation parameter estimates are quite similar if the plan choice model is estimated entirely on new enrollees (see Appendix F.3 of

[63]This method does a good job accounting for persistent preference heterogeneity captured by the (many) observed consumer factors in our data, including income, age, location, and health conditions. However, it may miss some persistent *unobserved* preference heterogeneity. See Pakes et al. (2021) for new methodological work seeking to distinguish state dependence vs. unobserved heterogeneity in the health insurance choice context.

his paper). Because our main structural simulations do not rely on inertia estimates — we simulate market outcomes with a consumer population of only new enrollees — we do not do further work to separate true inertia from persistent unobserved heterogeneity.

4. Plan dummy variables We include a large number of plan dummy variables and interactions to capture unobserved plan demand differences (e.g., due to insurer reputation) and to ensure proper identification of the premium coefficient. For each plan, we include separate dummies at the region-income group and region-year level, as well as interactions with age-sex groups and risk score deciles to allow unobserved quality to vary with medical risk. The CommCare program includes five regions (Boston, Central MA, Northern MA, Southern MA, and Western MA) and five income groups at which prices vary (0-100%, 150-200%, 200-250%, and 250-300% of poverty). After omitting empty cells where a plan is not available, there are 251 plan dummy variables/interactions in total.

Discussion of Identification

The specification of plan dummies is intended to aid in identifying the premium coefficients using only *within-plan* variation across income groups due to subsidies — just as in our reduced form difference-in-differences analysis in Section 4. Specifically, the plan-region-year dummies soak up any demand variation correlated with insurer pricing, which occurs at the plan-region-year level (or plan-year level from 2011-forward). The plan-region-income group dummies soak up any persistent plan preference differences across income groups within a region. The only remaining variation in premiums not soaked up by these dummies are *within-plan differences in premium changes* across income groups.

Appendix Figure A9 illustrates this identification by comparing changes in monthly market shares for price-paying (100-300% poverty) vs. zero-price (0-100% poverty) new enrollees when a plan changes its premium downward or upward. (This is exactly analogous to the event study regression estimates shown in Section 4 and Figure 5.) The graphs shows that there are parallel (and flat) pre-trends in market share prior to the premium change (at $t = 0$), and that the zero-price group shows no changes in market shares, consistent with zero or minimal changes in unobserved quality. For price-paying enrollees, market shares jump at $t = 0$ for price decreases (panel A) and fall for price increases (panel B). Premium coefficients in the demand model ($\alpha(Z_{it})$) are identified from these differential demand responses for price-paying vs. zero-price enrollees, separately for each set of consumer observables, Z_{it}.

Appendix Figure A9. Premium Coefficient Identification: New Enrollee Market Shares around Plan Price Changes

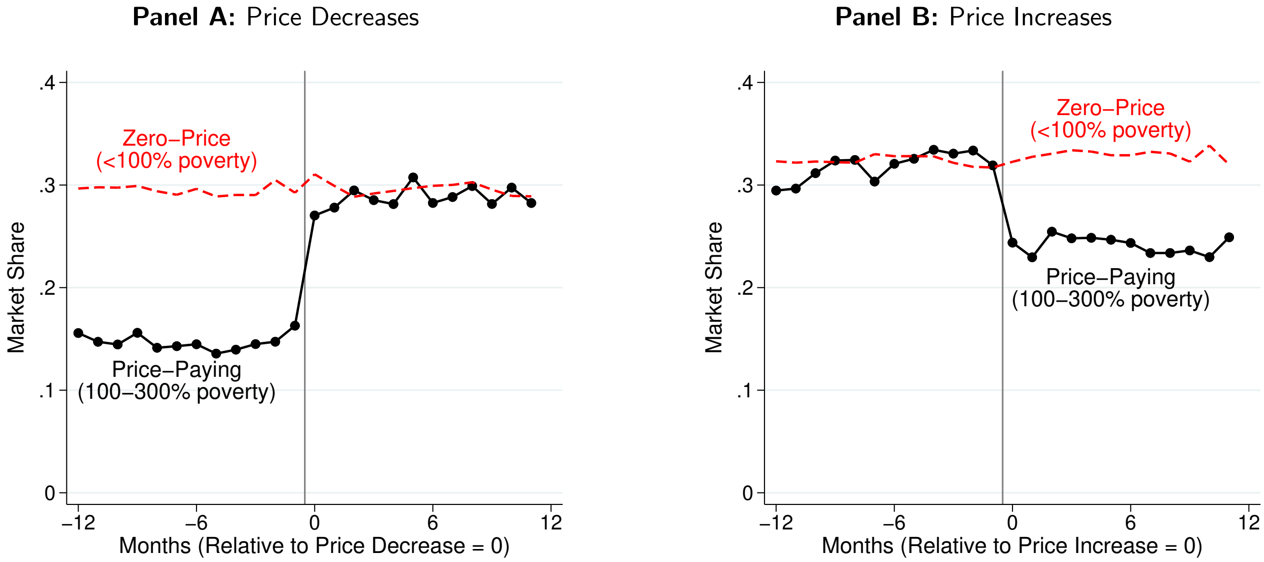

Note: These graphs show the source of identification for the premium coefficients in plan demand and test the key parallel trends assumption for the difference-in-differences approach. It replicates a figure in the appendix of Shepard (2022). Each graph shows average monthly plan market shares among new enrollees for plans that at time 0 decreased their prices (panel A) or increased their prices (panel B). Each point represents an average market share for an independent set of new enrollees. The identification comes from comparing demand changes for above-poverty price-paying enrollees (for whom premium changes at time 0) versus below-poverty zero-price enrollees (for whom premiums are unchanged at \$0). Consistent with the parallel trends assumption, trends in shares are flat and parallel for both groups at times other than the premium change but change sharply for price-payers only at the price change. The sample is limited to fiscal years 2008-2011. We drop 2012+ because below-poverty new enrollees became subject to a limited choice policy that required them to choose lower-price plans. In the demand estimates, we keep this sample but limit the choice set for this group accordingly.

Appendix Table A2. Insurance Plan Choice Model Estimates

Variable	Coeff.		Std. Error
Enrollee Premium (per $10/month): Avg. Coeff.	**-0.494**	***	**(0.005)**
Base Coeffs by Income: 100-150% poverty	**-0.807**	***	**(0.014)**
150-200% povery	**-0.525**	***	**(0.011)**
200-250% poverty	**-0.457**	***	**(0.010)**
250-300% poverty	**-0.409**	***	**(0.011)**
x High Cost Residual (>80th pctile)	0.085	***	(0.006)
x High Risk Score (>80th pctile)	0.072	***	(0.012)
x Any Chronic Illness	-0.009		(0.005)
x Cancer	0.029	***	(0.008)
x Immigrant	-0.095	***	(0.019)
x Age ≥45 years	0.073	***	(0.015)
Provider Network			
Network Utility (avg. coeff.)	**0.495**	***	**(0.006)**
x Income >100% poverty.	0.020		(0.011)
x High Risk Score (>80th pctile)	-0.236	***	(0.017)
x Any Chronic Illness	0.174	***	(0.006)
x Cancer	0.044	***	(0.012)
Share Prev Used Hosp. Covered (avg. coeff.)	**0.220**	***	**(0.017)**
x Income >100% poverty.	-0.244	***	(0.034)
x High Risk Score (>80th pctile)	0.186	***	(0.052)
x Any Chronic Illness	0.121	***	(0.033)
x Cancer	0.115		(0.060)
x Prev. Used Partners Hospitals	0.359	***	(0.032)
Inertia: Current Plan Dummy (avg. coeff.)	**4.631**	***	**(0.012)**
x Income >100% poverty.	-1.322	***	(0.025)
x High Risk Score (>80th pctile)	-0.113	*	(0.051)
x Any Chronic Illness	-0.055	**	(0.020)
x Cancer	0.012		(0.032)
x Immigrant	-0.238	***	(0.068)
Avg. Plan Dummies: BMC	*(normalized = 0)*		
CeltiCare	-0.796	***	(0.046)
Fallon	0.138	***	(0.036)
Neighborhood Health Plan (NHP)	-0.145	***	(0.016)
Network Health	0.067	***	(0.014)
Model Stats: Pseudo-R^2	0.536		
No. Choice Instances	863,052		
No. Unique Enrollees	383,891		

* $P < 5\%$, ** $P < 1\%$, *** $P < 0.1\%$

Note: This table shows main coefficient estimates for the multinomial logit plan choice model described in the appendix text. The table shows logit model estimates for key plan attributes, and their interactions with enrollee characteristics. Premiums are the amount paid by consumers after subsidies, in $10 per month; this varies by about $20-60 across plans. Network utility is the consumer-specific expected utility measure for a plan's hospital network, based on an estimated hospital demand model. Share previously used hospitals covered is the share of an enrollee's previously used hospitals that a plan covers. For most covariates, we report the average coefficient across all enrollees, as well as key interactions terms with consumer observables. Additional interaction terms are described in the appendix text.

E.2 Plan Cost Model

Our model of plan costs uses the individual-level cost data (from insurance claims), along with estimates of plan cost effects to predict the counterfactual spending of enrollee-year (i, t) in each possible plan j. The estimation approach closely follows that of Jaffe and Shepard (2020) in using enrollees who enroll in *different* plans across separate CommCare enrollment spells to infer plan effects on costs.

As noted in Section 5.2, we assume a model where observed costs for insurer j on enrollee i at time t is the product of an enrollee's potential cost in an average plan (Γ_{it}) times a factor capturing plan effects on costs $(\delta_{j,r})$ — capturing the impact of plan differences due to factors like provider networks — which we allow to vary by region r:

$$C^{obs}_{ijt} = \Gamma_{it} \times \delta_{j,r(i)}. \tag{33}$$

We then proceed in two steps. First, we estimate $\delta_{j,r}$. To do so, we leverage cases where the same individual enrolls in the market in two separate spells in which they choose different plans. As noted in Section 5.2, we estimate the following Poisson regression model with individual fixed effects:

$$E(C^{obs}_{ijt}|Z_{it}) = \exp\left(\alpha_i + \beta_t + Z_{it}\gamma + \lambda_{j,r}\right) \tag{34}$$

This specification controls for individual fixed effects (α_i), year fixed effects (β_t), and time-varying enrollee observables Z_{it} (age-sex bins, a spline in risk score, income group, and enrollee location). The $\lambda_{j,r}$ coefficients represent the plan-specific cost effects, which we allow to vary across regions r to account for differential cost structures based on a plan's regional provider network. The estimated multiplicative plan cost effect of interest is $\hat{\delta}_{j,r} = \exp(\hat{\lambda}_{j,r})$. We normalize the scale of these fixed effects so that $\hat{\delta}_{j,r}$ has an (enrollment-weighted) mean of 1.0 across all plans. The model assumes that plan cost effects are constant over time, though they can vary by region. This is reasonable only if the determinants of costs — in our setting, primarily networks — are stable, which is roughly true over our 2007-2011 sample period. Additionally, we have explored allowing plan effects to vary across years in the data and found similar results. These plan cost effects estimates let us do two things, discussed in turn.

1) Predicted Costs in Counterfactual Plans First, we can use these estimates to predict enrollee i's *counterfactual* costs had they enrolled in plan k as:

$$\hat{C}_{i,k,t} = C^{obs}_{i,j,t} * \left(\frac{\hat{\delta}_{k,r(i)}}{\hat{\delta}_{j,r(i)}}\right) \tag{35}$$

where j is the actual plan for enrollee-year (i, t) (who lives in region $r(i)$) and k is any available plan. For $k = j$, this simplifies to actual observed costs, $C^{obs}_{i,j,t}$. For $k \neq j$, this effectively adjusts their observed costs by the estimated ratio of plan cost effects for plan k vs. j. This implicitly assumes that plan effects take a constant multiplicative form (for each region r). While this is an assumption that rules out forms of "selection on moral hazard" (Einav et al., 2013), this level of detail for plan costs seems rich enough given our focus in this paper on general price competition.

2) Enrollee Unobserved Risk (cost residual) We also use the plan cost model estimates to generate estimates of enrollee-year-level "unobserved" risk, not captured by demographics and medical risk variables in our claims data. To do so, we start with choice-instance-level observations of individuals' monthly medical costs, averaged over each choice instance (i, t). We then adjust for plan-specific cost effects using the estimates of $\hat{\delta}_{j,r(i)}$ from the Poisson regression (34) above. For each individual-year i, t, the adjusted plan cost is equal to:

$$C_{i,t}^{adj} = C_{i,j,t}^{obs} / \hat{\delta}_{j,r(i)} \tag{36}$$

for the chosen plan j and represents the individual's cost if they had been in the average plan (recall that the average plan effect in each region, $\bar{\delta}_r$, is normalized to 1.0).

Next, we estimate a second Poisson regression of $C_{i,t}^{adj}$ on the other covariates in the demand model:

$$E\left(C_{i,t}^{adj} | Z_{i,t}\right) = \exp\left(a_{r(i),t} + Z_{i,t} a_1\right) \tag{37}$$

where $a_{r(i),t}$ is a region-year specific constant, and the covariates in $Z_{i,t}$ are: indicator variables for income group, age-sex groups (in 5-year bins), immigrant status, HCC risk score quantiles (five quintiles, plus a dummy for the top 5%), and dummies for any chronic disease and cancer. The regression is weighted by the number of monthly observations in each choice instance (i, t). We then compute the "cost residual" for each individual as the residual from this regression, computed as:

$$\omega_{i,t}^{Cost} = \frac{C_{i,t}^{adj}}{\hat{C}_{i,t}^{adj}} = \frac{C_{i,t}^{adj}}{\exp\left(\hat{a}_{r(i),t} + Z_{i,t} \hat{a}_1\right)} \tag{38}$$

where $\hat{C}_{i,t}^{adj}$ is the predicted cost for individual i from the Poisson regression. In the demand model, we allow individuals' price sensitivity coefficients to vary by deciles of these cost residuals.

E.3 Verifying Demand Model Fit (Figure 6)

To verify that our demand model can reproduce our reduced form observations, we compare our demand model predictions to the actual data, with results shown in Figure 6. To generate this figure, we start from our demand estimation sample at the choice instance-by-plan level, assigning predicted choice probabilities, plan-specific costs (see Section 5.2), and actual observed plan choices to each observation. We then extend each choice instance up to the month of the subsequent choice instance or when the enrollee exits the data. Next, we collapse the resulting data set to the plan-region-income group-month level, weighting each observation by either the observed plan choice $y_{ij} \in \{0, 1\}$ (resulting in observed shares and observed average costs) or the predicted share s_{ij} derived from the demand model (resulting in predicted shares and predicted average costs). Finally, we run event studies following the reduced form analysis, described previously in Section 4, for the observed vs. predicted shares and costs.

Additionally, we evaluate how the estimated price sensitivities of demand, $\alpha(Z_{it})$ in demand model equation (16), covary with enrollee costs, as discussed in Section 5.3. Appendix Figure A10 shows a binned scatter plot of this relationship. Individual costs are strongly negatively correlated with price-sensitivity coefficients, confirming that adverse selection is strong in this market.

Appendix Figure A10. Adverse Selection Correlation: Lower-Cost Types are More Price-Sensitive

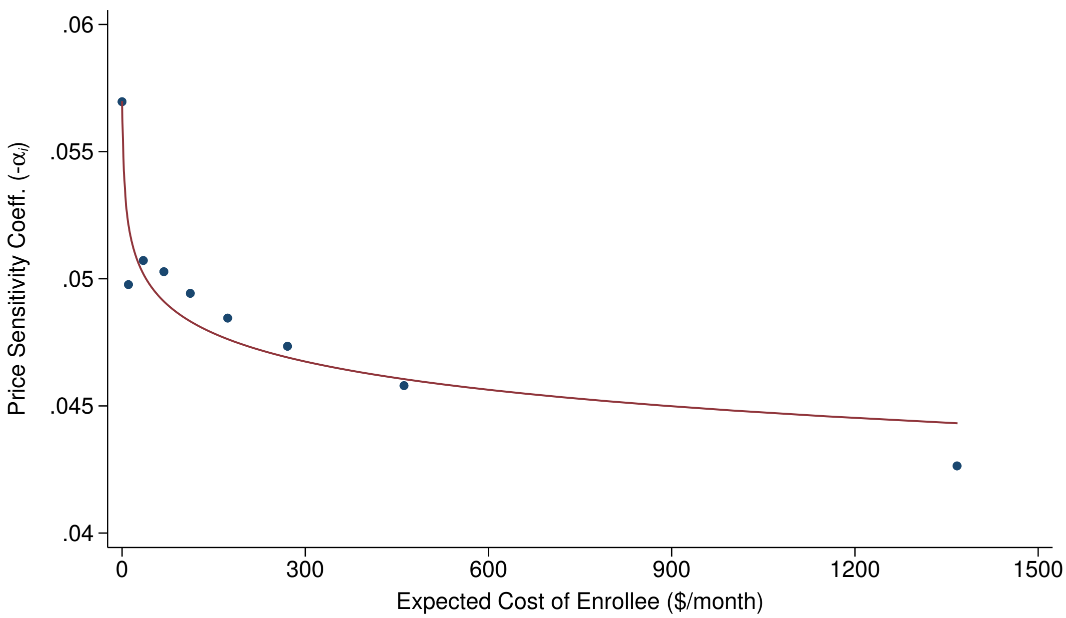

Notes: Figure shows a binned scatter plot of individual-level costs (for a given year t, and adjusted for plan effects) vs. individuals' price-sensitivity of demand coefficient, $\alpha(Z_i) = Z_{i,t}\alpha$, derived from the structural demand estimates of (16). Sample includes all new enrollees in the market from 2007-2014.

F Counterfactual Simulations: Detailed Methods

F.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 allowed 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 (where included) and by our maximum allowed price of $475 that binds only in the monopoly case.

When there are multiple competitors and the degree of adverse selection is high (e.g., without risk adjustment and price floors), we generally find no pure-strategy Nash pricing equilibrium. In such cases, we solve for mixed-strategy pricing equilibria, following the method in Subsection F.2 below. For all combinations of entrants, we are able to identify either a pure- or mixed-strategy pricing equilibrium, although the identified equilibrium does not always yield positive net profits to all firms.

For a given combination of firms, we first attempt to find a pure-strategy pricing equilibrium. To do this, we adopt the following grid search approach.[64] 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 $475. For 4 firms, there are $30^4 = 810{,}000$ possible price vectors. We then identify candidate price vectors that satisfy the following four conditions: (1) First, all firms' prices fall between the price floor (if applicable) and $475 (the maximum allowed price). (2) Second, all firms make positive profits. (3) Third, the FOCs are satisfied within a pre-defined tolerance.[65] (4) Fourth, no firm can deviate to a higher or lower price (within the bounds of the price floor and $475) and make higher profits. For each of the resulting candidate price vectors, we search continuously 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.[66]

To account for cases where firms' optimal prices are at a boundary (i.e., a price floor or ceiling), we systematically iterate through cases that allow subsets of firms to price at either a floor or ceiling. This process identifies pure-strategy pricing equilibria satisfying criterion (1) above.

This procedure delivers up to one candidate equilibrium price vector for each possible combination of firms.[67] Among these possible equilibria, we exclude combinations where an additional firm could enter and "break" the candidate equilibrium. The candidate equilibrium is broken if the additional

[64] Starting with a grid search helps ensure that we do not miss candidate solutions to the FOCs. Due to adverse selection, not all solutions to the FOCs will be global optima for the insurers' maximization problem. Indeed, some solutions to the FOCs are local minima for certain firms.

[65] The tolerance we use is N/J * .20, where N is the total number of individuals and J is the total number of plans. This effectively requires that marginal price changes not affect profits by more than $0.20 per enrollee per participating firm.

[66] For each firm, we check all possible global deviations in increments of $5 to verify that the firm cannot make greater profits by deviating to any other price.

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

firm can make positive profits in the resulting pure- or mixed-strategy pricing equilibrium that results, even if the resulting pricing equilibrium includes a different firm that makes negative profits. Thus, this so-called "breaking equilibrium" can break a candidate equilibrium without being a valid equilibrium itself. Appendix Section F.2 below describes how we find mixed-strategy breaking equilibria.

For example, consider a candidate equilibrium that includes BMC and Network competing in a duopoly. In some cases, Celticare finds it profitable to enter, resulting in a (potentially mixed) pricing equilibrium where either BMC or Network earns negative expected profits. This three-firm equilibrium with Celticare is not a valid equilibrium, since not all firms make positive profits. However, this "breaking equilibrium" still invalidates the original candidate equilibrium with BMC and Network. In these cases, the unique valid equilibrium is typically a Celticare monopoly.

In rare cases, no pure-strategy equilibrium satisfies Criterion (3). When this occurs, we look for a mixed strategy equilibrium (following Appendix Section F.2 below) by iterating through firm combinations where no pure-strategy pricing equilibrium was found.

F.2 Mixed-strategy pricing equilibria

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

Figure A11 below shows an two-firm example where pure-strategy equilibria do not exist. Starting from the top right of the figure, both firms undercut each other, but the low-cost firm (firm 2, Celticare) undercuts to a greater degree. Eventually, firm 4 (NHP) no longer earns positive profits, and there is a discontinuity in its best response curve, where instead of undercutting, firm 4 would rather price at the price ceiling to minimize its losses. This demonstrates how non-existence of pure-strategy equilibrium corresponds to discontinuities in one or both best-response curves.

Appendix Figure A11. Pricing Best-Response Curves for Firms 2 (CeltiCare) and 4 (NHP)

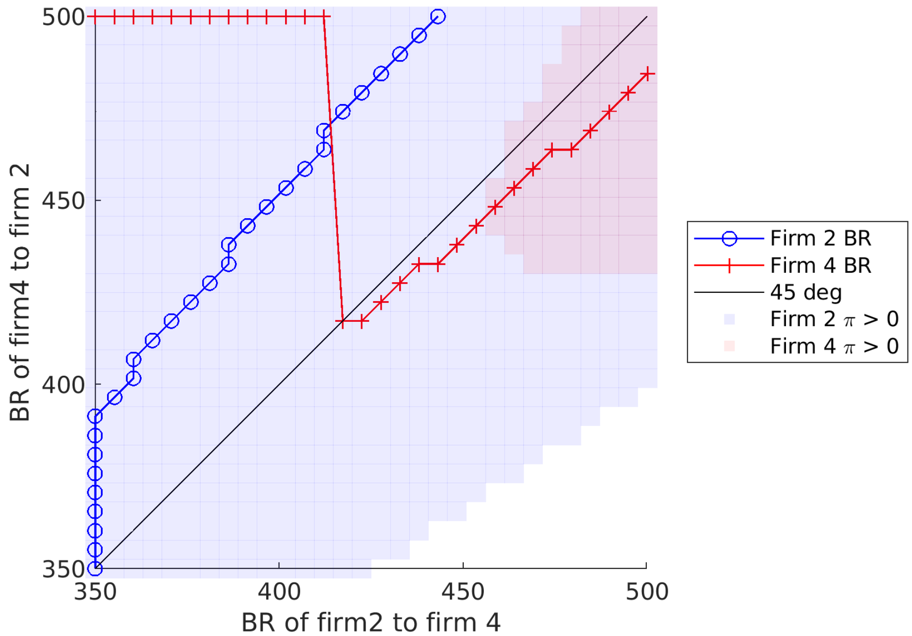

Notes: Figure shows best response curves for Celticare and NHP in the case with no risk adjustment. In both panels, the red curve labeled "Firm 4 BR" shows the optimal pre-subsidy premium of NHP on the Y-axis, given Celticare's pre-subsidy premium on the X-axis. The curve labeled "Firm 2 BR" shows the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

Although these cases do not permit pure-strategy equilibria, mixed-strategy equilibria do exist. In the case shown in Figure A11, the resulting mixed-strategy equilibrium involves a single price for firm 2 (Celticare) and firm 4 (NHP) mixing between a high and low price.

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

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

1. Cases where only one firm (denoted firm i) mixes over two prices p_{iL} and p_{iH}, where p_{iH} is set to the price ceiling of $475. Firm j sets a single price p_j. In this case, p_{iL} and p_j found by solving the corresponding pricing first order conditions.

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

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

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

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

G Analysis of Insurer Participation in the ACA Marketplaces

This appendix provides details for our analysis of insurer participation in the Affordable Care Act (ACA) Marketplaces. We start by describing the ACA datasets (section G.1) and the analysis used to construct Figure 10 showing trends in insurer participation and zero-premium plans (section G.2). The remainder of the appendix (sections G.3-on) describes an analysis of the effects of reinsurance on insurer participation in the ACA marketplaces, which we briefly described in Section 8 of the text.

G.1 ACA Marketplace Datasets

For both the analysis in Section 8 and our analysis below, we draw on publicly released data on the "federally facilitated" ACA Marketplaces whose states use the federal HealthCare.gov platform for enrollment. This sample includes between 32-39 states, depending on the year; the remaining "state-based" marketplaces coordinate their own enrollment and have less consistent data availability. For our analysis in Figure 10, we restrict to the balanced panel of 30 states that are consistently in this data from 2014-24. For the reinsurance analysis in Sections G.3 and following, we use the balanced panel of states in the data during the relevant period for that analysis, 2014-18.

The public ACA data include information on insurer participation, plan offerings and premiums, and aggregated data on enrollment by area (county markets) and income groups. The specific datasets we use are:

- **Qualified Health Plan (QHP) Landscape files:** This includes data on participating insurers and their plans and premiums (by age and family type) in each county-level market in each year. We use this data to measure insurer participation and premiums.[68]

- **Marketplace Open Enrollment Public Use Files:** This includes data on the number and composition (including income levels) of enrollees in each county who are enrolled in Marketplace coverage as of the annual open enrollment period. We use this data to measure total enrollment by county and to estimate income distributions for estimating subsidies and net premiums (in our zero-premium analysis).[69]

- **Issuer-level Enrollment Data:** This includes data on the number of ever enrolled and dis-enrolled individuals for each issuer-plan-year,[70] which we use to measure the number of ever-enrolled individuals for each issuer-year. We combine this with the reinsurance data below to calculate a per-enrollee reinsurance amount for each issuer-year. These datasets also include issuer-county-year level enrollment separately by sex, age group, income group, and tobacco use.[71]

[68]The QHP Landscape files are available at https://www.healthcare.gov/health-and-dental-plan-datasets-for-researchers-and-issuers/.

[69]The Open Enrollment Public Use files are available at https://www.cms.gov/data-research/statistics-trends-reports/marketplace-products/2024-marketplace-open-enrollment-period-public-use-files

[70]Note: an issuer is an insurer-state combination, such that each state has its own unique set of issuers.

[71]The issuer-level enrollment data are available at: https://www.cms.gov/marketplace/resources/data/issuer-level-enrollment-data

- **Reinsurance Data:** This contains reinsurance payment data at the insurer-state-year level for the 2014-16 period while the ACA's reinsurance program was in effect.[72]

G.2 Insurer Participation and Zero-Premium Plan Analysis

For Figure 10 in the text, we show insurer participation and estimated shares of enrollees with at least one zero-premium plan (after subsidies), capturing the relevance of the ACA's indirect price floor at $0. This section describes details of that analysis. As noted above, the sample is the 30 states that are consistently in the federally facilitated marketplace data in all years during the 2014-24 period we analyze.

Insurer Participation Starting from the QHP Landscape files, we calculate the number of unique insurers offering plans in each county market (the level at which participation occurs) in each year. Figure 10(a) reports the average number of participating insurers per county market and the share of counties for which there are just 1-2 insurers (monopoly or duopoly). All averages/shares are weighted by 2010 county population (of individuals under age 65), drawn from 2010 Census data.

Zero-Premium Shares We seek to estimate the share of ACA enrollees for whom at least one zero-premium gold/silver/bronze plan (or one zero-premium silver plan) is available. This calculation is complex because even with a given county-year market, pre-subsidy premiums vary by demographics (notably, age and family size) and subsidies vary by income and are based on the actual premium of the "benchmark" second-cheapest silver plan. To implement our calculation perfectly would require data at the county-year x demographics x income group level, which is unavailable. Therefore, we use a variety of approximations to estimate the zero-premium share, wherever possible making conservative assumptions that may somewhat understate this probability.

We start with data (from the QHP Landscape file) on premiums for the cheapest gold, silver, and bronze-tier plans, as well as the benchmark second-cheapest silver plan.[73] Premiums in the landscape file are only available for certain age and family size bins. For simplicity, we use the premiums for a 40-year-old individual, which is around the median age in the market.[74] We then calculate the premium subsidy available for each 1% of federal poverty level (FPL) bin in the relevant range (100-400% of FPL), using the difference between the benchmark premium and the expected monthly contribution specified by law.[75] People with incomes below 100% of FPL and above 400% of FPL are unsubsidized, so we assume they do not have access to any free plans.

[72] These data are available at https://www.cms.gov/marketplace/health-plans-issuers/premium-stabilization-programs

[73] We exclude Platinum plans, which are often unavailable and are very high-price (so unlikely to be free). We also exclude Catastrophic plans, which are only available to a subset of younger enrollees below age 30.

[74] Because of the way subsidies work, zero-premium prevalence is slightly higher for older groups and slightly lower for younger groups (once we exclude catastrophic plans), so using the median age should generate reasonable estimates.

[75] Individuals are expected to contribute between 0-9.5% of income (depending on their income and the year) for the benchmark silver plan. We thank David Anderson and Coleman Drake for sharing data on the expected monthly contributions schedule across all years.

Using the estimated subsidy, we then calculate the *net-of-subsidy* premium for the cheapest gold, silver, and bronze plan. If the subsidy amount exceeds the pre-subsidy premium, then the zero-premium floor kicks in, and the individual at that income bin has access to a zero-premium plan. We then collapse the data to the county-year level, using a weighted average across income bins using county-year-level data on the income distribution of enrollees (from the OE Public Use files).[76] With these county-year level estimates of the share of enrollees with zero-premium plans, we then calculate a national average, weighting by 2010 non-elderly population in the county.

The results are shown in Figure 10(b). We make several notes about the patterns in the figure. First, zero-premium *silver* plans are rare up to 2020, since the benchmark (second-lowest) silver plan by construction has a premium of at least 2% of income for everyone (and higher for higher-income groups). In these years, the cheapest silver plan is only free if the *gap* between the cheapest and second-cheapest silver plan exceeds the enrollee's benchmark contribution, which is rare. However, starting in 2021, the American Rescue Plan Act (ARPA) increased subsidies so that all enrollees with incomes below 150% of FPL qualified for two zero-premium silver plans. This accounts for the sharp spike in zero-premium silver plans starting in 2021.

Second, the prevalence of *any* zero-premium plans is largely driven by Bronze plans (with a few zero-premium Gold plans in later years). Bronze plans are free if the gap between their premium and the benchmark silver premium exceeds the enrollee's required contribution. This share jumps sharply in 2018 because of the introduction of "silver-loading" policies in which state regulators explicitly pushed insurers to raise silver-plan premiums to make up for lost cost-sharing reduction (CSR) subsidies that had been canceled by the Trump administration. This policy led to a sharp increase in silver plan premiums — and the associated subsidies linked to them — even relative to bronze and gold premiums. As a result, zero-premium bronze plans became extremely prevalent, with our estimates suggesting most consumers had access to at least one starting in 2018. This share further jumped in 2021 because of the expanded subsidies under ARPA. By 2024, over 70% of ACA enrollees have at least one zero-premium plan.

G.3 Overview: Analysis of Expiration of Reinsurance in 2017

Our model predicts that, among horizontally differentiated issuers, areas with steeper average cost curves (larger $\frac{dAC}{dP}$) should have less firm participation, but that risk adjustment and reinsurance can correct for this. Like risk adjustment, reinsurance may also affect $\frac{dAC}{dP}$. This will occur if very high-cost enrollees (who trigger reinsurance payments) are differentially sensitive to price relative to those who do not trigger payments or trigger lower payments. We provide an explicit characterization of the relationship between $\frac{dAC}{dP}$ and reinsurance below, showing that the effects of the removal of reinsurance on firm participation should be related to $\frac{dReins}{dP}$. We thus outline a method for estimating variation in $\frac{dReins}{dP}$ and then leveraging this variation, plus the end of reinsurance in the ACA in 2016, to

[76] This income distribution data is in coarse categories of 50-100% of FPL. For simplicity, we assume a uniform distribution within each category. This likely understates the zero-premium share because the ACA enrollee distribution skews lower-income, and zero-premium prevalence is also higher for lower-income groups.

demonstrate the relationship between $\frac{dAC}{dP}$ and insurer participation.

We find that areas with larger increases in $\frac{dAC}{dP}$ after the removal of reinsurance (i.e., areas where reinsurance was more important for flattening the average cost curve) have larger declines in participation, consistent with the predictions of our model.

G.4 Empirical Approach

G.4.1 Testing model predictions of firm entry: data and identification challenges

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

In terms of data availability, plan premiums are readily observed (QHP Landscape Files), but plan-specific enrollment is not released at the county level, which is the natural level at which to evaluate insurer participation, as insurers are allowed to selectively enter counties, despite only being able to set prices at the rating area level (Geddes, 2024; Ko and Fang, 2023). Similarly, average costs are not observable at the plan level, only at the issuer level.

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

Lastly, we note that quasi-random price variation alone is not sufficient: while this would allow us to obtain unbiased estimates of demand elasticities and average cost curve slopes, cross-market variation in these quantities may still be correlated with unobservable market characteristics. For instance, an association between larger price elasticities and lower insurer entry could reflect confounding due to unobservably lower income (and levels of demand).[77]

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

With these data and identification challenges in mind, we develop a test of our model's predictions regarding the slope of the average cost curve that leverages the nationwide removal of reinsurance

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

occurring after 2017 and pre-existing variation in the importance of reinsurance across states.[78] We describe the rationale of this approach in more detail in the following section.

G.4.2 End of reinsurance induces differential changes in the average cost curve slope

Reinsurance helps to reimburse firms for enrolling individuals with very high costs. Specifically, the ACA reinsurance program ran from 2014 to 2016 and reimbursed a fraction of insurers' costs for high-cost individuals. In 2014 and 2015, payments were triggered when an individual's annual spending passed \$45k. In 2016, payments began at \$90k annual spending. In 2014, reinsurance payments covered 100% of claims between \$45k and \$250k (though it was originally anticipated that they would only cover 80%). In 2015, reinsurance payments covered only 55% of claims between \$45k and \$250k. In 2016, reinsurance payments covered 53% of claims between \$90k and \$250k.[79]

The effect of reinsurance is to flatten the slope of the average cost curve. The degree of flattening can vary across markets. When reinsurance was removed, areas with more pre-period "flattening" experienced larger increases in the slopes of their average cost curves. It is this *heterogeneous increase in average cost curve slopes* that we use for identifying how the slope of the average cost curve affects participation. To show this in mathematical terms, we can follow Section 2 to write firms' *net* average cost per enrollee, now including reinsurance, as:

$$AC_j^{Net} = AC_j(P) - RE_j(P) \tag{39}$$

where $AC_j(P)$ denotes average claims cost per enrollee and and $RE_j(P)$ denotes the reinsurance payment given to firm j per enrollee. We can further break down the reinsurance term as:

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

where c is the share of costs reimbursed each year, $I(m)$ represents the set of enrollees in market m, $y_{ij}(P)$ denotes whether enrollee i chose plan j, and C_{ij} is the total annual cost of the enrollee to the insurer. Importantly, we can calculate $RE_j(P)$ at the state-by-insurer (i.e., "issuer") level, separately for 2014, 2015, and 2016. The crucial empirical object — the slope of the average cost curve — can then be obtained by differentiating AC_j^{Net} with respect to P_j:

$$\frac{\partial AC_j^{Net}(P)}{\partial P_j} = \frac{\partial AC_j(P)}{\partial P_j} - \frac{\partial RE_j}{\partial P_j}$$

As discussed in Section 2 of the main text, adverse selection in pricing occurs when $\frac{\partial AC_j}{\partial P_j}$ is positive. Here, the first term $\frac{\partial AC(P)_j}{\partial P_j}$ represents adverse selection in terms of claims costs. The second term,

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

[79] The timed phase-out of the reinsurance program after 2016 was determined in the original ACA legislation in 2010 and was common knowledge to all market participants. For more details, see https://www.kff.org/affordable-care-act/issue-brief/explaining-health-care-reform-risk-adjustment-reinsurance-and-risk-corridors/

$\frac{\partial RE_j(P)}{\partial P_j}$, encodes the ability of reinsurance to flatten the average cost curve, much like conventional risk adjustment. Taken together, the equation shows that an increase in P_j leads to an increase in claims costs (from enrolling disproportionately more high-cost individuals), but also an increase in reinsurance as the plan enrolls more individuals with costs above the attachment point.[80]

From the above equation, we can see that places where reinsurance payments respond strongly to price (large, positive values of $\frac{\partial RE_j(P)}{\partial P_j}$) are those where average cost curves will steepen more after the removal of reinsurance (after which we have $\frac{\partial AC_j^{Net}}{\partial P_j} = \frac{\partial AC_j(P)}{\partial P_j}$). Our model predicts that more insurer exit will occur in these markets.

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

G.5 Estimating reinsurance-premium elasticities at the state level

To compute how reinsurance responds to premium changes, we first need to compile data on premiums and reinsurance at the issuer-state-year level for 2014-2016.[81]

Premium data: For premiums, we start with the ACA landscape files, which are at the plan-county level. We first limit to 1 observation per plan per rating area (only silver plans), since premiums do not vary within counties in the same rating area. There are about 8 plans per issuer per rating area on average. We merge with plan-level enrollment data and keep the largest (participating) plan per issuer per rating area.[82]

Starting with the individual premium for each plan, we subtract the cheapest plan's premium for each rating area. Finally, we collapse these relative premiums across rating areas within a state by weighting each rating area by its below-65 population shares (using 2010 Census population data). This gives us a measure of P_{jst}, or the premium for issuer j in state s and year t.

Reinsurance data: We use reinsurance payment data at the insurer-state-year level from 2014-16, merged with the premium data above. We limit to 38 states with premium data and further to the 34 states with multiple years of data.[83] We divide by the number of ever-enrolled at the issuer-state-year

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

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

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

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

level. This gives us a measure of RE_{jst}, the per-enrollee reinsurance payment for issuer j in state s and year t. Note, both premiums and especially reinsurance are skewed variables with outliers, so we log transform both relative premiums P_{jst} and reinsurance RE_{jst}.[84]

Estimating the effect of premium changes on reinsurance: The baseline specification for estimating the reinsurance-premium elasticity at the state level is as follows:

$$\log(1 + RE_{jst}) = \delta_s \times \log(1 + P_{jst}) + \eta_j + \eta_s + \eta_t + \eta_{jst} \tag{40}$$

where δ_s is the coefficient of interest that describes the (state-specific) elasticity of per-enrollee reinsurance payments to changes in relative premiums.[85]The η_j, η_s, and η_t fixed effects control for fixed differences in reinsurance across issuers and states,[86] and account for the overall decline in reinsurance payments over time from 2014-2016.[87] Identification comes from changes in the relative premium *within* an issuer over time. We weight the above regression by the number of enrollees in each issuer-state-year observation, and use robust standard errors.[88]

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

$$\log(1 + RE_{jst}) = \delta \times \log(1 + P_{jst}) + \tilde{\eta}_j + \tilde{\eta}_s + \tilde{\eta}_t + \tilde{\eta}_{jst} \tag{41}$$

Here, we cluster standard errors at the state level.[89] A positive estimate of δ, if causally identified, implies that insurers raising their premiums can expect to have larger reinsurance payments (because higher premiums attract more higher-cost enrollees, some of whom will qualify for reinsurance payments). Finally, we perform Bayesian shrinkage to the ($n = 34$) state-specific slopes using the estimator of Morris (1983) and code from Chandra, Finkelstein, Sacarny, and Syverson (2015).

G.6 Firm participation regressions

We start with data on issuer participation at the county-year level from 2014-2018. We focus on issuers operating silver plans in each county. In the ACA marketplaces, counties are only included in the data if they have at least one silver plan. We restrict the data to counties present in the data for all years between 2014 and 2018. We consider several measures of participation, including the number of issuers per county, log number of issuers per county, number of counties with 5 or more firms, and the number of counties that have no more than 1, 2, or 3 firms. This gives us six participation outcome variables in total.

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

[85]Using the log(1+x) transform ameliorates the effect of outliers and cases where the relative premium is zero.

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

[87]The decline in payments is a function of the coinsurance rate and the attachment point.

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

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

We merge the county-level participation data with our state-level estimates of the reinsurance-premium elasticity δ_s and run the following regression:[90]

$$Y_{cst} = \beta \times Post_t \times \delta_s + \alpha_s + \alpha_t + \varepsilon_{cst} \tag{42}$$

where Y_{cst} represents the number of insurers in county c, state s, and year t. We examine participation at the *county* level, because that is the relevant geographic unit where firm participation decisions are allowed to vary.[91] $Post_t$ is an indicator variable for years 2017-18, and δ_s is the state-specific reinsurance-premium elasticity (which we also refer to as the state-specific reinsurance "slope") after Bayesian shrinkage is applied, as shown in Figure A15). The α_s and α_t terms are state and year fixed effects respectively, and ε_{cst} is the error term.[92] The coefficient of interest β encodes the effect of reinsurance on participation. We estimate the model via OLS and cluster standard errors at the state level.

Weighting: The choice of weights for the participation regression (equation 42) is independent from the weighting scheme used for estimating the state-specific slopes (equation 40). In the following results, we will consider both unweighted and weighted specifications. Our preferred specification is unweighted, treating each county equally. For the weighted version, we weight by each county's <65 year-old population from the 2010 Census. This gives more weight to participation changes for counties in larger states. Generally speaking, we find that the decreases in insurer participation associated with the end of reinsurance are driven by counties with smaller populations. This makes sense, as larger counties have more enrollees over which fixed costs can be spread.

G.7 Results

G.7.1 Insurer participation in the ACA over time

Figure A12 below shows how insurer participation has changed in the ACA marketplaces over time.[93] We can clearly see a large decline in participation in 2017-18, after reinsurance was removed after 2016.

During the period where reinsurance was in place (2014-2016), there were between 4-5 insurers participating in each county on average (Panel (a)) and very few markets were single-insurer monopolies (Panel (b)). After the end of reinsurance in 2017, the average number of insurers fell to below 3 per county, and the share of monopoly markets rose as high as 30%. In 2018, there were more than 1,500 monopoly counties, compared to fewer than 250 in 2016.

The timing of this wave of insurer exit is suggestive of a role for reinsurance for maintaining

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

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

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

[93] The estimates in the figure differ slightly from those in Figure 10 in the body text because of slightly different sample of states used for this exercise. Figure 10 uses states that are stably in the Healthcare.gov dataset from 2014-24, while the exercise in this appendix requires a balanced panel for only 2014-18.

competition. However, further analysis is required to argue that reinsurance (1) played a causal role and (2) exerted its effect via the channel identified in our model (i.e., through its effect on adverse selection and the slope of the average cost curve).[94]

Appendix Figure A12. Issuer Participation Over Time

(a) Average Number of Issuers (b) Share of Duopoly or Monopoly Counties

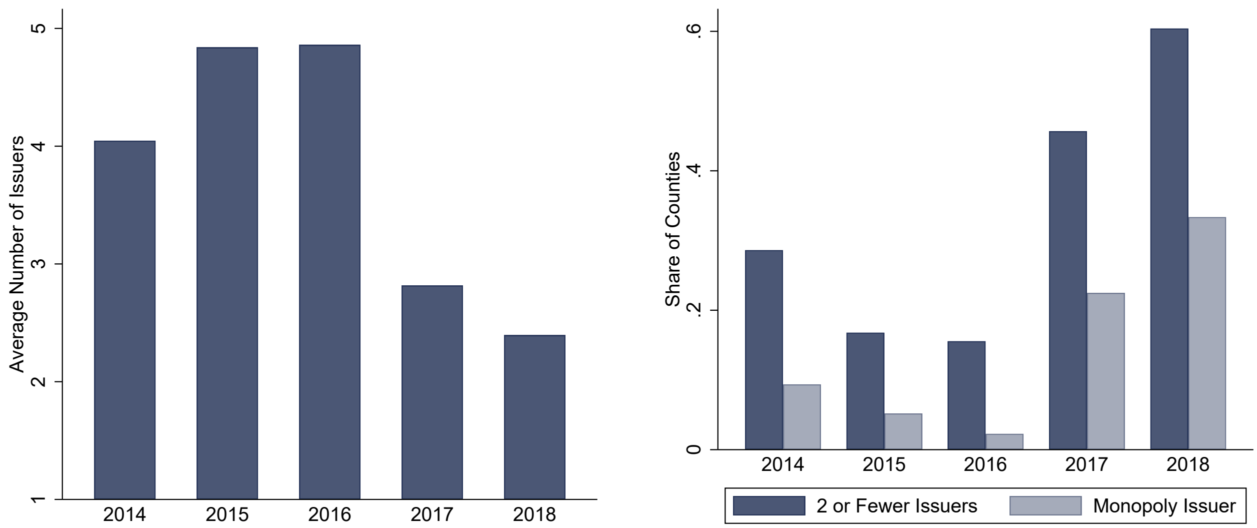

Notes: Figure shows how insurer participation has changed over time. Panel (a) shows the average number of issuers per county by year. Panel (b) shows the population-weighted share of counties with fewer than two issuers (duopoly) or just one issuer (monopoly) by year. Both panels show county-level data, with each county weighted by its total <65 year-old population in the 2010 Census.

G.7.2 Estimates of Reinsurance-Premium Elasticities

Figure A13 below shows the positive relationship between (log) reinsurance per enrollee and (log) relative premium. The left figure is a raw scatterplot, with no controls or fixed effects, where each circle represents an issuer-state-year observation, weighted by total enrollees. The right figure plots the same data as a binned scatterplot, where both (log) reinsurance and (log) relative premium have been residualized on issuer, state, and year fixed effects.

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

Appendix Figure A13. Reinsurance per enrollee vs Premiums

(a) Raw scatterplot, no controls

(b) Binned with controls

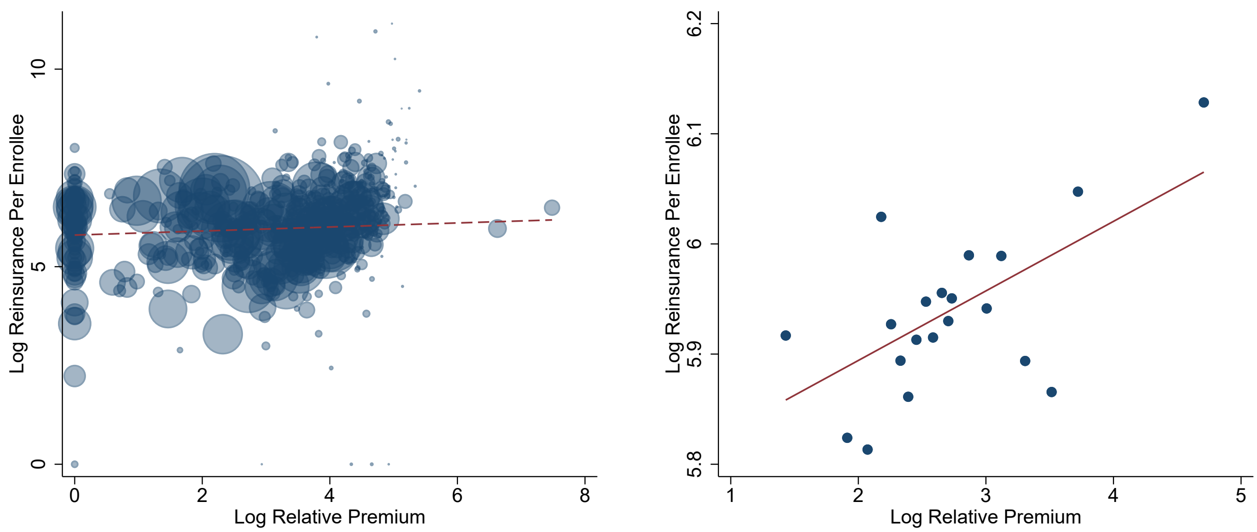

Notes: Figure shows a positive relationship between log(1 + reinsurance per enrollee) and log(1 + relative premium) at the issuer-state-year level, weighted by number of enrollees. Panel (a) shows a scatterplot of the raw data, with circle sizes reflecting the number of enrollees underlying each observation.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

Table A3 below shows estimates for the specification in equation (41), with different sets of fixed effects. Consistent with adverse selection, the sign of the reinsurance slope is positive across all specifications. In our preferred specification (column 5, with state, year, and issuer fixed effects), the elasticity of reinsurance with respect to premium is 0.063. Given an enrollee-weighted average relative premium of \$31 and average reinsurance per person of \$514, this elasticity corresponds to a \$1.06 increase in reinsurance for every \$1 increase in premiums. This suggests that reinsurance played a significant role in flattening the average cost curve in the ACA marketplaces prior to its expiration in 2016. Hence, removing reinsurance in 2016 resulted in a steepening of average cost curves; the magnitude of this effect varies from $\Delta dAC/dP = 0.86$ to 3.60, depending on the specification, with generally more modest effects when controlling for the full set of fixed effects.

Appendix Table A3. Estimates of the Aggregate Reinsurance Slope

	(1)	(2)	(3)	(4)	(5)
log Relative Premium	0.0511	0.135***	0.111***	0.215***	0.0631***
	(0.0320)	(0.0444)	(0.0323)	(0.0287)	(0.0218)
dReinsurance/dPremium	0.855	2.260	1.857	3.601	1.056
State FE		X		X	X
Year FE			X	X	X
Issuer FE					X
R-squared	0.007	0.136	0.402	0.561	0.951
Observations	604	604	604	604	542

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Table shows estimates of the aggregate reinsurance slope obtained using the specification in equation (41). Each observation represents a state-year-issuer. The coefficient on log Relative Premium can be interpreted as the elasticity of reinsurance per enrollee with respect to relative premiums. Standard errors are shown in parentheses. The row titled "dReinsurance/dPremium" re-expresses the elasticity estimates as the change in dollars of reinsurance per enrollee for every $1 increase in relative premiums. Column 5 shows our preferred specification, which controls for state, year, and issuer fixed effects. Regressions are weighted by the number of enrollees in each issuer-state-year.

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

G.7.3 State-specific reinsurance-premium elasticities:

In this section, we show estimates of Equation (40), focusing on specifications that include issuer, state, and year fixed effects. We do this to ensure that our identifying variation comes from *within-issuer* variation over time.

Figure A14 below plots the state-specific coefficients for weighted and unweighted specifications. The median premium elasticity of reinsurance is 0.066. These values are consistent with the results shown previously in Table A3. Reassuringly, we see that most of the state-specific slopes are positive.

Appendix Figure A14. State-specific slopes: reinsurance per enrollee vs premiums

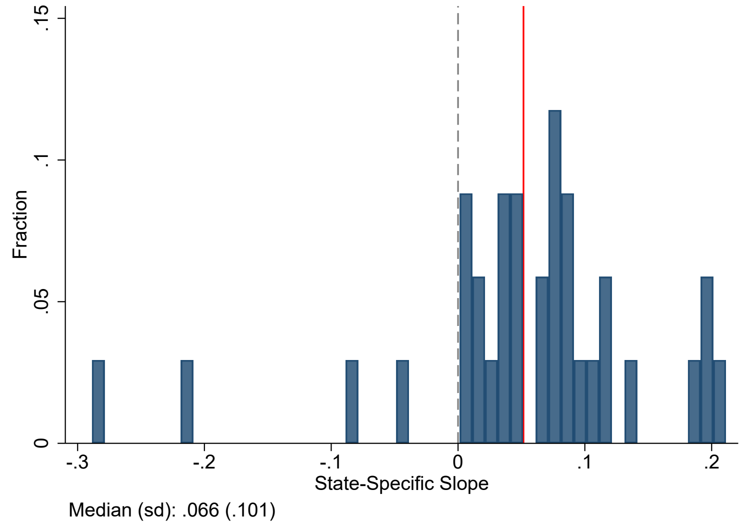

Notes: Figure shows the distribution of $\hat{\delta}_s$ from equation (40). Plotted estimates are for $n = 34$ states using Healthcare.gov over multiple years. Regressions use log(1 + reinsurance per enrollee) and log(1+ relative premium) with state, year, and issuer fixed effects. Results are weighted by the number of enrollees for each issuer-state-year observation with bin widths set to 0.01. Red vertical line indicates the median estimate.

G.7.4 Applying Bayesian shrinkage to the state-specific estimates:

Next, we apply Bayesian shrinkage to the estimates shown in Figure A14. While the original estimates are obtained using OLS, which is unbiased, certain small states may have extreme estimates as a result of sampling error, resulting in an overly dispersed distribution of reinsurance slope estimates. The shrinkage procedure selectively attenuates estimates towards the mean, based on the standard error of each estimate (estimates with larger standard errors are attenuated more).[95] Using the shrunk estimates in our participation regressions (equation (42)) also allows us to account for this differential sampling error.

Figure A15 below mirrors Figure A14, but shows the distribution of the estimates after shrinkage is applied. We see that dispersion is reduced after shrinkage. The following Figure A16 shows the effects of shrinkage on each state-specific estimate, where shrinkage is represented as a rotation from the 45-degree line to the X-axis.

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

Appendix Figure A15. State-specific slopes: after shrinkage procedure

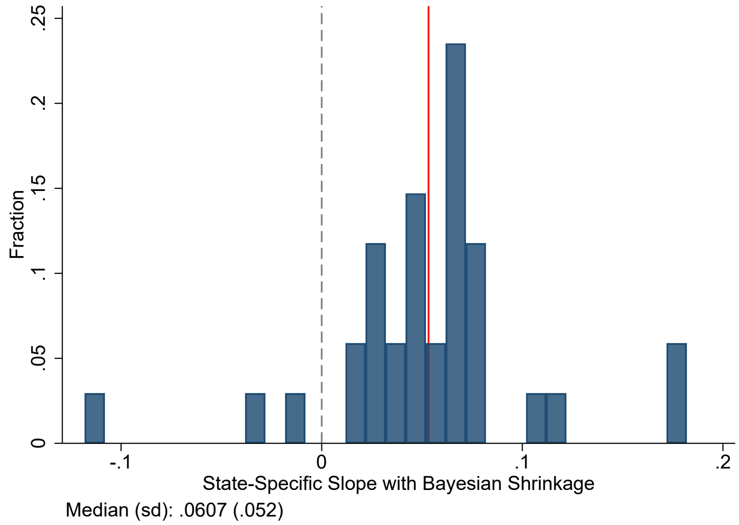

Notes: Figure shows the distribution of $\hat{\delta}_s$ from equation (40) after the Bayesian shrinkage procedure has been applied. Plotted estimates are for $n = 34$ states using Healthcare.gov over multiple years. Regressions use log(1 + reinsurance per enrollee) and log(1+ relative premium) with state, year, and issuer fixed effects. Regression results are weighted by the number of enrollees for each issuer-state-year observation with bin widths of 0.01. Red vertical line indicates the median estimate.

Appendix Figure A16. Visualizing the effects of the shrinkage procedure

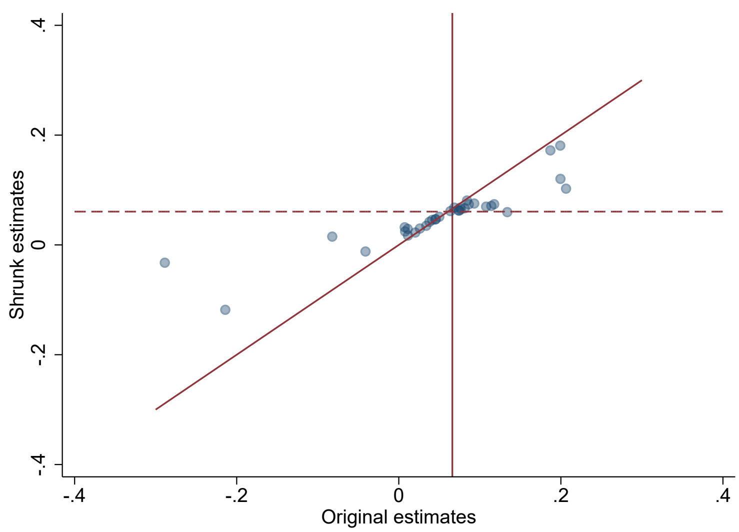

Notes: Figure plots shrunk estimates vs. original estimates of $\hat{\delta}_s$ from Equation (40).

G.8 Effects of Reinsurance on Firm Participation

We now return to the original test of our model: does increasing the slope of the average cost curve result in firm exit? Assuming that our state-specific estimates of the reinsurance-premium elasticity are well identified, we can evaluate this question using the difference-in-differences specification given in equation (42). While reinsurance expired nationwide after 2016, states with larger reinsurance-premium elasticities (reinsurance slopes) should experience more firm exit. Including state and year fixed effects ensures that we control for fixed differences in participation across states as well as aggregate time trends in participation over time.

As detailed in the following section, this is indeed what we find. Counties in states with larger reinsurance elasticities experienced larger declines in firm participation. This is consistent with our model, which ties firm participation to the slope of the average cost curve.

G.8.1 Firm participation results:

Figure A17 visualizes our model predictions. It shows that states with above-median reinsurance-premium elasticities had larger declines in insurer participation after the end of reinsurance after 2016, compared to below-median states. This is consistent with the idea that reinsurance played a larger role in flattening the average cost curve in these states in the pre-period (2014-2016). Before the end of reinsurance, above-median states had 1.42 more issuers per county. This difference dropped to 0.80 after the end of reinsurance, suggesting that the end of reinsurance was responsible for a 0.62 decline in the average number of issuers per county in above-median states. Overall, the average number of issuers declined by 1.9, from 4.3 in 2014-16 to 2.4 in 2017-18. Thus, the end of reinsurance is associated with $0.62/1.9 = 33\%$ of the overall decline in participation between these two time periods.

Appendix Figure A17. Issuer Participation Over Time by Reinsurance Slope

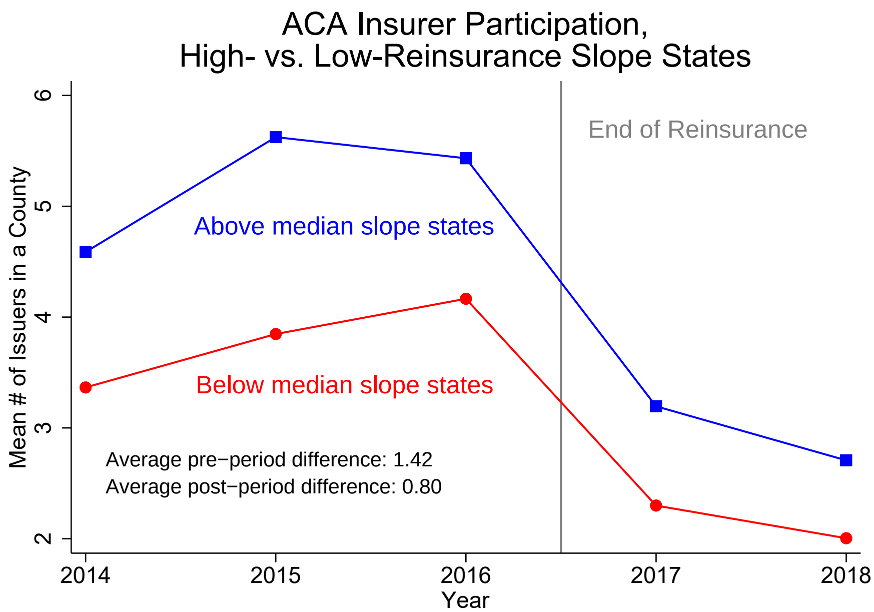

Notes: Figure shows the average number of issuers in counties with above- versus below-median values of the reinsurance-premium elasticity. The gray vertical line represents the end of the ACA reinsurance program after 2016. Counties are weighted by their <65 year-old population totals in the 2010 Census.

Table A4 below quantifies this effect by showing estimates from Equation (42). Column 2 shows our preferred specification. Columns 3-5 show robustness tests. Column 3 shows results weighted by county population. While no longer statistically significant, the sign of the effect is preserved. Also, this weighted specification is statistically significant in smaller counties (Appendix Table A5). Column 4 limits the post-period to 2017, to account for any effects of the 2016 US presidential election, as insurer participation decisions for 2017 were determined in 2016 prior to the results of the 2016 election. Our point estimate is largely unchanged, despite larger standard errors (as expected from halving the post-period sample size). Column 5 demonstrates that our results are robust to not applying Bayesian shrinkage to our estimates of δ_s.

Across all specifications, we see that larger reinsurance slopes are associated with greater exit after the expiration of reinsurance. The reported coefficients can be interpreted as the effect of suddenly increasing the reinsurance-premium elasticity by 1. In our preferred specification in Column 2, a 1 standard deviation increase in the reinsurance elasticity (a 0.0607 unit increase), corresponds to a decline of 0.27 firms after the end of reinsurance.

Appendix Table A5 shows that the effect of reinsurance on participation is more robust for counties with smaller population sizes, with much smaller standard errors for counties with below-median population sizes.

Appendix Table A4. Effect of Reinsurance Expiration on Insurer Participation in the ACA

	(1) First Stage	(2) Second Stage	(3) Second Stage: Weighted	(4) Only 2017 Post-Period	(5) Slope without Bayesian Shrinkage
log Premium	0.0631***				
	(0.0218)				
Post X Slope		-4.383***	-2.281	-4.176**	-2.492***
		(1.181)	(2.833)	(1.903)	(0.730)
Observations	542	12,470	12,470	9,976	12,470
R-squared	0.951	0.754	0.793	0.780	0.753
FEs	State/Year/Issuer	County/Year	County/Year	County/Year	County/Year
Pre-Period Mean	5.999	3.001	4.587	3.001	3.001

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Table shows the effect of reinsurance expiration on insurer participation in the ACA. Column 1 reproduces the first stage regression of reinsurance on premiums from Appendix Table A3. Columns 2-5 regress the number of insurers per county on the interaction between the first stage slope and an indicator for the post period (2017-2019) (Equation (42)). The coefficient on Post X Slope indicates the effect of a 1 unit increase in the reinsurance-premium elasticity δ_s.

Appendix Table A5. Effect of Reinsurance on Participation (By County Population)

(a) Above-Median County Population

	(1) First Stage	(2) Second Stage	(3) Second Stage: Weighted	(4) Only 2017 Post-Period	(5) Slope without Bayesian Shrinkage
log_premium	0.0631***				
	(0.0218)				
post_slope		-4.346**	-1.692	-3.306	-2.108*
		(1.736)	(2.964)	(2.878)	(1.080)
Observations	542	6,235	6,235	4,988	6,235
R-squared	0.951	0.776	0.792	0.799	0.774
FEs	State/Year/Issuer	County/Year	County/Year	County/Year	County/Year
Pre-Period Mean	5.999	3.428	4.731	3.428	3.428

(b) Below-Median County Population

	(1) First Stage	(2) Second Stage	(3) Second Stage: Weighted	(4) Only 2017 Post-Period	(5) Slope without Bayesian Shrinkage
log_premium	0.0631***				
	(0.0218)				
post_slope		-3.242***	-3.377***	-3.834***	-1.975**
		(1.064)	(1.025)	(1.206)	(0.732)
Observations	542	6,235	6,235	4,988	6,235
R-squared	0.951	0.686	0.710	0.701	0.686
FEs	State/Year/Issuer	County/Year	County/Year	County/Year	County/Year
Pre-Period Mean	5.999	2.574	2.623	2.574	2.574

Robust standard errors in parentheses
*** p<0.01, ** p<0.05, * p<0.1

Notes: Table shows the effect of reinsurance expiration on insurer participation in the ACA, separately by above- versus below-median county population from the 2010 Census. Panel (a) shows results for counties with above-median population. Panel (b) shows results for below-median counties. Column 1 reproduces the first stage regression of reinsurance on premiums from Appendix Table A3. Columns 2-5 regress the number of insurers per county on the interaction between the first stage slope and an indicator for the post period (2017-2019) (Equation 42). The coefficient on Post X Slope indicates the effect of a 1 unit increase in the reinsurance-premium elasticity δ_s.

The following Figure A18 visualizes the data underlying Column 2 of Table A4. The figure plots a binned scatterplot of firm participation against estimated reinsurance slopes, separately by time period (before and after the end of reinsurance). The raw (binned) data are shown without weights or fixed effects. Because we omit year fixed effects, we can see that firm participation drops substantially in the post-period (2017-18) relative to the pre-period (2014-16) across the board, but the largest decreases in participation are for counties that had higher reinsurance elasticities in the pre-period. Because we omit state fixed effects, we can see that the states with large pre-period reinsurance slopes actually had *more* robust insurer participation.

Appendix Figure A18. Insurer Participation in the ACA Before and After Removal of Reinsurance

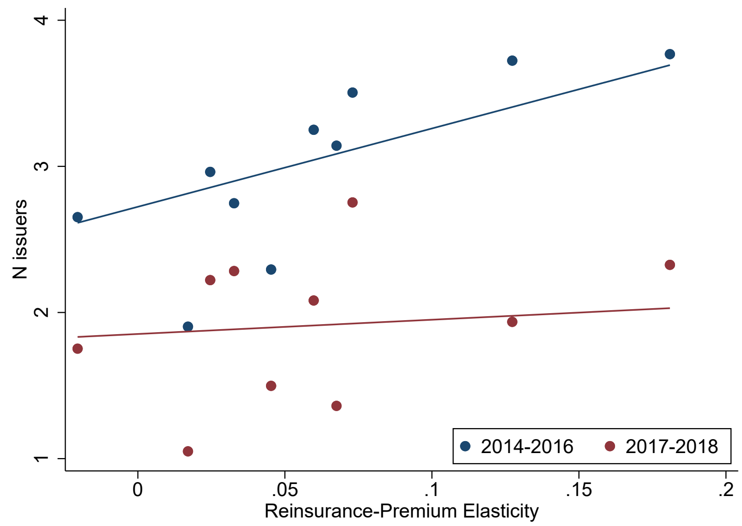

Notes: Figures show how county-level firm participation changes after reinsurance is removed from the ACA in 2016. The Y-axis plots the average number of issuers; the X-axis shows the magnitude of the state-specific reinsurance elasticity for each of $n = 34$ states (as shown in Figure A15).

H Additional Exhibits

Appendix Figure A19. CommCare Plans Pre-Subsidy Prices

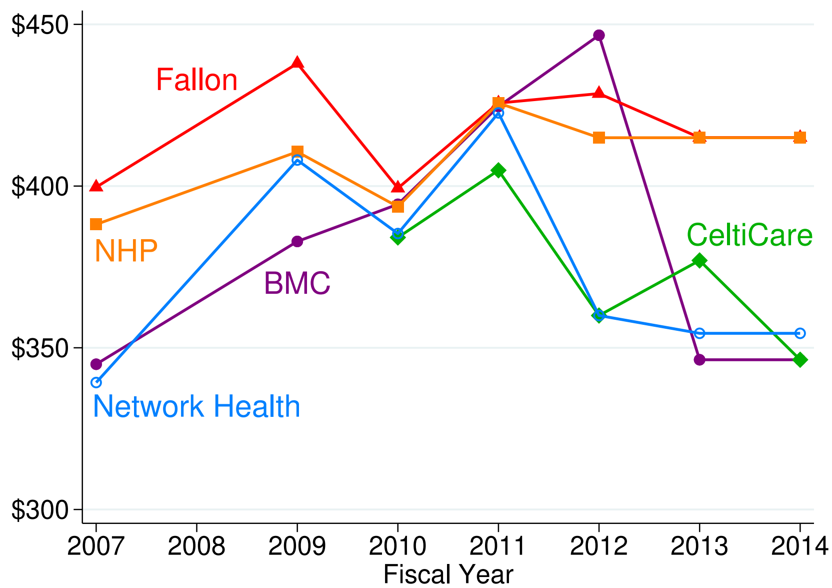

Note: The graphs show average pre-subsidy insurer prices for each insurer's plan in the CommCare market, by fiscal year. The five plans are shown in different colors and labeled. Values shown are averages for the plan's actual enrollees; underlying premiums and (in some years) prices vary by income group 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 A20. Premium Variation Example: Network Health (Boston region), 2010-13

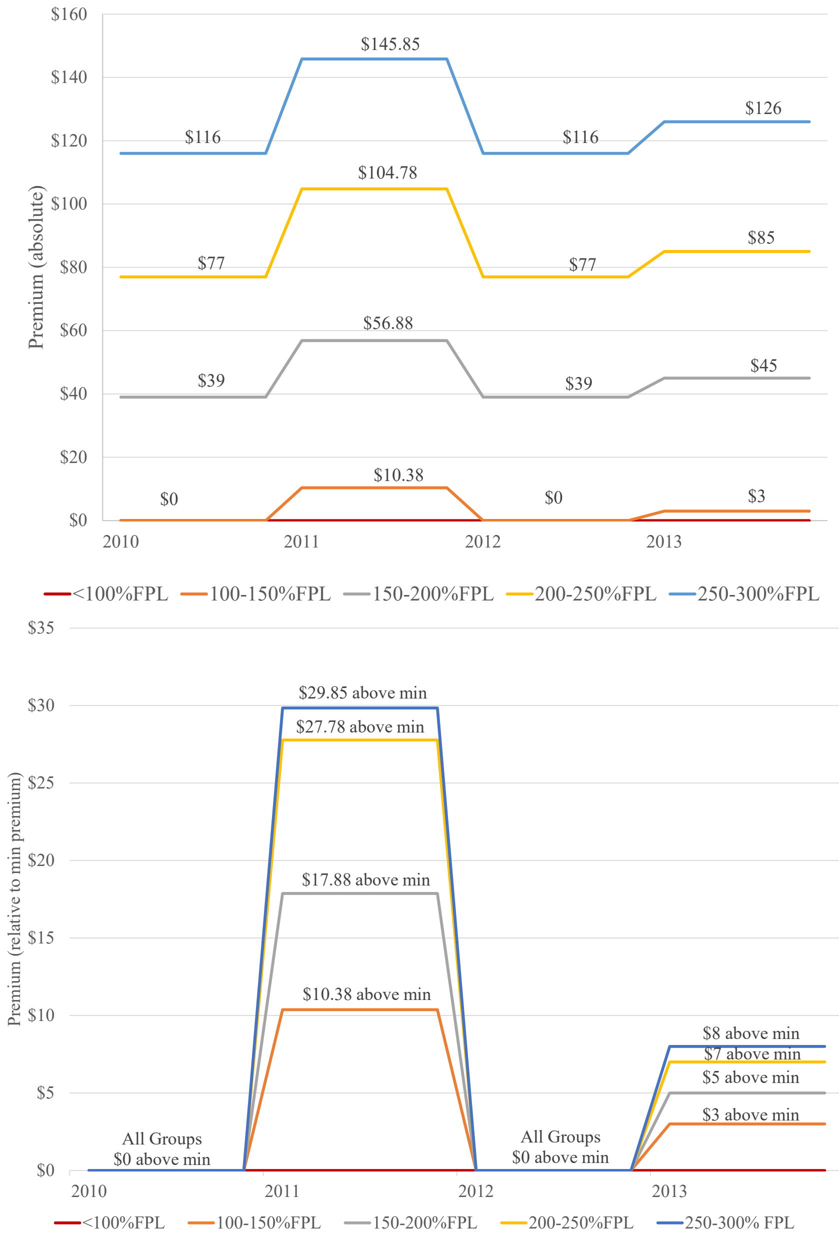

Note: The graphs shows the example of Network Health's (post-subsidy) enrollee premiums by income group over the 2010-2013 CommCare years. "FPL" refers to the federal poverty level. Pre-subsidy prices (and enrollee premiums) vary at the regional level in 2010, and the graph shows premiums specifically 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 A21. Hospital Coverage in Massachusetts Exchange Plans

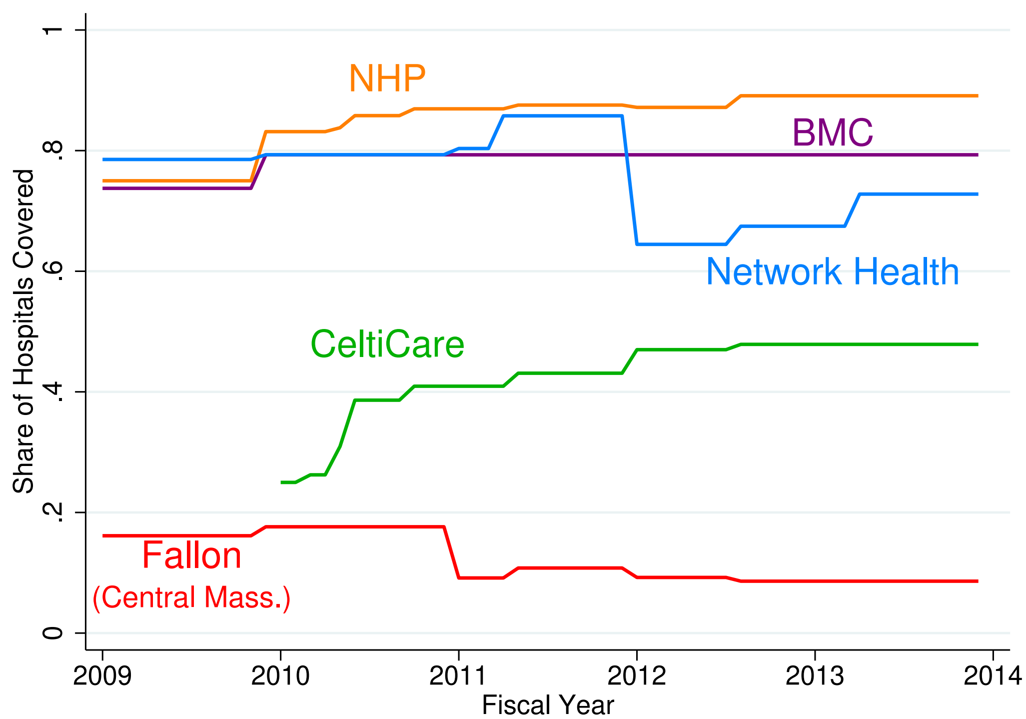

Note: The graph shows the shares of Massachusetts hospitals covered by each CommCare plan, where shares are weighted by hospital bed size in 2011. Fallon's hospital coverage share is much lower than other plans largely because it mainly operates in central Massachusetts and therefore does not have a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

H.1 Baseline Simulation Results and Robustness

Appendix Table A6. Counterfactual Simulation Results

(1) Policy Scenario	(2) Entrants	(3) Prices	(4) Average price	(5) Surplus	(6) Survives F= $10	(7) Survives F= $30	(8) Optimal Floor	(9) Surplus
Panel (a): Equilibria as a function of risk adjustment and price floors								
(1) Actual Risk Adj. (*baseline*)	**Monopoly** [*CeltiCare*]	[$475]	$475	$0	Yes	Yes	$393	$132
(2) No Risk Adj. ($\lambda = 0$)	**Monopoly** [*CeltiCare*]	[$475]	$475	$0	Yes	Yes	$393	$132
(3) High Risk Adj. ($\lambda = 0.7$)	**Three firms** [*BMC CeltiCare,Network*]	[$431, $371, $428]	$390	$119	No	No	$380	$122
(4) Perfect Risk Adj. ($\lambda = 1$)	**All four firms** [*BMC CeltiCare, NHP, Network*]	[$437, $375, $475, $431]	$397	$116	No	No	None	—
Panel (b): Robustness to alternative assumptions								
(5) No Celticare	**Two firms** [*BMC,Network*]	[$384, $389]	$386	$135	No	No	None	—
(6) No Plan Cost Diffs.	**Monopoly** [*CeltiCare*]	[$475]	$475	$0	Yes	Yes	$393	$131
(7) All Enrollees	**Two firms** [*BMC,CeltiCare*]	[(0.14*$395, 0.86*$429), $372]	$392	$133	Yes	Yes	$360	$160

Note: Table shows results from our counterfactual simulations across different levels of risk adjustment (Panel (a)) and alternative assumptions (Panel (b)). In all of these policy scenarios, we find a unique equilibrium that survives our refinement rule. Simulations in Rows 1-6 are conducted using a random sample of new enrollees (N=5053). Rows 1-4 vary the level of risk adjustment. Rows 5-7 represent modifications to our baseline specification in Row 1; these rows continue to assume actual risk adjustment. Row 5 ("No Celticare") removes Celticare from the market. Row 6 assumes that enrollee costs do not vary across plans. Row 7 uses a sample of all enrollees, which includes both new and current enrollees, which allows us to capture the effects of inertia in consumer demand (N=5026). Column 1 describes the policy scenario. Column 2 lists the equilibrium set of entrants. Column 3 lists the equilibrium prices for each firm (all values in $ per month), in order of the firms listed in column 2 (note: Row 7 shows a mixed strategy where BMC mixes between two prices). Column 4 shows the (share-weighted) average price. Column 5 gives the consumer surplus per enrollee in the market, relative to the surplus of the Celticare monopoly, which normalized to zero. Columns 6 and 7 denote whether the equilibrium can survive fixed cost magnitudes equivalent to $10 or $30 per enrollee in the market, split evenly across firms. Column 8 shows the optimal price floor (if one exists), and Column 9 shows the consumer surplus that would occur under the optimal price floor.

Appendix Figure A22. Market Outcomes across Risk Adjustment and Price Floors (All Outcomes)

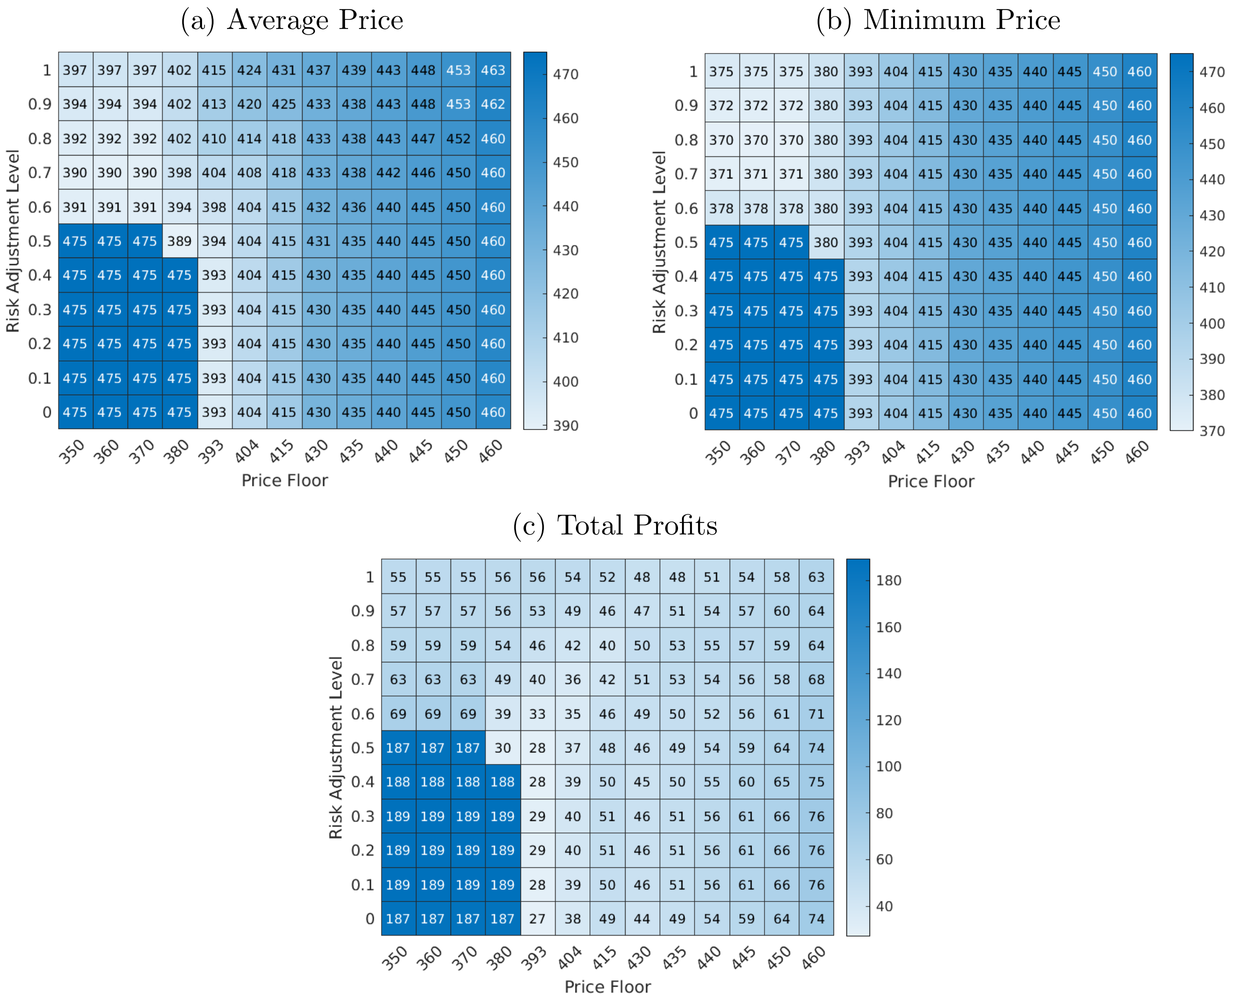

Notes: Figures show how equilibrium outcomes vary as a function of the level of risk adjustment $\lambda \in (0, 1)$ (shown on the Y-axis) and price floors (X-axis). This shows additional outcomes beyond the equilibrium number of firms and consumer surplus shown in Figure 9 in the body text. For each panel, the Z-axis (i.e., heatmap color) plots a different outcome of interest. Each cell corresponds to a unique equilibrium. Panel (a) shows the share-weighted average price; Panel (b) shows the minimum price of the cheapest plan; and Panel (c) shows total insurer profits (per enrollee). All values are in $ per enrollee-month.