

Subsidizing Quality for the Poor: Equity and Competition in Health Insurance Markets

Kai Shen Lim*

16 October 2026

[CLICK HERE FOR LATEST VERSION](#)

Abstract

For many essential goods and services, the poor are often priced out of higher-quality options. Governments respond by incrementally subsidizing these options to ensure affordability for the poor. This paper studies the trade-offs of incremental subsidies in the Massachusetts health insurance exchange, where the state paid part of the price gap between higher-quality (broad-network) plans and lower-quality plans for low-income enrollees. Using the removal of these subsidies in 2017, I show that the policy afforded low-income enrollees access to broad-network plans but weakened price competition among insurers. Equilibrium simulations show that 91 cents of every incremental subsidy dollar became insurer profits despite increasing the poor's enrollment in the broad-plan. The policy cannot be rationalized by a planner who maximizes consumer surplus net of public spending, and instead suggests a planner who places additional value on low-income enrollment in broad-network plans. A counterfactual fixed subsidy to enrollees choosing a broad-network plan yields higher enrollment at similar public cost, with lower insurer profits. Overall, for a planner who values equity in quality, subsidizing quality directly is optimal relative to subsidizing incremental prices.

1 Introduction

Many societies value the poor’s access to essential goods and services—medical care, schooling, and housing—beyond the redistribution of income (Tobin, 1970). A typical policy response is to make these goods and services affordable for the poor. Indeed, subsidies have been shown to raise the poor’s health insurance coverage (Frean et al., 2017; Finkelstein et al., 2019), college attendance and completion (Dynarski, 2003; Denning et al., 2019), and housing stability (Wood et al., 2008). These *level* subsidies provide equal amounts across all options, lowering the price of participating in these markets for the poor.

But these goods are often sold in differentiated markets, where higher-quality options yield lower mortality, greater educational attainment, and higher adult earnings (Abaluck et al., 2021; Deming et al., 2014; Chetty et al., 2016). Increasingly, government policy goes beyond subsidizing market participation to ensure the poor’s access to these higher-quality options: housing vouchers under ‘Small Area Fair Market Rents’ pay more in expensive neighborhoods (Collinson and Ganong, 2018), the Cal Grant pays more toward tuition at the University of California than at the California State University, and many state child care programs pay more to higher-rated providers. Despite the ubiquity of these equity-driven policies, the literature has paid less attention to the trade-offs of subsidizing higher-quality options for the poor.

In this paper, I study an *incremental* subsidy policy that Massachusetts’ ACA marketplace implemented to “afford meaningful choice”¹ for low-income enrollees. Specifically, the policy subsidizes a share of the price gap between higher-priced and cheaper plans. To understand the trade-offs behind the incremental subsidy policy, I provide a theory of subsidized insurance markets, where the social planner uses both a *level* and an *incremental* subsidy to achieve welfare goals. The model consists of a high-quality insurer that covers more hospitals and a low-quality insurer. Enrollees differ in income and health, and sicker enrollees value the high-quality plan more while costing more to insure.

Under perfect competition, adverse selection—when sicker enrollees value the high-quality plan more, raising average costs and prices—can justify the incremental subsidy on efficiency grounds. An incremental subsidy addresses this by lowering enrollees’ net price of the high-quality plan and keeping healthier (more price-sensitive) enrollees in the high-quality plan. However, this justification weakens if risk adjustment sufficiently addresses adverse selection by transferring funds from plans with healthier enrollees to plans with sicker enrollees. With imperfect competition, the subsidy introduces a competitive distortion. Subsidizing the price gap between the more expensive and cheaper option implies that enrollees now bear only a share of any price increase, and insurers are able to raise markups.

¹Massachusetts Connector ‘Health Connector Function Areas: Status Updates and Looking Ahead’ Presentation to the Board of Directors, November 10, 2022

Why subsidize the incremental price at all, when the planner could lower the poor's premium for every plan by the same amount using *level* subsidies and avoid the competitive distortion? The key difference is that the *incremental* subsidy incentivizes *enrollment in the high-quality plan*. A planner would choose it either to correct adverse selection, or because it values the poor's enrollment in the high-quality plan. Which motive the observed policy reveals is an empirical question, and I disentangle the two channels to understand its trade-offs.

I take theory to data in Massachusetts' ConnectorCare—the subsidized segment of the state's insurance marketplace created by the Affordable Care Act (ACA). Three features of this setting allow me to disentangle the planner's two motives. First, ConnectorCare implements both a *level* and an *incremental* subsidy for enrollees below 300 percent of the Federal Poverty Level (FPL). This is an important feature of the setting, since it allows me to separate the social planner's concern for the poor generally (level subsidies) from its concern for the poor's access to the high-quality plan (incremental subsidies). Second, insurers compete mainly on two dimensions—prices and hospital network breadth: ConnectorCare requires every insurance plan to cover the same set of services with limited enrollee cost-sharing. Third, in 2017 ConnectorCare removed incremental subsidies for the high-quality plan due to fiscal concerns. This policy variation identifies enrollees' price sensitivity for the high-quality plan and how much adverse selection worsens once the subsidy is removed.

Using administrative enrollment data from the Connector and medical claims from the Massachusetts All-Payer Claims Database, I show that the removal of the incremental subsidy in 2017 led to an exodus of enrollees from the high-quality plan to the lower-quality plans. The policy removal raised the net premium for the high-quality plan from \$0 to \$117 a month for enrollees below the poverty line. This also affected enrollees at higher incomes, though at a smaller gap increase from \$9 to \$122 at 100–150 percent of FPL and from \$43 to \$118 at 250–300 percent. Unsurprisingly, among below poverty enrollees, the high-quality plan's market share plummeted from 51 percent in 2016 to approximately 5 percent in 2017. Above poverty enrollees also switched away from the high-quality plan proportionally to the price increase: the high-quality plan's share fell from 23 to 9 percent at 100–150 percent of FPL and from 24 to 15 percent at 250–300 percent. Overall, the below poverty enrollees' response per dollar of price gap was about three times that of the other income groups.

The removal of the incremental subsidy worsened adverse selection in the high-quality plan, but risk adjustment compensated most of the cost increase. Within each income group, the high-quality plan's cost per enrollee rose relative to the narrow plans' after the removal, by 225 percent for below poverty enrollees and by 34 and 30 percent for above poverty enrollees. Risk adjustment covers 86 percent of the cost increase for below poverty enrollees and 60 to 68 percent for above poverty enrollees. This shows that risk adjustment weakens the efficiency case for the incremental subsidy. The below-poverty enrollees who switched

also lost access to care, as claims at the high-quality plan's exclusive hospitals fell by 98 percent, and office visits fell by 42 percent.

I then use the same policy removal to estimate a structural model of demand and costs for insurance plans. I use the model to simulate equilibria under alternative incremental subsidy schedules, mapping each schedule to equilibrium prices, enrollment, and welfare. Raising the incremental subsidy from zero to its observed level doubles low-income enrollment in the high-quality plan, from 2.4 to 5.0 percent of the market. The subsidy also draws in healthier enrollees, which mitigates adverse selection. However, the high-quality plan's average revenue increased from \$624 to \$776 per member-month, implying that 91 cents of every dollar the state spends on the incremental subsidy becomes insurer profits. Consumer surplus rises minimally since the subsidy is passed onto the high-quality plan's price rather than to enrollees.

No welfarist planner rationalizes the observed policy. I simulate welfare under two counterfactuals: removing the incremental subsidy, and reallocating its budget to the level subsidy as a lump sum to the same enrollees. An efficiency planner that weighs consumer surplus against the cost of public funds prefers removing incremental subsidies: the incremental subsidy costs \$35 per member-month and returns \$2 of consumer surplus, and the reallocated lump sum gains \$35 consumer surplus at a cost of \$45. A utilitarian planner with moderate inequality aversion ranks the three policies the same way. With strong inequality aversion, the planner prefers a larger level subsidy against incremental subsidies. Under every welfarist criterion, a level subsidy thus dominates the incremental subsidy.

Choosing the observed incremental subsidies requires valuing the poor's enrollment in the high-quality plan beyond consumer surplus. A planner that chooses the observed policy would have to value an additional low-income enrollee in the high-quality plan at \$3,821 per member-month, 4.9 times the plan's average revenue and 475 times enrollees' own willingness to pay for access. Such a large valuation could reflect benefits of broad-network care that enrollees do not internalize, a specific egalitarian preference for equal access in the sense of Tobin (1970), or Connector's preference in keeping a high-cost broad-network insurer in the subsidized market. Regardless of the interpretation, the policy reveals a planner that values enrollment in the high-quality plan.

Given that the planner has a strong preference for enrolling in the high-quality plan, what other policies could improve welfare without distorting competition? I consider a counterfactual policy of a quality-based subsidy, where enrollees are given a fixed subsidy for choosing the high-quality plan. With an approximately similar government cost, this policy would have yielded 26 percent less insurer profits and 67 percent more enrollment in the high-quality plan. This policy shows that subsidizing quality is a more cost-effective

policy relative to subsidizing incremental prices to ensure that high-quality options are affordable for the poor.

These implications go beyond health insurance. Housing vouchers under Small Area Fair Market Rents, the Cal Grant, and tiered child care subsidies among other policies implement incremental subsidies for the higher-quality option. In such settings, firms are able to raise markups, and a fixed payment for choosing high-quality options removes potential competitive distortions.

This policy takeaway however, rests on two important caveats. The first is that the planner has to define which aspect of quality to subsidize. Quality is often multidimensional, and subsidizing some dimensions but not others may lead firms to over-invest in subsidized dimensions but not others (Holmstrom and Milgrom, 1991). Second, the model holds input prices fixed and does not endogenize supply-chain (insurer-hospital) bargaining. In theory, the quality subsidy could raise or lower consumer welfare through this channel. A subsidy that expands the high-quality firm's output strengthens its position against its input suppliers, but it also raises the value of the inputs that define its quality, which strengthens the suppliers' position (Gowrisankaran et al., 2015; Ho and Lee, 2017). If suppliers capture part of the subsidy through higher input prices, the firm's costs and price rise and the quality subsidy's advantage over the incremental subsidy shrinks. Which effect dominates depends on how willing consumers are to substitute across inputs.

Related Literature. This paper contributes to two strands of the literature. The first studies subsidy design and competition in health insurance markets. Existing work shows that subsidies interact with market power: insurers game the low-income subsidy in Medicare Part D (Decarolis, 2015), capture nearly half of payment increases in Medicare Advantage (Cabral et al., 2018), and price against price-linked subsidy formulas (Jaffe and Shepard, 2020; Tebaldi, 2025). Finkelstein et al. (2019) estimate low-income willingness to pay for coverage in the same Massachusetts market. My contribution is in showing the trade-offs of an equity driven policy, and documenting a less distortionary policy alternative to incremental subsidies.

The second literature studies in-kind transfers and redistribution through allocation. This literature asks when in-kind provision beats cash (Nichols and Zeckhauser, 1982; Besley and Coate, 1991; Currie and Gahvari, 2008); the recent mechanism-design strand characterizes optimal allocation of quality when the planner controls prices directly (Dworczak et al., 2021; Akbarpour et al., 2024). I provide empirical evidence to this classic question, showing that a real-world policy is difficult to rationalize through standard welfare analyses.

The paper proceeds as follows. Section 2 presents the theory. Section 3 describes the setting and data. Section 4 presents the reduced-form evidence on enrollment, selection,

and care. Section 5 develops the structural model and estimation. Section 6 presents the counterfactual and welfare analysis. Section 7 concludes.

2 Theory

In this section, I present a model of an insurance market with two vertically differentiated insurers. The social planner chooses two subsidy instruments to maximize welfare. The first instrument is a *level* subsidy that lowers the price the poor pays for every plan. The second is an *incremental* subsidy that pays a share of the price gap between a high-quality (broad-network) plan and a low-quality (narrow-network) plan.

The key insight is that the incremental subsidy *changes the composition of the marginal enrollee* for the high-quality plan. Given this, there are two reasons a planner might use the incremental subsidy. The first is adverse selection. If subsidized poor enrollees are healthier and also price sensitive, making the high-quality plan affordable through incremental subsidies would decrease average costs for the high-quality plan, leading to possibly lower prices to consumers. However, this depends on the effectiveness of risk adjustment in correcting adverse selection. The second—and more likely explanation—is a direct preference for subsidized enrollees' access to the high-quality plan.

I begin in Section 2.1 by setting up the market and the two subsidy instruments. Next, Section 2.2 shows how adverse selection affects costs. Section 2.3 then derives the high-quality plan's price response under perfect competition and imperfect competition. Finally, Section 2.4 asks which planner would prefer the incremental subsidy.

2.1 Setup

Consider a market in which two insurers each sell a single plan to a mass-one population of low-income enrollees (indexed by i). Enrollees do not face an outside option². Insurer B sells a high-quality (broad-network) plan and insurer N a low-quality (narrow-network) plan, at gross premiums $p_B > p_N$, and I write $\Delta p \equiv p_B - p_N$ for the price gap between both plans. Each enrollee has an incremental value ΔV_i of the high-quality plan over the low-quality plan.

The key enrollee characteristic is health, $h_i \in \{H, S\}$ (healthy or sick), which affects both the value an enrollee has for the high-quality plan and what the enrollee costs the insurer. Although costs differ by health, insurers cannot price discriminate and charge a single premium for its plan. The market, therefore, features the asymmetric information that leads to adverse selection.

The planner subsidizes enrollees through two instruments. A level subsidy $\tau \geq 0$ pays the same amount toward either plan. An incremental subsidy pays a share $s \in [0, 1]$ of the price gap for enrollees who choose the high-quality plan. The net premium an enrollee

²Given that the model focuses on the enrollee's choice of a high-quality plan relative to a lower-quality plan, I focus the model exposition on the intensive margin of choice rather than the extensive margin of whether to participate in this market. My empirical setting also features a low uninsurance rate, given that Massachusetts imposed individual mandates for health insurance coverage.


pays for plan j is therefore:

$$\tilde{p}_j = \begin{cases} p_N - \tau & j = N \\ p_B - \underbrace{\tau}_{\text{level subsidy}} - \underbrace{s \cdot \Delta p}_{\text{incremental subsidy rate}} & j = B \end{cases} \quad (1)$$

An enrollee choosing between the plans responds to the difference in net premiums, $\Delta \tilde{p} \equiv \tilde{p}_B - \tilde{p}_N = (1 - s) \cdot \Delta p$, which I call the net price gap. Notably, the level subsidy enters both net premiums, whereas the incremental subsidy narrows the gap between the high-quality and low-quality plans to $(1 - s) \cdot \Delta p$. Thus, the incremental subsidy affects plan choice while the level subsidy does not. Every demand-side object below depends on prices and on the incremental subsidy only through $\Delta \tilde{p}$.

I visualize this difference in [Figure 1](#), which compares a level subsidy (Panel A) to an incremental subsidy (Panel B) for the same pair of gross premiums. In Panel A the state pays the same dollar toward either plan and enrollees face the full price gap; in Panel B it pays a larger share of the high-quality plan's premium and the gap narrows. This brings us to the first implication of the model:

Figure 1: Difference Between a Level and Incremental Subsidy



Notes: Green is the level subsidy τ the state pays toward either plan. Maroon is the incremental subsidy $s \cdot \Delta p$ the state pays toward the high-quality plan. The height of the dark blue bar is the enrollee's net premium. The level subsidy cancels from the gap enrollees face; the incremental subsidy narrows it to $(1 - s) \cdot \Delta p$.

Implication #1. *The level subsidy is equivalent to a lump-sum transfer. It lowers both net premiums by τ and does not affect plan choice.*

The planner and the insurers interact in a simple two-stage game with fixed plan attributes:

1. **Subsidy:** The planner commits to the level subsidy and the incremental subsidy rate, (τ, s) .
2. **Pricing:** The two insurers set gross premiums (p_B, p_N) simultaneously to maximize profits, and enrollees then choose plans.

2.2 Adverse Selection

Next, I show how adverse selection changes firm pricing incentives.

Demand and the marginal enrollee. An enrollee chooses the high-quality plan when the enrollee's incremental value exceeds the net price gap. Enrollment in the high-quality plan, $D_B(\Delta\tilde{p})$, is the share of enrollees that sort into the high-quality plan, and the remaining $1 - D_B$ enroll in the low-quality plan.³ Demand for the high-quality plan falls in its net premium, $D'_B(\Delta\tilde{p}) < 0$.

Notice that the two subsidy instruments affect this margin differently. The incremental subsidy lowers the high-quality plan's net premium by Δp per unit of s , which gives $\partial D_B / \partial s = -\Delta p \cdot D'_B(\Delta\tilde{p}) > 0$. By contrast, a \$1 increase in p_B raises the net premium by only $1 - s$, and $\partial D_B / \partial p_B = (1 - s) \cdot D'_B(\Delta\tilde{p}) < 0$. Effectively, a higher s shields the high-quality plan's enrollees from the full price increase.

I define the high-quality plan's semi-elasticity of demand with respect to the net price gap as $\tilde{\eta}_B(\Delta\tilde{p}) \equiv -D'_B(\Delta\tilde{p}) / D_B(\Delta\tilde{p})$. I define the high-quality plan's (own-price) semi-elasticity of demand as $\eta_B \equiv -(\partial D_B / \partial p_B) / D_B$. The two are related by $\eta_B = (1 - s) \cdot \tilde{\eta}_B(\Delta\tilde{p})$, since a dollar of gross premium affects the gap by $1 - s$.

Adverse selection arises because the sick, who expect more medical spending, get more out of the high-quality plan and so have a higher willingness to pay for it (Einav and Finkelstein, 2011). Formally, I assume that the sick share among the marginal enrollees rises with the price gap: as the high-quality plan becomes more expensive, the enrollees left at the margin are increasingly the sick.⁴

Costs and the selection wedge. Next, I turn to the high-quality plan's costs. All derivations are in [Appendix A](#). Covering an enrollee of health h costs the high-quality plan c_{hB} after risk-adjustment, with $c_{SB} > c_{HB}$. Briefly, risk adjustment compensates insurers for enrolling sicker people. This is to prevent the market from unraveling where expensive plans only attract the sick and cheaper plans attract the healthy. In practice, risk adjustment

³Formally, $D_B(\Delta\tilde{p}) = \Pr\{\Delta V_i \geq \Delta\tilde{p}\}$. I treat D_B as continuously differentiable in the gap with positive density at the margin, which lets me differentiate it below.

⁴This holds when incremental value is a health shift plus logistic taste, $\Delta V_i = V_{h_i} + \varepsilon_i$ with $V_S > V_H$, which is the form of the demand model in [Section 5](#), and it implies the covariance condition $\text{Cov}(\Delta V_i, \mathbf{1}[h_i = S]) > 0$ that Kong et al. (2024) show leads to adverse selection.

compensates the high-quality plan for only part of the cost difference. Under perfect risk adjustment $\delta_B = 0$, the uncompensated cost gap $\delta_B \equiv c_{SB} - c_{HB} \geq 0$ becomes zero.

The high-quality plan's average cost is then:

$$AC_B(\Delta\tilde{p}) = \underbrace{c_{HB}}_{\text{healthy enrollee's cost}} + \underbrace{\delta_B \cdot \phi_B(\Delta\tilde{p})}_{\text{gap times sick share of pool}} \quad (2)$$

where $\phi_B(\Delta\tilde{p}) \equiv D_{SB}(\Delta\tilde{p})/D_B(\Delta\tilde{p})$ is the sick share among the high-quality plan's enrollees and $D_{SB}(\Delta\tilde{p})$ is its sick enrollment. Notably, average cost depends on prices and on the incremental subsidy through the net price gap. The high-quality plan's average cost increases with its price, since marginal enrollees are healthier than inframarginal enrollees:

$$\frac{\partial AC_B}{\partial p_B} = (1 - s) \cdot AC'_B(\Delta\tilde{p}) \quad (3)$$

where the slope of average cost in the net price gap is

$$AC'_B(\Delta\tilde{p}) = \underbrace{\tilde{\eta}_B(\Delta\tilde{p})}_{\text{demand semi-elasticity}} \cdot \underbrace{\delta_B \cdot (\phi_B(\Delta\tilde{p}) - \phi_B^m(\Delta\tilde{p}))}_{\text{cost gap between pool and margin } (>0)} \quad (4)$$

and $\phi_B^m(\Delta\tilde{p}) \equiv D'_{SB}(\Delta\tilde{p})/D'_B(\Delta\tilde{p})$ is the sick share among the marginal enrollees, and $\phi_B^m < \phi_B$ under adverse selection ([Appendix A.1](#)). Equation (4) shows that average cost rises faster with price, the more marginal enrollees there are ($\tilde{\eta}_B$), the healthier they are than the pool ($\phi_B - \phi_B^m$), and the larger the uncompensated cost gap between the sick and healthy (δ_B). Under perfect risk adjustment the slope is zero $\delta_B = 0$.

2.3 Competition and Strategic Incentives

I am now ready to derive what adverse selection implies for the high-quality plan's price. First, I consider a perfectly competitive market in which insurers are price takers and prices equal average cost, $p_B = AC_B$ (the usual assumption in the risk-selection literature). Second, I extend the model to allow insurers to set prices and earn markups to maximize profit.

Perfect competition. In a perfectly competitive market, insurers are price-takers and earn zero profits. Enrollees differ in health, and insurers cannot price discriminate. Thus, prices are set to average costs (Einav and Finkelstein, 2011). For the low-quality plan, I assume perfect risk adjustment, $p_N = c_N$ and focus my analysis on the high-quality plan. The high-quality plan's price equates $p_B = AC_B$, which depends on the incremental subsidy.

Under zero-profit pricing, the incremental subsidy affects the high-quality plan's average cost. Given that marginal enrollees are healthier than the plan's existing pool of enrollees, s lowers AC_B . Holding prices fixed ([Appendix A.2](#)):

$$\frac{\partial AC_B}{\partial s} = -\Delta p \cdot AC'_B(\Delta \tilde{p}) = -\delta_B \cdot \underbrace{(\phi_B - \phi_B^m)}_{\substack{\text{health gap infra-} \\ \text{marginal vs. marginal } (>0)}} \cdot \underbrace{\frac{\partial D_B / \partial s}{D_B}}_{\substack{\text{proportional} \\ \text{enrollment gain } (>0)}} < 0 \quad (5)$$

The first term is the uncompensated cost gap for the broad plan. The second term is how much sicker the inframarginal enrollees are compared to the marginal enrollee. The third term captures how many switchers the subsidy attracts per existing enrollee. The second and third term are positive and thus, average cost decreases as the incremental subsidy rate s increases.

In equilibrium, price falls by more than this direct cost effect. Differentiating $p_B^*(s) = AC_B(\Delta \tilde{p})$ at $\Delta \tilde{p} = (1-s)(p_B^*(s) - p_N)$ gives

$$\frac{dp_B^*}{ds} = \frac{\overbrace{\frac{\partial AC_B}{\partial s}}^{\text{direct cost effect } (<0)}}{\underbrace{1 - (1-s) \cdot AC'_B(\Delta \tilde{p})}_{\in (0,1)}} \quad (6)$$

where the denominator lies between zero and one ([Appendix A.2](#)). In other words, the lower price attracts further healthy enrollees, which lowers average cost again, and each dollar of cost saving therefore passes through to price more than one for one. This leads us to the second implication:

Implication #2. *The incremental subsidy lowers the high-quality plan's price by lowering its average cost.*

Implication 2 is related to prior work on risk adjustment as a remedy for selection (Mahoney and Weyl, 2017), which shows a trade-off between mitigating selection and constraining market power. In the present model, the incremental subsidy and risk adjustment are direct substitutes on the cost margin. The take-away is that the selection case for the incremental subsidy is conditional on the effectiveness of risk adjustment.

While perfect competition is a useful benchmark, it abstracts from several real-world features of the health insurance market. Insurers are often price-setters, and earn markups. I turn to this feature below.

Imperfect competition. With imperfect competition, the high-quality insurer sets its price and earns

$$\pi_B = \underbrace{(p_B - AC_B(\Delta\tilde{p}))}_{\text{margin over average cost, } m_B} \cdot D_B(\Delta\tilde{p}) = p_B \cdot D_B(\Delta\tilde{p}) - TC_B(\Delta\tilde{p}) \quad (7)$$

where $TC_B = AC_B \cdot D_B$ is total cost and AC_B depends on the price through the net price gap $\Delta\tilde{p} = (1-s)(p_B - p_N)$ and the composition of enrollees in (2). The margin over average cost, $m_B \equiv p_B - AC_B$, is profit per enrollee. Writing $MC_B \equiv dTC_B/dD_B$ for the marginal enrollee's cost, the first-order condition is the Lerner condition on the marginal enrollee, $p_B - MC_B = 1/\eta_B$, where (See [Appendix A.3](#) and [Appendix A.1](#) for derivations):

$$MC_B(\Delta\tilde{p}) = \underbrace{c_{HB}}_{\text{healthy enrollee's cost}} + \underbrace{\delta_B \cdot \phi_B^m(\Delta\tilde{p})}_{\text{gap times marginal sick share}} \quad (8)$$

is the healthy enrollee's cost plus the uncompensated gap times the probability that the marginal enrollee is sick. Subtracting marginal cost from average cost yields the "selection wedge", $AC_B - MC_B = \delta_B \cdot (\phi_B - \phi_B^m) > 0$. Subtracting the selection wedge from both sides of the Lerner condition $p_B - MC_B = 1/\eta_B$ gives the margin over average cost:

$$\underbrace{m_B}_{\text{margin over average cost}} = \underbrace{\frac{1}{\eta_B}}_{\text{Lerner markup}} - \underbrace{(AC_B - MC_B)}_{\text{selection wedge}} \quad (9)$$

The Lerner term is the classic markup term. Adverse selection occurs if marginal enrollees are unhealthier (costlier) than the inframarginal average enrollee.

The incremental subsidy affects the high-quality plan's price through two channels. The cost channel is that of Implication 2: a higher s narrows the gap, the marginal enrollee is healthier, and MC_B falls. The markup channel is new to imperfect competition. A \$1 increase in p_B affects enrollees as $1-s$ dollars. This implies that the semi-elasticity for the insurer is $\eta_B = (1-s) \cdot \tilde{\eta}_B(\Delta\tilde{p})$, and the Lerner markup $1/\eta_B = 1/((1-s) \tilde{\eta}_B(\Delta\tilde{p}))$ rises in s . Differentiating the Lerner condition in s at fixed p_N yields

$$\frac{dp_B^*}{ds} = \frac{\overbrace{\delta_B \cdot \frac{\partial \phi_B^m}{\partial s}}^{\text{cost channel } (<0)} + \overbrace{\frac{\partial(1/\eta_B)}{\partial s}}^{\text{markup channel } (>0)}}{\underbrace{1 - \frac{\partial MC_B}{\partial p_B} - \frac{\partial(1/\eta_B)}{\partial p_B}}_{(>0)}} \quad (10)$$

where the partial derivatives hold p_B fixed. The denominator is positive (see [Appendix A.3](#)). The cost channel is negative under adverse selection and zero under perfect risk adjustment. The markup channel is positive for $s \in [0, 1]$ ([Appendix A.3](#)), and this brings us to the third implication.

Implication #3. *Under imperfect competition the incremental subsidy lowers the high-quality plan's price through the cost channel and raises it through the markup channel. Price falls if and only if the cost channel outweighs the markup channel.*

Which channel dominates depends on risk adjustment. If risk adjustment perfectly compensates the cost channel, the incremental subsidy introduces a purely distortionary competitive effect.

2.4 Welfare Objectives

I now turn to the social planner's problem to understand what type of welfare function would rationalize the incremental subsidy. The goal is to nest the standard welfarist planners and an additional term capturing the direct preference for access in one objective function, and to compare the two subsidy instruments at fixed prices. An enrollee in the low-quality plan pays $p_N - \tau$ toward a plan that is priced at p_N , and an enrollee in the high-quality plan pays $p_B - \tau - s \cdot \Delta p$ toward a plan priced at p_B . Summing over enrollees, government cost is:

$$G(\tau, s; \mathbf{p}) = \underbrace{\tau(D_B + D_N)}_{\text{level subsidy}} + \underbrace{s \cdot \Delta p \cdot D_B(\Delta \tilde{p})}_{\text{incremental subsidy}} \quad (11)$$

Notice that the level subsidy's cost is fixed by the mass of enrollees—which is $D_B + D_N \equiv 1$. By contrast, the incremental subsidy's cost depends on both the price gap and enrollment in the high-quality plan, each of which responds to s (at fixed prices through the switchers, and in equilibrium through the gap).

Welfare is a consumer welfare standard⁵ which consists of enrollees' consumer surplus, plus an access premium, less the fiscal cost, all measured in dollars of public funds:

$$W(\tau, s; \mathbf{p}) = \underbrace{\omega \cdot CS}_{\text{consumer surplus}} + \underbrace{\mu \cdot D_B(\Delta \tilde{p})}_{\text{access premium}} - \underbrace{\lambda \cdot G(\tau, s; \mathbf{p})}_{\text{fiscal cost}} \quad (12)$$

⁵I exclude insurer profits from welfare given that Connector implements multiple policies to constrain profits. First, Insurers in the Massachusetts individual market must meet the ACA's medical loss ratio floor, returning premiums to enrollees if administrative cost and profit exceed 20 percent of revenue. Second, ConnectorCare insurers enter through the Connector's Seal of Approval procurement, in which the Connector reviews each bid and can reject rates it finds too high. Third, the Division of Insurance reviews the same rates separately. A planner with these policies unlikely accounts for insurer profits in welfare objectives.

where CS is enrollee surplus, $\omega \geq 1$ is the social marginal welfare weight on enrollees, i.e., the value of a dollar to them relative to a dollar of public funds (Saez and Stantcheva, 2016), $\mu \geq 0$ is the planner's valuation in dollars, of an enrollee holding the high-quality plan, and $\lambda \geq 1$ is the marginal cost of public funds.

This objective nests three planners as specific cases. An *efficiency* planner values a dollar the same across all enrollees ($\omega = 1$) and places no value on access ($\mu = 0$). A *utilitarian* or *inequality-averse* planner values a dollar to enrollees more than a dollar of public funds ($\omega > 1$), and is indifferent to which plan they enroll in ($\mu = 0$). A *Tobin* planner values the poor's enrollment in the high-quality plan itself. ($\omega = 1$ and $\mu > 0$). In principle, one could specify alternate welfare functions of different functional forms. I mainly choose this linear approach for simplicity and transparency, acknowledging that the social planner may exhibit more complex preferences.

Now, consider first the level subsidy. It transfers every enrollee a dollar, valued at ω , at a fiscal cost of λ per dollar:

$$\frac{\partial W}{\partial \tau} = \omega - \lambda \quad (13)$$

Because τ affects neither the gap nor any cost, the level subsidy does not affect prices. (13) is the total effect of the level subsidy, at fixed prices and in equilibrium, $dW/d\tau = \partial W/\partial \tau$, since $\mathbf{p}^*$ does not depend on τ . For the incremental subsidy, at fixed prices it transfers Δp to each enrollee already in the high-quality plan, affords access for marginal enrollees who switch into the broad plan:

$$\frac{\partial W}{\partial s} = \underbrace{(\omega - \lambda) \cdot \Delta p \cdot D_B(\Delta \tilde{p})}_{\text{(a) transfer to enrollees already in B}} + \underbrace{\mu \cdot \frac{\partial D_B}{\partial s}}_{\text{(b) access value } (>0)} - \underbrace{\lambda \cdot s \cdot \Delta p \cdot \frac{\partial D_B}{\partial s}}_{\text{(c) subsidy on marginal enrollees } (>0)} \quad (14)$$

Term (a) is similar to the level subsidy term, except that it now scales with the price gap and enrollment in the high-quality plan. Term (b) is the access value of those who switch into the high-quality plan as a result of raising the incremental subsidies. Term (c) is the state's additional cost on the marginal enrollees. Comparing (13) and (14) per dollar of fiscal cost (Appendix A.4) gives the fourth implication:

Implication #4. *At fixed prices, the three planners rank the instruments as follows.*

- (i) *An efficiency planner subsidizes nothing: a dollar transferred is worth a dollar to enrollees and costs $\lambda \geq 1$.*
- (ii) *An inequality-averse planner prefers the level subsidy to the incremental subsidy: both transfer the same amount to consumer while the incremental subsidy has an additional cost on marginal enrollees.*

(iii) *A Tobin planner uses the incremental subsidy once the access value of the switchers exceeds the fiscal loss on enrollees already in the high-quality plan.*

It is important to note that Implication 4 holds prices fixed. In equilibrium, prices adjust, and from Implication 3 the incremental subsidy raises the high-quality plan's markup. The markup is a transfer from enrollees and the state to the insurer, and it is a loss to consumer welfare.⁶ Allowing prices to respond further strengthens points (i) and (ii) of Implication 4. The efficiency and inequality-averse planner would further prefer the level subsidy, as it does not affect insurer pricing decisions, and avoids competition distortions.

Parallel: cash versus in-kind transfers. The first part of Implication 4 is closely related to the classic case for cash over in-kind transfers (Currie and Gahvari, 2008). In that argument, a cash transfer allows the recipient to purchase any option—cash dominates unless the planner values the good itself or to address market failures (Tobin, 1970; Nichols and Zeckhauser, 1982). Both of those motives may exist in my setting—the incremental subsidies might address adverse selection, while a social planner that cares about the poor's enrollment in the broad plan may also adopt the incremental subsidy.

⁶If profit entered the welfare function, the markup would be a transfer to insurers rather than a pure loss. The incremental subsidy would also be justified with the margin the insurer earns on each switcher. Points (i) and (ii) would then hold only if that margin fell short of the state's payment per switcher.

3 Setting and Data

3.1 Massachusetts ConnectorCare Program

ConnectorCare is the subsidized segment of the Massachusetts Health Connector—the state’s insurance marketplace created under the federal Affordable Care Act (ACA). ConnectorCare opened in 2015 and replaced the state’s pre-ACA exchange established under a 2006 state reform. It serves enrollees with incomes up to 300 percent of the federal poverty line and covered more than 290,000 members as of 2026. The state layers its own premium subsidies on top of federal subsidies, and finances them from a dedicated state trust fund.

In ConnectorCare, enrollees choose among a small set of insurers that differ in hospital networks. Three insurers cover 94 percent of enrollment in my sample period. To preserve anonymity, I will refer to these plans as insurers N1, N2 and B, with B denoting the broad plan. Insurer B offers a broad network of academic hospitals and includes “star” hospitals (Shepard, 2022). Insurers N1 and N2 sell narrower-hospital network products. Insurer B’s network includes 68 of the 72 hospitals on the Connector’s 2015 network list and all six academic medical centers, against 57 and 50 hospitals and four and three academic medical centers for Insurers N1 and N2. The eight hospitals of the state’s largest system are exclusive in Insurer B’s network ([Table B1](#)).

One key feature of the ACA is that it restricts insurance companies to offer health plans with identical cost-sharing structures—such as set deductibles, copayments, and out-of-pocket maximums—within specific metal tiers. Within ConnectorCare, consumers only choose from ‘Silver’ plans. This is important, as the main quality differentiation lies on the network breadth of specific plans.

Each insurer files one statewide base gross premium with the Massachusetts Division of Insurance every year, and this single filed premium spans across both the subsidized and the unsubsidized segment of the marketplace. This base rate is then adjusted using a geographic rating factor, and a tobacco rating factor that is exogenous to the insurer pricing decision. The state then layers subsidies based on different income groups. These groups consist of those below the federal poverty line (FPL) through 250–300 percent of FPL. Enrollees may switch insurers each year at open enrollment during January. Special enrollment periods are also available for those who experience qualifying life events. Insurers choose whether to participate in ConnectorCare annually at the start of each year. Insurers also choose whether to participate in ConnectorCare in addition to the unsubsidized segment of the marketplace.

In practice, ConnectorCare uses a strong form of managed competition. Insurers file gross premiums with the Division of Insurance (DOI) reviews the filing against projected costs and medical loss ratio requirements—percentage of health insurance premium dollars an insurance company spends on direct medical care and quality improvements rather than

administrative costs and profits. This implies that insurers may not compete in a canonical Bertrand sense. Thus, I use ConnectorCare as the empirical setting for the subsidy design question in this paper rather than an accurate model of ConnectorCare.

Federal subsidies affect ConnectorCare enrollees and insurers through two separate policies. The Advance Premium Tax Credit (APTC) reduces enrollees' premiums by capping it at a fixed percentage of household income, with the federal government covering the difference between that amount and a benchmark. This benchmark is the second-cheapest Silver plan premium in the enrollees' choice set. The second policy is Cost-Sharing Reductions (CSR), which affects insurers' revenues. ConnectorCare plans are highly generous silver plans with minimal deductibles. This policy reimburses insurers for offering relatively more generous cost sharing features. Both are set at the federal level and are exogenous.

3.2 State Subsidy Policy and Removal of Incremental Subsidies in 2017

Connector implements both a *level* and an *incremental* subsidy across different income groups. For the level subsidy, Connector sets an "affordability schedule", which is the maximum monthly premium considered "affordable" for an individual or family at a given income. For example in 2016, the affordability schedule for a below poverty line single enrollees is 0 percent of monthly income—implying that this enrollee would receive health insurance for free. This schedule increases with income from 2.9 percent to 5 percent of income among those in the 250–300 percent FPL group.

In addition to the affordability schedule, Connector also decides how much of the gap between the broad and narrower network plans to subsidize. Connector's policy documents refer to this as "premium smoothing", but it resembles an incremental subsidy for high-quality plans. The purpose of the policy is "to ensure that the lowest income enrollees have *meaningful insurer choice*". This incremental subsidy is progressive, as wealthier enrollees have a smaller incremental subsidy relative to poor enrollees.

2017 Removal of Incremental Subsidies. [Figure 2](#) shows a four panel figure depicting how the removal of incremental subsidies in 2017 affected gross and net premiums for Insurer B and Insurer N1 (see [Figure B1](#) for a decomposition of level and incremental subsidies by income groups). From Panel A, the gross price gap between the broad and narrow plan was \$51 in 2016. When tracking the net prices across income groups between the broad and narrow plan in Panel B, the gap widens as income increases. For the below poverty income group, the net premium for the broad plan and narrow plan are both zero respectively. This widens as income increases, with \$167 and \$124 at the 250–300 percent FPL group.

In 2017, when Connector removed the incremental subsidies, the gap between Insurer B and Insurer N1 is now constant across income groups (Panel D). Notably, the gap across income groups varies between \$117 and \$122, which shows that ConnectorCare enrollees now face the full price gap between Insurer B and Insurer N1.


One noteworthy mention is that Insurer N1 and Insurer N2 are often the lowest gross premium carriers across regions. In 2016, Insurer N1 is the lowest across most regions, while this flipped to Insurer N2 in 2017. Insurer N1 and N2 have similar narrow networks, and therefore compete more intensely on price. To prevent either insurers from undercutting and leading to a death spiral where price sensitive consumers would pick the cheapest insurer, Connector provides a more generous incremental subsidy for the second-lowest priced insurer. For example, if gross premiums for Insurer N1 were \$300, while Insurer N2 was \$305, Connector would provide a larger incremental subsidy for Insurer N2 to ensure that prices are equal for the two insurers. This explains why the net premium gap between Insurer B and Insurer N1 often exceeds the gross premium gap in Panel C. I do not model this feature in the structural model. The paper instead focuses on the margin between the high-quality and low-quality plans. Modeling this rule where Connector equalized the premiums for the two cheapest plans would complicate the pricing game and potentially lead to equilibria non-existence.

How the subsidy is set in practice. Subsidy rates follow an annual timeline: the Connector’s Board votes the affordability schedule in the spring, insurers file rates in May, and the Board decides the ConnectorCare premium structure for the coming year in September. Importantly, Connector can adjust level subsidies and incremental subsidies after observing insurer pricing decisions. Connector can also return filed rates to insurers and negotiate before the September vote. Premiums under ConnectorCare therefore emerge from this exchange between the state and insurers rather than from a single round of simultaneous bids. The structural model abstracts from this sequence.

3.3 Data

I use individual-month enrollment data covering the universe of ConnectorCare enrollment between 2015 and 2019 from Connector’s administrative records. The data records each enrollee’s insurer, plan, income group, billed premium, age, gender, family structure, and ZIP code. I observe every subsidy component for each enrollee. This includes the rating area-risk adjusted gross premium, the federal premium tax credit, the state level and incremental subsidies, and the net premium. Insurer availability varies by region, and I observe each year-region menu. I restrict the sample to enrollees with individual coverage (single, unmarried), and exclude the American Indian and Alaska Native plan variants.

Figure 2: Gross and Net Premium Gaps for Insurer B (Broad) versus Insurer N1 (Narrow)



Notes: Panels A and C plot the enrollment-weighted gross premiums of the broad insurer and the first narrow insurer. Panels B and D plot the net premiums of both insurers in 2016 and 2017, respectively.

The panel is measured at the enrollee-year level and contains 699,596 enrollee-years over 2016 to 2019.

Each enrollee is classified into three categories. A ‘new enrollee’ is an enrollee who was never observed in ConnectorCare in any earlier year of the data. A continuing enrollee was enrolled in the prior year and remains in the program the following year. Of the 699,596 enrollee-years in the panel, 258,430 are new enrollees.

I use claims data from the Massachusetts All-Payer Claims Database (APCD) to observe ConnectorCare enrollee-episode costs and healthcare utilization. I take an annual January snapshot with one observation per member-year, restricted to the three insurers N1, N2 and B. The resulting panel contains 632,424 person-years over 2016 to 2019. APCD assigns its own person identifiers, and this sample cannot be merged to the Connector enrollment data at the individual level.

3.4 Data Construction

Connector enrollment panel. The enrollment analysis uses all enrollee-years in the Connector panel. This comprises new enrollees, continuing enrollees, and re-enrollees. Analyses built around the 2017 reform restrict to region-group cells where the broad-network insurer is offered in years between 2016 and 2019. Demand estimation uses only the new enrollees in this panel rather than the full set of all enrollees.

APCD sample. The selection analysis uses the APCD panel, where the outcome is the enrollee’s costs and risk score by insurer, and the care analysis uses APCD claims, where the outcomes are the broad-network share of care, office visits, and utilization counts. The APCD data further provides micro moments for the demand estimation. The demand sample comprises 219,970 new enrollees. APCD records income in three broader income categories rather than the five income groups. Given this, I split each member in the ‘100–200’ and ‘200–300’ percent FPL in proportion to their shares of that region-year’s enrollment in the Connector data. In addition, the APCD over-represents the broad-network insurer relative to true enrollment. Thus, I weight members within each region, year, and income range to match the insurer shares of new enrollees in the Connector data. The supply side uses APCD claims to measure each insurer’s insurer-paid cost per member-month, net of risk adjustment transfers and cost-sharing reductions.

Summary statistics. [Table 1](#) describes the two data sources. Enrollment skews toward the middle of the eligible income range, with the 150–200 percent group alone contributing 231,379 of the 699,596 enrollee-years. Lower income groups face larger shares of new enrollees (44.8 percent of enrollee-years below the poverty line). This share falls monotonically with income to 30.6 percent in the 250–300 group. In addition, insurer choice varies by in

Table 1: Summary Statistics

	<100% FPL	100–150% FPL	150–200% FPL	200–250% FPL	250–300% FPL
Panel A. Connector Enrollment Panel (2016–2019)					
Counts					
Enrollee-years	68,933	129,412	231,379	163,888	105,984
New enrollees (%)	44.8	41.6	37.7	33.0	30.6
Insurer chosen (%)					
Insurer N1	35.6	44.8	50.7	53.0	54.8
Insurer N2	46.8	40.8	33.1	31.2	28.2
Insurer B	14.7	9.2	10.1	9.9	10.9
Other insurers	2.9	5.2	6.1	6.0	6.1
Panel B. APCD Sample (2016–2019)					
	<100% FPL	100–200% FPL	200–300% FPL		
Counts					
Enrollee-years	50,053	311,271	271,100		
Demand sample, new enrollees	24,137	115,902	79,931		
Enrollee characteristics					
Age (years)	44.1	44.1	45.0		
Female (%)	60.6	57.2	54.4		
HCC risk score	0.89	0.90	0.95		
Dist. B-exclusive hospitals (mi)	13.7	19.4	20.4		

Notes: The table reports summary statistics for the two data sources over 2016 to 2019. Panel A comprises ConnectorCare enrollee-years from Health Connector administrative records, restricted to single, unmarried enrollees; columns correspond to the five subsidy income groups. Panel B comprises all enrollee-years in the Massachusetts All-Payer Claims Database sample. The demand sample row counts the new enrollees used in demand estimation. Risk score is the winsorized HCC risk score. Dist. ‘B-exclusive hospital’ is the enrollee’s distance to an Insurer B-exclusive hospital.

income systematically. Insurer N1’s share rises from 35.6 percent below the poverty line to 54.8 percent at 250–300 percent, while Insurer N2’s falls from 46.8 percent to 28.2 percent over the same range. This shows that the cheaper narrow insurer draws disproportionately most of the poor. Insurer B’s share is smaller from 14.7 percent below the poverty line to 10.9 percent.


Panel B describes the 632,424 enrollee-years in the APCD sample. The sample varies minimally across income groups. Age ranges from 44.1 years below poverty and at 100–200 percent of FPL to 45.0 at 200–300 percent. For risk score, from 0.89 to 0.95. Enrollees below the poverty line live closer to the broad network’s hospitals than higher-income enrollees—14 miles against 19 and 20 at 100–200 and 200–300 percent of FPL.

4 Reduced Form: Subsidy Removal on Enrollment and Selection

4.1 Descriptive Evidence

Among all ConnectorCare enrollees, below poverty enrollees faced the largest increase in Insurer B's price in 2017. A below poverty enrollee choosing Insurer B would have paid zero in 2016, while in 2017, a below poverty enrollee faced a \$117 price difference for Insurer B. Other income groups faced similar increases in Insurer B's price, though not as stark as a below poverty enrollee. In this section, I use the removal of the incremental subsidy as a natural experiment to quantify its effects on enrollment and health composition of enrollees in Insurer B.

Figure 3: Broad-Network Enrollment by Income Group, 2015–2019



Notes: Panel A plots the share of all enrollees by year within income groups; the dashed light gray line marks the 2017 removal of the incremental subsidy policy. Panel B plots the same shares across income group, by year.

Figure 3 shows the raw pattern on enrollment. Panel A plots Insurer B's share of all enrollees by income group over 2015 to 2019. Before the subsidy removal, Insurer B enrolled roughly half of below-poverty enrollees and about a quarter of every other group. After the subsidy removal, the share of below poverty enrollees fell substantially from approximately 50 percent to less than 5 percent. While not as large as the below poverty enrollees, this decline occurs across every income group. Notably, the ranking of income groups inverts. Below poverty enrollees—who had the highest Insurer B share of any group before the subsidy removal—have the lowest share of any group by 2019.

Panel B plots the same shares across income, by specific years. The figure shows the income gradient reversing after the subsidy removal. In 2016, Insurer B's share is highest among enrollees below poverty (51 percent) and roughly flat across the remaining groups (23 to 25 percent). By 2017, the gradient inverts: the share now rises with income, from

roughly 5 percent below poverty to 15 percent at 250–300 percent of poverty. The 2018 and 2019 lines preserve this upward slope while shifting downward, as Insurer B’s share continues to fall across years.

4.2 Empirical Strategy

Enrollment. The removal of the incremental subsidy increased the net price of Insurer B for every enrollee. I estimate the effects of the subsidy removal in a differences-in-differences design:

$$\text{Insurer B}_{igrt} = \sum_{g \neq 250-300} \beta_g (\mathbf{1}\{G_i = g\} \times \text{Post-2017}_t) + \gamma_{gr} + \tau_t + \varepsilon_{igrt} \quad (15)$$

where Insurer B_{igrt} is an indicator for choosing Insurer B, Post-2017_t equals one post-2017. Income groups g are the five income groups, while I specify the 250–300 percent group as the control group. γ_{gr} are income group by region fixed effects, and τ_t are year fixed effects. The group by region effects absorb group-region specific taste for Insurer B, while the year effects absorb statewide shocks. I estimate the specification separately for new enrollees and for continuing enrollees. Re-enrollees make up a small share of the sample, and I omit them. Standard errors are clustered at the income group by region by year level.

The identifying assumption is parallel trends—absent the subsidy removal, income groups within a region would have trended similarly in their Insurer B shares. Beyond the descriptive parallel trends across above poverty income groups in Panel A of [Figure 3](#), there are two other facts that support this identifying assumption. First, insurer menus and network breadth did not change between 2016 and 2017. Second, the removal did not cause enrollees to exit ConnectorCare. Enrollment rose in every income group in 2017, by 18 percent below poverty and by 4 to 21 percent above ([Appendix Figure C1](#) and [Table C1](#)). The share outcome therefore reflects switching among plans rather than enrollee entry or exit.

The subsidy removal occurred simultaneously when Insurer B raised its gross premiums. Gross premium increases are absorbed by year fixed effects, while the removal differentially affected the net premiums across income groups. Coefficients β_g thus quantifies the incremental subsidies effects on net premiums. The remaining threat is a shock in 2017 that shifted below-poverty enrollees’ taste for Insurer B relative to the 250–300 percent group for reasons other than price. However, ConnectorCare plans are standardized Silver plans with identical cost sharing within an income group, and insurer menus and networks did not change in 2017.

Selection. To understand if the removal of the subsidy led to healthier and cheaper enrollees leaving Insurer B, I estimate a triple difference-in-differences design. Using the

APCD claims panel data, for member i enrolled in insurer j in income group g , region r , and year t , I estimate

$$y_{igrt} = \sum_g \delta_g (\text{Insurer } B_{ij} \times \mathbf{1}\{G_i = g\} \times \text{Post-2017}_t) + \sum_g \lambda_g (\text{Insurer } B_{ij} \times \mathbf{1}\{G_i = g\}) + \sum_g \phi_g (\mathbf{1}\{G_i = g\} \times \text{Post-2017}_t) + \gamma_{gr} + \tau_t + \varepsilon_{igrt} \quad (16)$$

I specify three outcomes. The first is the enrollee's insurer cost per member-month. The second is cost net of a risk-adjustment transfer, $c_{it} - (\frac{h_{it}}{\bar{h}} - 1) \cdot \bar{P}$, where $\frac{h_{it}}{\bar{h}}$ is the enrollee's HCC risk score divided by the market mean in year t and $\bar{P}$ is the enrollment-weighted mean gross premium. The third outcome is the HCC risk score, which I show in Appendix Table C2. The claims panel uses a higher-level income group such that g consists of three groups: below poverty, 100–200 percent, and 200–300 percent. δ_g is the coefficient of interest. It represents a within-group difference-in-differences across insurers, and measures the post-2017 change in Insurer B's enrollees' outcome relative to the narrow insurers', for income group g . Comparing insurers within a group and year controls for group-level cost and health trends. The identifying assumption is that, absent the subsidy removal, the average cost of Insurer B's enrollees and the narrow insurers' enrollees would have trended similarly within each income group. Standard errors are clustered at the income group by region level.

4.3 Results

Enrollment. Figure 4 plots the β_g coefficients while Table 2 tabulates the coefficients. Panel A plots the income-group $\times$ post-2017 coefficients from Equation 15, separately for new and continuing enrollees. Each group's 2016 Insurer B share of new enrollees are bracketed below the FPL categories to give a sense of the magnitude. Among the below poverty new enrollees, enrollment in Insurer B fell sharply relative to the 250–300 percent FPL control group. This decline is estimated at -0.384 relative to a mean of 0.511, showing that the incremental subsidy removal led many to switch to Insurer N1 or N2 (Column 1, Table 2). The effect declines monotonically with income and is statistically indistinguishable from zero by 200–250 percent of poverty. Continuing enrollees trace the same downward-sloping pattern at uniformly smaller magnitude, consistent with switching costs dampening demand response.

Selection. Columns (3) and (4) of Table 2 tabulate the post-2017 change in Insurer B enrollees' cost per member-month relative to the narrow insurers. The results show that costs increase for every income group after the removal of the incremental subsidies. Insurer

Table 2: Effects of Removing Subsidies on Enrollment and Selection

Panel A. Enrollment in Insurer B	Binary		Panel B. Selection into Insurer B	\$ per member-month	
	New (1)	Continuing (2)		Cost (3)	Cost net of transfer (4)
Below 100% FPL $\times$ Post-2017	-0.384*** (0.050)	-0.285*** (0.041)	Below 100% FPL $\times$ Post-2017	866.8*** (184.3)	118.7*** (40.9)
100–150% FPL $\times$ Post-2017	-0.083** (0.031)	-0.043 (0.035)	100–200% FPL $\times$ Post-2017	170.6*** (24.9)	55.4*** (17.5)
150–200% FPL $\times$ Post-2017	-0.061** (0.027)	-0.027 (0.031)	200–300% FPL $\times$ Post-2017	167.9*** (36.7)	66.8*** (17.2)
200–250% FPL $\times$ Post-2017	-0.004 (0.022)	-0.008 (0.031)			
2016 Insurer B share:			2016 Insurer B mean:		
Below 100% FPL	0.511	0.508	Below 100% FPL	386	356
100–150% FPL	0.233	0.267	100–200% FPL	496	430
150–200% FPL	0.223	0.265	200–300% FPL	553	442
200–250% FPL	0.170	0.257			
250–300% FPL	0.184	0.257			
Income group $\times$ region FE	Yes	Yes	Income group $\times$ region FE	Yes	Yes
Year FE	Yes	Yes	Year FE	Yes	Yes
			Insurer $\times$ income group levels	Yes	Yes
Observations	219,228	369,146	Observations	601,866	601,866

Notes: Panel A tabulates the subsidy removal's effect on choosing Insurer B among new and continuing enrollees. Columns (1) and (2) tabulate equation (15). Panel B tabulates the triple difference-in-differences of equation (16). Column (3) is the insurer's cost per member-month. Column (4) nets out a risk-adjustment transfers. Both panels restrict to region-group cells where Insurer B is offered in both 2016 and 2019. Panel B is weighted by member-months. All columns include income-group-region and year fixed effects. Columns (3) and (4) additionally include the insurer $\times$ income group level terms of equation (16). Standard errors, in parentheses, are clustered at the income-group $\times$ region $\times$ year level in Panel A and at the income-group $\times$ region level in Panel B. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.


B's cost per enrollee rose by \$171 per member-month relative to the narrow insurers' at 100–200 percent of poverty and by \$168 at 200–300 percent. For below poverty enrollees, the increase is \$867 on a 2016 mean of \$386. The enrollees who left Insurer B were therefore disproportionately cheap and healthier.

However, risk adjustment compensates most of the adverse selection effects. Column (4) shows that the transfer covers 86 percent of the cost increase for below poverty enrollees and similarly for above poverty enrollees, leaving Insurer B with only a cost increase of \$119, \$55, and \$67 per member-month. This shows that while that adverse selection against the Insurer B is substantial, risk adjustment mostly offsets the justification for using incremental subsidies to mitigate selection.

4.4 Are the Poor Disproportionately Affected?

The results thus far show that removing the incremental subsidies led to below poverty enrollees switching out of Insurer B. However, if below poverty enrollees simply substitute away from Insurer B's hospital networks to Insurer N1/N2, this implies that losses may be mitigated and there might not be an efficiency rationale to subsidize Insurer B for this income group.

Figure 4: Effects of Removing Subsidies on Enrollment and Selection



Notes: Panel A plots the income group $\times$ post-2017 coefficients from equation (15), estimated separately for new enrollees and continuing enrollees, with 250–300 percent FPL enrollees as the control group. The sample is restricted to region–group cells where Insurer B is offered in both 2016 and 2019. Bracketed values below each income group show the group’s 2016 Insurer B share among new enrollees. Panel B plots the post-2017 change in Insurer B enrollees’ cost per member-month relative to the narrow insurers’, by income group, from the triple difference-in-differences of equation (16). Bars are 95 percent confidence intervals.

To estimate the effects of Insurer B on below poverty enrollees’ health care use, the ideal experiment would assign the loss of Insurer B at random among below-poverty enrollees in 2016. In reality, those who left Insurer B were more price sensitive and less attached to Insurer B’s covered hospitals. A comparison of switchers against enrollees who kept Insurer B would therefore confound the loss of the network with factors that might also affect health care utilization.


Instead, I compare switchers against below-poverty enrollees already in Insurer N1 or N2 in 2016. This group faced the same narrow networks and the same post-2017 shocks to care without a transition to a new insurance plan. The identifying assumption is that, absent the switch, switchers’ care would have followed the control group’s trend. For enrollee i in year t , I estimate

$$y_{it} = \sum_{t' \neq 2016} \beta_{t'} (\text{Switcher}_i \times \mathbf{1}\{t = t'\}) + \alpha_i + \tau_t + \varepsilon_{it} \quad (17)$$

where Switcher_i indicates the treated group, the control group is the omitted category, 2016 is the reference year, and α_i and τ_t are person and year fixed effects; the person effects absorb time-invariant differences in care use. The coefficients $\beta_{t'}$ trace switchers’ care relative to their own 2016 level and to the control group’s trajectory. The outcomes are

office-visit claims, claims at Insurer B-exclusive hospitals, inpatient claims, and emergency department claims, each per member-month (Figure 5). Standard errors are clustered at the region by year level.

Figure 5: Care Disruption Among Enrollees Priced Out of Insurer B



Notes: Each panel plots the event-study coefficients β_t of equation (17) for below-poverty enrollees who held Insurer B in 2016 and left it after the removal of the incremental subsidy, relative to below-poverty enrollees observed in Insurer N1 or N2 in 2016, with 2016 as the reference year. Outcomes are claim counts per member-month. Shaded bands are 95 percent confidence intervals from standard errors clustered at the region by year level. The vertical line marks the removal of the incremental subsidy. Switchers' 2016 means are 0.843 office-visit claims, 0.266 claims at Insurer B-exclusive hospitals, 0.011 inpatient claims, and 0.229 emergency department claims per member-month.

Figure 5 plots the results. Switchers' office-visit claims fall by 0.354 per month in 2017 against a 2016 mean of 0.843, and the gap widens to 0.441 by 2019. Their claims at Insurer B-exclusive hospitals fall by 0.262 per month in 2017 against a 2016 mean of 0.266. Inpatient claims do not appear to be affected by switching away from Insurer B. Emergency department claims fall by 0.101 per month in 2017 and return to the control trajectory by 2019.

This decline in office visits and utilization in Insurer B-exclusive hospitals may partially be explained by below poverty enrollees' location. Below poverty enrollees tend to live closer (13.7 miles) to Insurer B-exclusive hospitals compared to above poverty enrollees

(19.7 and 20.5 miles respectively for 100–200 and 200–300 percent FPL groups). This implies that below poverty enrollees are more likely to face health care disruptions upon losing Insurer B’s hospitals networks [Table 1](#). Overall, the below poverty enrollees who left Insurer B faced disruptions to their health care, concentrated in primary care and in Insurer B-exclusive hospitals with effects persisting post-2017.

5 Structural Model and Estimation

5.1 Overview

In this section, I describe and estimate a structural model that comprises of demand for insurance plans, and costs from enrollee-level claims. Insurers compete by setting gross premium rates after observing the level and incremental subsidy rates set by the social planner. I use this model in [Section 6](#) to understand how incremental subsidies affect insurers' strategic incentives, and simulate counterfactual subsidies that would improve welfare.

5.2 Demand

Specification. I model enrollees' plan choices as a function of net premiums, health, and distance to Insurer B exclusive hospitals. I restrict my sample to only new enrollees: consumers have multiple "choice instances" during enrollment periods when they can choose different plans. New enrollees are enrollees that newly enrolls in the marketplace—in contrast to continuing enrollees that have the option to switch plans during the annual open enrollment period. My sample period ranges between 2016 and 2019. A market is a region, year, and income group, such that every enrollee in a market chooses from the same insurers at the same net premiums. Enrollee i in income group g , region r , and year t chooses insurer j to maximize

$$\begin{aligned}
 u_{ijgrt} = & \underbrace{\alpha_g \tilde{p}_{jgrt}}_{\text{net premium, by income}} + \underbrace{\beta \mathbf{1}\{b_{jt} \psi_{jrt} = \min_k b_{kt} \psi_{krt}\}}_{\text{lowest-price insurer in the region}} + \underbrace{\zeta_{jr} + \zeta_{jg}}_{\text{insurer quality, by region and income}} + \underbrace{\theta_h h_i \mathbf{1}\{j = B\}}_{\text{health sorting into Insurer B}} \\
 & + \underbrace{\theta_d d_i \mathbf{1}\{j = B\}}_{\text{distance to Insurer B exclusive hospital}} + \underbrace{\zeta_{i,n(j)} + (1 - \rho) \varepsilon_{ijgrt}}_{\text{correlated taste nest, correlation } \rho}
 \end{aligned} \tag{18}$$

where $\tilde{p}_{jgrt}$ is the net premium after level and incremental subsidies are implemented for income group g for insurer j , and $b_{jt} \psi_{jrt}$ is insurer j 's gross premium in region r : the statewide base rate b_{jt} scaled by the region's rating factor ψ_{jrt} . ζ_{jr} and ζ_{jg} are unobserved insurer quality terms that vary by region and by income group, h_i is the enrollee's risk score, and d_i is distance to an Insurer B-exclusive hospital. The health and distance terms enter only through the Insurer B-exclusive hospital choice.

I add an additional 'lowest-price insurer in the region' term to the utility function. Without this term, the model fit for equilibrium premiums far exceeds the observed data. This occurs given that enrollees choose the lowest priced insurer more than a standalong

price sensitivity parameter predicts. Insurer N1 and N2 both make up close to 80 percent of market share, and when Insurer N1 (or N2) becomes the cheapest insurer, the cheaper insurer receives a much higher market share.

The error term follows a nested structure. Insurer N1 and Insurer N2 form a correlated taste nest, and $\zeta_{i,n(j)}$ is a taste draw common to insurers in nest $n(j)$, with ρ governing its weight and ε type-one extreme value. I choose Insurer N1 and Insurer N2 as the correlated nest because the data show that enrollees substitute between the two narrow-network insurers far more frequently than between either and Insurer B. I do not model un-insurance, as the incremental subsidy policy mostly affects the intensive margin of choice rather than the extensive margin of participating in insurance markets.

Estimation and identification. I estimate demand using BLP with micro moments (Berry et al., 1995, 2004; Conlon and Gortmaker, 2020). Markets consists of 160 region, year, and income group cells. The demand model implicitly captures two sources of variation. The first is the 2017 removal of the incremental subsidy and the second is the regional variation in which insurer is cheapest. From the claims data, I use two additional moments: the mean risk score and the mean distance among enrollees who choose Insurer B.

Each parameter is identified by a distinct source of variation. The nest correlation ρ is identified by how enrollees choose Insurer N1 over Insurer N2 in response to their net premium gap across years. The premium on the lowest-price insurer β is identified by regions in which the identity of the cheapest insurer changed across years. The risk-score and distance moments identify the health and distance terms. The price coefficients α_g are identified by the 2017 reform.

Price endogeneity is a standard concern that arises. Specifically, an insurer that expects enrollees in a region or an income group to favor it might raise its rate, and net premium then rises with the taste shock ζ , which biases α_g toward zero. The fixed effects remove most of this channel. Insurer-by-region effects absorb a region's persistent taste for an insurer, insurer-by-income-group effects absorb an income group's, and market fixed effects absorb anything common to all insurers in a region, year, and income group. The remaining threat is an unobserved region-income-year specific shock to an insurer's appeal.

Given this threat, I instrument the net premium with the insurer's statewide gross rate interacted with income group and year. This instrument fulfills the exclusion restriction, since a statewide rate cannot respond to a shock confined to a region or an income group.

Within a year the instrument has no cross-group variation: the gross rate and the subsidy schedule are both statewide and collinear with the income-group effects. The variation that identifies α_g is therefore the 2017 reform. Insurer B raised one statewide rate for 2017, and the removal of the incremental subsidy passed that single increase through to net premiums by a different amount in each group. The differences across groups thus

Table 3: Demand Estimates

<i>Panel A. Estimated Parameters</i>			
Price sensitivity α_g (per \$100 per month)			
Below 100% FPL			−2.72 (0.10)
100–150% FPL			−1.42 (0.06)
150–200% FPL			−1.23 (0.05)
200–250% FPL			−1.11 (0.05)
250–300% FPL			−0.93 (0.05)
Lowest-price insurer premium β			0.054 (0.017)
HCC Risk Score $\times$ Insurer B θ_h			0.200 (0.040)
Distance to Insurer B hospital θ_d			−0.265 (0.125)
Correlated taste nest ρ			0.80 (0.061)
<i>Panel B. Own-price semi-elasticity of enrollment (percent per \$1 of net premium)</i>			
	Insurer N1	Insurer B	Insurer N2
100–150% FPL	−1.80	−1.05	−4.55
150–200% FPL	−1.62	−0.93	−4.23
200–250% FPL	−1.17	−0.91	−4.06
250–300% FPL	−0.87	−0.75	−3.50
Enrollment-weighted	−1.45	−0.93	−4.16

Notes: Panel A tabulates the parameters of equation (18). ρ is the correlation within the correlated taste nest of Insurer N1 and Insurer N2. Robust standard errors in parentheses. Panel B reports the percent change in each insurer's enrollment per one dollar increase in its monthly net premium, at observed 2016 prices, with the other insurers' premiums held fixed.

comes only from the subsidy schedule. Each α_g is identified by how that group's share of Insurer B responded to its own change in net premium.

Results. Table 3 tabulates the estimated parameters and the implied own-price semi-elasticities of demand for each insurer by income group. Price sensitivity declines in income: the coefficient is −2.72 per \$100 below the poverty line and −0.93 at 250 to 300 percent of poverty. Panel B translates these coefficients into insurer-specific enrollment responses. A one dollar increase in Insurer B's monthly net premium lowers its enrollment by 0.93 percent when weighted across income groups, and the response ranges from 1.05 percent at 100 to 150 percent of poverty to 0.75 percent at 250 to 300 percent. Insurer B's enrollment is the least price sensitive of the three because it stands alone in its nest. Insurer N1 loses 1.45 percent of its enrollment per dollar and Insurer N2 loses 4.16 percent. The difference reflects Insurer N2's smaller share within the nest of the two narrower insurers.

5.3 Costs and Risk Adjustment

A key component of the model is the degree of adverse selection given equilibrium net-of-subsidy prices. To measure this, I require each enrollee's potential cost at each insurer. I specify an enrollee's i potential cost to insurer j as

$$c_{ij} = \underbrace{\hat{c}_i}_{\text{predicted cost in the base insurer}} \times \underbrace{\delta_j}_{\text{insurer cost effect}} \quad (19)$$

where $\hat{c}_i$ is person i 's expected cost given observables and δ_j is insurer j 's multiplicative cost effect. To estimate $\hat{c}_i$, I estimate individual expected cost in a regression on the full claims panel, without person fixed effects, as a function of age, sex, income tier, and the risk score, with region and year effects and insurer indicators. This prediction removes the insurer component and re-centers it at its enrollment-weighted mean. Insurer differences enter equation (19) only through the cost effects δ_j . Appendix [Table D1](#) reports the model's coefficient.

The cost effects across insurers δ_j are the key parameters that capture insurer j 's specific cost structure that affect an individual's predicted cost. A simple comparison of means across insurers confounds insurer costs with enrollee health. I identify δ_j from enrollees who switch insurers, comparing the same person's costs before and after the switch. I estimate:

$$\mathbb{E}[c_{it} \mid j(i, t), h_{it}, i, t] = \exp(\beta_B \text{Insurer B}_{it} + \beta_{N2} \text{Insurer N2}_{it} + \varrho h_{it} + v_i + \kappa_t) \quad (20)$$

using Poisson pseudo-maximum likelihood, where c_{it} is person i 's realized monthly spending in year t , Insurer B_{it} and Insurer $N2_{it}$ indicate enrollment in each insurer with Insurer N1 as the omitted baseline, and β_B and β_{N2} are the log cost effects of enrolling in each relative to Insurer N1. h_{it} is the risk score in year t with coefficient ϱ , v_i is a person fixed effect, and κ_t is a year effect. Standard errors are clustered at the individual level.

Panel A of [Table 4](#) tabulates the observed cost and risk scores of enrollees in 2016 while Panel B shows the δ_j estimates. In 2016, Insurer B paid \$552.5 per member-month, nearly twice the \$277.8 and \$299.2 that Insurers N1 and N2 paid. This is reflected in the risk score differences, where Insurer B has a mean risk score of enrollees at 1.36 against 0.96 and 0.94. Switching from Insurer N1 to Insurer B raises insurer costs by 50 percent, and moving to Insurer N2 by 33 percent.

Table 4: Insurer Costs and Estimated Insurer Cost Effects

	(1) Insurer N1	(2) Insurer B	(3) Insurer N2
<i>Panel A. Observed costs 2016</i>			
Average cost (\$/month)	277.8	552.5	299.2
Mean risk score of enrollees	0.96	1.36	0.94
Person-years	62,901	28,432	24,362
<i>Panel B. Insurer cost effects δ_j</i>			
Cost effect (2016–2019)	baseline	1.50 (0.046)	1.33 (0.121)

Notes: Panel A reports insurer costs per member-month for single-member policies in 2016, enrollment month-weighted, from the all-payer claims database. Panel B reports within-person Poisson estimates of each insurer's multiplicative cost effect δ_j in the cost model $c_{ij} = \hat{c}_i \delta_j$, from enrollees who switch insurers across adjacent years over 2016–2019, with person and year fixed effects and a concurrent risk-score control. Standard errors are in parentheses.

Risk adjustment. ConnectorCare transfers money between insurers based on the health of their enrollees. Insurer j receives, for enrollee i ,

$$T_i = \left(\frac{h_i}{\bar{h}} - 1 \right) \bar{P} \quad (21)$$

which prices the enrollee's excess risk, $h_i/\bar{h} - 1$, at the statewide enrollment-weighted mean gross premium $\bar{P}$. Transfers are rescaled to sum to zero across insurers. Insurers take $\bar{P}$ as given, such that the transfer is invariant to any one insurer's rate. This simplifies the game considerably, and follows Tebaldi (2025), who models the ACA transfer formula and then holds risk adjustment fixed throughout his analysis. However, risk adjustment is endogenous to the insurer's health pool of enrollees $h_i/\bar{h} - 1$. A healthier enrollee pool yields lower transfers and vice versa for sicker enrollee pools.

5.4 Insurer Profits

Given demand and costs, insurer j 's profit is defined as:

$$\pi_j = \sum_i \mathcal{S}_{ij} \left[\underbrace{b_j \psi_{jr(i)}}_{\text{gross premium: base rate} \times \text{rating factors}} \left(1 + \underbrace{\text{CSR}_{g(i)}}_{\text{cost-sharing reduction payments}} \right) - \underbrace{c_{ij}}_{\text{cost}} + \underbrace{T_i}_{\text{risk-adjustment transfer}} \right] \quad (22)$$

where $\mathcal{S}_{ij}$ is person i 's choice probability, b_j is the insurer's statewide gross premium base rate, ψ_{jr} collects the rating factors that scales the gross premium based on regional factors, CSR_g is the cost-sharing reduction payments for income group g , c_{ij} is the cost of equa-

tion (19), and T_i is the risk-adjustment transfer of equation (21). Recall that ConnectorCare plans carry lower copays and deductibles than the underlying silver plans, and the state and federal government reimburse insurers for the generous cost-sharing structure. In the data, the payment is proportional to the premium within each income group, between 31 to 36 percent across income groups.

5.5 Model Fit and Discussion

Table 5 compares the model’s equilibrium under the observed 2016 incremental subsidy policy with the actual data. Costs replicate the observed data. The model reproduces Insurer N1’s average cost within one percent and Insurer B’s within two percent, while Insurer N2’s model cost is 19 percent above its observed cost.

The model also comes close to reproducing the average revenue of the two narrow-network insurers, \$421 against \$374 per member-month for Insurer N1 and \$422 against \$399 for Insurer N2, gaps of 13 and 6 percent. However, the model does not come close to predicting Insurer B’s revenue. Its equilibrium average revenue is \$774 against the \$472 observed. At the observed revenue Insurer B’s margin after risk adjustment on this segment is negative, \$32 per member-month below cost.

There are three reasons why the model overpredicts Insurer B’s revenue. First, the model has no medical loss ratio constraint. Massachusetts requires insurers in this market to spend at least 88 percent of premium revenue on claims. At the model’s revenue of \$774 and cost of \$563, Insurer B’s loss ratio would be 73 percent. Imposing the floor would cap its revenue near \$640, and a federal 80 percent floor near \$704. The constraint therefore closes part of the gap but not fully.

Second, Insurer B is vertically integrated with the hospital system whose hospitals are exclusive to its network. An integrated insurer values an enrollee not only for marginal revenue from premiums but also from health care utilization. A premium below cost can be rational if utilization margins covers the difference.

Third, Insurer B was operating on a loss during the sample period, and in this sense the model is accurate. A Bertrand-Nash insurer facing Insurer B’s demand would not price at \$472: that revenue is below any static best response, and an insurer would either exit or raise its price. In reality, Insurer B did raise its base rate significantly after 2016 and reducing its ConnectorCare market share.

Table 5: Model Fit at the Observed Incremental Subsidy Policy

	Data	Model
<i>Average revenue (\$/member-month)</i>		
Insurer N1	373.6	421 [13%]
Insurer N2	399.0	422 [6%]
Insurer B	471.8	774 [64%]
<i>Average cost (\$/member-month)</i>		
Insurer N1	277.8	275.4 [−1%]
Insurer N2	299.2	356.8 [19%]
Insurer B	552.5	562.9 [2%]
<i>Market shares</i>		
Insurer N1	0.577	0.584 [1%]
Insurer N2	0.202	0.261 [29%]
Insurer B	0.221	0.154 [−30%]

Notes: The ‘Data’ column reports 2016 observed quantities. The ‘Model’ column is the solved equilibrium under the observed 2016 schedule. Average revenue is premium plus cost-sharing reduction payments per member-month, gross of risk adjustment. Average cost is the model’s cost of each insurer’s observed 2016 enrollees, inclusive of the insurer cost effect and gross of risk adjustment.

6 Welfare and Counterfactuals

I use the demand and cost estimates to simulate equilibrium pricing under different incremental subsidy rates and its effects on insurers' incentives. I use the model to first examine how varying the intensity of the incremental subsidy policy affects insurers' incentives. Next, I simulate consumer surplus and the fiscal cost of financing the subsidy policies to infer the social planner's welfare objectives. Finally, I simulate a counterfactual policy of subsidizing quality directly (instead of incremental prices).

6.1 Subsidy Policies and Simulation Assumptions

As mentioned in Section 5, insurers set gross premiums in Bertrand competition, taking both the level and incremental subsidies as given, and enrollees choose plans at net-of-subsidy premiums.

Sample I simulate the market for new enrollees at 100–300 percent of poverty and treat the below-poverty group as a fixed segment that does not respond to equilibrium simulations. This is because below poverty enrollees faced zero prices for every plan in 2016, and Insurer B's best response would be to price at the ceiling.


Varying the Incremental Subsidy I vary the incremental subsidy policy for below poverty enrollees and income groups between 100 and 200 percent of poverty. The observed 2016 incremental subsidy rates are 1 for below poverty enrollees, 0.73 and 0.59 in the 100–150 and 150–200 percent FPL group, while being close to zero for above 200 percent of poverty (0.07 and 0.002). Specifically, let each policy be denoted by $\bar{s}$, the average rate for 100–200 percent enrollees. I then fit the two rates on the straight line as an average $\bar{s}$:

$$s_g(\bar{s}) = \max \left\{ 0, 1 - (1 - \bar{s}) \frac{y_g - y_1}{\bar{y} - y_1} \right\}, \quad g \in \{100-150, 150-200\} \quad (23)$$

where y_g is the income assigned to group g , y_1 the income assigned to the below-poverty group, and $\bar{y}$ the mean of the two low-income groups' incomes. A higher $\bar{s}$ subsidizes more of the price gap for the 100–150 and 150–200 percent FPL group. At $\bar{s} = 0$ the two low-income groups receive zero incremental subsidy, while as $\bar{s}$ approaches one, both groups approach the below-poverty group's full subsidy. Figure 6 illustrates the design with two examples. The red line ($\bar{s} = 0.87$) subsidizes 90 percent of the gap at 100–150 percent of poverty and 83 percent at 150–200 percent; the blue line ($\bar{s} = 0.33$) subsidizes 50 and 17 percent. I solve the equilibrium on a grid from $\bar{s} = 0$ to 0.7 in steps of 0.1. At $\bar{s} = 0.8$ Insurer B's prices at the ceiling, and I censor equilibrium simulations beyond that.

The incremental subsidy applies to plans that are not the lowest-price carrier in a region. In the data, the incremental subsidy is applied to both Insurer N2 and Insurer B.

Figure 6: Illustrating Alternate Incremental Subsidy Schedules




Notes: The black line is the observed 2016 incremental subsidy schedule. The red and blue lines are two illustrative members of the family of schedules in equation (23), labeled by $\bar{s}$, the average rate across the 100–150 and 150–200 percent groups

Endogenous Level Subsidies I treat the amount of level subsidies as endogenous. Connector fixes the affordability schedule, which sets the net premium of the cheapest plan for enrollees in each income group. While I treat the affordability schedule as exogenous, this implies that level subsidies would endogenously adjust based on insurer gross premiums. In other words, the state’s expenditure on level subsidies rises and falls with insurers’ prices. This implies that the cheapest insurer’s enrollees pay the same amount regardless of insurer prices, since τ_{gr} rises one for one with the lowest gross premium. This is the price-linked design of Jaffe and Shepard (2020).

6.2 The Incidence of Incremental Subsidies

Figure 7 is a four panel figure which shows how varying the incremental subsidy affects insurers’ incentives, welfare and enrollment. Appendix Table D4 tabulates these equilibrium outcomes.

Figure 7: Incidence of the Incremental Subsidy



Notes: Each point is an equilibrium simulation for a value of $\bar{s}$. The dashed line marks the policy with the observed average rate, $\bar{s} = 0.657$. Panel A depicts Insurer B's average revenue (premium plus cost-sharing reduction payments), average cost, marginal cost (the cost of the enrollees at the margin of its demand), and average cost net of the risk-adjustment, measured at per member-month of its enrollees. Panel B does the same for Insurer N2. Panel C shows changes from the equilibrium at $\bar{s} = 0$ in government cost, insurer profit, and consumer surplus, per member-month. Panel D depicts Insurer B's share of each income group's enrollees.

Insurer B. At $\bar{s} = 0$, Insurer B's average revenue is \$624 against an average cost of \$628. With risk adjustment, average costs are reduced from \$628 to \$496. In addition, as $\bar{s}$ increases, both average and marginal cost reduce, showing the positive effect of incremental subsidies on adverse selection: Insurer B's average cost falls from \$628 to \$594, and the cost of its marginal enrollee from \$593 to \$575. However, average revenue increases much more than the reduction in average costs, as Insurer B's average revenue rises significantly, from \$624 to \$776 per member-month at the observed schedule. This simulation shows that while incremental subsidies do address adverse selection, Insurer B is able to raise markups significantly.

Insurer N2. Panel B shows average revenue and costs for Insurer N2. As $\bar{s}$ increases, Insurer N2's revenue rises from \$396 to \$433 per member-month. This implies that while the incremental subsidy targets Insurer B, the narrow insurers are also able to raise their markups; when a narrow insurer raises its price, existing enrollees in the narrow insurer pay the same amount (given that level subsidies fix the affordability schedule). However, the gap between Insurer B and the cheapest plan shrinks, and Insurer B becomes relatively cheaper. Increasing $\bar{s}$ in this case implies that enrollees only face a fraction $1 - s_g$ of that change, as the narrow insurers' demand become more inelastic.

Incidence. Panel C shows how the incremental subsidy is divided among government, insurers, and enrollees. Specifically, the figure plots each welfare component relative to the equilibrium with zero incremental subsidy ($\bar{s} = 0$). Government cost G is the state's spending on the level subsidy, the incremental subsidy, and state cost sharing,

$$G = \sum_i \left[\underbrace{\tau_{g(i)r(i)}}_{\text{level subsidy}} + \underbrace{s_{g(i)}(p_{j(i)r(i)} - \underline{p}_{r(i)})}_{\text{incremental subsidy}} + \underbrace{\text{CSR}_i^{\text{state}}}_{\text{state cost sharing}} \right] \quad (24)$$

where p_{jr} is the gross premium of enrollee i 's plan and $\underline{p}_r$ the cheapest gross premium in the region. Federal premium tax credits and federal cost-sharing payments are held fixed at their observed levels and excluded. Moving from zero incremental subsidy to the observed policy raises government cost by \$38 per member-month. Insurer profit rises by \$34, or 91 percent of the added spending, and consumer surplus by a modest \$2, or 5 percent.

Insurer B Enrollment. Panel D shows that the incremental subsidy does afford access for the poor into Insurer B. I find that Insurer B's share of the 100–150 percent group rises from 3.9 to 12.4 percent between zero incremental subsidy and the observed schedule. Its share of the 150–200 percent group also rises from 5.1 to 8.9 percent. However, its share of the two groups above 200 percent of poverty falls (from 5.7 to 2.6 percent and from 7.1 to 3.5 percent). This is because those enrollees face the full increase of Insurer B's prices.

6.3 Which Social Planner Chooses the Incremental Subsidy?

Given these trade-offs, I ask which planner would choose the observed incremental subsidy schedule. I simulate three policies: (i) the observed level and incremental subsidy, (ii) the level subsidy alone, and (iii) reallocating the incremental subsidy budget to a larger level subsidy, and compare them under three consumer welfare objectives:

$$W = \sum_g \omega_g CS_g + \mu S_B - 1.3 G \quad (25)$$

where CS_g is consumer surplus in income group g , ω_g is the planner's weight on that group's surplus, S_B is the specific egalitarian preference for the poor. Specifically, it is the share of 100–200 percent of poverty enrollees who enroll in Insurer B. G is the state's spending of equation (24), and 1.3 is the marginal cost of public funds.⁷

I define three planners as specific cases of equation (25):

1. *Efficiency*: $\omega_g = 1$ for every group and $\mu = 0$. Every dollar of consumer surplus is worth the same in each income group and enrollment in Insurer B has no inherent value to the planner.
2. *Utilitarian (Inequality-averse)*: $\omega_g > 1$ and $\mu = 0$, with each group's surplus weighted by its marginal utility of income. Under constant relative risk aversion,

$$\omega_g = \left(\frac{y_{\text{med}}}{y_g} \right)^\gamma \quad (26)$$

normalized to one for the 150–200 percent group, which contains the median enrollee, with γ the degree of inequality aversion. At $\gamma = 0$ the utilitarian planner is the efficiency planner.

3. *Tobin*: $\omega_g = 1$ and $\mu > 0$, so that the planner values the enrollment of the poor in Insurer B beyond consumer surplus.

Table 6 tabulates the results. Column (1) is the equilibrium at the observed incremental subsidy rates in addition to the level subsidy. Column (2) removes the incremental subsidy for the two low-income groups while the level subsidy remains. The difference in government cost between the two columns is the incremental subsidy's fiscal cost, \$37.7 per member-month. Column (3) distributes that \$37.7 to enrollees at 100–200 percent of poverty as an equal lump sum subsidy per enrollee.

⁷A weight of 1.3 on government spending is typical for analyses of subsidized insurance markets, including Decarolis et al. (2020) for Medicare Part D and Saltzman et al. (2021) for the ACA exchanges.

Table 6: Which Planner Chooses the Incremental Subsidy?

	Observed policy (level + incremental) (1)	Level only (2)	Increased level subsidies (3)
<i>Welfare by planner (\$ per member-month)</i>			
Efficiency	64.4	107.8	97.4
Utilitarian ($\gamma = 1.4$)	40.8	82.9	79.7
Utilitarian ($\gamma = 5$)	270.0	306.8	348.8
Tobin ($\mu = 3,821$)	255.6	197.8	187.4

Notes: All columns are simulated equilibria. Column (1) is the observed level and incremental subsidy. There is multiple equilibria at the observed schedule, which differ by \$7 in the base rates of Insurers N1 and N2. I report the average of the two in Column (1). Column (2) removes the incremental subsidy. Column (3) adds to column (2) additional level subsidy for enrollees at 100–200 percent of poverty equal to the difference in government cost between columns (1) and (2)—\$37.7 per member-month. Appendix Table D5 reports the consumer surplus, government cost, insurer profits, and Insurer B share behind each column, and Appendix D.2 derives μ .

Efficiency. I find that the efficiency planner prefers the level subsidy, with welfare of \$107.8 per member-month against \$64.4 under the observed policy (Row 1 of Table 6). The incremental subsidy costs \$34.7 per member-month and yields just under \$2 of consumer surplus, while \$30.6 of it becomes insurer profit (see Table D5). The planner also prefers a lower level subsidy (comparing columns 2 versus 3), since a lump-sum transfer still costs \$1.30 per dollar to the social planner.

Utilitarian. I find that the utilitarian planner also rejects the observed incremental subsidy policy at any degree of inequality aversion. At $\gamma = 1.4$, the planner prefers the level subsidy over the incremental subsidy (82.9 against 40.8)⁸. At an extremely high rate of inequality aversion $\gamma = 5$, the planner prefers increased level subsidies instead (348.8 against 270.0). Thus, an inequality averse planner would prefer the level subsidy over the incremental subsidy to improve welfare for the poor. This is because the level subsidy leaves every income group better off than the observed incremental policy at the same government cost.

Tobin. Suppose a planner values the enrollment of the poor in the broad network itself beyond welfarist components. I specify the parameter μ as the planner's member-month valuation. Assuming the observed incremental subsidy policy ($\bar{s} = 0.657$) is the planner's optimal policy, the first-order condition $dW/ds = 0$ implies

$$\mu = - \frac{dCS/d\bar{s} - 1.3 dG/d\bar{s}}{dS_B/d\bar{s}} \quad (27)$$


⁸1.4 is the value the Office of Management and Budget uses for federal benefit-cost analysis.

which I evaluate numerically (Appendix D.2). I find a valuation of \$3,821 per low-income Insurer B member-month. At this valuation the Tobin planner chooses the observed policy: its welfare is 255.6 against 197.8 under the level subsidy alone and 187.4 under increased level subsidies (Row 4 of Table 6). This is because the share of the poor increase from 0.024 to 0.05, and the planner's valuation of this outweighs the surplus it forgoes. The valuation is large: it is 4.9 times Insurer B's revenue per enrollee of \$776, and about 475 times the \$8 at which the average low-income enrollee values access to Insurer B.

6.4 Counterfactual Policy for the Revealed Planner

If the objective is to ensure low-income enrollment in Insurer B, the observed policy is an expensive policy. Instead, I consider a counterfactual policy that subsidizes Insurer B directly rather than the price gap. This involves removing the incremental subsidy in every income group above poverty and paying a fixed credit of v dollars per member-month toward Insurer B's premium for enrollees at 100–200 percent of poverty. Prices are no longer scaled by $1 - s_g$, and enrollees face the full margin of every insurer's price increase. I solve the equilibrium under a \$100 credit and compare it with the observed policy in Figure 8; Appendix Table D6 provides a detailed tabulation of the components.

Figure 8: A Broad Plan Credit Compared With the Incremental Subsidy



Notes: Each panel compares two solved equilibria of the model, the observed incremental subsidy and a \$100 credit per member-month toward Insurer B's premium for enrollees at 100–200 percent of poverty.

I find that a \$100 credit raises low-income enrollment in Insurer B from 0.051 to 0.084 of the market, which amounts to two-thirds more than the observed schedule. This comes at a lower government cost (\$289 against \$310 per member-month) and slightly higher consumer surplus (\$467.3 against \$463.5). Importantly, insurer profits fall by 26 percent,

from \$104 to \$77. This suggests that subsidizing quality instead of incremental prices increases enrollment in the broad-plan while reducing government costs and insurer profits.

There are two important limitations to consider when extrapolating to other markets. The first is that quality is defined in this exercise as the broad-network plan. In reality, quality has many dimensions. A subsidy tied to one dimension may lead firms to over-invest in the subsidized dimension and under-invest in the rest (Holmstrom and Milgrom, 1991). Relatedly, subsidizing incremental prices instead of quality is a more convenient approach, as it does not require the planner to define quality dimensions.

Second, the model holds hospital prices fixed, and a credit tied to hospital networks changes the bargaining between insurers and those hospitals. In a bargaining model, the price an insurer pays a hospital depends on each side's outside option: the insurer loses the enrollees who value that hospital, and the hospital loses the volume it gets from the insurer's enrollees (Gowrisankaran et al., 2015; Ho and Lee, 2017). The credit's effect on this bargain is ambiguous. On one hand, Insurer B gains market share and with it bargaining power (Ho and Lee, 2017; Choi, 2025). On the other hand, the credit makes the named hospitals worth more to the insurer, which strengthens the hospitals' hand. If the credit draws other insurers to add those hospitals in order to qualify, the hospitals become a must-have for every plan (Ho and Lee, 2019).

7 Conclusion

This paper studies trade-offs of subsidizing the incremental price of a high-quality health insurance plan for the poor in Massachusetts. Theory shows that adverse selection and concern for the poor's enrollment in the high-quality plan can justify the observed incremental policy. The structural model shows that, empirically, adverse selection is dominated by insurer markups. The high-quality insurer is able to raise markups, given that enrollees now face only a fraction of any price increase. A counterfactual policy of subsidizing quality instead of incremental prices achieves higher enrollment in the high-quality plan while lowering both insurer profits and government costs.

This policy has direct relevance for many markets. Governments routinely narrow quality gaps for the poor by paying the incremental price of the higher-quality option. The 'Cal Grant'—California's main need-based college grant—covers systemwide tuition and fees at the University of California and California State University. The subsidy amount increases one for one with tuition rate and is larger at higher tuition UC colleges. Small Area Fair Market Rents set the Housing Choice Voucher's payment standard at the ZIP code's rent rather than the metropolitan area's. The government then pays the rent gap between a cheap neighborhood and an expensive zip code. Overall, my findings show that a fixed subsidy for choosing the higher-quality option affords access without distorting competition.

References

- Abaluck, J., Caceres Bravo, M., Hull, P., and Starc, A. (2021). Mortality effects and choice across private health insurance plans. *The Quarterly Journal of Economics*, 136(3):1557–1610.
- Akbarpour, M., Dworczak, P., and Kominers, S. D. (2024). Redistributive allocation mechanisms. *Journal of Political Economy*, 132(6):1831–1875.
- Berry, S., Levinsohn, J., and Pakes, A. (1995). Automobile prices in market equilibrium. *Econometrica*, 63(4):841–890.
- Berry, S., Levinsohn, J., and Pakes, A. (2004). Differentiated products demand systems from a combination of micro and macro data: The new car market. *Journal of Political Economy*, 112(1):68–105.
- Besley, T. and Coate, S. (1991). Public provision of private goods and the redistribution of income. *American Economic Review*, 81(4):979–984.
- Cabral, M., Geruso, M., and Mahoney, N. (2018). Do larger health insurance subsidies benefit patients or producers? evidence from medicare advantage. *American Economic Review*, 108(8):2048–2087.
- Chetty, R., Hendren, N., and Katz, L. F. (2016). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902.
- Choi, J.-H. (2025). Insurer size and negotiated hospital prices: Insights from the Affordable Care Act in Arkansas. *Health Economics*, 34(11):2097–2113.
- Collinson, R. and Ganong, P. (2018). How do changes in housing voucher design affect rent and neighborhood quality? *American Economic Journal: Economic Policy*, 10(2):62–89.
- Conlon, C. and Gortmaker, J. (2020). Best practices for differentiated products demand estimation with pyblp. *The RAND Journal of Economics*, 51(4):1108–1161.
- Currie, J. and Gahvari, F. (2008). Transfers in cash and in-kind: Theory meets the data. *Journal of Economic Literature*, 46(2):333–383.
- Decarolis, F. (2015). Medicare part d: Are insurers gaming the low income subsidy design? *American Economic Review*, 105(4):1547–1580.
- Decarolis, F., Polyakova, M., and Ryan, S. P. (2020). Subsidy design in privately provided social insurance: Lessons from medicare part d. *Journal of Political Economy*, 128(5):1712–1752.
- Deming, D. J., Hastings, J. S., Kane, T. J., and Staiger, D. O. (2014). School choice, school quality, and postsecondary attainment. *American Economic Review*, 104(3):991–1013.
- Denning, J. T., Marx, B. M., and Turner, L. J. (2019). Propelled: The effects of grants on graduation, earnings, and welfare. *American Economic Journal: Applied Economics*, 11(3):193–224.

- Dworczak, P., Kominers, S. D., and Akbarpour, M. (2021). Redistribution through markets. *Econometrica*, 89(4):1665–1698.
- Dynarski, S. M. (2003). Does aid matter? measuring the effect of student aid on college attendance and completion. *American Economic Review*.
- Einav, L. and Finkelstein, A. (2011). Selection in insurance markets: Theory and empirics in pictures. *Journal of Economic perspectives*, 25(1):115–138.
- Finkelstein, A., Hendren, N., and Shepard, M. (2019). Subsidizing health insurance for low-income adults: Evidence from massachusetts. *American Economic Review*, 109(4):1530–1567.
- Frean, M., Gruber, J., and Sommers, B. D. (2017). Premium subsidies, the mandate, and Medicaid expansion: Coverage effects of the Affordable Care Act. *Journal of Health Economics*, 53:72–86.
- Gowrisankaran, G., Nevo, A., and Town, R. (2015). Mergers when prices are negotiated: Evidence from the hospital industry. *American Economic Review*, 105(1):172–203.
- Ho, K. and Lee, R. S. (2017). Insurer competition in health care markets. *Econometrica*, 85(2):379–417.
- Ho, K. and Lee, R. S. (2019). Equilibrium provider networks: Bargaining and exclusion in health care markets. *American Economic Review*, 109(2):473–522.
- Holmstrom, B. and Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, and Organization*, 7(Special Issue):24–52.
- Jaffe, S. and Shepard, M. (2020). Price-linked subsidies and imperfect competition in health insurance. *American Economic Journal: Economic Policy*, 12(3):279–311.
- Kong, E., Layton, T., and Shepard, M. (2024). Adverse selection and (un)natural monopoly in insurance markets. Working Paper 33187, National Bureau of Economic Research.
- Mahoney, N. and Weyl, E. G. (2017). Imperfect competition in selection markets. *Review of Economics and Statistics*, 99(4):637–651.
- Nichols, A. L. and Zeckhauser, R. J. (1982). Targeting transfers through restrictions on recipients. *American Economic Review, Papers and Proceedings*, 72(2):372–377.
- Saez, E. and Stantcheva, S. (2016). Generalized social marginal welfare weights for optimal tax theory. *American Economic Review*, 106(1):24–45.
- Saltzman, E., Swanson, A., and Polsky, D. (2021). Inertia, market power, and adverse selection in health insurance: Evidence from the ACA exchanges. Working Paper 29097, National Bureau of Economic Research.
- Shepard, M. (2022). Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange. *American Economic Review*, 112(2):578–615.

- Tebaldi, P. (2025). Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca. *The Review of Economic Studies*, 92(1):586–620.
- Tobin, J. (1970). On limiting the domain of inequality. *Journal of Law and Economics*, 13(2):263–277.
- Wood, M., Turnham, J., and Mills, G. (2008). Housing affordability and family well-being: Results from the housing voucher evaluation. *Housing Policy Debate*, 19(2):367–412.

Appendix Contents

A Theory Appendix	49
A.1 Average Cost, Marginal Cost, and the Selection Wedge	49
A.2 Perfect Competition	50
A.3 Strategic Pricing	50
A.4 The Planner	51
B Appendix: Setting and Data	52
C Appendix: Reduced Form	54
D Appendix: Structural Model and Welfare Simulations	56
D.1 Private Willingness to Pay for Broad-Network Access	57
D.2 Recovering the Access Value	58

Appendix Tables

B1 Insurer Networks and Quality Ratings	53
C1 ConnectorCare Enrollment by Income Group, 2016 to 2019	54
C2 Selection into Insurer B: HCC Risk Score	54
C3 Dist. to B-exclusive hospitals by Income Group	55
D1 The Expected-Cost Model	56
D2 Private Willingness to Pay for Broad-Network Access	57
D3 Incremental Subsidy Policies and Equilibrium Base Prices	59
D4 Equilibrium Outcomes and Pricing Along the Family of Policies	60
D5 Simulated Outcomes and Welfare Behind the Planner Comparison	61
D6 A Broad Plan Credit Compared With the Incremental Subsidy: Equilibrium Outcomes	62

Appendix Figures

B1 Premium Decomposition by Income Group (% FPL), Insurer A vs. Insurer B, 2016–2017	52
C1 ConnectorCare Enrollment, 2016 to 2019	55
C2 Insurer B’s Below-Poverty Enrollees in the Claims Panel, 2016 to 2019	55

A Theory Appendix

This appendix derives the equations in Section 2. Throughout, $\Delta\tilde{p} \equiv (1-s) \cdot \Delta p$ is the price gap enrollees face. Enrollment D_B and sick enrollment D_{SB} depend on s and p_B only through $\Delta\tilde{p}$, with $\partial\Delta\tilde{p}/\partial s = -\Delta p$ and $\partial\Delta\tilde{p}/\partial p_B = 1-s$. Every derivative below is a derivative in $\Delta\tilde{p}$ times one of these factors.

In particular, $\phi_B^m = (\partial D_{SB}/\partial p_B)/(\partial D_B/\partial p_B) = D'_{SB}(\Delta\tilde{p})/D'_B(\Delta\tilde{p})$, and $\eta_B = (1-s) \cdot \tilde{\eta}_B(\Delta\tilde{p})$ with $\tilde{\eta}_B(\Delta\tilde{p}) \equiv -D'_B(\Delta\tilde{p})/D_B(\Delta\tilde{p})$. As a reminder, adverse selection implies that the sick share among marginal enrollees, $\phi_B^m(\Delta\tilde{p})$, increases with the price gap.

A.1 Average Cost, Marginal Cost, and the Selection Wedge

Deriving Equation (2). With healthy enrollment $D_{HB} = D_B - D_{SB}$, $AC_B = (c_{HB} \cdot D_{HB} + c_{SB} \cdot D_{SB})/D_B = c_{HB} + \delta_B \cdot \phi_B$, which is (2). Total cost is $TC_B = AC_B \cdot D_B = c_{HB} \cdot D_B + \delta_B \cdot D_{SB}$.

Deriving Equation (8). Differentiating total cost with respect to price and dividing by marginal enrollment:

$$\begin{aligned} MC_B &= \frac{\partial TC_B / \partial p_B}{\partial D_B / \partial p_B} \\ &= \underbrace{c_{HB}}_{\text{healthy cost}} + \underbrace{\delta_B \cdot \phi_B^m(\Delta\tilde{p})}_{\text{marginal selection cost}} \end{aligned} \quad (\text{A1})$$

which is (8).

Deriving Equations (3) and (4). Differentiate $AC_B = TC_B/D_B$ with the quotient rule and use $\partial TC_B / \partial p_B = MC_B \cdot \partial D_B / \partial p_B$:

$$\begin{aligned} \frac{\partial AC_B}{\partial p_B} &= \frac{\partial D_B / \partial p_B}{D_B} \cdot (MC_B - AC_B) \\ &= (1-s) \cdot \tilde{\eta}_B(\Delta\tilde{p}) \cdot \underbrace{\delta_B \cdot (\phi_B - \phi_B^m)}_{\text{selection wedge}} \end{aligned} \quad (\text{A2})$$

where the second line uses $\partial D_B / \partial p_B = -\eta_B \cdot D_B = -(1-s) \cdot \tilde{\eta}_B \cdot D_B$ and $AC_B - MC_B = \delta_B \cdot (\phi_B - \phi_B^m)$ from (A1). This is (3), and dividing by $1-s$ gives (4).

Sign of the selection wedge. Given that the pool's sick share ϕ_B is the enrollment-weighted average of ϕ_B^m over all gaps above the current one, an increasing ϕ_B^m implies $\phi_B^m(\Delta\tilde{p}) < \phi_B(\Delta\tilde{p})$: the marginal enrollee is healthier than the pool. The wedge is positive under adverse selection, since $\phi_B^m < \phi_B$, and so is the slope in (A2).

A.2 Perfect Competition

Deriving Equation (5). Differentiating $\phi_B = D_{SB}/D_B$ with respect to s at fixed prices,

$$\frac{\partial \phi_B}{\partial s} = \underbrace{(\phi_B^m - \phi_B)}_{<0: \text{marginal healthier than pool}} \cdot \underbrace{\frac{\partial D_B / \partial s}{D_B}}_{>0: \text{enrollment rises}} < 0 \quad (\text{A3})$$

since the enrollees the subsidy draws in are healthier than the pool. Equation (5) is $\partial AC_B / \partial s = \delta_B \cdot \partial \phi_B / \partial s$ with (A3) substituted, or, by the chain rule, $-\Delta p \cdot AC'_B(\Delta \tilde{p})$.

Sign of the pass-through denominator. Under zero-profit pricing $p_B = AC_B$ and $p_N = c_N$. The denominator of the pass-through in the main text, $1 - (1-s) \cdot AC'_B(\Delta \tilde{p})$, is below one because the slope is positive by (A2), and above zero when the zero-profit equilibrium is stable, the average cost curve crossing the price line from below. It then lies in $(0, 1)$ and $|dp_B^* / ds| > |\partial AC_B / \partial s|$.

A.3 Strategic Pricing

Deriving Equation (9). Profit is $\pi_B = (p_B - AC_B) \cdot D_B$. Differentiating and substituting (A2),

$$\begin{aligned} \frac{\partial \pi_B}{\partial p_B} &= \underbrace{D_B}_{\text{revenue from stayers}} + \underbrace{(p_B - AC_B) \cdot \frac{\partial D_B}{\partial p_B}}_{\text{margin lost on leavers}} - \underbrace{D_B \cdot \frac{\partial AC_B}{\partial p_B}}_{\text{composition effect on } AC_B} \\ &= \eta_B \cdot D_B \cdot \left[\underbrace{\frac{1}{\eta_B}}_{\text{Lerner markup}} - \underbrace{(p_B - MC_B)}_{\text{margin}} \right] \end{aligned} \quad (\text{A4})$$

where the second line collects the last two terms into $(p_B - MC_B) \cdot \partial D_B / \partial p_B$ and uses $\partial D_B / \partial p_B = -\eta_B \cdot D_B$. Setting (A4) to zero gives the Lerner condition $p_B - MC_B = 1/\eta_B$. Subtracting the selection wedge $AC_B - MC_B$ from both sides gives (9).

Deriving Equation (10). Differentiating (A4) once more at $p_B - MC_B = 1/\eta_B$, the second-order condition is

$$\frac{\partial^2 \pi_B}{\partial p_B^2} = -\eta_B \cdot D_B \cdot \left[1 - \underbrace{\frac{\partial MC_B}{\partial p_B}}_{\text{marginal cost slope}} - \underbrace{\frac{\partial(1/\eta_B)}{\partial p_B}}_{\text{markup slope}} \right] < 0 \quad (\text{A5})$$

Differentiating the Lerner condition $p_B - MC_B = 1/\eta_B$ with respect to s at fixed p_N and solving for dp_B^* / ds gives (10), of which the denominator is the bracket in (A5). The cost channel is $\delta_B \cdot \partial \phi_B^m / \partial s$, and since a unit of s lowers the gap by Δp at fixed prices, $\partial \phi_B^m / \partial s = -\Delta p \cdot \phi_B^{m'}(\Delta \tilde{p}) < 0$ by adverse selection: the subsidy makes the marginal enrollee healthier.

For the markup channel, $1/\eta_B = 1/[(1-s) \cdot \tilde{\eta}_B(\Delta\tilde{p})]$, so at fixed prices

$$\frac{\partial(1/\eta_B)}{\partial s} = \frac{1}{(1-s)} \cdot \left[\underbrace{\frac{1}{(1-s) \cdot \tilde{\eta}_B}}_{\text{shielding enrollees}} + \underbrace{\Delta p \cdot \frac{\tilde{\eta}'_B(\Delta\tilde{p})}{\tilde{\eta}_B^2}}_{\text{change in semi-elasticity}} \right] \quad (\text{A6})$$

The first term captures the shielding of enrollees from any price increase. The second is the change in the net-premium semi-elasticity as the gap shrinks, and is non-negative whenever $\tilde{\eta}_B$ does not fall as the gap widens, i.e., whenever D_B is log-concave in the gap, as it is for logistic demand. Under that condition the markup channel is positive, and it grows without bound as $s \rightarrow 1$ through the first term.

A.4 The Planner

Deriving Equations (13) and (14). Mean enrollee surplus is $CS = \mathbb{E}_i[\max\{V_{iN} - \tilde{p}_N, V_{iB} - \tilde{p}_B\}]$, where V_{ij} is enrollee i 's value of plan j and $\Delta V_i = V_{iB} - V_{iN}$. By the envelope theorem, $\partial CS/\partial \tau = 1$, since both net premiums fall by one dollar; $\partial CS/\partial s = \Delta p \cdot D_B$, since the high-quality plan's net premium falls by Δp for the inframarginal D_B enrollees. Differentiating (12) with respect to τ and using (19) gives (13), because D_B does not depend on τ . Differentiating with respect to s at fixed prices gives (14).


Deriving Implication 4. For the comparison per dollar of fiscal cost, the marginal fiscal cost of the incremental subsidy is $\partial G/\partial s = \Delta p \cdot D_B + s \cdot \Delta p \cdot \partial D_B/\partial s$, the transfer to enrollees already in the plan plus the cost on induced enrollment. With $\mu = 0$:

$$\underbrace{\frac{\partial W/\partial s}{\partial G/\partial s}}_{\text{incremental}} - \underbrace{(\omega - \lambda)}_{\text{level}} = -\omega \cdot \underbrace{\frac{s \cdot \Delta p \cdot \partial D_B/\partial s}{\partial G/\partial s}}_{\text{induced share of cost}} < 0 \quad \text{for } s > 0 \quad (\text{A7})$$

which is the ranking in Implication 4: per dollar of fiscal cost the incremental subsidy delivers less than the level subsidy whenever $s > 0$ and exactly as much at $s = 0$. The Tobin planner's condition follows from setting (14) to zero at $s = 0$ with $\omega = 1$.

B Appendix: Setting and Data

Figure B1: Premium Decomposition by Income Group (% FPL), Insurer A vs. Insurer B, 2016–2017



Notes: Each bar decomposes the gross premium into four components, stacked bottom to top. Gray: APTC (Advance Premium Tax Credit). Green: State-level subsidy. Blue: Net premium. Red: State-Incremental subsidy. Top row shows 2016; bottom row shows 2017. Income groups are expressed as % of the Federal Poverty Line (FPL).

Table B1: Insurer Networks and Quality Ratings

	Insurer B	Insurer N1	Insurer N2
<i>Panel A: Hospital network, ConnectorCare (2015)</i>			
Hospitals in network (of 72)	68	57	50
Share of listed hospitals (percent)	94	79	69
Academic medical centers (of 6)	6	4	3
Specialty hospitals (of 3)	3	2	1
Hospitals of the largest system (of 8)	8	0	0
Hospitals in no other ConnectorCare network	8	0	0
<i>Panel B: NCQA Health Insurance Plan Ratings, 2019–2020</i>			
Overall rating (out of 5)	4.5	4.5	4.0
Patient experience	4.5	3.5	2.5
Prevention	4.5	4.5	4.0
Treatment	4.0	4.5	3.5

Notes: The source for Panel A is the Massachusetts Health Connector, *ConnectorCare Hospital Network List*, effective January 1, 2015, which reports for each of 72 Massachusetts hospitals whether it participates in each ConnectorCare insurer's network. Academic medical centers are the six hospitals the state's Center for Health Information and Analysis classifies as such. Specialty hospitals are the state's pediatric, cancer, and eye and ear hospitals. The largest system row counts the eight hospitals of the state's largest hospital system. The last row counts hospitals that appear in that insurer's network and in no other ConnectorCare insurer's network. The source for Panel B is the National Committee for Quality Assurance, *Health Insurance Plan Ratings 2019–2020*, Medicaid product line. The ratings are for each insurer's Medicaid plan, which shares the insurer's provider network. .

C Appendix: Reduced Form

Table C1: ConnectorCare Enrollment by Income Group, 2016 to 2019

	2016	2017	2018	2019	Change 2016 to 2017 (percent)
<i>Panel A. All enrollees</i>					
Below 100% FPL	15,007	17,711	18,117	18,220	18.0
100–150% FPL	28,885	30,066	32,621	38,096	4.1
150–200% FPL	52,803	58,045	59,048	62,086	9.9
200–250% FPL	33,768	39,438	43,149	47,954	16.8
250–300% FPL	21,386	25,887	27,641	31,377	21.0
All	151,849	171,147	180,576	197,733	12.7
<i>Panel B. New enrollees</i>					
Below 100% FPL	7,480	8,317	8,077	7,051	11.2
100–150% FPL	13,501	12,856	13,329	14,276	–4.8
150–200% FPL	25,695	22,034	19,794	19,822	–14.2
200–250% FPL	14,926	13,234	12,555	13,521	–11.3
250–300% FPL	9,032	8,094	7,343	7,994	–10.4
All	70,634	64,535	61,098	62,664	–8.6


Notes: The table counts enrollee-years in the enrollment panel by income group and year. New enrollees are those entering the marketplace in that year. The last column is the percent change from 2016 to 2017, the year the incremental subsidy was removed. The panel begins in 2016.

Table C2: Selection into Insurer B: HCC Risk Score

	Below 100% FPL (1)	100–200% FPL (2)	200–300% FPL (3)
Insurer B $\times$ Post-2017	0.693*** (0.142)	0.266*** (0.066)	0.243*** (0.060)
2016 mean, Insurer B	1.16	1.29	1.45
2016 mean, narrow insurers	0.97	0.92	0.99
Observations		601,866	


Notes: The table reports the triple difference-in-differences of equation (16) with the HCC risk score as the outcome, on the same sample as Panel B of Table 2. Each column is the post-2017 change in Insurer B enrollees' risk score relative to the narrow insurers' within that income group. The regression is unweighted and includes income-group–region and year fixed effects and the insurer $\times$ income group level terms. Standard errors in parentheses are clustered at the income-group–region level. * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$.

Figure C1: ConnectorCare Enrollment, 2016 to 2019



Notes: Panel A plots total ConnectorCare enrollment by year. Panel B plots enrollment by income group. The dashed line marks the removal of the incremental subsidy. Counts are enrollee-years in the enrollment panel and match [Table C1](#).

Figure C2: Insurer B's Below-Poverty Enrollees in the Claims Panel, 2016 to 2019



Notes: Panel A plots below-poverty person-years in the claims cost panel by insurer and year. Panel B plots the mean and median insurer cost per member-month among those enrollees.

Table C3: Dist. to B-exclusive hospitals by Income Group

	(1) <100% FPL	(2) 100–200% FPL	(3) 200–300% FPL
<i>Dist. to nearest B-exclusive hospital (miles)</i>			
Mean	13.70	19.74	20.52
Median	8.16	18.00	18.43
Standard deviation	12.61	15.54	15.60
Observations	10,869	64,970	50,626

Notes: The sample is ConnectorCare enrollees in 2016. Distance is the driving distance from the enrollee's ZIP code to the nearest hospital that only Insurer B covers.

D Appendix: Structural Model and Welfare Simulations

Table D1: The Expected-Cost Model

	Insurer cost (\$/member-month)
Log risk score	1.122 (0.014)
Female	0.129 (0.009)
Age	0.022 (0.002)
Age squared	−0.00022 (0.00002)
Income tier 3	0.018 (0.010)
Region and year effects	Yes
Insurer indicators	Yes
Pseudo R^2	0.41
Observations	582,309

Notes: The table reports the Poisson pseudo-maximum-likelihood regression behind the expected-cost function $\hat{c}_i$ of equation (19). The sample is the claims panel of above poverty enrollees from 2016 to 2019. The dependent variable is insurer cost per member-month. Log risk score is $\ln(1 + h_i)$. Income tier 3 is 200 to 300 percent of poverty. The omitted tier is 100 to 200 percent. Standard errors in parentheses are clustered by region.

Table D2: Private Willingness to Pay for Broad-Network Access

	Access value (\$/member-month)
100–150% FPL	8.86
150–200% FPL	7.62
200–250% FPL	2.23
250–300% FPL	3.70
Mean, 100–200% FPL	8.05

Notes: Each value is equation (D2) averaged within the income group with enrollment weights, in dollars per member-month.

D.1 Private Willingness to Pay for Broad-Network Access

This appendix measures enrollees' willingness to pay for Insurer B. I calculate this through measuring the compensating variation of removing Insurer B from the choice set at the observed-policy equilibrium. Holding all other premiums fixed, I convert each enrollee's expected-utility loss into dollars. Fix the observed-policy equilibrium with base prices b^* and the implied net premiums p_{ij} . For enrollee i in income group g with choice set $\mathcal{J}$, expected utility is the two-level logsum

$$V_i(\mathcal{J}) = \log \left(\sum_n \left[\sum_{j \in \mathcal{J}_n} \exp \left(\frac{u_{ij}}{1-\rho_n} \right) \right]^{1-\rho_n} \right), \quad (\text{D1})$$

where u_{ij} is the indirect utility of equation (18). Nests n partition the alternatives; Insurer N1 and Insurer N2 share the correlated taste nest with parameter ρ , and Insurer B stands alone. Let $\mathcal{J}_{-B}$ denote the same choice set with Insurer B removed. The access value in dollars per member-month is

$$AV_i = \frac{V_i(\mathcal{J}) - V_i(\mathcal{J}_{-B})}{|\alpha_g|}, \quad (\text{D2})$$

the expected-utility loss divided by the income group's marginal utility of income, and I aggregate to the group with enrollment weights, $AV_g = \sum_{i \in g} w_i AV_i / \sum_{i \in g} w_i$.

Results. Table D2 reports the estimates. The average enrollee values Insurer B's presence at \$8.86 and \$7.62 per member-month in the two groups that receive the incremental subsidy and at \$2.23 and \$3.70 in the two groups that do not, where Insurer B's equilibrium net premium is \$300 and \$359. The revealed planner valuation, \$3,821 (Appendix D.2), is about 475 times the mean private access value of the 100 to 200 percent groups, \$8.05.

The comparison is conservative. μ is a valuation per enrolled Insurer B member-month, and the access value averages over all group members, including the majority who never choose Insurer B. The two objects nonetheless price the same margin, Insurer B's presence for the group, and the multiple states the wedge between the planner's valuation and the group's mean private valuation. No welfarist planner whose objective scales the private valuations in Table D2 by plausible weights generates a wedge of this size.

D.2 Recovering the Access Value

This appendix details the recovery of the Tobin planner's valuation μ in Section 6.3 and the two thresholds in Panel C of Table D5. Throughout, $\mathcal{S}_B$ is low-income Insurer B enrollment per market enrollee: the number of enrollees at 100–200 percent of poverty who hold Insurer B, divided by all enrollees in the market.

Why the fitted policy. The recovery requires a one-parameter rationalization of the observed incremental subsidy rates. The observed 2016 policy is not exactly a member of the family in Section 6.1, since its two low-income rates, 0.727 and 0.587, do not lie on one line through the full subsidy below poverty. The member with the same average rate, 0.657, has rates 0.743 and 0.571, within two percentage points of each observed rate. I treat this fitted policy as the state's choice.

The condition. Along the family (incremental subsidy rates across income groups), the Tobin planner's welfare is $W(\bar{s}) = CS(\bar{s}) + \mu \mathcal{S}_B(\bar{s}) - 1.3 G(\bar{s})$, with $\omega_g = 1$. If the fitted policy is the planner's optimum, the first-order condition $dW/d\bar{s} = 0$ holds, and solving for μ gives equation (27): μ is the fall in $CS - 1.3 G$ per unit of $\mathcal{S}_B$ gained along the family. The ratio does not depend on how the family is indexed: reindexing by the rate at 100 to 150 percent of poverty, or by the slope of the line, multiplies numerator and denominator by the same factor.

Implementation. Along the family, the Tobin planner's welfare is $W(\bar{s}) = CS(\bar{s}) - 1.3 G(\bar{s}) + \mu \mathcal{S}_B(\bar{s})$. If the fitted policy $\bar{s}^*$ is the planner's optimum, then

$$\left. \frac{dW}{d\bar{s}} \right|_{\bar{s}^*} = \frac{d(CS - 1.3 G)}{d\bar{s}} + \mu \frac{d\mathcal{S}_B}{d\bar{s}} = 0 \implies \mu = - \frac{d(CS - 1.3 G)/d\bar{s}}{d\mathcal{S}_B/d\bar{s}}. \quad (D3)$$

I evaluate the two derivatives as central differences between the equilibria at average rates 0.600 and 0.700, the fitted policy's neighbours on the grid of Table D4:

$$\mu = - \frac{(CS - 1.3 G)|_{0.700} - (CS - 1.3 G)|_{0.600}}{\mathcal{S}_B|_{0.700} - \mathcal{S}_B|_{0.600}}. \quad (D4)$$

Between them, $CS - 1.3 G$ falls by \$12.25 per member-month (consumer surplus is unchanged to the cent and government cost rises by about \$9.4), while low-income Insurer B enrollment per market enrollee rises by 0.00321; the ratio in equation (D4) gives $\mu = 3,821$.

Two baselines. The recovery of μ uses the fitted member of the family, while the thresholds use the equilibria at the actual observed policy, averaged as in column (1) of Table D5, because that column is the one the main text reports. The two differ by \$3.8 of $CS - 1.3 G$ and 0.0009 of $\mathcal{S}_B$. Evaluated at the fitted member instead, the threshold against the level subsidy alone is \$1,723 rather than \$1,638, and the threshold against the larger level subsidy is \$1,310 rather than \$1,246. Both remain below \$3,821.

Table D3: Incremental Subsidy Policies and Equilibrium Base Prices

Incremental subsidy rate			Base price (\$ per month)		
100–150%	150–200%	Average	Insurer N1	Insurer B	Insurer N2
0.000	0.000	0.000	258	390	258
0.200	0.000	0.100	267	392	267
0.400	0.000	0.200	266	398	266
0.475	0.125	0.300	268	406	268
0.550	0.250	0.400	270	417	270
0.625	0.375	0.500	273	434	270
0.700	0.500	0.600	275	461	273
0.743	0.571	0.657	278	486	278
0.775	0.625	0.700	279	513	279
0.850	0.750	0.800	282	548 [†]	282
<i>Observed policy, two equilibria</i>					
0.727	0.587	0.657	273	484	273
0.727	0.587	0.657	280	484	280

Notes: Each row is a solved equilibrium on the family of Section 6.1. The two low-income rates lie on a line through the full subsidy below poverty, with the rate at 150 to 200 percent floored at zero. The rates above 200 percent of poverty are fixed at 0.065 and 0.002 in every row. The average $\bar{s}$ is the simple mean of the two low-income rates. The row at $\bar{s} = 0.657$ is the fitted policy of Appendix D.2. The final two rows use the observed policy with its region-specific rates, at which the solver finds two equilibria. Base prices are statewide, with all three re-solved.

[†]Insurer B's price sits at the solver's upper bound. I exclude this row from every calculation.

Table D4: Equilibrium Outcomes and Pricing Along the Family of Policies

<i>Panel A. Outcomes by policy</i>				
Average incremental rate	Consumer surplus	Insurer profits	Government cost	Low-income Insurer B
0.000	461.6	70.1	272.2	0.024
0.100	462.1	80.9	284.6	0.029
0.200	462.3	81.0	284.0	0.031
0.300	462.6	83.8	287.3	0.035
0.400	463.0	87.7	291.7	0.040
0.500	463.3	92.6	297.2	0.044
0.600	463.3	97.6	303.5	0.049
0.657	463.5	104.5	309.9	0.051
0.700	463.3	106.5	312.9	0.052
0.800 [†]	464.4	116.4	325.0	0.064
<i>Observed policy, two equilibria</i>				
0.657 (273, 484, 273)	463.2	96.3	302.4	0.049
0.657 (280, 484, 280)	463.4	105.0	311.4	0.051

Notes: Panel A reports outcomes for the schedules of [Table D3](#), indexed by the average incremental subsidy rate across the two low-income groups. Consumer surplus, insurer profits, and government cost are per member-month over all enrollees at 100 to 300 percent of poverty. Low-income Insurer B is S_B , low-income Insurer B enrollment per market enrollee. [†]Censored row, see [Table D3](#). Panel B reports pricing at the first row of Panel A and at the fitted policy.

Table D5: Simulated Outcomes and Welfare Behind the Planner Comparison

	Observed policy (level + incremental) (1)	Level only (2)	Increased level subsidies (3)
<i>Panel A. Simulated outcomes (\$ per member-month)</i>			
Consumer surplus	463.3	461.6	496.3
of which, 100–200 percent groups	264.0	261.2	295.8
Government cost	306.9	272.2	306.9
Insurer profits	100.7	70.1	70.1
Low-income Insurer B enrollment per market enrollee (S_B)	0.0500	0.0236	0.0236
<i>Panel B. Welfare by planner (\$ per member-month)</i>			
Efficiency	64.4	107.8	97.4
Utilitarian ($\gamma = 1.4$)	40.8	82.9	79.7
Utilitarian ($\gamma = 5$)	270.0	306.8	348.8
Tobin ($\mu = 3,821$)	255.6	197.8	187.4

Notes: All entries are dollars per member-month of the simulated population at 100 to 300 percent of poverty. Column (1) averages the two equilibria at the observed policy. Column (2) sets the incremental subsidy rates below 200 percent of poverty to zero and lets insurers reprice. Column (3) adds to column (2) a lump sum of \$34.7 per member-month—the estimated budget for incremental subsidies—to enrollees at 100 to 200 percent of poverty. I hold federal premium tax credits and cost-sharing payments at their observed levels. Panel B applies equation (25) to Panel A. The efficiency row is consumer surplus less 1.3 times government cost. The utilitarian rows weight each group’s surplus by $(y_{\text{med}}/y_g)^\gamma$, with the 150 to 200 percent group as baseline and group midpoints as y_g . The weights are 1.60, 1.00, 0.70, and 0.53 at $\gamma = 1.4$ and 5.38, 1.00, 0.28, and 0.10 at $\gamma = 5$. The Tobin row adds μS_B to the efficiency row at $\mu = 3,821$.

Table D6: A Broad Plan Credit Compared With the Incremental Subsidy: Equilibrium Outcomes

	Level only (1)	Observed policy (level + incremental) (2)	\$100 credit (level + credit) (3)
<i>Panel A. Base rates (\$ per month)</i>			
Insurer N1	258	279	261
Insurer B	390	484	381
Insurer N2	258	279	261
<i>Panel B. Outcomes (\$ per member-month)</i>			
Consumer surplus	461.6	463.5	467.3
Insurer profits	70.1	104.4	77.3
Government cost	272.2	309.9	288.5
Low-income Insurer B enrollment	0.0236	0.0506	0.0845
<i>Panel C. Welfare (\$ per member-month)</i>			
Efficiency	107.8	60.6	92.2
Tobin ($\mu = 3,821$)	197.8	253.9	414.9
<i>Panel D. Cost of low-income Insurer B enrollment, against column (1)</i>			
Government cost per unit of S_B (\$)		1,394	268

Notes: All columns are solved equilibria. Column (1) is the level subsidy alone, the $\bar{s} = 0$ equilibrium of Table D4. Column (2) is the observed policy, at the second of the two observed-policy equilibria in Table D3. Column (3) removes the incremental subsidy in every income group above poverty and pays \$100 per member-month toward Insurer B's premium for enrollees at 100 to 200 percent of poverty, with net premiums floored at zero and no cash paid. Consumer surplus, profits, and government cost are per member-month over all enrollees at 100 to 300 percent of poverty. S_B is low-income Insurer B enrollment per market enrollee. Efficiency welfare is consumer surplus less 1.3 times government cost. Tobin welfare adds μS_B at $\mu = 3,821$. Panel D divides the change in government cost from column (1) by the change in S_B from column (1).