Online Appendix for:
The Two Margin Problem in Insurance Markets

A Analysis in a General Model (Relaxing Vertical Assumptions)

In this appendix, we present a formal mathematical analysis of the equilibrium impacts of tuning the parameters governing the two main policies discussed in Section 3: the mandate penalty and risk adjustment. We implement this analysis in a general model that does not invoke the vertical assumptions used for our graphical approach. This lets us show how the vertical assumptions interact with the model's main predictions.

Horizontal differentiation allows for an additional margin of substitution, between H and U, that the vertical model shuts down. As we show below, this adds additional terms to the comparative statics defining the policy effects on prices and market shares. But as long as these H-U substitution terms are not too large—e.g., as long as when M increases, most of the newly insured buy the cheaper L plan, not H—then they do not reverse the sign of the vertical model predictions. Thus, our results are not a knife-edge case driven by the assumption of pure vertical differentiation. Rather, as long as vertical differentiation is the "main" way that H and L compete, the model provides a useful approximation. This is consistent with the findings of our empirical robustness check that allows for horizontal differentiation in Appendix D.4.1.

A.1 Model Setup

The setup is identical to that of Section 2, with two plans H and L and $P = \{P_H, P_L\}$ denoting insurer prices. Let $G = \{S_H, S_L, M\}$ denote plan-specific government subsidies (S_j) and the mandate penalty (M). Throughout this section (as in Section 2), we assume $S_H = S_L = S$, though the framework would generalize if this were not true. Nominal consumer prices equal $P_j^{cons} = P_j - S$ for $j = \{L, H\}$ and $P_U^{cons} = M$.

Unlike in the vertical model, we will not assume that W_H and W_L are perfectly correlated. Instead, we allow consumers to vary along both willingness to pay dimensions. Each consumer type is characterized by an ordered pair $s = (s_H, s_L)$, where s_H indexes WTP for H and s_L indexes WTP for L. We once again normalize $W_U \equiv 0$. Note that a single s-index is no longer sufficient to characterize consumer willingness-to-pay. Without loss of generality, the s index takes a bivariate uniform distribution, so it represents an index of the percentile of the WTP distribution for H and L.

The set of consumers who choose a given option $j \in \{H, L, U\}$ is defined as $A_j(P, G) = \{s : W_j(s) - P_j^{cons} \geq W_k(s) - P_k^{cons}\ \forall k\}$. Demand is defined as the size of this group: $D_j(P, G) = \int_{A_j(P,G)} ds$.

For each "WTP-type," we once again have a plan-specific expected cost $C_j(s)$. We again make the adverse selection assumption that costs in a given plan are increasing in WTP for that plan. Hence $\partial C_j(s_H, s_L)/\partial s_j < 0$ for plan j. Average costs for plan $j \in \{L, H\}$ equal the average of $C_j(s)$ over the enrolling set of consumers:

$$AC_j(P; G) = \frac{1}{D_j(P; G)} \int_{A_j(P,G)} C_j(s) ds \tag{6}$$

Similarly, we can define the average risk score functions:

$$\overline{R}_j(P;G) = \frac{1}{D_j(P;G)} \int_{A_j(P,G)} R(s)ds \tag{7}$$

where $R(s)$ is the average risk score among type-s consumers. The baseline per-enrollee risk adjustment transfer from L to H is a function of these average risk scores, the (share-weighted) average risk score in the market ($\equiv \overline{R}(P;G)$) and the (share-weighted) average price in the market ($\equiv \overline{P}(P;G)$):

$$T(P;G) = \left(\frac{\overline{R}_H(P;G)}{\overline{R}(P;G)} - 1 \right) \overline{P}(P;G). \tag{8}$$

Finally we introduce a parameter $\alpha \in (0,1)$ that multiplies the transfer, $\alpha \cdot T(P;G)$, allowing us to vary the strength of risk adjustment by scaling the transfers up or down such that $\alpha = 0$ represents no risk adjustment, $\alpha \in (0,1)$ is partial risk adjustment, $\alpha = 1$ is full-strength risk adjustment, and $\alpha > 1$ is over-adjustment.

We define equilibrium as prices equal average costs net of risk adjustment transfers:

$$\begin{aligned} P_H &= AC_H\left(P;G\right) - \alpha T(P;G) \equiv AC_H^{RA}(P;G,\alpha) \\ P_L &= AC_L\left(P;G\right) + \alpha T(P;G) \equiv AC_L^{RA}(P;G,\alpha) \end{aligned} \tag{9}$$

where $AC_j^{RA}(P;G,\alpha)$ are risk-adjusted costs for plan $j = \{L,H\}$.

A.2 Approach and Assumptions on Signs of Demand/Cost Curve Slopes

We now consider the equilibrium response to an increase in the uninsurance penalty M and an increase in α, i.e. the strength of the risk adjustment transfers. Our goal is to understand the cross-margin interactions—the effect of M on demand for H and the effect of risk adjustment on the share uninsured. To do so, we use the equilibrium conditions to derive the relevant comparative statics, $\frac{dD_H}{dM}$ and $\frac{dD_U}{d\alpha}$. The comparative statics take account of both direct effects—denoted with partial derivatives below (e.g., $\frac{\partial AC_H}{\partial P_H}$)—and equilibrium effects on market prices—denoted with total derivatives (e.g., $\frac{dP_H}{dM}$). These comparative statics allow us to show the features of demand and cost that determine the sign and magnitude of the cross-margin effects.

In analyzing these comparative statics, we will assume a *stable equilibrium* that is characterized by *adverse selection*. These assumptions let us sign the slopes of several demand/cost curves that enter the equations. In particular, we assume:

- Equilibrium stability, which requires that $1 - \frac{\partial AC_j}{\partial P_j} > 0$ for $j = \{H,L\}$ locally to the equilibrium point.
- Adverse selection, which requires that (on average) the highest-cost types buy H, middle-cost types buy L, and the lowest-cost choose U. More specifically, we assume:
    1. The marginal H consumer is lower-cost than the average H consumer and higher-cost than the average L consumer—which implies that $\frac{\partial AC_H}{\partial P_H} > 0$ and $\frac{\partial AC_L}{\partial P_H} > 0$.

2. The consumer on the margin of H and L is lower-cost than the average H consumer—so $\frac{\partial AC_H}{\partial P_L} < 0$
3. The marginal uninsured consumers are lower-cost than the average consumer of H or L, so $\frac{\partial AC_H}{\partial M} \leq 0$ and $\frac{\partial AC_L}{\partial M} \leq 0$.

For the analysis of risk adjustment, we also assume that the analogous stability and adverse selection conditions hold for *risk-adjusted* average costs AC_H^{RA} and AC_L^{RA}. This is true in our empirical simulations, where we find that risk adjustment is imperfect, so risk-adjusted cost curves are characterized by adverse selection.

Further, while we do not impose the vertical model, it is useful to note its implications for several relevant partial derivatives:

- Vertical model assumes that no consumers are on the H-U margin, which implies that $\frac{\partial D_H}{\partial M} = \frac{\partial AC_H}{\partial M} = \frac{\partial D_U}{\partial P_H} = 0$.

In the analysis below, we color in red the terms that are zero under the vertical model. This lets readers see where relaxing the vertical assumptions adds additional terms to the comparative statics.

A.3 Increase in Uninsurance Penalty (M)

We derive comparative statics for enrollment in H in response to a change in the uninsurance penalty M. Throughout this section, we assume that there is no risk adjustment in place, which simplifies the math.

We start by analyzing $\frac{dD_H}{dM}$, the cross-margin effect of a mandate penalty on enrollment in H. This comparative static is comprised of two parts. First, in red is the direct enrollment change in H for a change in M, holding fixed P_H and P_L. In the vertical model, this $\frac{\partial D_H}{\partial M}$ term would be zero. The second term is the indirect effect on D_H through the change in relative prices of H and L. Formally:

$$\frac{dD_H}{dM} = \underbrace{\frac{\partial D_H}{\partial M}}_{\text{HU margin}} + \underbrace{\underbrace{\frac{\partial D_H}{\partial \Delta P_{HL}}}_{(-)} \cdot \left(\frac{dP_H}{dM} - \frac{dP_L}{dM} \right)}_{\text{HL margin}}. \tag{10}$$

In the vertical model, $\frac{\partial D_H}{\partial M} = 0$, so under the vertical assumption the sign of $\frac{\partial D_H}{\partial M}$ would be fully determined by the change in the incremental price of H vs. L caused by an increase in M. If an increase in M leads to an increase in $\Delta P_{HL} = P_H - P_L$, then an increase in M will lead to lower demand for H. This positive relationship between M and ΔP_{HL} would occur under our assumptions about adverse selection because an increase in M would induce a fall in P_L as the consumers on the margin between L and U who are induced to purchase L are relatively healthy. If the vertical model does not hold, $\frac{\partial D_H}{\partial M} > 0$, which would partly offset the decrease in D_H but not fully do so as long as it is small in magnitude.

Thus, to sign the cross-margin effect, we need to show that $\frac{dP_H}{dM} - \frac{dP_L}{dM} > 0$. We now fully differentiate P_H and P_L with respect to M to characterize this relationship more explicitly.

$$\begin{aligned} \frac{dP_H}{dM} &= \frac{\partial AC_H}{\partial M} + \frac{\partial AC_H}{\partial P_H}\frac{dP_H}{dM} + \frac{\partial AC_H}{\partial P_L}\frac{dP_L}{dM} \\ \frac{dP_L}{dM} &= \frac{\partial AC_L}{\partial M} + \frac{\partial AC_L}{\partial P_H}\frac{dP_H}{dM} + \frac{\partial AC_L}{\partial P_L}\frac{dP_L}{dM} \end{aligned} \tag{11}$$

Notice, that unlike under the purely vertical model, a change in M impacts direct costs for both H and L. Solving this system of equations again for $\frac{dP_H}{dM}$, we get the expression below.

$$\frac{dP_H}{dM} = \left[\frac{\partial AC_H}{\partial M} + \frac{\partial AC_L}{\partial M}\frac{\partial AC_H}{\partial P_L}(1 - \frac{\partial AC_L}{\partial P_L})^{-1}\right] \times \Phi_H^{-1} \tag{12}$$

where $\Phi_H = \{1 - \frac{\partial AC_H}{dP_H} - \frac{\partial AC_H}{\partial P_L}\frac{\partial AC_L}{\partial P_H}(1 - \frac{\partial AC_L}{\partial P_L})^{-1}\}$.

We now can sign $\frac{dP_H}{dM}$ as follows:

$$\frac{dP_H}{dM} = \left[\underbrace{\frac{\partial AC_H}{\partial M}}_{\text{Ext. Margin Selection}(\leq 0)} + \underbrace{\underbrace{\frac{\partial AC_L}{\partial M}}_{(-)} \cdot \underbrace{\frac{\partial AC_H}{\partial P_L}}_{(-)} \underbrace{\left(1 - \frac{\partial AC_L}{\partial P_L}\right)^{-1}}_{(+)}}_{\text{Substitution to } L\ (+)}\right] \times \underbrace{\Phi_H^{-1}}_{(+)} \tag{13}$$

and $\Phi_H = \underbrace{(1 - \frac{\partial AC_H}{dP_H})}_{(+)} - \underbrace{\frac{\partial AC_H}{\partial P_L}\frac{\partial AC_L}{\partial P_H}}_{(-)}\underbrace{(1 - \frac{\partial AC_L}{\partial P_L})^{-1}}_{(+)} > 0$, where all signs are determined by the adverse selection and stability assumptions laid out above.

Therefore, we can sign $\frac{dP_H}{dM} > 0$ under the vertical model. The intuition is as we have already described: the mandate penalty lowers P_L, leading relatively healthy H consumers to leave H and substitute to L, which raises AC_H and therefore P_H. When the vertical model does not hold, extensive margin selection of consumers on the HU margin into H ($\frac{\partial AC_H}{\partial M} < 0$) pushes in the other direction. But as long as extensive margin substitution is not too large, the main effect of substitution to L will dominate.

We derive the expression for $\frac{dP_L}{dM}$ in a similar way:

$$\frac{dP_L}{dM} = \left[\underbrace{\frac{\partial AC_L}{\partial M}}_{\text{Ext. Margin Selection}(-)} + \underbrace{\underbrace{\frac{\partial AC_H}{\partial M}}_{(\leq 0)} \cdot \underbrace{\frac{\partial AC_L}{\partial P_H}}_{(+)} \underbrace{\left(1 - \frac{\partial AC_H}{\partial P_H}\right)^{-1}}_{(+)}}_{\text{Substitution to } H(\leq 0)}\right] \times \underbrace{\Phi_L^{-1}}_{(+)} \tag{14}$$

where $\Phi_L = \{1 - \frac{\partial AC_L}{dP_L} - \frac{\partial AC_L}{\partial P_H}\frac{\partial AC_H}{\partial P_L}(1 - \frac{\partial AC_H}{\partial P_H})^{-1}\} > 0$ as with Φ_H above.

Thus, under the vertical model where $\frac{\partial AC_H}{\partial M} = 0$, we can unambiguously say that P_L falls with a higher mandate penalty ($\frac{dP_L}{dM} < 0$). This conclusion also holds when we relax the vertical model (as shown by the negative substitution term), as any extensive margin substitution into H acts to lower the price of H, drawing the sickest consumers away from L and pushing L's costs and price even further down.

Returning now to $\frac{dD_H}{dM}$, we observe under the vertical model that $\left(\frac{dP_H}{dM} - \frac{dP_L}{dM}\right) < 0$, which implies that $\frac{dD_H}{dM} > 0$. In other words, the "unintended consequence" of decreasing enrollment in H should always occur under the vertical model. When we relax the vertical model, this result will also hold as long substitution on the HU margin is not too large.

A.4 Increasing the Strength of Risk Adjustment (α)

We now consider in our more general model the effect of a small increase in the α parameter on the share of the population that is uninsured. As in the previous section, we color in red the terms that are zero under the vertical model. This lets readers see where relaxing the vertical assumptions adds additional terms to the comparative statics.

The change in the share of the uninsured population given a change in α is comprised of two parts: changes in enrollment from the HU margin (in red) and LU margin (in black). Under the vertical model assumptions, the HU margin is not present.

$$\frac{dD_U}{d\alpha} = \underbrace{\underbrace{\frac{\partial D_U}{\partial \Delta P_{HU}}}_{\geq 0} \frac{d\Delta P_{HU}}{d\alpha}}_{\text{HU margin}} + \underbrace{\underbrace{\frac{\partial D_U}{\partial \Delta P_{LU}}}_{(+)} \frac{d\Delta P_{LU}}{d\alpha}}_{\text{LU margin}} \tag{15}$$

where $\Delta P_{HU} = P_H - S - M$ and $\Delta P_{LU} = P_L - S - M$ are the net prices of H and L relative to uninsurance.

By the law of demand, $\frac{\partial D_U}{\partial P_H} \geq 0$, $\frac{\partial D_U}{\partial P_L} > 0$. Under the vertical model, $\frac{\partial D_U}{\partial P_H} = 0$, so the cross-margin effect of risk adjustment on uninsurance is entirely determined by the sign of the LU margin. We now consider the impact of a change in α on ΔP_{HU} and ΔP_{LU}. The change in prices depends on the nature of subsidies. With subsidies linked to the price of L, ΔP_{LU} $(= P_L - S - M)$ is fixed by construction. Therefore, the LU margin of substitution is shut down. In the vertical model, we will have $\frac{dD_U}{d\alpha} = 0$.

Let us now consider the case where there is a fixed subsidy and therefore prices can be affected by the level of transfers. We fully differentiate (9) and rearrange to get a system of equations. These are identical under both the horizontal and vertical model.

$$\frac{dP_H}{d\alpha} = \underbrace{T(.)}_{(+)} \times \left[\underbrace{-1}_{\text{Direct}(-)} + \underbrace{\frac{\partial AC_H^{RA}}{\partial P_L} \left(1 - \frac{\partial AC_L^{RA}}{\partial P_L}\right)^{-1}}_{\text{Substitution from L }(-)} \right] \times (\Phi_H^{RA})^{-1} < 0$$

where $\Phi_H^{RA} \equiv 1 - \frac{\partial AC_H^{RA}}{\partial P_H} - \frac{\partial AC_L^{RA}}{\partial P_H} \frac{\partial AC_H^{RA}}{\partial P_L} (1 - \frac{\partial AC_L^{RA}}{\partial P_L})^{-1}$. As in the mandate section above, this Φ_H^{RA} term must be positive under the assumptions on stability and adverse selection we have made.

The term in brackets is composed of two effects. First, there is a direct effect of stronger risk adjustment transferring money to H, which tends to lower P_H. Second, there is an indirect substitution effect, arising from substitution of relatively healthy consumers on the margin between H and L opting for H and lowering H's average cost and thus its price. Thus, $\frac{dP_H}{d\alpha} < 0$ because both the direct and indirect effects push P_H down.

Doing the same for $\frac{dP_L}{d\alpha}$ gives

$$\frac{dP_L}{d\alpha} = \underbrace{T(.)}_{(+)} \times \left[\underbrace{1}_{\text{Direct}(+)} + \underbrace{\left(-\frac{\partial AC_L^{RA}}{\partial P_H}\right)\left(1 - \frac{\partial AC_H^{RA}}{\partial P_H}\right)^{-1}}_{\text{Substitution to H }(-)} \right] \times \underbrace{(\Phi_L^{RA})^{-1}}_{(+)}$$

where $\Phi_L^{RA} \equiv 1 - \frac{\partial AC_L^{RA}}{\partial P_L} - \frac{\partial AC_H^{RA}}{\partial P_L}\frac{\partial AC_L^{RA}}{\partial P_H}(1 - \frac{\partial AC_H^{RA}}{\partial P_H})^{-1}$, which must be positive under the stability and adverse selection assumptions.

Here, the direct effect is positive because larger transfers take money from L, driving up the price of L. However, the indirect substitution effect is negative—since $\frac{\partial AC_L^{RA}}{\partial P_H} > 0$ by adverse selection. Intuitively, stronger risk adjustment transfers increase the price of L, causing consumers on the H-L margin to opt for H instead of L. These consumers are the highest-cost L enrollees, implying that their exit from L will lower L's average cost and thus its price. Therefore, the indirect substitution effects will mute (or even fully offset) the direct effect of risk adjustment on P_L. Because of this direct and indirect effect, it is ambiguous whether P_L will increase or decrease, and in general, any change in P_L will be smaller than one would expect from the direct effect alone.

Further, the question of whether the direct or indirect effect dominates depends on whether the substitution term is greater than or less than 1 in absolute value. If it is greater than 1, then the substitution term will dominate. This will occur if $\frac{\partial AC_L^{RA}}{\partial P_H} > 1 - \frac{\partial AC_H^{RA}}{\partial P_H}$. This will tend to occur when intensive margin adverse selection is very strong (even after risk adjustment) so that both $\frac{\partial AC_L^{RA}}{\partial P_H}$ and $\frac{\partial AC_H^{RA}}{\partial P_H}$ are large. Conversely, if adverse selection is weak, the direct effect will dominate.

This expression also tells us how the size of any cost advantage for L may affect the effects of increasing α. When L has no cost advantage over H (the cream-skimmer case), the only reason L gets any demand is intensive margin adverse selection. When adverse selection is strong in the cream-skimmer case, L exists but the substitution effect is also large, muting the direct effect of risk adjustment. When adverse selection is weak in the cream-skimmer case, L fails to exist. Thus, it is more likely that increasing α will have little or no (or possibly negative) effect on P_L in the case where L has no cost advantage than in the case where L has a cost advantage.

To summarize the case with fixed subsidies, $\frac{dD_U}{d\alpha}$ is ambiguous even under the vertical model because we cannot theoretically sign the change in P_L when when α increases. If the direct effect dominates, then P_L will increase with α and uninsurance will rise under the vertical model. If the substitution to H dominates, then P_L will fall and uninsurance will also fall.

When we relax the vertical assumptions, the potential for stronger risk adjustment to increase uninsurance is further mitigated by the presence of the HU extensive margin. The term $\frac{\partial D_U}{\partial P_H}\frac{dP_H}{d\alpha}$ in equation (15) will be positive. Because $\frac{dP_H}{d\alpha} < 0$, consumers on the HU margin will tend to become insured (in H) when risk adjustment is strengthened. This may offset any rise in uninsurance along the LU margin if P_L rises, as more consumers leave uninsurance to buy H.

B Appendix: Extensions to the Graphical Model

B.1 Graphical Analysis of Perfect Risk Adjustment

In this section, we illustrate how our graphical model can be used to show the effects of perfect risk adjustment on equilibrium prices and market shares. Under *perfect* risk adjustment, transfers perfectly capture all variation in C_L across consumer types. The graphical representation of the role of risk adjustment in the two margin problem is complicated by the fact that risk adjustment transfers cause RAC_H (the risk-adjusted cost curve) to become an equilibrium object rather than a stable market primitive (like AC_H), as any effects of selection into the market are at least partially shared between L and H due to the risk-based transfers.

To simplify exposition, we assume that the causal cost difference between H and L equals a constant value of δ for all consumer types s. We define perfect risk adjustment as transfers such that the average cost in H net of risk adjustment always equals the average cost in L net of risk adjustment plus δ: $RAC_H(P) = RAC_L(P) + \delta$. Under perfect risk adjustment, the average risk-adjusted cost in H and L does not depend on consumer sorting between H and L. Instead, the average cost of both plans depends only on consumer sorting between insurance and uninsurance. If new healthy consumers join the market (buying the L plan), the risk transfers share the improved risk pool equally between H and L, maintaining the δ difference between their average costs. The important simplifying feature of *perfect* risk adjustment is that when it comes to average costs, there is only one relevant margin of adjustment: the extensive margin. With *imperfect* risk adjustment, residual intensive margin selection that is not compensated by risk adjustment remains relevant, complicating the graphical analysis.

We depict the perfect risk adjustment case in Figure A1. Note that here we do not assume that L is a pure cream-skimmer but instead that L has a cost advantage equal to δ. Risk adjustment affects the curves in a number of ways. First, as depicted in panel (b), risk adjustment causes the average cost curve for L to shift upward and rotate slightly to make it parallel with the original, unadjusted average cost curve for H. This shift reflects the risk transfer away from L (and to H) that raises L's effective costs. $RAC_L(s_{LU})$ still slopes down because of extensive margin adverse selection, but it is now a fixed curve that does not depend on the price of H or sorting between H and L.[45] The new, higher average cost curve for L, RAC_L implies a new, higher equilibrium price for L, $\hat{P}^e_L$. This higher price of L implies a new demand curve for H, shifted upward from the previous demand curve and depicted in panel (c) of Figure A1. This higher demand curve for H reflects the fact that the higher price of L makes L less attractive relative to H.

Panel (d) of Figure A1 illustrates the second direct effect of risk adjustment. For the H plan, risk adjustment causes the average cost curve, $RAC_H(s_{HL})$, to be *rotated downward* relative to the unadjusted curve, $AC_H(s_{HL})$. RAC_H is now a flat line, since sorting between plans (i.e., the value of s_{HL}) does not affect average costs. The level of RAC_H equals $AC_H(s_{LU})$—the average cost if the entire population up to the extensive margin type s_{LU} were to enroll in H.

Figure A2 shows how this shift in H's average cost curve combines with the shift in H's demand curve to produce a new lower equilibrium price of H, $\hat{P}^e_H$ and a higher quantity of consumers enrolling in H.

[45] One can show that RAC_L is parallel to the old AC_H since it is capturing the overall average costs of everyone from $s = 0$ up to a given s_{LU} cutoff.

Figure A1: Equilibrium under Perfect Risk Adjustment

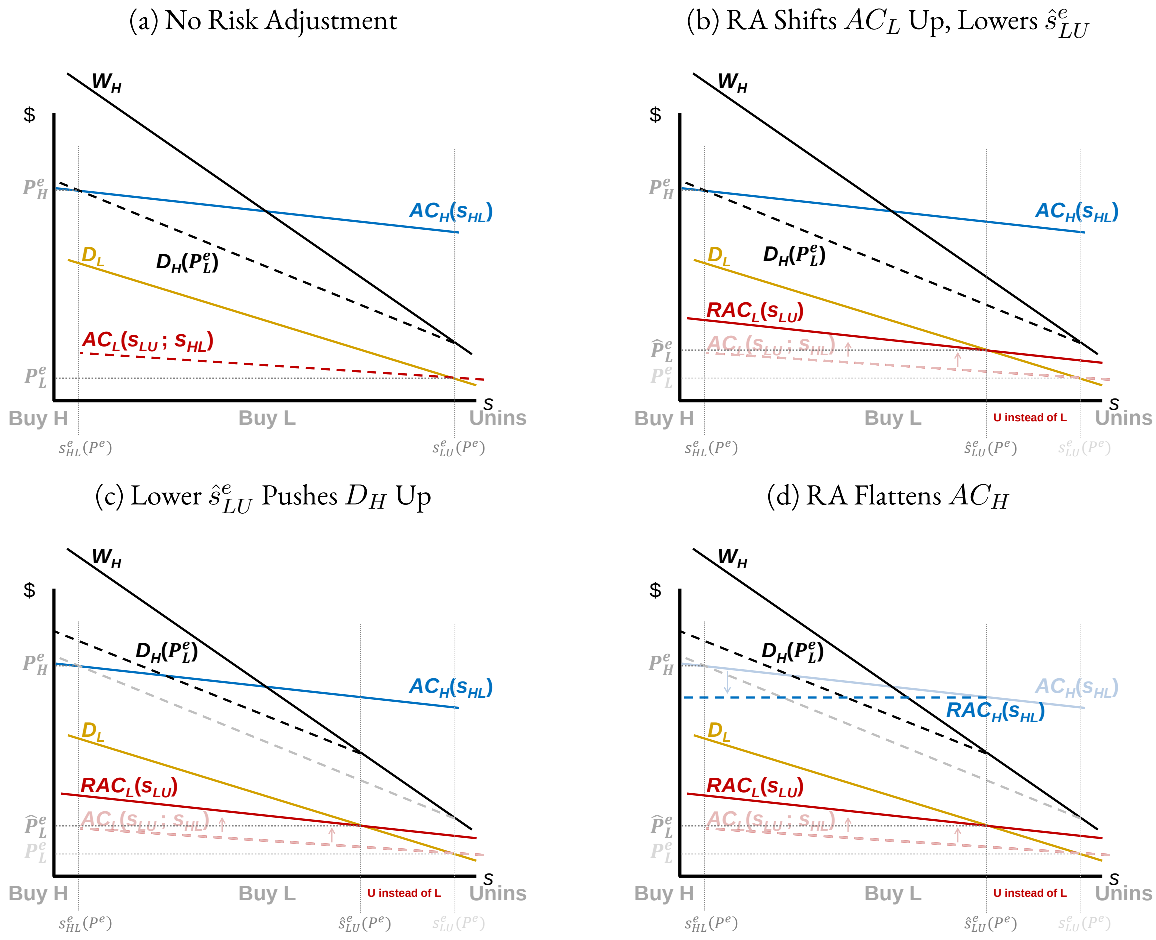

Notes: Starting from equilibrium in panel (a) and introducing perfect risk adjustment in panel (c), perfect risk adjustment shifts up the average cost of L from $AC_L(s_{LU})$ to $RAC_L(s_{LU})$, reflecting the transfer away from L to H. Unlike AC_L, the risk adjusted RAC_L only depends on the extensive margin S_{LU}, not on the allocation across plans (s_{HL}). The risk adjusted curve $RAC_L(s_{LU})$ intersects D_L at a lower point, shifting out the extensive margin from s^e_{LU} to $\hat{s}^e_{LU}$. Next, in panel (c) we see that this lower extensive margin-type $\hat{s}^e_{LU}$ shifts up D_H. Finally, in panel (d) we see that risk adjustment flattens the risk adjusted average cost of H, RAC_H, which like RAC_L no longer varies depending on sorting between the two plans, s_{HL}.

Figure A2: Equilibrium under Perfect Risk Adjustment

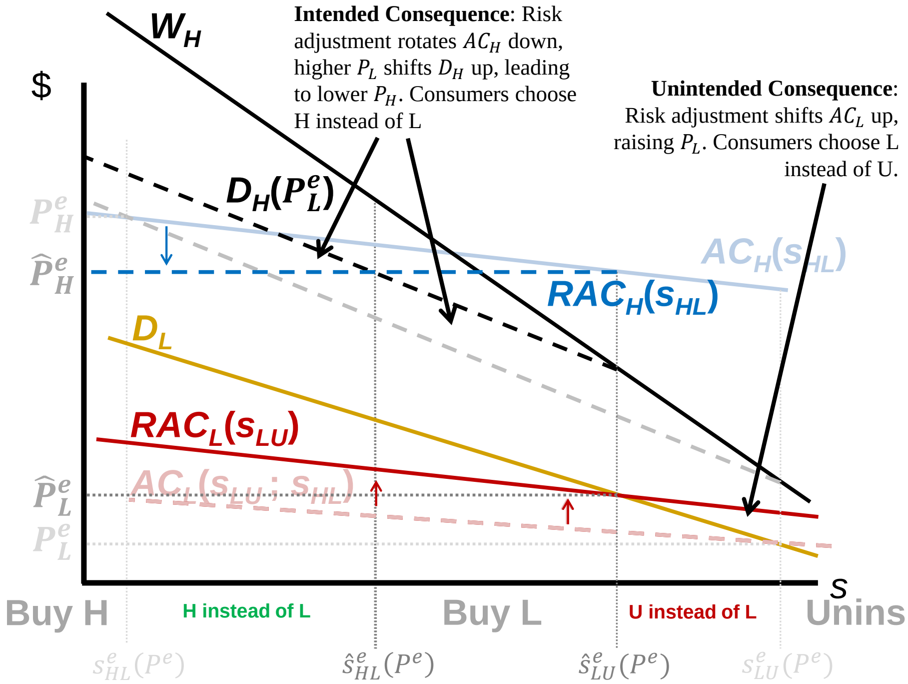

Notes: Under perfect risk adjustment, the risk-adjusted average cost curve for H is completely flat for a given s_{LU}. Equilibrium occurs at s_{HL} and s_{LU} values such that RAC_H intersects D_H and RAC_L intersects D_L.

In summary, perfect risk adjustment has two effects. First, it causes the average cost curve for H to rotate downward until it is flat. This rotation of the cost curve causes s_{HL} to shift right, indicating a shift of consumers from L to H. This is the intended effect of risk adjustment, and it is caused by a transfer from L to H to compensate H for the externality imposed on it by intensive margin selection from L. Second, it causes the average cost curve for L to both rotate and shift up.[46] This change in AC_L causes s_{LU} to shift left, indicating a shift of consumers from L to U, increasing uninsurance. This is the unintended effect of risk adjustment. It occurs because the transfer to H comes from L, resulting in an increase in L's costs and price, forcing some consumers out of the market. In Section 3 we also provide a graphical description of the welfare consequences of risk adjustment, both perfect and imperfect.

In Appendix A and Appendix D.4.1 we also explore (both theoretically and empirically) how the effects of risk adjustment are affected by the relaxation of our vertical model assumption, finding that the presence of consumers with non-vertical preferences can act to weaken the unintended effects of risk adjustment on the extensive margin.

Finally we note that if risk adjustment is perfect—as assumed in this subsection—it will often lead to countervailing effects with some consumers opting for H instead of L and other consumers opting for U instead of L. With imperfect risk adjustment, in contrast, the unintended extensive margin effect may or may not occur, depending on the relative sizes of the direct and indirect effects.

[46]The curve remains downward-sloping because perfect risk adjustment only addresses intensive margin selection, leaving selection on the extensive margin in place.

B.2 Extension: Medicare Advantage + Traditional Medicare

Our graphical model can be extended to other cases beyond the baseline $H/L/U$ setup modeled on the ACA Marketplaces. One setting of particular policy interest is the Medicare Advantage (MA) market, in which plans of varying quality compete with an outside option of Traditional Medicare (TM). A key difference for the MA-TM setting is that the inside-option plans are *advantageously* selected relative to the outside option. Unlike the ACA case where the outside option of uninsurance attracts the lowest-cost consumers, TM has historically attracted the sickest and highest-cost enrollees.[47] We show in this section how our graphical model can capture the MA-TM case under the maintained assumption of vertical differentiation. (For non-vertical differentiation, a 2-D graphical approach is not feasible, but see the math in Appendix A that captures the general case.)

The MA-TM extension works as follows. We start by setting up a model with three vertically ranked plans: (1) TM, the most preferred option; (2) H, a high-quality MA plan (middle option); and (3) L, a lower-quality MA plan (least preferred). We think of TM as representing Traditional Medicare bundled with a generous Medigap plan so that it is the most generous option for both cost sharing and provider network. H could be a broad-network MA plan (e.g., a PPO), while L could be a narrow-network MA plan (e.g., an HMO). Importantly, we assume that TM is the outside option whose price is set *exogenously* by policymakers (e.g., via the Part B premium and rules for MA subsidies/benchmarks), while the prices of H and L are determined in equilibrium. We note, of course, that the real-world MA-TM market is much more complicated than this setup and that vertical differentiation is an approximation. Our model should be seen as an approximation, and the caveats discussed for our baseline model also apply here.

To capture advantageous selection with respect to TM, we *reorder the plan sorting* along our maintained "s type" x-axis. Rather than have the lowest-WTP types choose the outside option (of uninsurance) as in our baseline model, the highest-WTP types now choose the outside option of TM. Middle-WTP types choose the H MA plan, and the lowest-WTP types choose the L MA plan. We will continue to assume that WTP correlates with sickness (cost), so the sickest types choose TM, middle types choose H, and the healthiest types choose L. This reordering lets us define demand and costs curves and competitive equilibrium in a similar manner as in our baseline H/L/U model. We note that this reordering is different from the EFC-graph approach to advantageous selection, which instead uses upward sloping curves corresponding to a market where consumer preference for more generous coverage itself is negatively correlated with costs.

Formally, we maintain the vertical model assumptions of Section 2 with labeling changes. We normalize $W_{i,L} \equiv 0$ and make the following two assumptions:

Assumption 3. *Vertical ranking:* $W_{i,TM} > W_{i,H} > W_{i,L} \equiv 0$ *for all* i

Assumption 4. *Single dimension of WTP heterogeneity: There is a single index* $s \sim U[0,1]$ *that orders consumers based on declining WTP, such that* $W'_H(s) < 0$ *and* $W'_{TM}(s) - W'_H(s) < 0$ *for all* s.

We assume that the consumer price of TM, P_{TM}, is set exogenously. The prices of the H and L MA plans are set competitively to equal their average costs:

$$P_H = AC_H(P) \qquad \text{and} \qquad P_L = AC_L(P) \tag{16}$$

[47] There is evidence that in recent years, improved risk adjustment has offset some of these differences. The model in this section should be seen as illustrative of the traditional case where MA was still advantageously selected.

As in the baseline model, there could be non-uniqueness, and we limit attention to equilibria that meet the requirements of the Riley Equilibrium (RE) notion (see Appendix C.3). For the graphical presentation, we focus on the case of monotonic adverse selection in which higher-WTP correlates with higher costs. For graphical simplicity, we also focus on the pure cream-skimming case where $C_H(s) = C_L(s)$ for all s. The more general case would be similar but would involve plotting two separate type-specific cost curves. Finally, we depict the case with positive demand for all contracts, though in principle the model allows one more contracts to unravel.

Figure A3 shows equilibrium in the MA-TM case under these assumptions. The graph is similar to equilibrium in the baseline H/L/U case (see Figure 4) but with a few differences. First, the price of *TM* is exogenous and there is therefore no need to show the average cost of *TM*. Second, all demand and average cost curves are now equilibrium objects that depend on P_H or P_L; it is no longer possible to define AC_H and D_L based on primitives alone. This makes the setup slightly more complex to describe, but the basic concepts and cross-margin policy effects are similar.

Walking through Figure A3, suppose we start with an exogenous P_{TM} (set by policymakers) and an initial guess for P_H and P_L. The demand curve that determines sorting between TM and H is $D_{TM}(s) = W_{TM}(s) - W_H(s) + P_H$. The type indifferent between these two options is $s^*_{TM,H}$, defined by $D_{TM}(s^*_{TM,H}) = P_{TM}$. Types to the left of this point ($D_{TM}(s) > P_{TM}$) choose TM, while types to the right of this point choose H or L. Sorting between H vs. L is determined by the yellow curve $D_H(s) \equiv W_H(s) + P_L$, with indifferent type $s^*_{H,L}$ defined by $W_H(s^*_{H,L}) + P_L = P_H$. Types to the left of $s^*_{H,L}$ choose H (since $D_H(s) > P_H$), while types to the right choose L (since $D_H(s) < P_H$). Notice that both the dashed black and yellow curves equal WTP (for TM and H) shifted upward by P_L. This is similar to the way that the mandate penalty (price of the lowest-quality option) shifted upward WTP for insurance plans in our baseline H/L/U model, but in this case the price of L is endogenous.

Turning to costs, the pink curve is the type-specific cost curve, $C_H(s) = C_L(s)$, for this pure cream-skimming case (though this would be easy to generalize). The average cost curve for H starts at $s^*_{TM,H}$ and slopes downward to the right (lying above the $C_H(s)$ curve), capturing the average costs of all individuals choosing H (i.e., $s \in [s^*_{TM,H}, s_{HL}]$). In equilibrium, $AC_H(s)$ intersects $D_H(s)$ at $s = s^*_{H,L}$ so that $AC_H(s^*_{H,L}) = P_H$. For the L plan, the average cost curve starts at this $s^*_{H,L}$ type and slopes downward to the right (lying above the $C_L(s)$ curve). Since all $s \in [s^*_{H,L}, 1]$ choose L, the final average costs of L equals the value of $AC_L(s)$ at $s = 1$. In equilibrium, $AC_L(1) = P_L$.

This model can also be used to think about cross-margin policy effects. For instance, suppose the government decreases the price of *TM*, intending to get more consumers to choose the higher-quality *TM* option. Some consumers then shift from H into *TM* at the $s^*_{TM,H}$ margin, captured by a movement along the D_{TM} curve. These people leaving H are its highest-cost consumers, so the AC_H curve shifts downward, resulting in a lower P_H and a shift from L to H on the intensive margin. Therefore, a change in the extensive margin price (P_{TM}) results in a demand shift on the intensive margin from the L to the H plan. Notice, however, that unlike the H/L/U case, the cross-margin effects *reinforce* the original policy's goal of getting consumers into higher-quality plans. In addition to the intended shift from H to TM (higher quality), there is a cross-margin shift from L to H (also higher quality). In words, lowering the price of TM results in more MA enrollees choosing higher-quality MA options.[48]

[48] Similar analysis could also be applied to study the cross-margin impact of a risk adjustment transfer from L to H, which might lower the price of H and draw consumers into H from TM. In reality in the MA-TM market, risk adjustment applies across all three options, making the analysis somewhat different.

Figure A3: Equilibrium in Medicare Advantage + Traditional Medicare Case

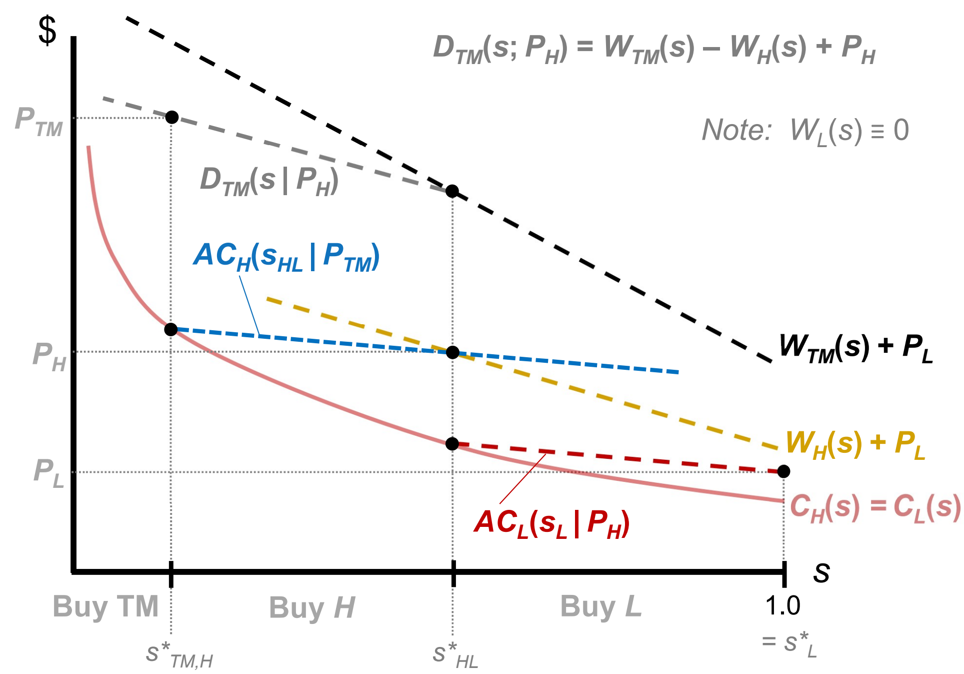

Notes: The graph shows equilibrium in the Medicare Advantage (MA) + Traditional Medicare (TM) case, as described in the appendix text. Assumptions, curve setups, and equilibrium are similar to our baseline H/L/U model, but with sorting reordered so that the highest-WTP (furthest left) types choose the outside option of *TM*, middle-WTP types choose the higher-quality MA plan (*H*), and the lowest-WTP types choose the lower-quality MA plan (*L*).

B.3 Formal Social Welfare Function

In this appendix, we derive a formal expression for welfare, building on the graphical presentation in the body text. We allow for cases where C_U is non-zero—e.g., if the outside option involves social costs like uncompensated care.

We define social welfare as:

$$\widehat{SW}(P) = \int_{0}^{s_{HL}(P)} (W_H(s) - C_H(s))\,ds + \int_{s_{HL}(P)}^{s_{LU}(P)} (W_L(s) - C_L(s))\,ds - \int_{s_{LU}(P)}^{1} C_U(s)\,ds \quad (17)$$

Recall that the level of utility was normalized above by setting $W_U = 0$. As in the figures, we can express welfare in terms of three curves and two areas (integrals) if we make the following transformations. First, add a constant equal to total potential cost of U, defining $SW = \widehat{SW} + \int_0^1 C_U(s)\,ds$. Second, define "net costs" of L (in excess of C_U) as $C_L^{Net}(s) \equiv C_L(s) - C_U(s)$. Rearranging and simplifying, this yields the following expression for social welfare:

$$SW = \underbrace{\int_0^{s_{HL}(P)} \left(W_H^{Net}(s) - W_L(s)\right) ds}_{\text{Intensive Margin Surplus from } H \text{ vs. } L} + \underbrace{\int_0^{s_{LU}P)} \left(W_L(s) - C_L^{Net}(s)\right) ds}_{\text{Extensive Margin Surplus from } L \text{ vs. } U} \quad (18)$$

The first term is the intensive margin surplus (H vs. L) for consumers who buy H, $s \in [0, s_{HL}]$. Notice that $W_H^{Net}(s) - W_L(s) = \Delta W_{HL} - \Delta C_{HL}$, so this is indeed capturing the intensive margin surplus. The second term is the extensive margin surplus from insurance (in L) relative to uninsurance, which applies to everyone who buys insurance, $s \in [0, s_{LU}]$. Equation (18) shows that it is straightforward to calculate welfare even when $C_U \neq 0$, as long as the researcher has information about C_U.

C Appendix: Simulation Method Details

C.1 Constructing Demand and Cost Curves

As discussed in section 4, we draw on separate demand and cost estimates for both low-income subsidized consumers from Finkelstein, Hendren and Shepard (2019) (abbreviated "FHS") and high-income unsubsidized consumers from Hackmann, Kolstad and Kowalski (2015) (abbreviated "HKK"). We describe how each respective paper produced its primitives as well as our modifications below.

C.1.1 Low-Income Demand and Costs: FHS (2019)

FHS Primitives

- Population: FHS estimate insurance demand in Massachusetts' pre-ACA subsidized health insurance exchange, known as "CommCare." CommCare was an insurance exchange created under the state's 2006 "Romneycare" reform to offer subsidized coverage to low-income non-elderly adults (below 300% of poverty) without access to other health insurance (from an employer, Medicare, Medicaid, or another public program). This population was similar, though somewhat poorer, than the subsidy-eligible population under the ACA.

- Market structure: CommCare participation was voluntary: consumers could choose to remain uninsured and pay a (small) penalty. As FHS show, a large portion of consumers (about 37% overall) choose the outside option of uninsurance, despite the penalty and large subsidies. The CommCare market featured competing insurers, which offered plans with standardized (state-specified) cost sharing rules but which differed on their provider networks. In 2011, the main year that FHS estimate demand, the market featured a convenient vertical structure among the competing plans. Four insurers had relatively broad provider networks and charged nearly identical prices just below a binding price ceiling imposed by the exchange. One insurer (CeltiCare) had a smaller provider network and charged a lower price. FHS pool the four high-price, broad network plans into a single "H option"—technically defined as each consumer's preferred choice among the four plans—and treat CeltiCare as a vertically lower-ranked "L option." FHS present evidence that this vertical ranking is a reasonable characterization of the CommCare market in 2011.

- FHS Estimation: To estimate demand and costs, FHS use a regression discontinuity design leveraging discontinuous cutoffs in subsidy amounts based on household income. Because subsidies vary across income thresholds, there is exogenous net price variation that can transparently identify demand and cost curves with minimal parametric assumptions. FHS leverage discontinuous changes in net-of-subsidy premiums at 150% FPL, 200% FPL, and 250% FPL arising from CommCare's subsidy rules. They estimate consumer willingness-to-pay for the lowest-cost plan (L) and incremental consumer willingness-to-pay for the other plans (H) relative to that plan.[49] This method provides estimates of the demand curve for particular ranges of s. The same variation is used to estimate $AC_H(s)$ and $C_H(s)$, the average and marginal cost curves for H. Our goal is not to innovate on these estimates but rather to apply them as primitives in our policy simulations to understand the empirical relevance of our conceptual framework.

Our Modifications to FHS Primitives

- Extrapolating to extremes of s distribution: The FHS strategy provides four points of the $W_L(s)$ curve and four points of the $W_{HL}(s) = W_H(s) - W_L(s)$ curve. As shown in Figure 10 from FHS, for the W_L curve these points span from $s = 0.36$ to $s = 0.94$ and for the W_{HL} curve these points span from $s = 0.31$ to $s = 0.80$. Because our model allows for the possibility of zero enrollment in either L or H or both, we need to modify the curves, extrapolating to the full range of consumers, $s \in [0, 1]$. We start by extrapolating linearly, and then we "enhance" demand for H among the highest WTP consumers, as we view this as more realistic than a linear extrapolation. (We explore the sensitivity of our empirical results to alternative assumptions about this WTP enhancement in Appendix D.4.2) We then smooth the enhanced demand curves to eliminate artificial kinks produced by the estimation and extrapolation.

  (1) Linear demand: For the linear demand curves, we extrapolate the curves linearly to $s = 0$ and $s = 1.0$. Call these curves $W_L^{lin}(s)$ and $W_H^{lin}(s)$, with incremental WTP defined as $W_{HL}^{lin} = W_H^{lin} - W_L^{lin}(s)$.

  (2) Enhanced demand: For the enhanced demand curves $\left(W_L^{enh}(s) \text{ and } W_H^{enh}(s)\right)$, we inflate consumers' relative demand for H vs. L in the extrapolated region, relative to a linear extrapolation. We implement enhanced demand in an *ad hoc* but transparent way: We first generate $W_L^{enh}(s) = W_L^{lin}(s)$ for all s. For all $s >= 0.31$ (the boundary of the "in-sample" region of $W_{HL}(s)$), we likewise set $W_{HL}^{enh}(s) = W_{HL}^{lin}(s)$. For $s = 0$, we set $W_{HL}^{enh}(s = 0) = 3W_{HL}^{lin}(s = 0)$, so that the maximum enhanced incremental willingness-to-pay is three times the value suggested by the primitives. We then linearly connect the incremental willingness to pay between s=0 and s=0.31, setting $W_{HL}^{enh}(s < 0.31) = W_{HL}^{lin}(s) + 3 \times \frac{(0.31-s)}{0.31} \times W_{HL}^{lin}(0)$ so that the enhanced curve is equal to the linear curve for $s >= 0.31$, equal to three times the linear curve at $s = 0$, and linear between $s = 0.31$ and $s = 0$. This approach assumes that there exists a group of (relatively sick) consumers who exhibit very strong demand for H relative to L, which seems likely to be true in the real world. Thus,

$$W_{HL}^{enh}(s) = \begin{cases} W_{HL}^{lin}(s) & \text{for } s \in [0.31, 1] \\ W_{HL}^{lin}(s) + 3 \times \frac{(0.31-s)}{0.31} \times W_{HL}^{lin}(0) & \text{for } s \in [0, 0.31) \end{cases} \tag{19}$$

[49] Because the base subsidy for L and the incremental subsidy for H change discontinuously at the income cutoffs, there is exogenous variation in both the price of L and the incremental price of H.

and

$$W_H^{enh}(s) = W_L^{lin}(s) + W_{HL}^{enh}(s). \tag{20}$$

Both the linear and the enhanced WTP curves are shown in the top panel of Figure A4.

- Cost of L plan: We need to produce estimates of $C_L(s)$ to complete the model. FHS provide suggestive evidence that $C_L(s)$ is quite similar to $C_H(s)$—i.e., that for a given enrollee, L does not save money relative to H. We conducted further analyses to provide additional evidence on this question (leveraging entry of the L plan in some areas but not others, leveraging additional price variation for L vs. H, etc.), consistently finding a lack of evidence of any cost advantage for L among the enrollees marginal to these sources of variation. While L may indeed be a pure cream-skimmer in this setting, the assumption that $C_H(s) = C_L(s)$ for all s seems unlikely to hold in many other settings. Thus, we consider both the setting where L has a 15% cost advantage so that $C_L(s) = 0.85C_H(s)$ and the setting where, consistent with the empirical evidence, L is a pure cream-skimmer, i.e. $C_L(s) = C_H(s)$.

- Smoothing primitives: Because they were estimated using a regression discontinuity design, the primitives above all have discrete "kink points" at which the slope of the curve with respect to the share of the population enrolled changes discretely. In these regions, equilibrium allocations are extremely sensitive to small changes in policy parameters. To avoid this unrealistic sensitivity, we smooth the cost curves as well as the enhanced demand curves using a fourth degree polynomial. Specifically, for primitive $Y(s)$, we run the following regression.

$$Y = \hat{\beta}_0 + \hat{\beta}_1 s + \hat{\beta}_2 s^2 + \hat{\beta}_3 s^3 + \hat{\beta}_4 s^4 + \epsilon$$

  Using the fitted coefficients, we then use the predicted value $\hat{Y}$,

$$\hat{Y} = \hat{\beta}_0 + \hat{\beta}_1 s + \hat{\beta}_2 s^2 + \hat{\beta}_3 s^3 + \hat{\beta}_4 s^4$$

  This "smoothing" process was done on both the WTP curves as well as the cost curve primitives.

C.1.2 High-Income Demand and Costs: HKK (2015)

For our simulations, we also consider demand of higher-income groups, which allows us to simulate policies closer to the ACA. Under the ACA, low-income households receive subsidies to purchase insurance while high-income households do not. We construct WTP curves for high-income households using estimates of the demand curve for individual-market health insurance coverage in Massachusetts from Hackmann, Kolstad and Kowalski (2015) ("HKK").

HKK Primitives

- Population: HKK estimate demand in the unsubsidized pre-ACA individual health insurance market in Massachusetts, which is for individuals with incomes above 300% of poverty (too high to qualify for CommCare).

- Estimation: HKK use the introduction of the state's individual mandate in 2007-08 as a source of exogenous variation to identify the insurance demand and cost curves. HKK only estimate demand for a single L plan.

Our Modifications to HKK Primitives

- Constructing $W_L^{HI}(s)$: We start by constructing $W_L^{HI}(s)$, based on the estimates from Hackmann, Kolstad and Kowalski (2015). The superscript HI refers to high income. The HKK demand curve takes the following form:

$$W_{HKK}(s) = -\$9,276.81 * s + \$12,498.68 \tag{21}$$

  This demand curve is "in-sample" in the range of $0.70 < s < 0.97$. As with the low-income, subsidized consumers, we linearly extrapolate $W_{HKK}(s)$ out-of-sample to construct $W_L^{HI,lin}(s)$. Specifically, we let $W_L^{HI,lin}(s) = W_{HKK}(s)$ for all s.

- Constructing $W_H^{HI,lin}(s)$ and $W_H^{HI,enh}(s)$: HKK only estimate demand for a single L plan. Similar to FHS, we start by estimating a linearly extrapolated WTP for H, $W_H^{HI,lin}(s)$, and then "enhance" demand for H among the highest WTP types, $W_H^{HI,enh}(s)$, using the W_{HL}^{lin} and W_{HL}^{enh} as constructed for the low-income population above (i.e. we assume that extensive margin WTP for insurance is different between the high-income and low-income groups, but intensive margin WTP for H vs. L is the same):

$$W_H^{HI,lin}(s) = W_L^{HI} + W_{HL}^{lin}(s)$$
$$W_H^{HI,enh}(s) = W_L^{HI} + W_{HL}^{enh}(s)$$

- Constructing $C_L^{HI}(s), C_H^{HI}(s)$: We assume that the cost curves for this group are equivalent to the cost curves of the subsidized population, Thus,

$$C_H^{HI}(s) = C_H(s)$$
$$C_L^{HI}(s) = C_L(s)$$

  where $C_H(s)$ is drawn from FHS and $C_L(s)$ is the curve as constructed in the previous section. We note that these assumptions imply that the high-income consumers have a level shift in WTP with no difference in the extent of intensive or extensive margin selection from the low-income consumers.

- Smoothing primitives: Similar to above, we also smooth primitives.

We thus have two demand systems: one for low-income consumers and one for high-income consumers. Both exhibit WTP for H that is "enhanced" for the highest WTP types beyond what a simple linear extrapolation would imply. We combine these systems to form one set of demand and cost curves, by assuming that 60% of the market is low-income and 40% of the market is high-income, consistent with the population in the ACA Health Insurance Marketplaces.

Figure A4: WTP Curves for H and L

(a) Low-Income

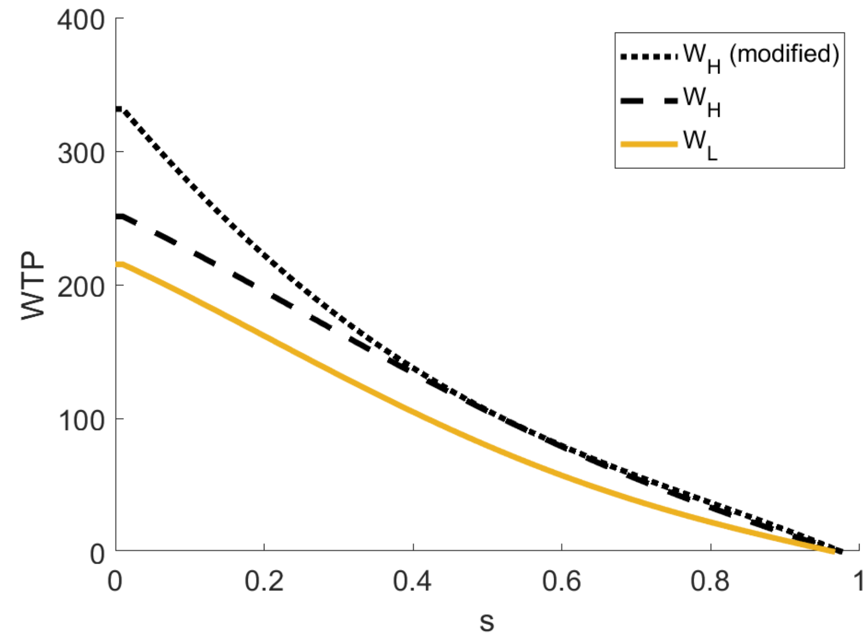

(b) High-Income

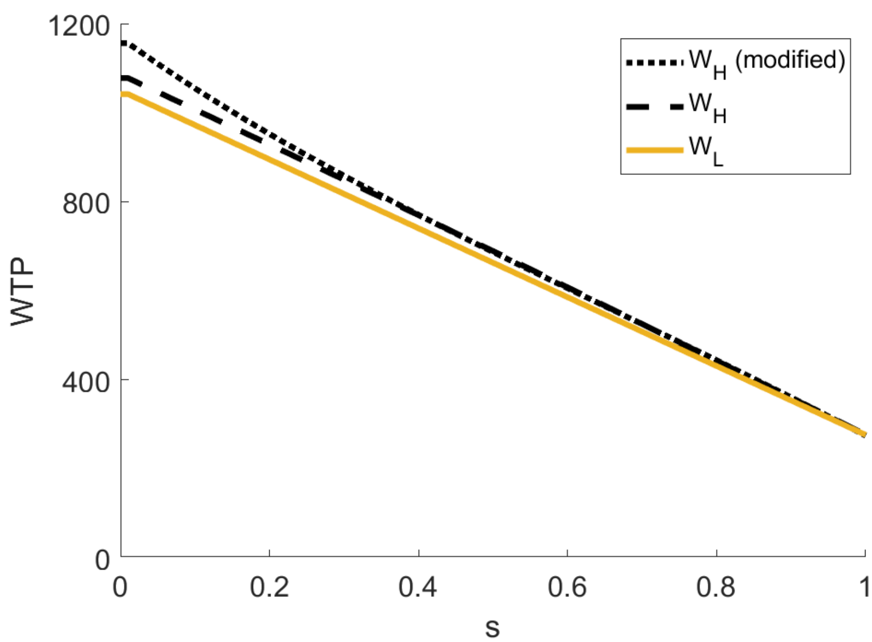

Notes: Figure shows WTP Curves for H and L, $W_H(s)$ and $W_L(s)$. The top panel shows curves for low-income group which come from (Finkelstein, Hendren and Shepard, 2019). The bottom panel shows curves for high-income group which come from (Hackmann, Kolstad and Kowalski, 2015). Linear curves extrapolate linearly over the out-of-sample range [0,0.31]. Modified (i.e. "enhanced") curves assume that the lowest s-types have very high incremental WTP for H.

C.2 Estimation of Risk Score Curve

Like WTP and costs, we use FHS's regression discontinuity approach to estimate a risk adjustment function for each s-type, $R(s)$. This function characterizes the *expected* cost of each s-type, as predicted by the actual risk scores of each enrollee, RA_i^{HCC}. To calculate this, we first compute these scores for each individual in our data, based on diagnosis codes present in the individual-level claims. All risk scores are computed using the Hierarchical Condition Categories (HCC), a risk adjustment model used by the Centers for Medicare and Medicaid Services for the ACA Marketplaces.[50]

Once we have a risk score for each individual in the data A_i^{HCC}, the risk score curve $R(s)$ was identified off of the same premium discontinuities as used to identify the demand curve in FHS. We then connect and smooth segments in a similar fashion to our construction of the cost and WTP curves to generate the $R(s)$ we use in our analysis. Similar to our assumption that the cost curve $C_H(s)$ estimated on the subsidized population applies to the un-subsidized population, we assume that this $R(s)$ curve estimated on the un-subsidized population also applies to the subsidized population.

Figure A5 shows a measure of risk-adjusted costs for the H plan in comparison to raw costs $C_H(s)$. It plots $C_H(s)$ and $C_H(s)/R(s)$; the latter would be constant in s under perfect risk adjustment. Consistent with risk adjustment being meaningful but imperfect, the risk-adjusted cost curve is much flatter than raw costs but still downward sloping. Over the $s \in [0, 1]$ interval, the risk-adjusted cost curve falls by about \$130, compared to a fall of \$367 in raw costs. Thus, by this measure, risk scores net out about 35% of the cost variation along the marginal cost curve for H. Since this simulation exclusively uses cost and risk score primitives from the subsidized population of pre-ACA Massachusetts, this finding should not necessarily be seen as generalizable to the entire ACA exchange population.

C.3 Riley Equilibrium Concept

We consider equilibria that meet the requirements of the Riley Equilibrium (RE) notion. In words, a price vector of $P = (P_H, P_L)$ is a Riley Equilibrium if there is no profitable deviation for which there is no "safe" (i.e. weakly profitable) reaction that would make the deviating firm incur losses. We slightly modify the definition presented in Handel, Hendel and Whinston (2015) below

DEFINITION 1: A *Riley Equilibrium* is a set of break-even price offers $P \in \mathcal{P}^{BE}$ for which there exists no Riley Deviation P'. A Riley Deviation (P') is a set of offers such that $P' \cup P$ is closed and $P' \cap P = \emptyset$. This P' is a Riley deviation if the following criteria are satisfied.

1. The Riley Deviation plan P' is weakly profitable and garners non-zero enrollment when the original prices are also offered: $P'_j \geq AC_j(P'_j)$ when $P \cup P'$ is offered and $P'_j \neq P_j$ (Note that this deviates from Handel, Hendel and Whinston (2015), which requires that the Riley Deviation is *strictly* profitable)
2. No "Safe Response" (P'') exists

We define a safe response as a set of price offers P'' such that $P \cup P' \cup P''$ is closed and P'' is disjoint from $P \cup P' \cup P''$ such that

[50] In practice, the methodology involves grouping diagnoses into different conditions, such as diabetes, etc. Individuals are then assigned risk scores based on the weighted value of all of their conditions. CMS publishes its weights annually on its website (https://www.cms.gov/medicare/health-plans/medicareadvtgspecratestats/risk-adjustors.html)

Figure A5: Raw Costs (C_H) versus Risk-Adjusted Costs

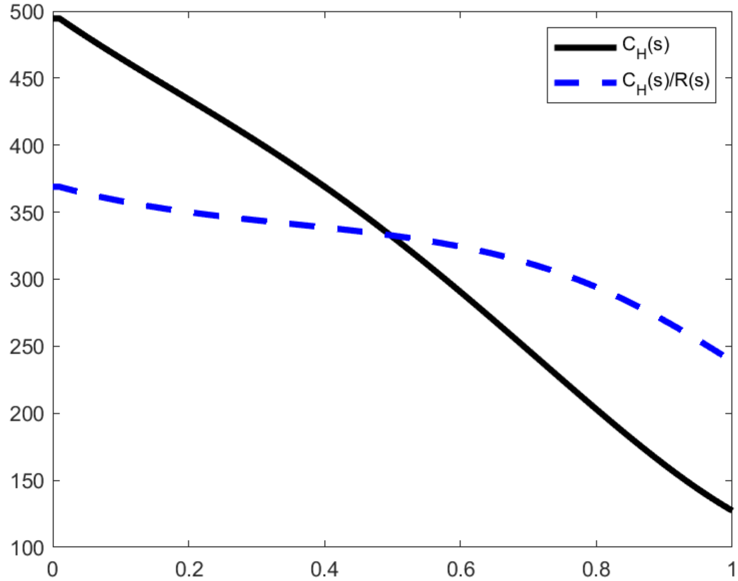

Notes: Figure shows raw C_H (black, continuous line) and risk-score normalized C_H (blue, dashed). While the risk score is able to flatten out the cost curve somewhat, not all risk is captured by the score, leaving some slope.

1. P' incurs losses when $P \cup P' \cup P''$ is offered

2. P'' does not incur losses when any market offering $\hat{P}$ containing $P \cup P' \cup P''$ is offered

It is straightforward to show that in our setting no price vector that earns positive profits for either L or H is a RE (see Handel, Hendel and Whinston, 2015 for a proof). This limits potential REs to the price vectors that cause L and H to earn zero profits. We refer to these price vectors as "breakeven" vectors. This set consists of the following potential vectors:

1. No Plan Enrollment: Prices are so high that no consumer enrolls in H or L

2. L-only: P_H is high enough that no consumer enrolls in H while P_L is set such that P_L equals the average cost of the consumers who choose L.

3. H-only: P_L is high enough that no consumer enrolls in L while P_H is set such that P_H equals the average cost of the consumers who choose H.

4. H and L: P_L and P_H are set such that both L and H have positive enrollment and P_L is equal to the average cost of the consumers who choose L and P_H is equal to the average cost of the consumers who choose H.

To simplify exposition, in Section 2 we assume that there is a unique RE such that there is positive enrollment in both H and L. However, we note that under certain conditions the competitive equilibrium will instead consist of positive enrollment in only one of the two plan options. We allow for these possibilities in the empirical portion of the paper and are able to find an unique RE where at least one plan has non-zero enrollment for every setting tested. See Appendix C.4 for details on the algorithm.

C.4 Reaction Function Approach to Finding Equilibrium

Evaluating demand, profits: For each uninsurance penalty, risk adjustment strength, L-plan cost advantage, and subsidy type setting, we find the equilibrium price configuration (P_H, P_L) using the following grid-search method. We construct a grid of P_H, P_L price combinations, with H on the vertical axis and L on the horizontal axis. For most simulations, we use a coarse grid with $1 units. For each pair, we evaluate H and L profits using the demand, cost, and risk-adjustment equations as detailed in the body of the paper. For insurance types H, L and uninsurance U we evaluate demand by finding the "indifference points"–the first and the last points in the s distribution such that each type of insurance's enrollment conditions are satisfied. Because of the vertical model, we can attribute all intermediate points of the s distribution between these indifference points to a given plan. If no points on the s vector satisfy the plan's enrollment conditions, the plan has zero enrollment. We have indifference points s_{HL}, s_{LU} if both H and L have non-zero enrollment and s_{HU}, s_{LU} if L or H has zero enrollment, respectively. If there is non-zero demand for both H and L, we calculate the average risk of those enrolled in each plan and construct transfers from the less risky plan to the more risky plan, per the ACA risk adjustment formula (see equation 4). In some counterfactual policy simulations, the transfer is multiplied by α. Finally, average costs are calculated for each plan with non-zero enrollment. The function returns the H, L profit grids Π^H, Π^L with which we can then evaluate equilibrium.

Finding equilibrium: For a given grid coarseness, we set a tolerance value T equal to the increment between grid points. A plan is considered to have zero profits if its profits are between $-T$ and T. Potential equilibria are all price pairs where (1) only H has non-zero enrollment and is making zero profits (2) only L has non-zero enrollment and is making zero profits (3) both H and L have non-zero enrollment and are both making zero profits. Given the coarseness of the grid, there are usually multiple potential equilibria of each type. We use the following process to refine this set down to the final equilibrium point according to our concept of the Riley Equilibrium.

- Single plan equilibria: First, we refine our $L-$only and $H-$only equilibria. For the remainder of this paragraph, we will refer to the potential $L-$only equilibria, but an analagous methodology also applies to refining potential $H-$only equilibria. Let $\mathcal{P}^{L-only}$ be the set of potential $L - only$ equilibria. Price vector $(P_H, P_L) \in \mathcal{P}^{L-only}$ iff. at (P_H, P_L)

  1. $\Pi^L(P_H, P_L) \in [-T, T]$
  2. L has nonzero enrollment
  3. H has zero enrollment.

  Given the curved nature of the primitives, for some settings, especially those where L has a large cost advantage, there are multiple unique P_L that are potential L-only equilibrium vectors.

  Further, for each potential L-only P_L, $\exists\ P_H^{min}\ s.t.\ \forall\ P_H > P_H^{min}(P_H, P_L) \in \mathcal{P}^{L-only}$[51] For each potential L-only equilibrium price P_L, we evaluate whether the conditions of a Riley Equilibrium are satisfied at (P_H^{min}, P_L). We need only evaluate P_H^{min} since any potential deviations from

[51] If at (P_H, P_L), L has non-zero enrollment and earns zero profits and H gets zero enrollment, then if H increases its price to $P'_H > P_H$, enrollment allocations will remain exactly the same and L will continue to make zero-profits.

(P_H^{min}, P_L) would also be deviations from (P_H, P_L), $P_H > P_H^{min}$.

To test an L-only equilibrium for an H-deviation, the process is as follows: Starting with the lowest $P_L \in \mathcal{P}^{L-only}$, the Riley Equilibrium refinement algorithm evaluates whether a Riley Deviation exists for a given potential L-only P_L using three nested loops. For L-only equilibria (P_L, P_H^{min}), we consider H-only Riley Deviations (P'_H, P_L) where $P'_H < P_H^{min}$.

1. *Find Potential Riley Deviations:* The outer loop evaluates each $P'_H < P_H^{min}$ to identify whether $\Pi^H(P'_H, P_L) > T$ (i.e. H makes positive profits). If no such potential H-deviations are found, (P_H^{min}, P_L) is considered a RE. If a potential H-deviation is found, the second loop is called.
2. *Find Potential Retaliations:* This loop evaluates each grid point (P'_H, P'_L), $P'_L < P_L$ to identify potential L-retaliations where $\Pi^L(P'_H, P'_L) > -T, \Pi^H(P'_H, P'_L) < -T$ (i.e. L makes weakly positive profits and H makes negative profits. If no such potential retaliations are found for a given potential H-deviation, then (P_H^{min}, P_L) is not a Riley Equilibrium (since there exists a Riley Deviation with no retaliation).
3. *Determine if Retaliation is "Safe":* If a potential retaliation is found, a third loop is activated to evaluate if there is any point (P''_H, P'_L), $P''_H < P'_H$ that makes a given retaliation "unsafe" where unsafe is defined as $\Pi^L < -T$ (i.e. L makes negative profits). If no such "unsafe" point exists, then the retaliation point is safe and the potential deviation would not succeed.

If no retaliation-proof deviation exists for a given (P_L, P_H^{min}), then the point is a RE. If a deviation does exist, the next larger $(P'_L, P_H min') \in \mathcal{P}^{L-only}$ is tested.

- H-L equilibria: Because of the coarseness of the grid, there are usually multiple connected points where both H and L have enrollees and are making zero profits. We pick the point with the lowest P_L to evaluate. For each potential HL equilibrium, we test if any single-plan deviations exist. This consists of checking whether any Riley Deviations that change P_H holding fixed P_L or change P_L holding fixed P_H exist, using the same set of RE loops described in the previous paragraph. If either type of deviation is found, the HL equilibrium is not an RE.

We apply this algorithm to every cost, risk adjustment, mandate penalty, and subsiyd type setting and in every case are able to find an unique equilibrium that satisfies our Riley Equilibrium conditions.

D Appendix: Additional Simulation Results

D.1 Simulation Results for Mandate/Uninsurance Penalty

Tables A1 and A2 Show additional outcomes for the mandate/uninsurance penalty simulations discussed in Section 5 and shown in Figure 9. In all cases, the welfare measure represents the social surplus under the particular policy setting as a percent of the difference between minimum possible social surplus and maximum possible social surplus achieved.

Table A1: Varying Mandate Penalty

(a) ACA-like subsidy, L cream-skimmer

mandate	0	15	30	45	60
price H	382	374	371	360	349
price L	352	344	337	325	313
share H	.42	.42	.3	.26	.23
share L	.31	.37	.55	.67	.77
share U	.27	.21	.15	.069	0
subsidy	297	289	282	270	258
welfare	.91	.76	.49	.24	0

(b) Fixed $275, L cream-skimmer

mandate	0	15	30	45	60
price H	387	381	373	349	349
price L	357	351	341	313	313
share H	.42	.42	.37	.23	.23
share L	.24	.3	.44	.77	.77
share U	.35	.28	.19	0	0
subsidy	275	275	275	275	275
welfare	.93	.79	.56	0	0

(c) ACA-like subsidy, L cost advantage

mandate	0	15	30	45	60
price H	414	409	404	399	.
price L	307	300	292	283	273
share H	.021	.017	.013	.0065	0
share L	.73	.79	.86	.93	1
share U	.25	.19	.13	.067	0
subsidy	252	245	237	228	218
welfare	.95	.75	.52	.27	0

(d) Fixed $250 , L cost advantage

mandate	0	15	30	45	60
price H	415	404	.	.	.
price L	307	294	273	273	273
share H	.019	.016	0	0	0
share L	.73	.84	1	1	1
share U	.26	.15	0	0	0
subsidy	250	250	250	250	250
welfare	.27	.16	0	0	0

Notes: Table A1 contains equilibrium prices, market shares, subsidy levels and relative welfare under varying levels of mandate penalties. Panels (a) and (b) are results for when L is a cream-skimmer ($\Delta C_{HL} = 0$) while panels (c) and (d) are for when L has a 15% cost advantage. In panels (a) and (c), the market has a price-linked subsidy while in panels (b) and (d), the market has a fixed subsidy. Relative welfare is calculated as $\frac{welfare - min(welfare)}{max(welfare) - min(welfare)}$ where max and min are taken over integer mandate penalty values 0 to 60 under the panel's same L cost advantage, subsidy scheme.

D.2 Simulations of Benefit Regulation

Tables A3 and A4 characterize equilibrium results with and without an L-plan offered when the L-plan is a pure cream-skimmer and when L has a 15% cost advantage. For a given setting, the welfare loss is reported in dollars and represents loss relative to welfare under the optimal allocation.

The results indicate that for the ACA-like price-linked subsidies, removing L from the choice set always (weakly) improves welfare. This is because removing L results in a higher subsidy and more people entering the market. In the fixed subsidy cases, we find that removing L often causes both an increase in H's market share and an increase in the uninsurance rate (especially when L has a 15% cost advantage). However, we find that in all cases, benefit regulation improves welfare, implying that the welfare losses from more people being uninsured are more than offset by welfare gains from more people enrolling in H.

Table A2: Varying Risk Adjustment (α)

(a) ACA-like subsidy, L cream-skimmer

α	0	.5	1	1.5	2
price H	.	437	382	362	362
price L	372	362	352	.	.
share H	0	.082	.42	.78	.78
share L	.72	.64	.31	0	0
share U	.28	.28	.27	.22	.22
subsidy	317	307	297	307	307
welfare	.46	.59	.91	.91	.91

(b) Fixed $275, L cream-skimmer

α	0	.5	1	1.5	2
price H	495	438	387	377	377
price L	381	369	357	.	.
share H	.0095	.097	.42	.66	.66
share L	.57	.52	.24	0	0
share U	.42	.38	.35	.34	.34
subsidy	275	275	275	275	275
welfare	.68	.73	.93	1	1

(c) ACA-like subsidy, L cost advantage

α	0	.5	1	1.5	2
price H	.	.	414	361	362
price L	308	308	307	313	.
share H	0	0	.021	.16	.78
share L	.75	.75	.73	.59	0
share U	.25	.25	.25	.25	.22
subsidy	253	253	252	258	307
welfare	.93	.93	.95	.99	.58

(d) Fixed $250, L cost advantage

α	0	.5	1	1.5	2
price H	.	.	415	365	381
price L	309	309	307	316	.
share H	0	0	.019	.16	.6
share L	.74	.74	.73	.56	0
share U	.26	.26	.26	.29	.4
subsidy	250	250	250	250	250
welfare	.24	.24	.27	.48	1

Notes: Table A2 contains equilibrium prices, market shares, subsidy levels and relative welfare under varying strengths of risk adjustment α. Panels (a) and (b) are results for when L is a cream-skimmer ($\Delta C = 0$) while panels (c) and (d) are for when L has a 15% cost advantage. In panels (a) and (c), the market has a price-linked subsidy while in panels (b) and (d), the market has a fixed subsidy. relative welfare is reported as $\frac{welfare - min(welfare)}{max(welfare) - min(welfare)}$ where max and min are taken over integer mandate penalty values 0 to 60 under the panel's same L cost advantage, subsidy scheme.

D.3 Additional Welfare Results from Simulations

D.3.1 Graphical Illustration of Welfare Consequences of an Uninsurance Penalty

In this appendix we show how to estimate the welfare consequences of an uninsurance penalty with our graphical model. This exercise corresponds to the similar exercise analyzing the welfare consequences of risk adjustment in the main text. Panel (a) of Figure A6 plots the empirical analogs to our welfare figure from Section 2 for the case where L is a pure cream-skimmer. Instead of plotting C_L, we plot $C_L^{Net} = C_L - C_U$, as in Eq. (18) to account for the fact that $C_U \neq 0$. We indicate the equilibrium s cutoffs for the baseline ACA setting, where subsidies are linked to the price of the lowest-priced plan, $\alpha = 1$, and there is no uninsurance penalty. The intensive margin equilibrium cutoff is s^e_{HL} and the extensive margin cutoff is s^e_{LU}. Thus, consumers with $s < s^e_{HL}$ enroll in H, consumers with $s^e_{HL} < s < s^e_{LU}$ enroll in L, and consumers with $s > s^e_{LU}$ remain uninsured.

It is apparent that, from a social surplus perspective, no consumer should be in L because $W_H - (C_H - C_L)$ is everywhere above W_L. This is because L is a pure cream-skimmer: All consumers value H more than L and L has no cost advantage over H. In addition, in this setting some consumers (those with

Table A3: Benefit Regulation : L-plan Cream Skimmer

	ACA-like all sub		Fixed = Avg. Cost		Fixed = 300		Fixed = 275		Fixed =2 50	
	L offered	No L	L offered	No L	L offered	No L	L offered	No L	L offered	No L
price H	382	362	353	390	429	429	448	448	461	461
price L	352	.	308	.	.	.	.	.	.	.
share H	.42	.78	.29	.65	.43	.43	.31	.31	.22	.22
share L	.31	0	.71	0	0	0	0	0	0	0
share U	.27	.22	0	.35	.57	.57	.69	.69	.78	.78
subsidy	297	307	322	322	300	300	275	275	250	250
welfare	-229	-225	-266	-213	-211	-211	-219	-219	-228	-228

Notes: Table A3 contains equilibrium prices, market shares, subsidy levels and welfare for various subsidy settings with and without the L plan offered. All results are for a setting where L is a cream-skimmer ($\Delta C_{HL} = 0$). The first two columns contain results for ACA-like price-linked subsidies. The following columns are for various fixed subsidies. Welfare is calculated under the baseline assumption, $C_U(s) = 0.64C_H(s) - 97$.

Table A4: Benefit Regulation : L-plan 15% cost advantage

	ACA-like all sub		Fixed = Avg. Cost		Fixed = 300		Fixed = 275		Fixed =2 50	
	L offered	No L	L offered	No L	L offered	No L	L offered	No L	L offered	No L
price H	414	362	.	390	.	429	441	448	462	461
price L	307	.	273	.	273	.	345	.	373	.
share H	.021	.78	0	.65	0	.43	.066	.31	.088	.22
share L	.73	0	1	0	1	0	.47	0	.25	0
share U	.25	.22	0	.35	0	.57	.46	.69	.67	.78
subsidy	252	307	322	322	300	300	275	275	250	250
welfare	-406	-236	-469	-224	-469	-222	-345	-230	-298	-239

Notes: Table A4 contains equilibrium prices, market shares, subsidy levels and welfare for various subsidy settings with and without the L plan offered. All results are for a setting where L has a 15% cost advantage. The first two columns contain results for ACA-like price-linked subsidies. The following columns are for various fixed subsidies. Welfare is calculated under the baseline assumption, $C_U(s) = 0.64C_H(s) - 97$.

Figure A6: Empirical Estimates of Foregone Surplus

(a) Baseline Foregone Surplus (b) Impose \$60 Penalty

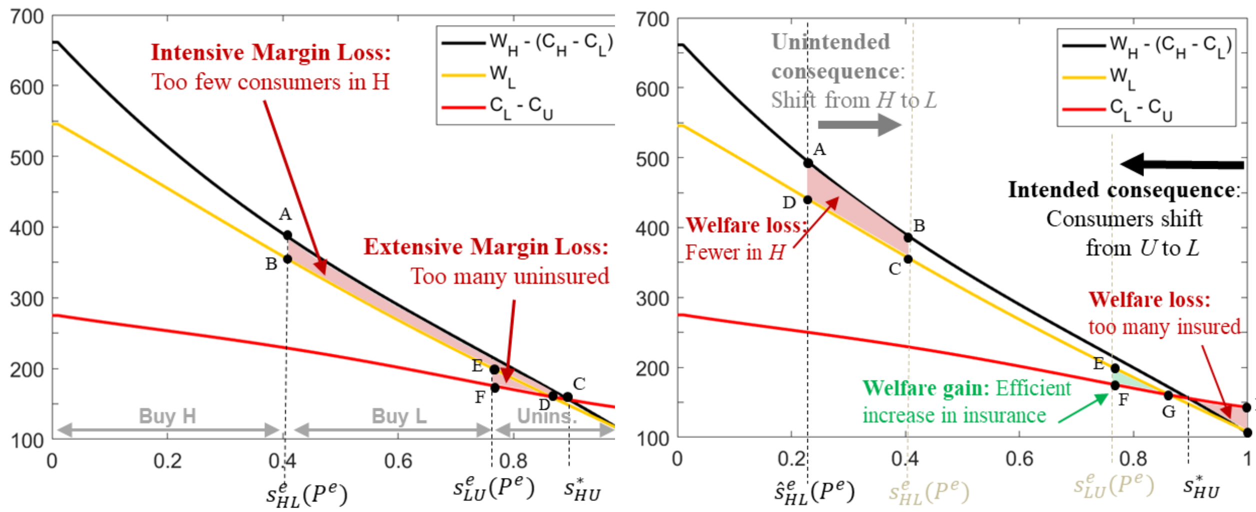

Notes: Panels (a) and (b) show welfare losses under ACA-like subsidies relative to efficient sorting, when L is a cream-skimmer and when L has a 15% cost advantage over H, respectively. In both settings, 60% of the population is low-income and 40% of the population is high-income, so WTP curves are weighted sums of both types. Efficient cutoffs are indicated with a * while equilibrium outcomes are denoted with an e superscript. For both panel (a) and (b), we assume $C_U(s) = 0.64C_H(s) - 97$.

$s > s^*_{HU}$) should not be insured at all. These consumers do not value either H or L more than the (net) cost of enrolling them, making it inefficient for them to be insured. In the figure, we depict the foregone surplus in the baseline ACA setting with shaded areas. The foregone intensive margin surplus in panel (a) (lost surplus due to consumers choosing L instead of H) is described by the area between W_H^{Net} and W_L for the consumers not enrolled in H, $ACDB$. This area represents a welfare loss of \$41.92. The foregone extensive margin surplus (lost surplus due to consumers choosing U instead of L) is given by the area between W_L and C_L^{Net} for the consumers who are not enrolled in insurance but should be, EDF. This area represents a welfare loss of \$16.58. The total foregone surplus in the baseline ACA setting in panel (a) of Figure A6 is \$58.50.

Panel (b) of Figure A6 shows how we estimate the welfare consequences of adding an uninsurance penalty of \$60 per month to the baseline case from Panel (a). Recall from the top-left panel of Figure 9 that the imposition of a \$60 mandate (1) induces all previously uninsured consumers to purchase insurance and (2) causes a shift of 19% of the market from H to L. Effect (1) is the intended consequence of the penalty, and it implies both welfare gains and losses. Welfare gains occur among those consumers who value L more than $C_L^{Net} = C_L - C_U$ and who newly enroll in L (green welfare triangle EFG). Welfare losses occur among those consumers who value L less than C_L^{Net} and who newly enroll in L (red welfare triangle GHI). Together, the intended consequence of the penalty, inducing all consumers to purchase insurance, implies a net welfare gain of \$16.59. Effect (2) is the unintended consequence of the penalty, shifting consumers from H to L. Here, it implies a welfare loss of \$57.83, which arises because H and L have similar costs but all consumers value H more than L. Overall a \$60 uninsurance penalty leads to a welfare loss of \$41.25 in this setting.

We report welfare impacts of a mandate in other market settings in Appendix D.3.2. Those results, which correspond to the cases in Figures 9, show that it is common for an uninsurance penalty to negatively affect welfare. Given the demand and cost primitives we consider, the unintended consequence of shifting consumers from H to L often more than offsets welfare gains from inducing some consumers who value insurance more than its cost to become insured. This is true both when L is a cream-skimmer and when L has a cost advantage. However, it is not clear that this result would generalize to other settings with different consumer willingness-to-pay for H vs. L.

D.3.2 Additional Welfare Estimates Corresponding to Market Share Simulations

Figures A7 and A8 present welfare results corresponding to the market shares in Figures 9 and 10. For a given parameter setting k, we report here welfare normalized as follows: $W_k = \frac{welfare - min(welfare)}{max(welfare) - min(welfare)}$. We characterize welfare under three different assumptions of the cost of uninsured individuals. The first baseline assumption is the same as in the body of the text:

$$C_U(s) = \frac{(1-d)C_H(s)}{1+\phi} + \omega,$$

where the share of total uninsured health care costs that the uninsured pay out of pocket is $d = 0.2$, the assumed moral hazard from insurance is $\phi = 0.25$, and the fixed cost of uninsurance is $\omega = -97$. In addition to this baseline specification, we also show welfare results where we assume uninsured individuals to have the same cost as they would in H ($C_U = C_H$) and where uninsured individuals have no cost $C_U = 0$.

When the cost of the uninsured is high ($C_U = C_H$), a stronger mandate is generally optimal in all settings. When the uninsured are less costly, however, lower mandates and higher risk adjustment are generally optimal.

D.3.3 Optimality under Interacting Policies, Further Results

In Figure A9, we present welfare results under interacting extensive margin (mandate) and intensive margin (risk adjustment α parameter) policies for all settings studied in Figures 9 and 10 in the main text. These results are similar to the results we report in Section 6 but correspond to different market and policy settings. We see that the optimal mandate and risk adjustment combination depends on both the subsidy as well as the cost structure. When the L plan is a cream-skimmer, moderate to strong risk adjustment is preferable in order to induce more consumers to enroll in H vs. L. When L has a cost advantage, however, weaker risk adjustment is preferable. Further, when L is a cream-skimmer, the optimal mandate for a given level of risk adjustment also varies, with ACA-like subsidies warranting a lower mandate compared to the fixed subsidy case.

D.4 Empirical Robustness: Varying Simulation Model Assumptions

D.4.1 Empirical Robustness: Relaxing the Vertical Model

The demand primitives from Finkelstein, Hendren and Shepard (2019) were estimated in a setting where insurance options could be clearly ranked from most to least desirable for all consumers and where WTP was assumed to vary along a single dimension of heterogeneity. As a result, these primitives are consistent

Figure A7: Welfare with Varying Mandate Penalty (M)

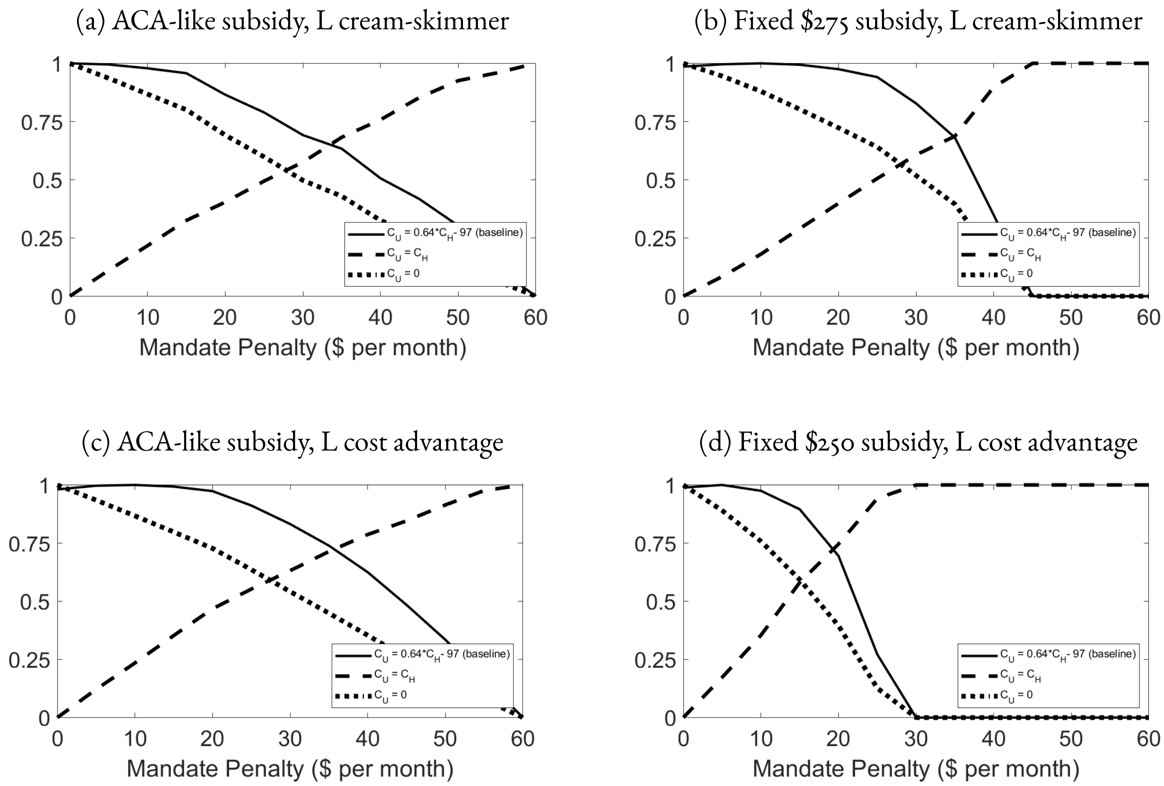

Notes: Figure A7 depicts equilibrium relative welfare under varying levels of the mandate penalty. The simulations are the same as in figure 9. Panels (a) and (b) are results for when L is a cream-skimmer ($\Delta C_{HL} = 0$) while panels (c) and (d) are for when L has a 15% cost advantage. In panels (a) and (c), the market has a price-linked subsidy while in panels (b) and (d), the market has a fixed subsidy. For each set of simulations, we present relative welfare under three different assumptions about the social cost of uninsurance. Relative welfare is calculated as $\frac{welfare - min(welfare)}{max(welfare) - min(welfare)}$ where max and min are taken over the possible mandate penalties within a set of simulations and C_U assumptions.

with a vertical demand structure. In effect, this means that throughout our main simulations, individuals are only on the margin between H and L or L and U, never on the margin between H and U (except in cases where the market "upravels" and nobody chooses L). As the theoretical analysis in Appendix A shows, allowing for an HU substitution margin that would be present with horizontal differentiation adds additional terms to the comparative statics defining cross-margin policy effects.

We can investigate how robust our empirical results are to the vertical model by assuming some portion of the population does not value L at all and is thus solely on the margin between H and U. To do this, we perform the following exercise:

Simulation modifications

- From our standard population comprising 60% subsidized low income types and 40% unsubsidized high income types, we assume γ percent of each type do not value L so that they may only choose between H and U

Figure A8: Welfare with Varying Strength of Risk Adjustment (α)

(a) ACA-like subsidy, L cream-skimmer

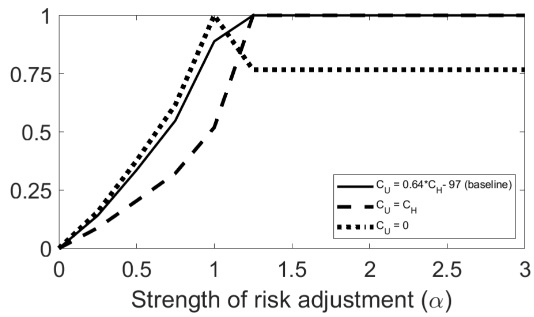

(b) Fixed $ 275 subsidy, L cream-skimmer

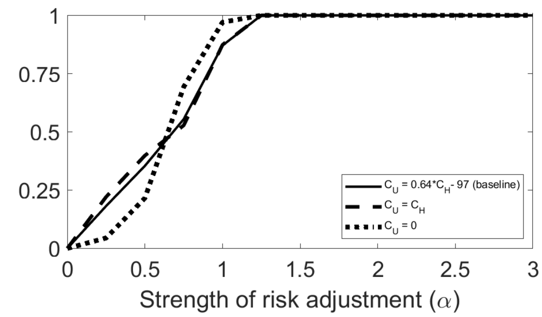

(c) ACA-like subsidy, L cost advantage

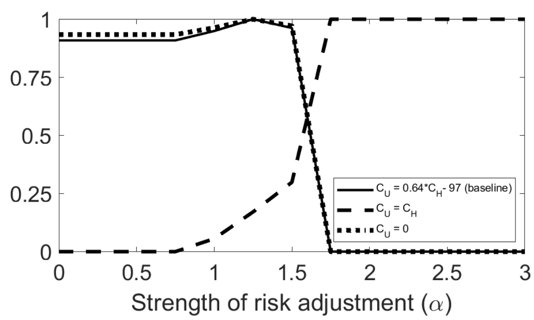

(d) Fixed $ 250 subsidy, L cost advantage

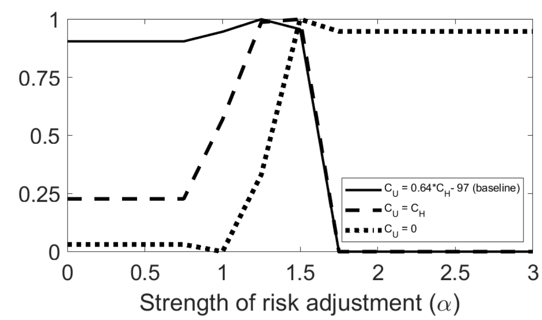

Notes: Figure A7 depicts equilibrium relative welfare under varying strengths of risk adjustment α. The simulations are the same as in figure 10. Panels (a) and (b) are results for when L is a cream-skimmer ($\Delta C_{HL} = 0$) while panels (c) and (d) are for when L has a 15% cost advantage. In panels (a) and (c), the market has a price-linked subsidy while in panels (b) and (d), the market has a fixed subsidy. For each set of simulations, we present relative welfare under three different assumptions about the social cost of uninsurance. Relative welfare is calculated as $\frac{welfare - min(welfare)}{max(welfare) - min(welfare)}$ where max and min are taken over the possible α values within a set of simulations and C_U assumptions.

- We assume that this γ portion has the standard $W_H(s)$ and $W_H^{HI}(s)$ curves and same s distribution as in our baseline simulations
- The remaining $1-\gamma$ portion of the population has the standard demand primitives and may choose between H, L, and U as normal
- For a given price bid, P_H and P_L, and subsidy, we allow both types to choose plans, estimating profits and equilibrium in the typical way

Impact of HU margin types on mandate results

In panel (a) of Figure A10 we estimate demand shares with ACA-like subsidies where the L plan is a pure cream-skimmer and with increasingly larger values of γ (i.e., increasing proportions of HU margin types) from 0% up to 20%. For every mandate penalty level, the market allocation to H is everywhere higher with larger shares of HU margin types. As the uninsurance penalty increases, consumers move from U to L *and* from U to H. There is still an unintended shifting of consumers from H to L as

Figure A9: Welfare under Interacting Extensive and Intensive Margin Policies

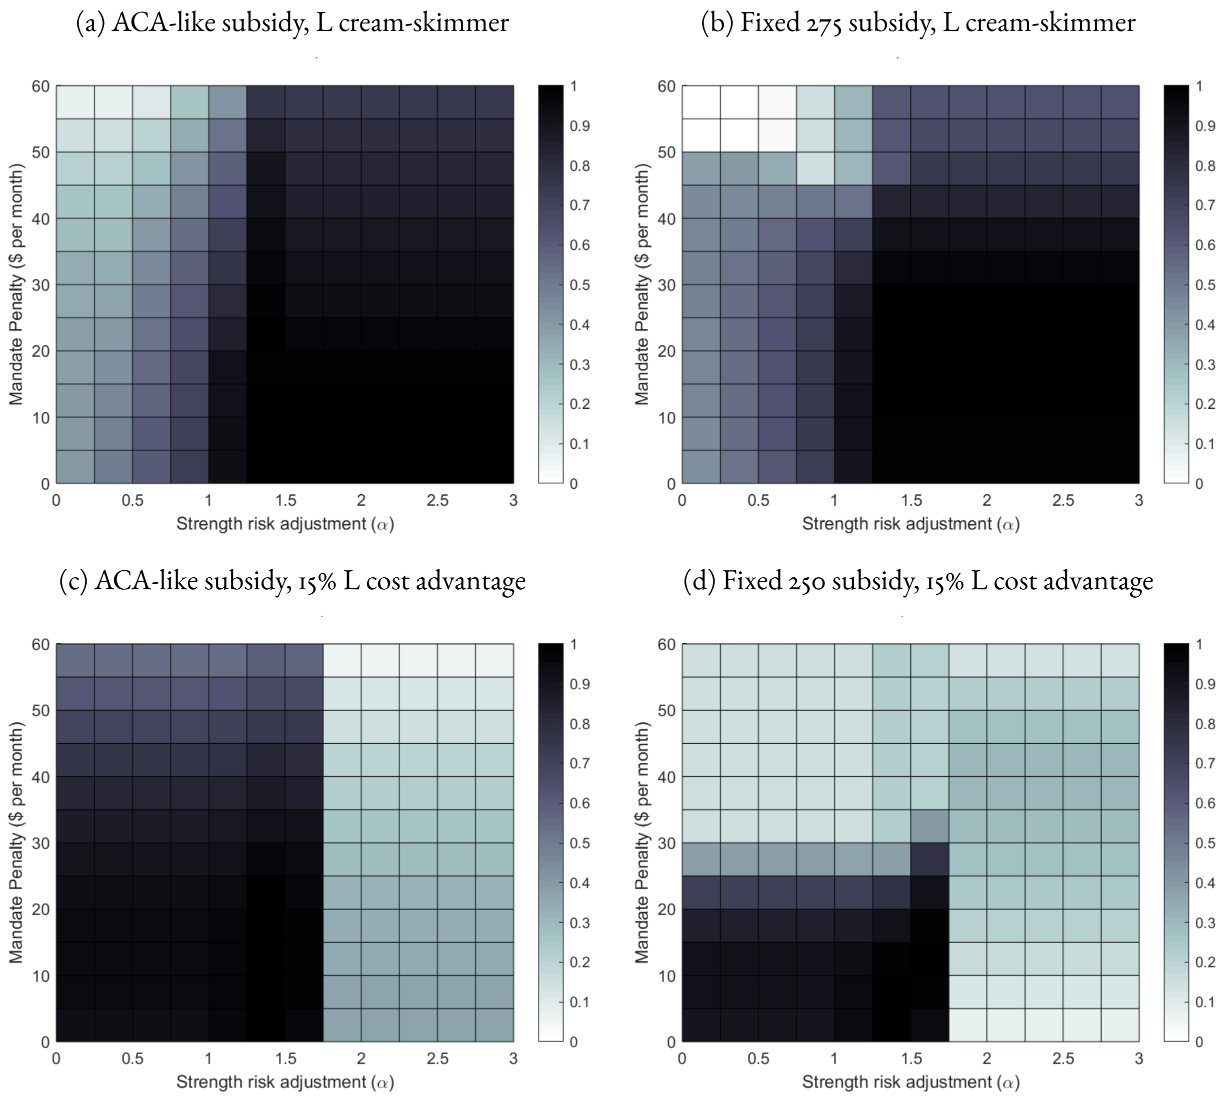

Notes: Figure A9 depicts equilibrium relative welfare under varying levels of the mandate penalty and strength of risk adjustment α. Panels (a) and (b) are results for when L is a cream-skimmer ($\Delta C_{HL} = 0$) while panels (c) and (d) are for when L has a 15% cost advantage. In panels (a) and (c), the market has a price-linked subsidy while in panels (b) and (d), the market has a fixed subsidy. Relative welfare is calculated as $\frac{welfare - min(welfare)}{max(welfare) - min(welfare)}$ where max and min are taken over all the possible mandate penalties and risk adjustment strengths within a subsidy and cost setting. For all simulations, we use our baseline assumption of the social cost of uninsurance, $C_U = 0.64C_H - 97$.

highlighted in Section 5 of the paper, but there are countervailing forces, composed of (1) the shifting of consumers from U to H, and (2) the fact that the presence of some lower-cost HU margin types in H lowers the price of H and the price differential between H and L.

On net, D_H still declines with a stronger mandate with a γ of 10% or 20%. This shows that the empirical "unintended" effect of the mandate on D_H is robust to some horizontal differentiation. However, the net decline is increasingly muted as γ increases, and a level of γ much larger than 20% would eventually result in D_H being flat or increasing with the mandate penalty.

Impact of HU margin types on risk adjustment results

Next, in panel (b) of Figure A10 we estimate demand shares as we vary risk adjustment strength for the case of fixed subsidies when L has a 15% cost advantage. Recall that this is the risk adjustment simulation where we saw a trade-off between extensive and intensive margin selection: Stronger risk adjustment induced consumers to move from L to H but it also induced some consumers to exit the market and opt for U.

Similar to our mandate simulations allowing for some consumers to be on the HU margin, we see that the initial allocations to H absent risk adjustment are higher when we have more HU margin types compared to our baseline setting. Because lower cost HU margin types will enroll in H compared to our baseline types, the cost differential between the two plans is lower with larger shares of HU margin types. Consequently, the size of risk adjustment transfers for a given α are lower. However, the level of α that causes the market to "upravel" to H is the same for all levels of γ. Further, the uninsurance rate also depends very little on γ, with the U market share at any given level of α being similar across levels of γ. This indicates that our result that under certain conditions risk adjustment can unintentionally increase the uninsurance rate while simultaneously shifting consumers from L to H is largely robust to our vertical model assumption for the market primitives we examine.

D.4.2 Empirical Robustness: Varying ΔW_{HL}

Demand for H critically depends on the incremental willingness to pay for H relative to L, $\Delta W_{HL} = W_H(s) - W_L(s)$. Below, we see how sensitive our results are to variations in this incremental willingness to pay. Specifically, we estimate equilibrium under simulations where we hold fixed $W_L(s)$ at baseline but scale $\Delta W_{HL}(s)$ by a multiplier $\rho \in [0.25, 4]$:

$$\Delta W_{HL}^{adj}(s) = \Delta W_{HL}(s)^{raw} * \rho$$

$$W_H^{adj}(s) = W_L(s) + \Delta W_{HL}^{adj}(s)$$

This scaling changes both the level and slope of $W_H(s)$, as seen in Figure A11.

Using our typical counterfactual process, we estimate equilibrium market shares under these modified primitives for varying levels of the mandate penalty and risk adjustment strength. Simulation results are presented in Figure A12. We find that under both increased and decreased incremental willingness to pay (i.e. higher and lower ρ), the general patterns of our counterfactual exercises do not change.

Panel (a) shows that demand for H declines with a larger mandate penalty, except at the very high scalar $\rho = 4$. When $\rho = 4$, the marginal willingness to pay for H relative to L is sufficiently high that an incrementally higher mandate penalty induces individuals to enter the market and then choose H over L. As a result, demand for H is weakly increasing in the mandate penalty throughout the range of penalties

Figure A10: Relaxing vertical model

(a) Mandate penalty

ACA-like subsidy, Cream-skimming L plan

(b) Risk adjustment

Fixed subsidy = $250, 15% cost advantage L plan

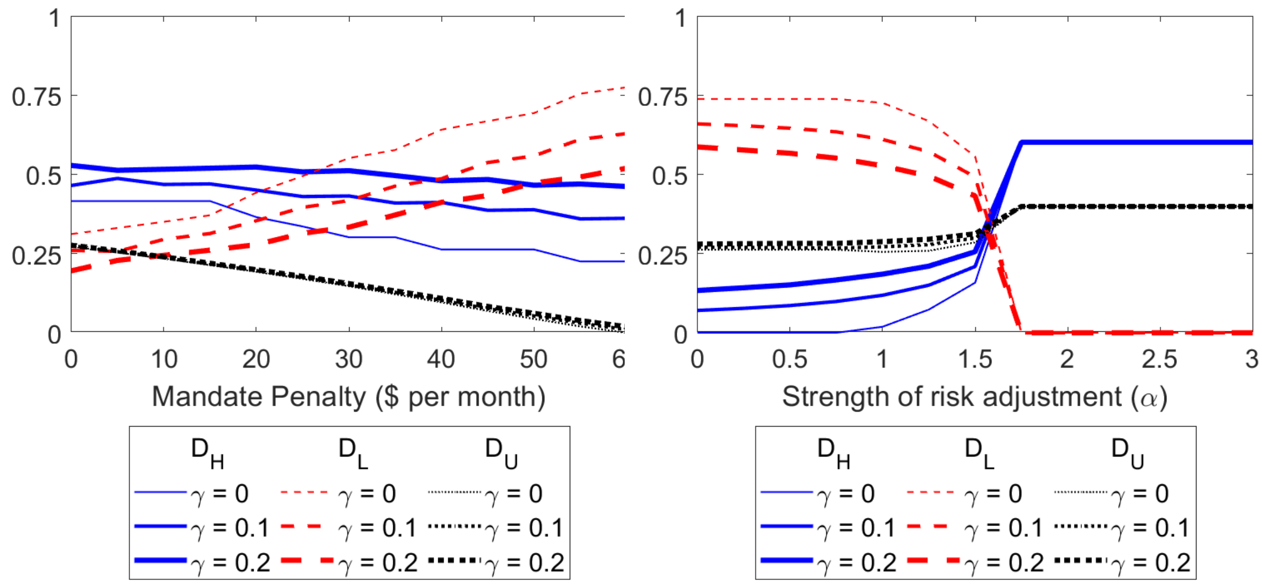

Notes: Panels (a) and (b) of Figure A10 depicts equilibrium market shares of H,L, and uninsurance under varying levels of the mandate penalty and risk adjustment strength (α), respectively. Three separate simulations are presented. The thinnest line is our baseline simulation where no individuals are on the margin between H and uninsurance ($\gamma = 0$) while the thickest lines correspond to when 20% of individuals do not consider L and are thus on the margin between H and U ($\gamma = 0.2$). All simulations in panel (a) are for a cream-skimming L plan and ACA-like price linked subsidy and all simulations in panel (b) are for an L plan with a 15% cost advantage and fixed subsidy of $250 for both plans.

tested while demand for L only rises for high levels of the mandate. The rise in L only occurs in the range of mandate penalties where the individuals induced to enter the market are of sufficiently low marginal willingness to pay that some choose L instead of H. Because this is a relatively small group, the cost differential between H and L remains small.

Panel (b) shows that increasing the strength of risk adjustment has similar effects at all levels of ρ. Initially, stronger risk adjustment induces consumers to choose H instead of L. But in all cases, there is also eventually an unintended increase in the uninsurance rate. The effect of modifying ρ is that the shifts in market share (both from L to H and from L to U) occur at different levels of α with shifts occurring at lower levels of α for higher levels of ρ. That is, when marginal willingness to pay for H relative to L is higher, a lower level of risk adjustment is needed to induce changes in market shares.

Figure A11: Scaled WTP_H

(a) Low income demand

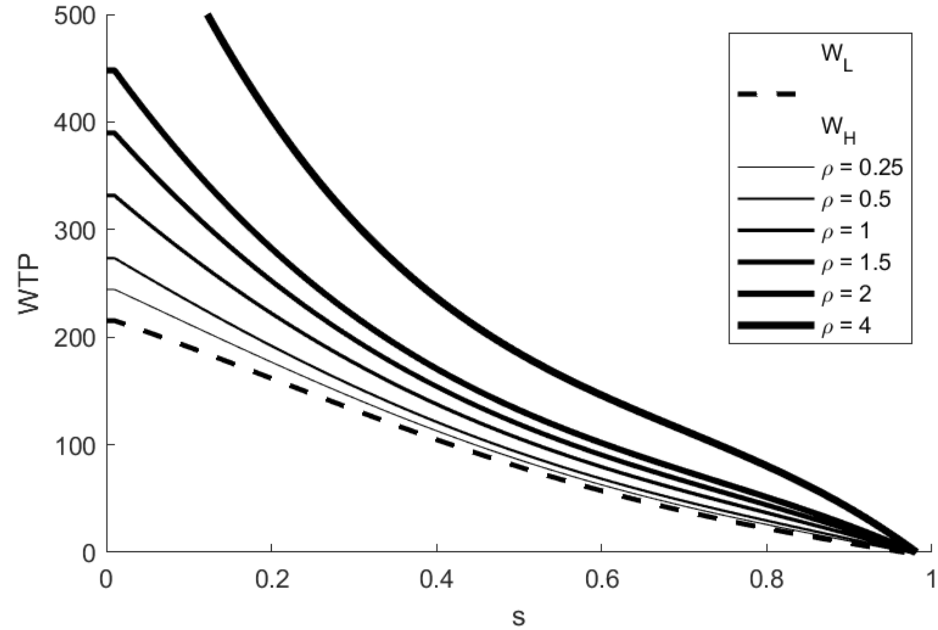

(b) High income demand

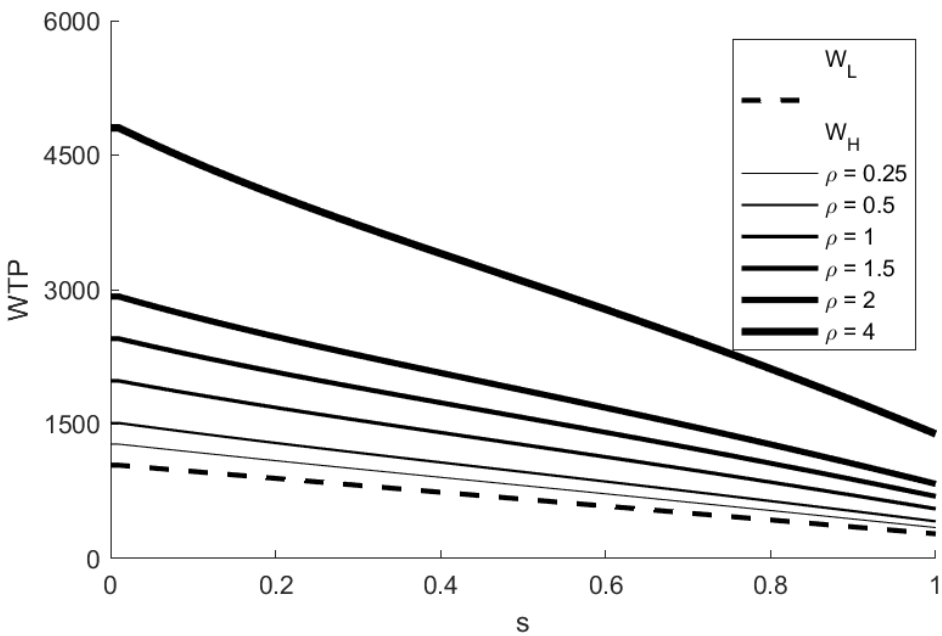

Notes: Panels (a) and (b) of A11 depicts willingness to pay curves for high and low-income consumers, respectively, under various scaling factors ρ of $\Delta W_{HL}^{adj} = \rho \Delta W_{HL}$. The thickest lines are for high marginal WTP for H relative to L. Baseline is for $\rho = 1$. Willingness to pay for L is the dashed line and remains unmodified.

Figure A12: Scaling ΔWTP

(a) Mandate penalty

ACA-like subsidy, Cream-skimming L plan

(b) Risk adjustment

Fixed subsidy = $250, 15 % cost advantage L plan

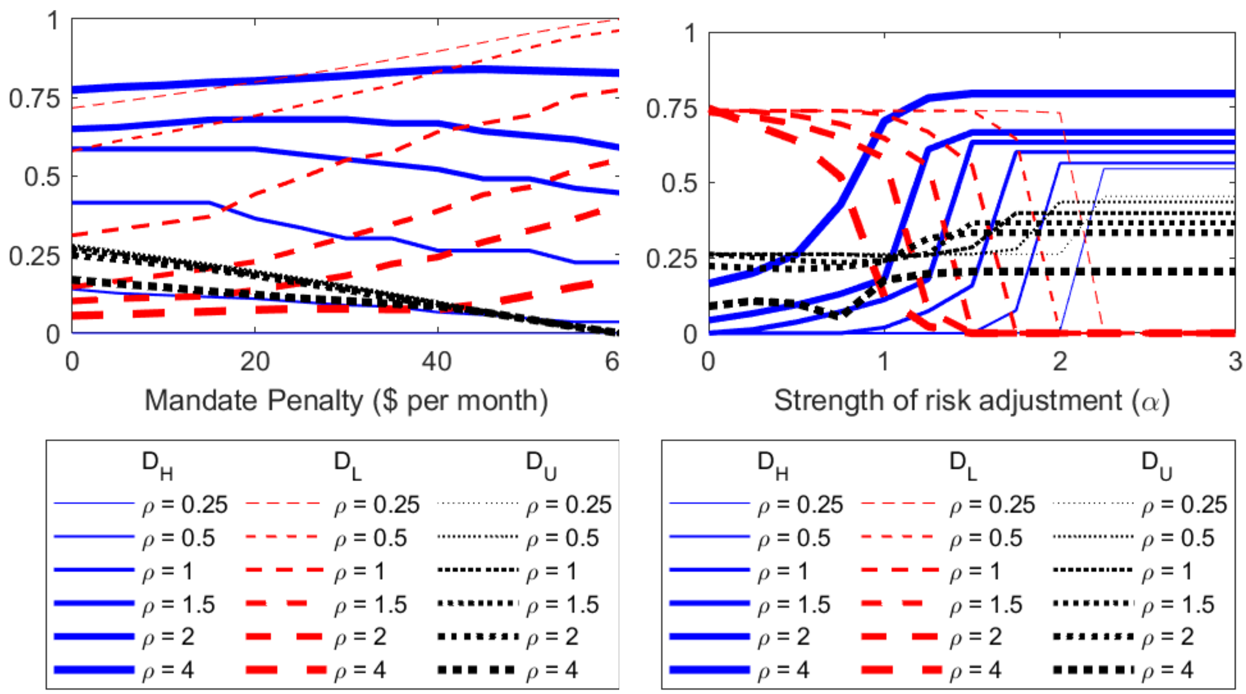

Notes: Figure A12 shows market shares for H,L, and uninsurance under the different scaled ΔWTP curves depicted in figure A11. Panel (a) depicts shares for different mandate penalties under an ACA-like price-linked subsidy and cream-skimming L plan ($\Delta C_{HL} = 0$). Panel (b) depicts shares for different strengths of risk adjustment (α) under a fixed subsidy and a 15% L plan cost advantage. As in figure A11, thicker lines correspond to market shares when marginal willingness to pay for H relative to L is set higher (higher ρ).