The Two Margin Problem in Insurance Markets[*]

Michael Geruso[†]   Timothy J. Layton[‡]   Grace McCormack[§]   Mark Shepard[¶]

May 3, 2019

Abstract

Insurance markets often feature consumer sorting along both an extensive margin (whether to buy) and an intensive margin (which plan to buy), but most research considers just one margin or the other in isolation. We present a graphical theoretical framework that incorporates both selection margins and allows us to illustrate the often surprising equilibrium and welfare implications that arise. A key finding is that standard policies often involve a trade-off between ameliorating intensive vs. extensive margin adverse selection. While a larger penalty for opting to remain uninsured reduces the uninsurance rate, it also tends to lead to unraveling of generous coverage because the newly insured are healthier and sort into less generous plans, driving down the relative prices of those plans. While risk adjustment transfers shift enrollment from lower- to higher-generosity plans, they also sometimes increase the uninsurance rate by raising the prices of less generous plans, which are the entry points into the market. We illustrate these trade-offs in an empirical sufficient statistics approach that is tightly linked to the graphical framework. Using data from Massachusetts, we show that in many policy environments these trade-offs can be empirically meaningful and can cause these policies to have unexpected consequences for social welfare.

[†]University of Texas at Austin and NBER. Email: mike.geruso@gmail.com

[‡]Harvard University and NBER. Email: layton@hcp.med.harvard.edu

[§]Harvard University. Email: gamccormack@g.harvard.edu

[¶]Harvard University and NBER. Email: Mark_Shepard@hks.harvard.edu

1 Introduction

Adverse selection is an important and persistent problem in insurance markets that motivates significant policy intervention. By distorting prices, selection can cause some consumers to inefficiently remain uninsured and cause others to select inefficient levels of coverage. Much of the prior literature aimed at understanding equilibria and optimal selection-related policy in insurance markets has largely considered these two forms of selection in isolation, either assuming that all consumers choose a contract and focusing on the intensive margin of choice between plans (e.g., Handel, Hendel and Whinston, 2015), or assuming that all contracts in the market are effectively identical and focusing on the extensive margin of gaining insurance (e.g., Hackmann, Kolstad and Kowalski, 2015).

By ignoring one margin or the other, the selection problem is usefully simplified. In empirical work, it becomes amenable to a sufficient statistics approach based on demand and cost curves defined in reference to a single price—either the price of gaining insurance or the price difference between a more and a less generous plan (Einav, Finkelstein and Cullen, 2010). But this simplification does not allow for potential *interactions* between these two margins of selection. For example, an insurance mandate—which aims to correct extensive margin selection—can negatively impact the generosity of coverage among consumers for whom the mandate does not bind (Azevedo and Gottlieb, 2017). Intuitively, the marginal consumers induced to enroll by the mandate are likely among the healthiest in the market. If these healthy consumers select the lowest-price (and lowest-quality) plans, these plans' risk profiles will get even healthier, allowing the plans to lower prices and siphon more consumers from higher-quality plans on the intensive margin.

Interactions between intensive and extensive margins of risk selection are relevant in a variety of insurance settings. For instance, in the Affordable Care Act exchanges, consumers choose both whether to buy insurance (the extensive margin) and which plan to buy among a variety of generosity levels (the intensive margin). In Medicare Advantage, enrollees choose both among competing private plans (intensive margin) and whether to stick with traditional Medicare (extensive margin). Similar dynamics may also be relevant in other settings with both a participation and plan choice decision, including employer programs featuring plan choice (e.g., CalPERS) and national health insurance systems with an opt-out (e.g., Germany).

Our first goal in this paper is to show how policy interventions that counteract selection on one margin can interact with the other. We show that a mandate's impact on plan generosity is, in fact, an

instance of a broader phenomenon that encapsulates many relevant policy interventions currently in place in insurance markets. These include plan benefits requirements, network adequacy rules, risk adjustment, reinsurance, subsidies, and behavioral interventions like plan choice architectures.

To see this, consider risk adjustment, a policy targeted at intensive margin selection. Risk adjustment typically enforces transfers from less generous plans with lower-cost enrollees to more generous plans with higher-cost enrollees. The transfers bring down the price of the more generous plans in equilibrium, leading more consumers to buy them. This shift towards generous coverage is the intended effect of the policy. However, because risk adjustment tends to lower the prices of more generous plans at the expense of raising the prices of less generous plans, some low-cost consumers who would have been nearly indifferent between less generous coverage and uninsurance decide to exit the market altogether as a result of risk adjustment. This is the unintended cross-margin effect: countering adverse selection across plans *within* a market has exacerbated adverse selection *into* the market.

Our second goal in this paper is to present a graphical framework that allows readers to visualize the cross-margin policy interactions in simple and transparent demand and cost curves. Recent complementary work has pointed to the theoretical possibility of cross-margin interactions (Azevedo and Gottlieb, 2017) or allowed for both margins of selection within a structural model of an insurance market (Domurat, 2018; Saltzman, 2017). We provide intuition for the two margin problem in a series of figures that parallel the simple graphical framework of Einav and Finkelstein (2011). As in Einav, Finkelstein and Cullen (2010), there is a tight link between our model and the estimation of sufficient statistics used to characterize the market and generate counterfactuals. We illustrate how equilibrium prices, allocations, and (critically) welfare can be recovered from demand and cost curves. Econometric identification is analogous to that in Einav, Finkelstein and Cullen (2010), though here, exogenous price variation along two margins is required—for example, independent variation in the price of a skimpy plan and in the price of a generous plan.[1] After we develop the core ideas in the context of perfect competition and a vertical model, we discuss how the key insights are affected by various extensions, including horizontal plan differentiation and irrational or ill-informed consumers.

With the intuition and price theory in place, we analyze the model's insights empirically using demand and cost estimates from Massachusetts' CommCare program, a precursor to the state's ACA

---

[1]Or alternatively, variation in a market-wide subsidy for selecting any plan and independent variation in the price difference between bare bones and generous plans.

health insurance Marketplace. CommCare was introduced in 2006 to provide subsidized health insurance coverage to low-income Massachusetts residents who did not qualify for Medicaid. In this setting, Finkelstein, Hendren and Shepard (2017) document significant adverse selection both into the market and within the market between a narrow-network, lower-quality option and a set of wider-network, higher-quality plans. In a regression discontinuity design that exploits discontinuities in the income-based premium subsidy scheme, they construct demand and cost curves for the lower and higher quality plans. We use these demand and cost curves in a number of illustrative counter-factual exercises that examine enrollment and prices as we vary benefit design rules, mandates and penalties, and risk adjustment strength. With the additional neoclassical assumption that revealed consumer preferences reflect underlying valuations, and with an estimate of the social cost of uninsurance, optimal policy can be evaluated graphically on the basis of social surplus. The overarching insight—both theoretically, and empirically—is that there is an interaction between extensive and intensive margin selection, and that this interaction often implies a policy trade-off between selection on one margin and selection on the other.

The empirical exercise, beyond demonstrating how our framework can be used, generates several substantive policy insights. The size of the unintended cross margin effects can be large enough to imply significant impacts on the allocation of consumers across contracts. We find that a strong mandate sufficient to move all consumers into insurance—increasing enrollment by around 25 percentage points in our setting—can cause the market share of more generous plans to shrink by more than 15 percentage points, or 35% of baseline market share. In the other direction, strengthening risk adjustment transfers to the point where the market "upravels" to include only generous coverage (or, equivalently, enforcing minimum coverage generosity to eliminate low-quality options) can substantially reduce market-level consumer participation—in our setting by as much as 15 percentage points or 60% of the baseline uninsurance rate. We find that the cross-margin welfare impacts can be similarly large (and often first-order), under a range of assumptions about the external social cost of the uninsured.

Further, we show that in some settings, cross-margin interactions are critical for determining optimal policy: When extensive margin policies (such as a mandate) are weak, it is optimal to also have weak intensive margin policies (such as risk adjustment). But when extensive margin policies are strong, on the other hand, it is optimal to also have strong intensive margin policies. These results

show that in these markets, regulators are indeed operating in a world of the second-best and must consider interactions between the two margins of selection in order to determine optimal policy. The previous literature that has focused on either the intensive *or* the extensive margin in isolation, can offer no guidance on this trade-off.

This trade-off is not just an academic curiosity. Today, states are being given increasing regulatory flexibility over selection-related policies in a way that outpaces insights from the prior research. For example, federal regulations now allow states to weaken Essential Health Benefits in their individual insurance Marketplaces in order to reduce uninsurance rates.[2] In so doing they are relaxing an intensive-margin selection policy in order to affect extensive margin selection in a way that our model captures. Similarly, states are being given flexibility to weaken risk adjustment transfers in order to address concerns that gross prices of the lowest-priced plans in the market are "too high" for many unsubsidized consumers. The goal is that reducing the risk adjustment payments these plans are making to higher-price plans will lead to lower prices for the lowest-price plans and induce entry of currently uninsured, healthy consumers. More broadly, much attention is currently being paid to policy proposals whose intention is to increase the number of covered lives in the individual health insurance market (Domurat, Menashe and Yin, 2018). Our work makes it clear that such policies might involve a significant tradeoff in terms of intensive margin selection, expanding and generalizing an insight from Azevedo and Gottlieb (2017).

Related Literature   Our work contributes to the literature on adverse selection in insurance markets. The closest paper to ours is Azevedo and Gottlieb (2017) who develop a theoretical model of insurance market equilibrium in a perfectly competitive market.[3] Our work also relates to several other parts of the insurance market literature. Ericson and Starc (2015), Handel, Hendel and Whinston (2015), Tebaldi (2017), Saltzman (2017), Domurat (2018), and other empirical, policy-oriented studies have analyzed equilibria and optimal policy in an insurance exchange. A related, but distinct

---

[2]Beginning in 2019, all states will have the option to relax Essential Health Benefits rules and to scale down the size of risk adjustment transfers. See Code of Federal Regulations Vol. 83, No. 74.

[3]The key insight of Azevedo and Gottlieb (2017) is that the literatures studying the effects of selection in a fixed contracts setting and an endogenous contracts setting can be united by allowing for a large contract space and allowing consumer preferences and pricing dynamics to determine which contracts survive in equilibrium. Empirical simulations of their model reveal an interaction between the mandate and intensive margin sorting of consumers. We show here that this mandate insight can be shown in the simpler setting with 2 vertically differentiated plans plus the outside option of uninsurance, a setting which is amenable to graphical representation. In our simpler setting, a graphical sufficient statistics approach can be used to characterize the impacts of a variety of selection-related regulatory interventions using reduced-form estimates of demand and cost curves.

literature on adverse selection on contract design has shown how socially efficient contracts fail to arise in equilibrium in selection markets (Rothschild and Stiglitz, 1976; Glazer and McGuire, 2000; Veiga and Weyl, 2016). There is empirical evidence that this occurs in practice in the context of health insurance (Carey, 2017a,b; Lavetti and Simon, 2016; Shepard, 2016; Geruso, Layton and Prinz, 2018). However, this literature has not previously incorporated the outside option of uninsurance. The lesson here is that the outside option can be important when considering the policies typically used to combat this type of selection problem: risk adjustment, reinsurance, and benefit design regulations. Indeed, our paper can be thought of as the simplest possible way to unite the literatures focusing on consequences of adverse selection in fixed vs. endogenous contracts settings.

With respect to risk adjustment in particular, the prior literature has considered only impacts along the margin of selection targeted by the risk adjustment policy (see, e.g., Layton, 2017 and Mahoney and Weyl, 2017). Two exceptions are Newhouse (2017) and Einav, Finkelstein and Tebaldi (2018). Newhouse (2017) notes that risk adjustment within a market can impact the composition of the population opting into the market, but does not pursue the cross-margin effects described here. Einav, Finkelstein and Tebaldi (2018) compare risk adjustment and subsidies as tools for dealing with extensive margin selection problems, not considering the intensive margin effects of the policies. The phenomenon we highlight provides a new lens through which to understand the potential impacts of these ubiquitous regulatory tools. It also opens new avenues of research. In particular, Medicare markets can be described by extensive selection margin between the private Medicare Advantage segment and the Traditional Medicare public option, as well as an intensive selection margin across the various private Medicare Advantage plans. Prior work has explicitly or implicitly treated risk adjustment as affecting only selection into the private plan segment (Brown et al., 2014; Newhouse et al., 2015; Geruso and Layton, 2018). Some insights here may likewise carry to state Medicaid programs in which a private managed care plans compete with each other and with the state in a highly regulated managed competition setting.

We note that our work also generates new implications for considering behavioral interventions and decision aids to address choice frictions such as inertia (Handel, 2013; Polyakova, 2016), misinformation (Kling et al., 2012; Handel and Kolstad, 2015), or complexity (Ericson and Starc, 2016; Ketcham, Kuminoff and Powers, 2016) in insurance markets. Any intervention that successfully impacts plan choice within the market can in principle affect rates of uninsurance as plan prices (and

therefore the entry point into the market) endogenously respond to the altered risk sorting. Similarly, interventions targeted at improving take-up of insurance (Domurat, Menashe and Yin, 2018) can in principle affect selection across plans within the market, similar to the effects of an insurance mandate or uninsurance penalty. Importantly, because price and enrollment predictions in our model are derived from demand and cost curves without making assumptions about deeper preference or belief parameters underlying these, our model delivers policy-relevant predictions regarding rates of uninsurance and plan prices even in the presence of behavioral choice frictions even though welfare estimation is more challenging.

Finally, our model and insights contribute to the broader literature on selection markets, including other types of insurance and credit markets. For example, our model offers insights for considering the effects of auto insurance mandates on selection across plans with varying deductibles within the market. With respect to markets for long-term care insurance, prior work has shown evidence that public coverage via Medicaid reduces take-up of insurance (Brown, Finkelstein and Coe, 2007). Our model suggests that if there is selection in these "crowd-out" effects, the presence of Medicaid may not only affect take-up but also the allocation of consumers across products within the market for long-term care insurance. In consumer credit markets, there is evidence suggesting both selection into a market and selection across products differentiated by down payments within the market Adams, Einav and Levin (2009); Einav, Jenkins and Levin (2012). In such a setting, subsidies encouraging the purchase of automobiles may not only result in more automobile purchases but they may also shift the allocation of consumers across loan products within the market.

2 Model

Our goal in this section is to develop a theoretical and graphical model that depict insurance market equilibrium and welfare in the spirit of Einav, Finkelstein and Cullen (2010) ("EFC"), while allowing for the possibility that interventions affecting selection on one margin may affect selection on another. This requires an insurance plan choice set with at least three options. Consider two fixed contracts, $j = \{H, L\}$, where H is more generous than L on some metric, and an outside option, U. In the focal application of our model to the ACA's individual markets, U represents uninsurance.

Each plan $j \in \{H, L\}$ sets a single community-rated price P_j that (along with any risk adjustment transfers—see below) must cover its costs. Consumers make choices based on these prices and on

the price of the outside option, $P_U = M$.[4] In our focal example, M is a mandate penalty. The distinguishing feature of U is that its price is exogenously determined; it does not adjust based on the consumers who select into it. This is natural for the case where U is uninsurance or a public plan like Traditional Medicare. $P = \{P_H, P_L, P_U\}$ is the vector of prices in the market.

In the most general formulation, demand in this market cannot be easily depicted in two-dimensional figures. To make the cross-margin effects of interest clearer, we impose a *vertical model* of demand, which assumes contracts are identically preference-ranked across consumers. Although the strict vertical assumption is not necessary for many of our key insights to hold, it captures the key features of the issues raised by simultaneous selection on two margins in a simple way that allows for graphical representation. In the next subsections, we present the vertical model, then add the cost curves, and finally show how to find equilibrium and welfare. In the appendix, we discuss the implications of relaxing the vertical demand assumption.

2.1 Demand

The model's demand primitives are consumers' willingness-to-pay (WTP) for each plan. Let $W_{i,H}$ be WTP of consumer i for plan H, and $W_{i,L}$ be WTP for L, both defined as WTP relative to U ($W_{i,U} \equiv 0$). We make the following two assumptions on demand:

Assumption 1. *Vertical ranking:* $W_{i,H} > W_{i,L}$ *for all i*

Assumption 2. *Single dimension of WTP heterogeneity: There is a single index $s \sim U[0,1]$ that orders consumers based on declining WTP, such that $W_L'(s) < 0$ and $W_H'(s) - W_L'(s) < 0$ for all s.*

These assumptions, which are a slight generalization of the textbook vertical model,[5] involve two substantive restrictions on the nature of demand. First, the products are vertically ranked: all consumers would choose H over L if their prices were equal. This is a statement about the *type of setting* to which our model applies. The vertical model works best when plan rankings are clear — e.g., a low- vs. high-deductible plan, or a narrow vs. complete provider network. Importantly, these are precisely the settings where intensive margin risk selection is most relevant. When plans are hor-

---

[4]Below, we allow that consumers may receive a subsidy, S, so that choices are based on post-subsidy prices, $P_j^{cons} = P_j - S$.

[5]Our vertical model follows the format of Finkelstein, Hendren and Shepard (2017). It is a generalization of the textbook vertical model in which products differ on quality (Q_j) and consumers differ on taste for quality (β_i), so that WTP equals: $W_{i,j} = \beta_i Q_j$ and utility equals $U_{i,j} = W_{i,j} - P_j = \beta_i Q_j - P_j$.

izontally differentiated (such as in the Covered California market (Tebaldi, 2017; Einav, Finkelstein and Tebaldi, 2018; Saltzman, 2017)), it is less likely that high-risk consumers will heavily select into a single plan or type of plan. In this case, the existing EFC framework captures the main way risk selection matters: in vs. out of the market (the extensive margin). Our model is designed to study the additional issues that arise when *both* intensive and extensive margins matter simultaneously.[6]

Second, consumers' WTP for H and L—which in general could vary arbitrarily over two dimensions—are assumed to collapse to a single-dimensional index, $s \in [0, 1]$. Higher s types have both lower W_L *and* a smaller gap between W_H and W_L. Lower-s types both care more about having insurance (L vs. U) and more about the generosity of coverage (H vs. L). This assumption is natural in many cases; indeed it holds exactly in a model where plans differ purely in their coinsurance rate (see, e.g., Azevedo and Gottlieb, 2017).

Substantively, Assumption 2 restricts consumer *sorting and substitution patterns* among options when prices change. Under prices at which all options are chosen, consumers sort into plans with the highest-WTP types choosing H, intermediate types choosing L, and low types choosing U. Consumers are only on the margin between adjacent-generosity options—between H and L or between L and U. If the price of U (the mandate penalty) increases modestly, the newly insured all buy L (the cheaper plan), not H. This restriction captures in a strong way the general (and testable) idea that these are the *main* ways consumers substitute in response to price changes. Weakening this assumption—allowing an H-U margin—does not change the key implications of the model (see Appendix B).

Figure 1 plots a simple linear example of $W_H(s)$ and $W_L(s)$ curves that satisfy these assumptions. The x-axis is the WTP index s, so WTP declines from left to right as usual. Let $s_{LU}(P)$ be the extensive-marginal type who is indifferent between L and U at a given set of prices P. Assuming for now that $P_U \equiv M = 0$, this cutoff type is defined by the intersection of L's WTP curve W_L and L's price:

$$W_L\left(s_{LU}\right) = P_L. \tag{1}$$

Consumers to the right of s_{LU} go uninsured. Those to the left buy insurance. Therefore, $W_L(s)$

---

[6]However, we note that researchers should recognize that an apparent lack of vertical differentiation in a market may sometimes be an equilibrium outcome, where the market has already unraveled to a low-quality set of horizontally differentiated plans, rather than a fixed market characteristic, suggesting that selection on the intensive margin between vertically differentiated options may be relevant even if it is not empirically observable.

Figure 1: Consumer Sorting under Vertical Model
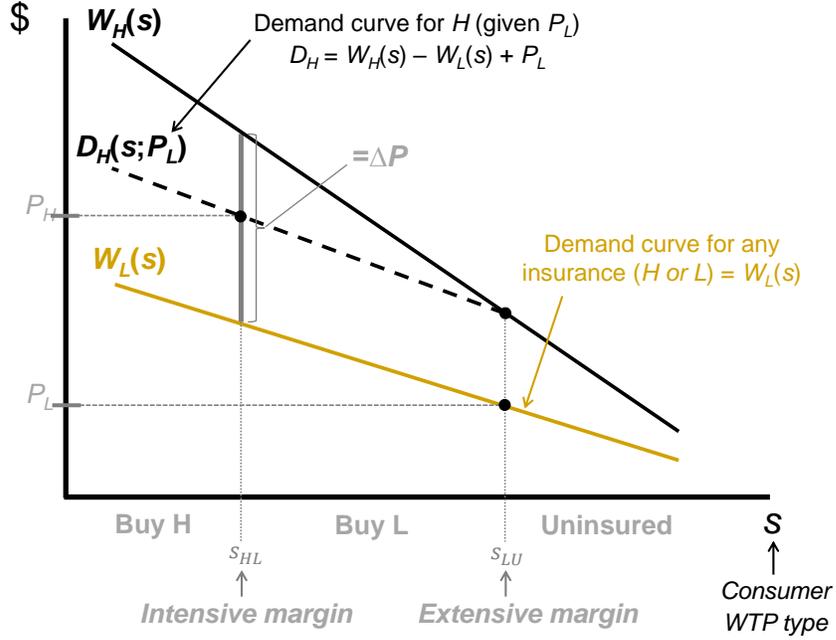
represents the *(inverse) demand curve for any formal insurance* (*H* or *L*). [7]

Let $s_{HL}(P)$ be the intensive-marginal type who is just indifferent between *H* and *L*. This cutoff type is defined by:

$$\Delta W_{HL}(s_{HL}) \equiv W_H(s_{HL}) - W_L(s_{HL}) = P_H - P_L \qquad (2)$$

Consumers to the left of s_{HL} buy *H* because their incremental WTP for *H* over *L*—which we label ΔW_{HL}—exceeds the incremental price. With demand for *H* and for *H* + *L* thus determined by Equations (1) and (2), demand for *L* equals the difference between the two.[8]

---

[7]In the more general case where consumers receive subsidies for purchasing insurance or pay a penalty when choosing *U*, $W_L(s)$ and the (inverse) demand curve for insurance will diverge. Specifically, $D_L(s) = W_L(s) + S + M$. For simplicity, we ignore the subsidy and penalty terms here but fully incorporate consumer subsidies when we use the model to study the effects of common policies (Section 3) as well as in the empirical exercise (Section 5).

[8]Formally, the demand functions for the general case where $M \neq 0$ are defined by the following equations, where $\Delta P \equiv P_H - P_L$:

$$D_H(P) = s_{HL}(\Delta P)$$
$$D_L(P) = s_{LU}(P_L - M) - s_{HL}(\Delta P)$$
$$D_U(P) = 1 - s_{LU}(P_L - M)$$

Rearranging equation (2) yields the (inverse) demand for H, given a fixed P_L:

$$D_H(s; P_L) \equiv W_H(s) - W_L(s) + P_L \tag{3}$$

Figure 1 shows $D_H(s; P_L)$ with a dashed line. One can draw D_H by noting that it intersects the W_H curve at the cutoff type s_{LU} (since $W_L(s_{LU}) = P_L$).[9] It then proceeds leftward at a slope equal to that of ΔW_{HL}, and its intersection with P_H determines s_{HL}. $D_H(s; P_L)$ is flatter than W_H because its slope equals that of $\Delta W_{HL}(s)$. This will tend to make D_H (and therefore intensive margin sorting) relatively more price elastic than D_L (extensive margin sorting).

Most importantly, $D_H(s; P_L)$ is not a pure primitive that could be identified off of exogenous price variation, but instead depends on both WTP primitives (W_H, W_L) and, critically, on P_L. Because demand for H depends on the price of L, policies targeted at altering the allocation of consumers on the extensive margin of insurance/uninsurance can have unintended effects on the sorting of consumers across the intensive H/L margin if these policies affect the price of L. The dependency of demand for H on the price of L generates an interaction between the intensive and extensive margins, a key theme of this paper.

2.2 Costs

The model's cost primitives are expected insurer costs for consumers of type s in each plan j.[10] These "type-specific costs" are defined as:

$$C_j(s) = E\left[C_{ij} \mid s_i = s\right] \tag{4}$$

$C_j(s)$ is analogous to "marginal cost" in the EFC model—so called because it refers to consumers on the margin of purchasing at at a given price. However, to avoid confusion in our model where there are two margins of adjustment, we refer to $C_j(s)$ as type-specific costs, or simply costs. In addition, we

---

[9]D_H is not defined to the right of s_{LU}, since if P_H falls further than its level at this point, nobody buys L. As a result, the demand curve for H thereafter equals $W_H(s)$.

[10]A key insight of the EFC model is that – while costs may vary widely across consumers of a given WTP type – it is sufficient for welfare to consider the cost of the *typical* consumer of each type. The reason is that with community rated pricing, consumers sort into plans based only on WTP. There is no way to segregate consumers more finely than WTP type, and since insurers are risk-neutral, only the expected cost within type matters. We note, however, that this argument breaks down when leaving the world of community rated prices, as pointed out by Bundorf, Levin and Mahoney (2012), Geruso (2017), and Layton et al. (2017). Our model (like the model of EFC) thus cannot be used to assess the welfare consequences of policies that allow for consumer risk-rating.

define $C_U(s)$ as the expected costs of uncompensated care of type-s consumers if uninsured. Along with adverse selection, external uncompensated care costs motivate subsidy and mandate policies.

Figure 2: Cost Curves under Vertical Model
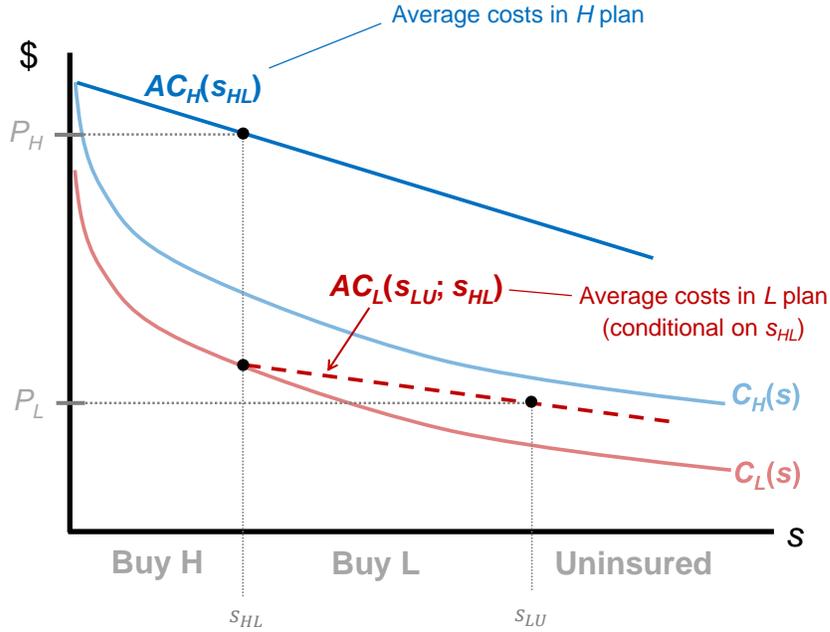
Plan-specific average costs, which are important in determining the competitive equilibrium, are defined as the average of $C_j(s)$ for all types who buy plan j at a given set of prices:

$$AC_j(P) = \frac{1}{D_j(P)} \int_{s \in D_j(P)} C_H(s)ds \tag{5}$$

where (abusing notation slightly) $s \in D_j(P)$ refers to s-types who buy plan j at prices P.

We illustrate the construction of these cost curves in Figure 2. We show a case where cost curves, C_H and C_L, are downward sloping, indicating adverse selection, though the framework could also be applied to advantageous selection. The gap between the two curves for a given s-type describes the difference in plan spending if the s-type consumer enrolls in H vs. L. We refer to this gap as the "causal" plan effect, since it reflects the true difference in insurer spending for a given set of people.[11]

We start by deriving $AC_H(P)$, the average cost curve for the H plan. To avoid ambiguity later,

---

[11]As in EFC, the causal plan effect reflects both a difference in coverage (e.g., lower cost sharing) conditional on behavior, and any behavioral effect (or moral hazard) of the plans.

it is helpful to redefine the argument of AC_H as the marginal type that buys H at price P, $s_{HL}(P)$. We use this notation in Figure 2. AC_H integrates over individual costs (C_H) from the left: For $s_{HL} = 0$, the only consumers enrolled in H are the very sickest consumers. For these consumers, $s = 0$, implying that $AC_H(s = 0) = C_H(s = 0)$. Then, as s_{HL} increases, moving right along the horizontal axis, H includes increasingly less-costly consumers, resulting in a downward sloping average cost curve. Eventually, when $s_{HL} = 1$ and all consumers are enrolled in H, $AC_H(s_{HL} = 1)$ is equal to the average cost in H across *all* potential consumers. Because H only has one marginal consumer type (the intensive margin), the derivation of $AC_H(s_{HL})$ is identical to that of the average cost curve in EFC. For each value of s_{HL}, there is only one possible value of AC_H. This implies that the curve can be calculated directly from a market primitive (by integrating over $C_H(s)$) and is not an equilibrium object.

The average cost curve for L, on the other hand, is more complicated because it is an average over a range of consumers, $s \in [s_{HL}, s_{LU}]$, with two endogenous margins. For each value of s_{LU}, there are many possible values of AC_L, depending on consumer sorting between H and L as determined by $s_{HL}(P)$. This fact makes it impossible to plot a single fixed AC_L curve as we did with AC_H. Nonetheless, it is possible to plot $AC_L(s_{LU})$ *conditional* on $s_{HL}(P)$. We denote this curve $AC_L(s_{LU}; s_{HL})$ and illustrate it with a dashed line in Figure 2. There are many such iso-s_{HL} plots of AC_L (not pictured) that hold P_H fixed at various levels. The leftmost point of the AC_L curve depends on the s_{HL} cutoff type determined by P_H. Higher values of s_{HL} imply that $AC_L(s_{LU}; s_{HL})$ starts from a higher point. Just as AC_H equals C_H at $s = 0$, AC_L equals C_L at $s = s_{HL}$. Moving rightward from $s = s_{HL}$, L adds more relatively healthy consumers, resulting in a downward sloping average cost curve.

In summary, While AC_H is fixed and does not depend on the price of L, AC_L is an equilibrium object in that it changes as the the cutoff s_{HL} changes. Because s_{HL} is determined by P_H, this implies that the average cost of L and thus the price of L depends on the price of H. This implies another dependency between extensive and intensive margin selection, the focus of this paper. For example, a subsidy targeted to H that results in a lower P_H and a larger (rightward-shifted) s_{HL} in Figure 2 would cause the leftmost point on AC_L to shift down and rightward and would cause the curve to have a less-steep slope. In a competitive market, this would likely result in a lower P_L, causing additional consumers to enter the market.

2.3 Competitive Equilibrium

We consider competitive equilibria where plan prices, P^*, exactly equal their average costs:[12]

$$P_H = AC_H(P)$$

$$P_L = AC_L(P) \tag{6}$$

In some settings, there will be multiple price vectors that satisfy this definition of equilibrium, including vectors that result in no enrollment in one of the plans or no enrollment in either plan. Because of this, we follow Handel, Hendel and Whinston (2015) and only consider equilibria that meet the requirements of the Riley Equilibrium (RE) notion. We discuss these requirements and provide an algorithm for empirically identifying the RE in Appendix A.1.

With the outside option of uninsurance, the equilibration process for the prices of H and L differs somewhat from the more familiar settings explored by EFC and Handel, Hendel and Whinston (2015). In those settings, it is assumed that all consumers choose either H or L. Assuming full insurance conveniently simplifies the equilibrium condition from two expressions to one: Namely, that the differential average cost must be set equal to the differential price.

To provide intuition for determining the equilibrium in our more complex setting, we build up from the classic case considered by EFC, which includes only H and U as plan options.[13] The EFC equilibrium can be seen in Panel A of Figure 3. It is defined by the intersection of W_H and AC_H, which determines the competitive equilibrium price. Any s-type whose WTP for H exceeds the price of H will buy H and all other s-types will opt to remain uninsured.

We now add L to the choice set. To illustrate the equilibrium, we proceed in four steps, corresponding to the four panels in Figure 3. Panels A and B show how P_H is determined, given a fixed price of L. Panel A shows that the fixed P_L implies a given extensive margin cutoff, s_{LU}. Panel B shows that this in turn implies an H plan demand curve, $D_H(P_L)$ (shown in dashed black), whose intersection with H's average cost curve determines P_H (and the intensive margin cutoff s_{HL}). This process determines the reaction function $P_H^*(P_L)$, which describes the breakeven price of H for a given

---

[12]We note that this definition of equilibrium prices differs slightly from the definition of Einav, Finkelstein and Cullen (2010) who consider a "top-up" insurance policy where only the price of H is required to be equal to its average cost, while the price of L is fixed. It is consistent, however, with the definition of Handel, Hendel and Whinston (2015)

[13]The correct analogy from EFC to our framework considers the choice between H and U rather than between H and L because the distinguishing feature of U is that its prices is exogenously determined, like the lower coverage option in the EFC setting.

fixed price of L.

Figure 3: Equilibrating Process with H, L, and Outside Option

(b) Panel B

(a) Panel A
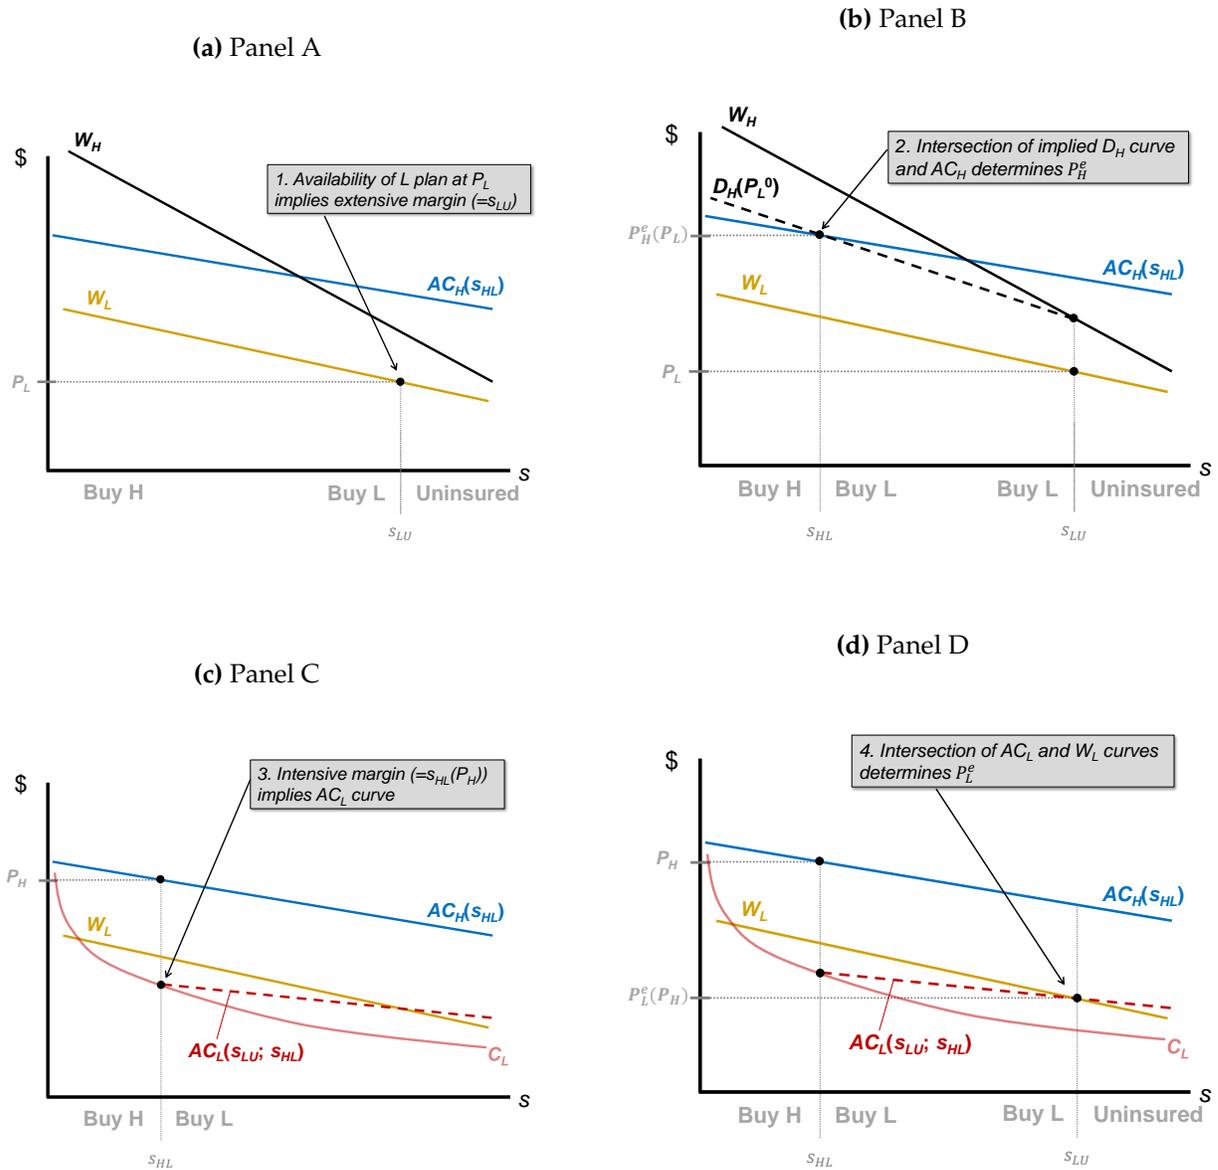
Panels C-D of Figure 3 show how P_L is determined, given a fixed P_H. Panel C shows that the fixed P_H implies a given intensive margin cutoff (s_{HL}), which in turn fixes the AC_L curve. Panel D shows how the intersection of AC_L with W_L determines P_L (and the extensive margin cutoff s_{LU}). This process determines the reaction function $P_L^e(P_H)$, which gives the breakeven price of L for a given fixed price of H.

In equilibrium, the reaction functions must equal each other: $P_H = P_H^e(P_L)$ and $P_L = P_L^e(P_H)$. Figure 4 depicts the equilibrium, including the AC_L and D_H curves as dashed lines. These dashed

lines are themselves equilibrium outcomes, even holding fixed consumer preferences and costs. In other words, there were many possible iso-s_{HL} AC_L curves and many possible iso-P_L D_H curves. The equilibrium vector of prices are the prices at which demand for L generates $D_H(P_L)$ *and* this demand for H simultaneously implies the AC_L curve by its intersection with AC_H.

Figure 4: Equilibrium
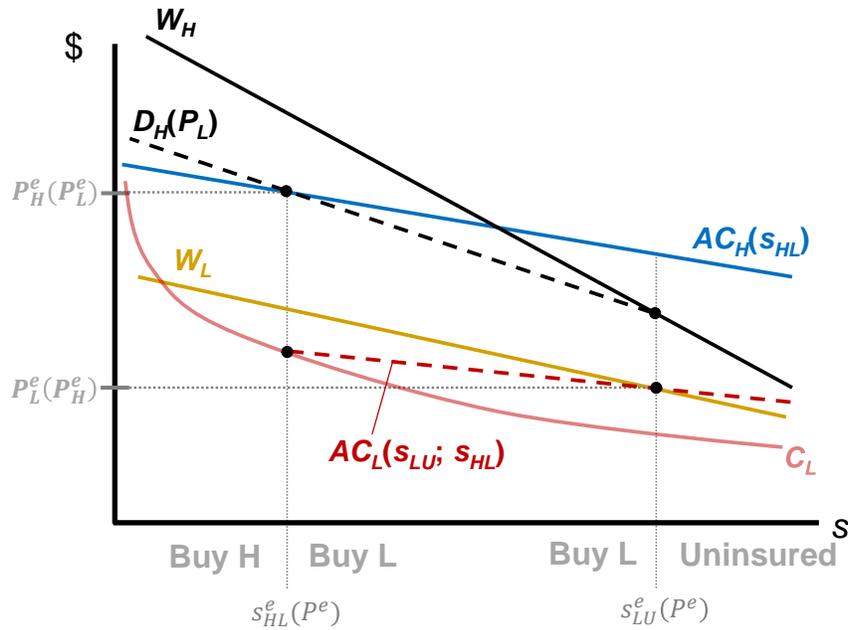
2.4 Social Welfare

We now show how our framework can be used to assess the welfare consequences of different policies. We define social welfare in the conventional way, as total social surplus. We provide a formal definition below, but we start by showing what we mean graphically. In order to make the figures simpler and more intuitive, we set C_U, the social cost of uninsurance, equal to zero. We nonetheless allow for a positive social cost of uninsurance in our empirical application below.

To build intuition, We start in Panel a of Figure 5 by illustrating the case where L is a pure cream-skimmer. That is, L has low costs on average because it attracts low cost individuals, but has no causal effect on costs, so $C_L = C_H$. For this case, given W_H, W_L, and $C_L = C_H$ we can find total social surplus for any allocation of consumers across plans described by the equilibrium cutoff values s_{HL}^e and s_{LU}^e.

15

Figure 5: Welfare

(a) Welfare when L is a Pure Cream-Skimmer
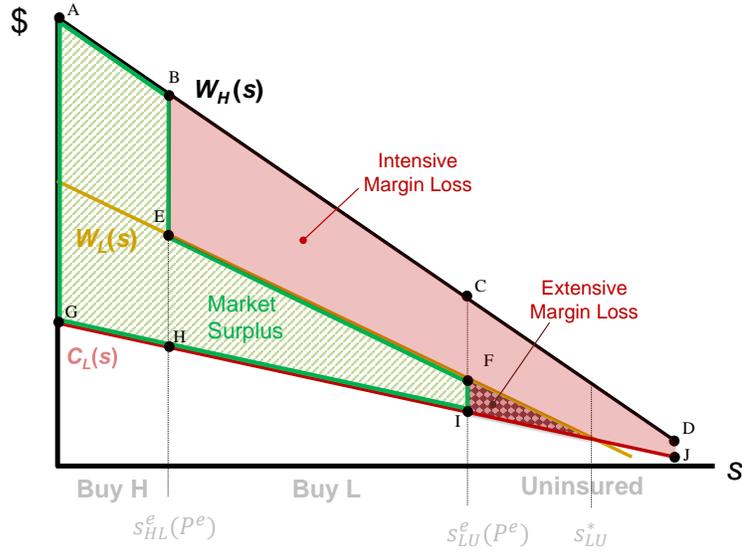
(b) Welfare when L has a Cost Advantage
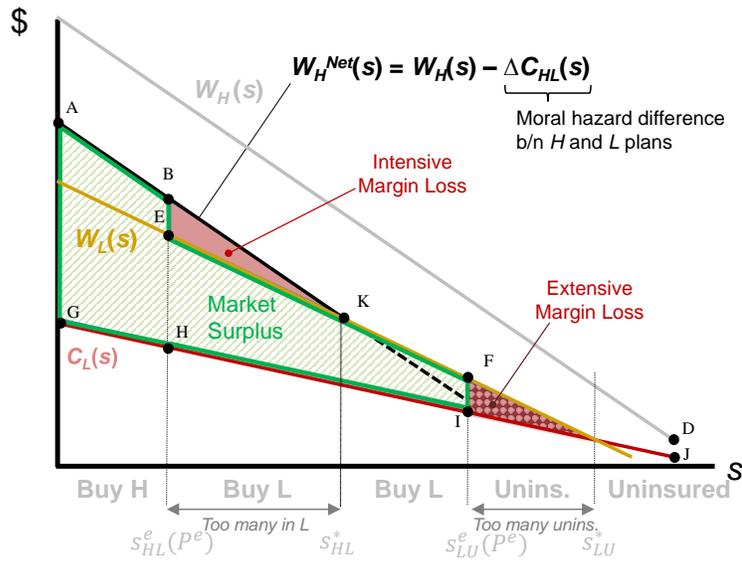
Panel a of Figure 5 shows that social surplus consists of two pieces. The first piece is the social surplus for consumers purchasing H, given by the area between W_H and $C_L = C_H$ for consumers with $s < s_{HL}$ ($ABHG$). The second piece is the social surplus for consumers purchasing L, given by the area between W_L and $C_L = C_H$ for consumers with $s_{HL} < s < s_{LU}$ ($EFIH$)). Panel a of Figure 5 also illustrates foregone surplus for this particular allocation of consumers across plans. Here, the foregone surplus consists of three components. The first component is the foregone surplus due to the

fact that consumers with $s_{HL} < s < s_{LU}$ purchased L when they would have generated more surplus by purchasing H, and it is described by the area between W_H and W_L for these consumers ($BCFE$). The second component is the foregone surplus due to the fact that consumers with $s > s_{LU}$ did not purchase insurance when they would have generated positive surplus by purchasing H, and it is described by the area between W_H and $\max\{W_L, C_L\}$ ($CDJF$). We refer to these two components as "intensive margin loss". The third component is the foregone surplus due to the fact that consumers with $s_{LU} < s < s_{LU}^*$ did not purchase insurance when they would have generated positive surplus by purchasing L, and it is described by the area between W_L and C_L for those consumers.

The figure thus shows how our graphical framework can be used to estimate welfare for any allocation of consumers across H, L, and U. It also shows that, like EFC, our framework accommodates the possibility that it is more efficient for some consumers (i.e. with $s > s_{LU}^*$) to go uninsured rather than purchase L. Further, the framework, makes it easy to determine the optimal allocation of consumers between insurance and uninsurance and between H and L. In the case of the particular demand and cost primitives drawn, it is straightforward to see that the optimal allocation of consumers across plans is for all consumers to be in H. If H were not available, however, the optimal allocation of consumers across L and U would consist of all consumers with $s < s_{LU}^*$ purchasing L and all other consumers remaining uninsured.

In Panel b of Figure 5, we show how our framework can also accommodate the possibility that it is more efficient for some consumers to be enrolled in L vs. H. To do this, we change the assumption that L is a pure cream-skimmer and instead assume that costs in H are higher than in L for each consumer and that the cost gap is constant across consumers: $C_H(s) - C_L(s) = \delta > 0$. It is convenient to define a new curve $W_H^{Net}(s) = W_H(s) - (C_H(s) - C_L(s))$, or WTP for H net of the incremental cost of H vs. L. Under the assumption that δ is constant, $W_H^{Net}(s)$ will be parallel to and below W_H. This is shown in Panel b of Figure 5: As L's cost advantage over H increases, W_H^{Net} shifts further down.[14]

Given this new W_H^{Net} curve, social welfare is still fully characterized by the three curves, W_H^{Net}, W_L, and C_L, and the social surplus and foregone surplus are defined in a similar manner to Panel a. Social surplus still consists of two components. The first is the surplus generated by the consumers enrolled in H, and it is characterized by the area between W_H^{Net} and C_L for consumers with $s < s_{HL}$ ($ABHG$). To see this, note that this gap is equal to $W_H^{Net}(s) - C_L(s) = W_H(s) - (C_H(s) - C_L(s)) -$

---

[14]Heterogeneity in L's cost advantage across s types could also be accommodated and would result in W_H^{Net} not being parallel to W_H.

$C_L(s) = W_H(s) - C_H(s)$. Note that this component is smaller than it was in Panel a due to the fact that now H has higher costs than L. In panel b it is thus less socially advantageous for these consumers to be enrolled in H vs. L. The second component is the surplus generated by the consumers enrolled in L, and it is characterized exactly as before by the area between W_L and C_L for consumers with $s_{HL}^e < s < s_{LU}^e$ (*EFIH*). Foregone surplus is also illustrated in the figure and in panel b and consists of two components. The first is the foregone intensive margin surplus due to the fact that consumers with $s_{HL}^e < s < s_{HL}^*$ are enrolled in L but would generate more surplus if they were enrolled in H. It is characterized by the area between W_H^{Net} and W_L for these consumers (*BKE*). (Unlike in Panel a, with H's higher costs it is now inefficient for any consumer with $s > s_{HL}^*$ to enroll in H.) The second component represents the extensive margin foregone surplus, and it is identical to the extensive margin foregone surplus in Panel a. In summary, our model can accommodate settings where it is not socially efficient for all consumers to be enrolled in H or even in L, such as settings where there is moral hazard, administrative costs, etc.

More formally, for cases where $C_U \neq 0$ we define social welfare as:

$$\hat{SW}(P) = \int_0^{s_{HL}(P)} (W_H(s) - C_H(s)) \, ds + \int_{s_{HL}(P)}^{s_{LU}(P)} (W_L(s) - C_L(s)) \, ds - \int_{s_{LU}(P)}^1 C_U(s) \, ds \qquad (7)$$

Recall that the level of utility was normalized above by setting $W_U = 0$. It is convenient to renormalize social welfare by adding a constant equal to total potential uncompensated care, defining $SW = \hat{SW} + \int_0^1 C_U(s) \, ds$. Rearranging and simplifying, this yields the following expression:

$$SW = \underbrace{\int_0^{s_{LU}(P)} \left(W_L(s) - C_L^{net}(s)\right) ds}_{\text{Net Surplus from Insurance in } L} + \underbrace{\int_0^{s_{HL}(P)} \left(\Delta W_{HL}(s) - \Delta C_{HL}(s)\right) ds}_{\text{Extra Surplus from } H} \qquad (8)$$

where $\Delta C_{HL}(s) \equiv C_H(s) - C_L(s)$ and $C_L^{net}(s) \equiv C_L(s) - C_U(s)$. Social welfare equals the sum of two terms. The first is the net surplus from insurance (in L) relative to uninsurance, which applies to all types who buy insurance, $s \in [0, s_{LU}]$. The second is the extra surplus from H for the subset of enrollees who buy H, $s \in [0, s_{HL}]$. Equation 8 shows that it is straightforward to calculate welfare even when $C_U \neq 0$, as long as the researcher has information about C_U.

3 Two-Margin Impacts of Risk Selection Policies

In this section, we use our model to assess the consequences of three policies commonly used to combat adverse selection in insurance markets: benefit regulation, the mandate penalty on uninsurance, and risk adjustment transfers. Each of these policies is targeted at one margin of adverse selection, but our model shows how they affect the other. We discuss each policy in turn and provide graphical illustrations for their consequences. We conclude with a discussion of other policies where cross-margin impacts on selection may be relevant, including behavioral interventions targeting take-up.

3.1 Benefit Regulation

We start by examining benefit regulation. In Figure 6, we consider a rule that eliminates L plans from the market. This thought experiment captures a variety of policies that set a binding floor on plan quality – e.g., network adequacy rules, caps on out-of-pocket limits, and the ACA's "essential health benefits." These policies seek to address *intensive margin* adverse selection problems by eliminating low-quality, cream-skimming plans. But they can also have unintended *extensive margin* consequences.

Panel (a) of Figure 6 shows the baseline equilibrium with both H and L plans, while Panel (b) shows equilibrium with L plans eliminated, which reduces to the classic EFC equilibrium. Panel (c) shows the welfare impact of benefit regulation. This involves two competing effects: Some consumers formerly in L shift to H (the intended consequence), and some consumers formerly in L become uninsured (the unintended consequence).

In the textbook cream-skimming case – where H is the socially efficient plan for everyone (though L is still better than uninsurance for most consumers) – these two effects have opposite welfare consequences. The first (intended) effect *improves welfare* by shifting people out of L – an inefficient plan that exists only by cream-skimming – and into H. The second (unintended) effect, however, *lowers welfare* by shifting some L consumers into uninsurance. Thus, even in this textbook case where the L plan is an inefficient cream-skimmer, banning it has ambiguous welfare consequences.[15]

What explains this counter-intuitive result? This can be thought of as an example of "theory of the second best"-style interactions that emerge with two margins of selection. Regulation that bans a

---

[15]The net welfare impact depends on the market primitives (W_H, W_L, C_H, C_L) and the social cost of uninsurance, C_U. Section 2 presents the framework for how these can be measured and the net welfare impact quantified.

Figure 6: Impact of Benefit Regulation

(a) No benefit regulation

(b) Plan L eliminated
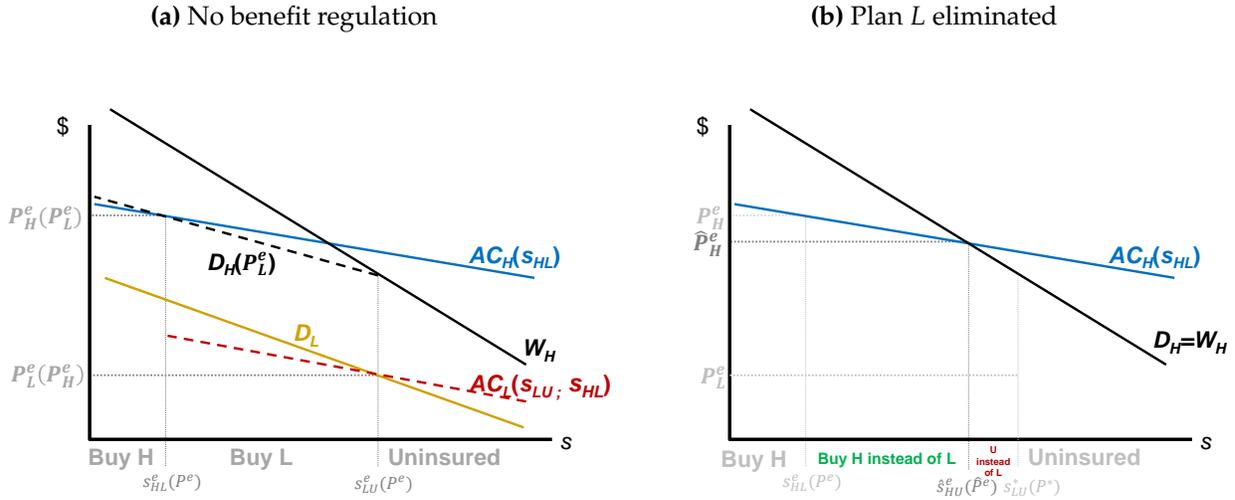
(c) Net enrollment and welfare impacts (B-A)
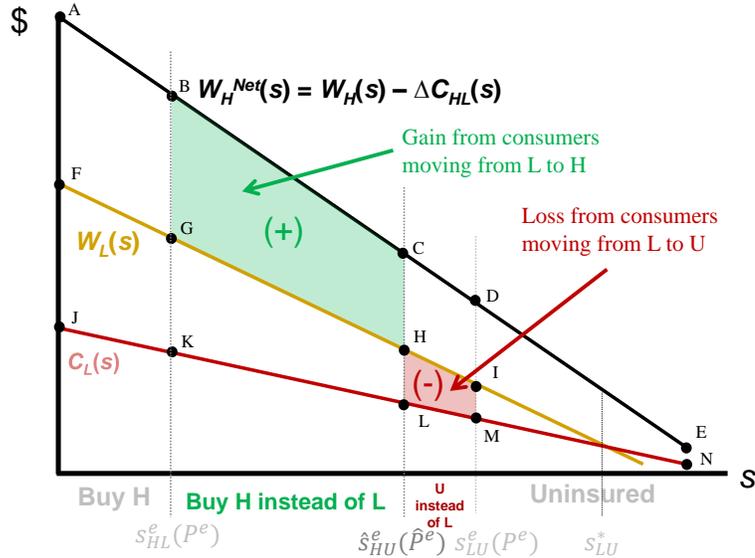
Notes: The figure shows the impact on equilibrium (panels a and b) and welfare (panel c) of a benefit regulation that eliminates the L plan. This thought experiment captures a variety of policies that set a binding floor on plan quality, thus eliminating low-quality plans. For welfare impacts, we show the textbook case where H is the efficient plan for all consumers and L is more efficient than U.

pure cream-skimming L plan addresses an intensive margin selection problem. But it has the unintended side effect of worsening the extensive margin selection problem of too much uninsurance. Put differently, a pure cream-skimming L plan adds no social value *within* the market, but by segmenting the healthiest people into a low-price plan, it can improve welfare by bringing new consumers *into* the market.[16]

---

[16]Of course, this reasoning depends on the market stabilizing to a separating equilibrium where both H and L survive.

3.2 Mandate Penalty on Uninsurance

Next we consider the consequences of a mandate penalty on uninsurance (U). The analysis is also applicable for the effect of larger insurance subsidies, which likewise reduce consumers' net price of buying insurance relative to remaining uninsured.

The mandate penalty has both a direct effect and an indirect effect through equilibrium price adjustments. The direct effect of a mandate penalty is to increase the demand for insurance. Panel (a) of Figure 7 shows this via an upward shift in W_L and W_H by \$$M$, reflecting that both are now cheaper relative to U (whose utility and price are normalized to zero). As a result of this shift, some people who were previously uninsured buy insurance in the L plan. This is the intended effect of the penalty.

Panel (b) depicts the unintended, equilibrium effects of the penalty. By definition under extensive margin adverse selection, the newly insured individuals are relatively healthy. Because they buy the low-price L plan, they lower L's average costs (i.e., a movement down the AC_L curve, not a shift in the AC_L curve) and therefore its price. The lower P_L leads some consumers to shift on the *intensive margin* from H to L – as captured by the downward shift in H's demand curve, $D_H(P_L)$. This is the main unintended effect of the penalty: although it is intended to reduce uninsurance, the penalty also shifts people toward lower-quality plans on the intensive margin.[17]

There is a second equilibrium effect from this shift in consumers from H to L. The consumers who shift are high-cost relative to L's previous customers, pushing up its average costs. In panel (b), this is depicted via an upward shift in the $AC_L(P_H)$ curve – which has to occur because of the higher P_H and the leftward shift in the marginal s_{HL} type. The higher average costs in L partly offset the fall in P_L due to the mandate and dampens the impact of the mandate on the price of L. Thus our model shows how and why cross-margin effects may make a mandate less effective than one would predict from its direct effects alone: The penalty induces healthy people to enter the market but also induces relatively sick people to move from H to L. Nonetheless, as long as the original equilibrium is stable, one can show that on net, a larger penalty decreases P_L and uninsurance (see Appendix B

If the market unravels to the L plan, insurance coverage will typically not be higher: the price of L will not be low (since it attracts all consumers), and because the quality of L is lower, uninsurance will typically be higher than in an H-only equilibrium where L is banned. Whether the market stabilizes to a separating equilibrium or unravels to L/upravels to H depends on the market primitives.

[17]We show in our simulations and in Appendix B that this prediction is largely robust to relaxing the vertical model. It is driven by two properties: (1) that the newly uninsured are relatively healthy (extensive margin adverse selection), and (2) that the newly insured mostly choose the low-priced L plan.

Figure 7: Impact of Mandate Penalty on Uninsurance

(c) Welfare Effects

Notes: The figure shows the impact of a mandate penalty in our framework. Panel (a) shows the direct effect: higher demand for insurance. Panel (b) shows the unintended equilibrium effect: an intensive margin shift from H to L. Panel (c) shows the welfare effects in the textbook case where H is the efficient plan for all consumers and L is more efficient than U.

for a proof).

Panel (c) of Figure 7 shows the welfare effects in the textbook case where H is the efficient plan for all consumers. There are again competing effects: (intended) welfare gains from newly insured consumers and (unintended) welfare losses from consumers moving from H to the lower-quality L plan. Thus, the interaction of the two margins of selection makes the welfare impact of a mandate ambiguous even in this textbook case. An uninsurance penalty seeks to limit the effects of adverse selection on the extensive margin by inducing healthy consumers to enroll in insurance. However,

by bringing healthy consumers into L from uninsurance, the penalty may also lead some consumers previously enrolled in high-quality coverage to shift to lower-quality coverage. In the extreme, a penalty could even lead to a market where high-quality contracts are unavailable to consumers (i.e. market unraveling to L).

3.3 Risk Adjustment Transfers

Of the three policies we consider, risk adjustment is the most difficult to illustrate graphically because the policy adds new risk-adjusted cost curves (for both L and H) that crowd the figure. Additionally, risk adjustment transfers cause AC_H to be affected by both margins instead of just one, as any effects of selection into the market are at least partially shared between L and H due to the transfers. Nonetheless, risk adjustment is an important policy lever used to combat intensive margin selection, so we show how our graphical model works with a simplified version of it. Specifically, we graph how *perfect* risk adjustment, where transfers perfectly capture all variation in C_L across consumer types, affect equilibrium outcomes. We then describe the effects of magnifying *imperfect* risk adjustment transfers, the real-world policy option currently available to states and the policy we study empirically in Section 5. For that, we rely on comparative statics derived in Appendix B.

Perfect Risk Adjustment   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, intensive margin selection remains relevant, complicating the graphical analysis.

Figure 8: Equilibrium under Perfect Risk Adjustment

(a) Effect of Perfect Risk Adjustment on Prices and Market Shares

(b) Welfare Effects of Perfect Risk Adjustment

We depict the perfect risk adjustment case in Panel (a) of Figure 8. 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 average cost curves whose intersection with demand determine equilibrium – shifting from the old unadjusted AC_H and AC_L (shown with semi-transparent lines) to risk-adjusted RAC_H and RAC_L. The risk-adjusted curves differ from the unadjusted curves in the following ways. For the L plan, the risk-adjusted curve, $RAC_L(s_{LU})$, is *shifted upward* relative to the unadjusted $AC_L(s_{LU})$, reflecting 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.[18]

For the H plan, $RAC_H(s_{HL})$ is *rotated downward* versus unadjusted $AC_H(s_{HL})$. 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. In addition to these shifts in H's cost curve, the higher P_L under risk adjustment shifts up the demand curve for H, further increasing the number of people who buy H instead of L.

Therefore, perfect risk adjustment has two effects. First, it narrows the average cost and therefore price gap between H and L, leading consumers to shift on the intensive margin towards the H plan. This is the intended purpose. Second, it pushes up the average cost and therefore the price of L. This results in some consumers choosing to be uninsured who would have chosen L in the absence of risk adjustment. This is the unintended, cross-margin consequence of risk adjustment.

Panel (b) of Figure 8 depicts the welfare effects of these changes in sorting, again in the textbook case where H is the efficient plan for all (even though L now has a cost advantage equal to δ, incorporated into $W_H^{Net} = W_H - \delta$). As with benefit regulation and the mandate penalty, there are opposing effects: a welfare gain from the intensive margin shift from L to H and a welfare loss from the extensive margin shift from L to uninsurance. (There is also a welfare gain on the extensive margin due to the fact that some of the people induced to choose uninsurance instead of L generate negative social surplus when enrolled in L) This suggests that, like the other policies, the welfare effects of risk adjustment are theoretically ambiguous. Our model provides a simple framework for estimating the net welfare effects given the relevant curves.

---

[18]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.

Imperfect Risk Adjustment While perfect risk adjustment is a useful thought experiment, most markets include an imperfect form of risk adjustment where transfers are based on individual risk scores computed from diagnoses appearing in health insurance claims. (See Geruso and Layton, 2015 for an overview.) For instance, in the ACA Marketplaces, the per-enrollee transfer from L to H is determined by the following formula:[19]

$$T(P) = \left(\frac{\overline{R}_H(P) - \overline{R}_L(P)}{\overline{R}(P)} \right) \cdot \overline{P}(P) \tag{9}$$

where $\overline{R}_j(P)$ is the average risk score of the consumers enrolling in plan j given price vector P, $\overline{R}(P)$ is the (share-weighted) average risk score among all consumers purchasing insurance, and $\overline{P}(P)$ is the (share-weighted) average price in the market. Note that the transfer is positive as long as H's average risk score is larger than L's average risk score.

In Appendix B we introduce a parameter α and define the transfer from H to L as $\alpha \cdot T(P)$ so that α describes the strength of risk adjustment with $\alpha = 0$ implying no risk adjustment, $\alpha = 1$ implying ACA risk adjustment, $\alpha = 2$ implying transfers twice as large as ACA transfers, and so on. We then derive some comparative statics describing the effect of an increase in α (i.e. a magnification of the imperfect transfers) on P_H and P_L. These comparative statics mimic the simulations we perform in the empirical section where we simulate equilibria under no risk adjustment and with increasingly large risk adjustment transfers (i.e. increasingly large values for α).

The comparative statics reveal that larger values of α (i.e. stronger transfers) unambiguously lower the price of H, as in the perfect risk adjustment case above. The effect of an increase in α on the price of L, however, is ambiguous. In addition to risk adjustment's direct effect to push up L's average costs (which drove the results under perfect risk adjustment), there is a second indirect effect. The consumers who shift from L to H tend to be L's most expensive enrollees, even net of imperfect risk adjustment. This lowers L's risk-adjusted average costs, pushing the price of L downward. This indirect effect will be larger when intensive margin adverse selection is severe (even after risk adjustment) and when consumers are quite price elastic on the intensive margin. Indeed, we find in some of our simulations (which are built on well-identified estimates of the relevant elasticities) that the indirect effect is large, and risk adjustment has minimal effects or even decreases P_L.[20]

---

[19]The actual formula used in the Marketplaces is a more complicated version of this formula that adjusts for geography, actuarial value, age, and other factors. Our insights hold with or without these adjustments, so we omit them for simplicity.

[20]This is particularly likely to happen when one allows for L to have a large cost advantage over H.

In summary, our model provides clear predictions for the effects of risk adjustment in a setting where consumers choose between H, L, and the outside option U. If risk adjustment is perfect, 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, the unintended extensive margin effect may or may not occur, depending on the relative sizes of the direct (transfer) effect and the indirect (substitution) effect.

3.4 Other Policies

The same price theory can be applied to other policies not explicitly discussed above. The key insight is that anything that affects selection on one margin has the potential to affect selection on the other margin, as firms adjust prices in equilibrium to compensate for the changing consumer risk pools. For example, the stylized benefit regulation case in Section 3.1 nests many specific interventions to address intensive margin selection, including network adequacy rules, Essential Health Benefits, and actuarial value requirements.

Our framework also informs on the potential impacts of reinsurance, a federal policy in place from 2014 to 2016 in the ACA Marketplaces. Reinsurance has gained research attention for desirable market stabilization and incentive properties (Geruso and McGuire, 2016; Layton, McGuire and Sinaiko, 2016) and has been adopted in various forms by some states since the federal program expired.[21] To the extent reinsurance is implemented as a system of budget-neutral enforced transfers based on insurer losses for specific conditions, it generates effects similar to those we document for risk adjustment. To the extent that reinsurance is implemented as an external subsidy into the market by fees assessed on plans outside of the market (as in the ACA), it shares properties of both the mandate penalty (by providing an overall insurance subsidy, making both H and L cheaper) and risk adjustment (by targeting the subsidy to higher-cost enrollees more likely to be in H than in L), resulting in simultaneous extensive and intensive margin effects that would be difficult to assess in models focusing only on one margin or the other.[22]

It is important to understand that the cross margin effects are relevant not only for policies that *aim* to address selection, but also for policies for which selection impacts are incidental or a nui-

---

[21]In policy practice, the term "reinsurance" is used to describe a wide gamut of regulatory interventions. see Harrington (2017) for a typology.

[22]In particular, it has the same kind of "indirect" effect of the mandate, which is to move the net cost curves—here because of payments to plans, rather than because of risk pool composition shifts.

sance. Handel (2013), for example, shows how addressing inertia through "nudging" can exacerbate intensive margin selection in an employer-sponsored plan setting. Our model implies that in other market settings, where uninsurance is a more empirically-relevant concern, there is a further effect of nudging: Worsening risk selection on the intensive margin through behavioral nudges may improve risk selection on the extensive margin, potentially counterbalancing the welfare harm documented in Handel (2013). Similar insights apply to any behavioral intervention that even *incidentally* impacts plan choice in a way that affects the sorting of consumer risks (expected costs) across plans. Similarly, behavioral interventions intended to increase take-up of insurance, such as information interventions or simplified enrollment pathways, may have important intensive margin consequences similar to the effects of a mandate.

4 Simulations: Methods

Any set of reduced form estimates of demand and cost functions could be used to demonstrate how our model can be applied empirically. Here, we draw on estimates of demand and costs from the Massachusetts pre-ACA subsidized health insurance exchange, known as Commonwealth Care or "CommCare," from Finkelstein, Hendren and Shepard (2017) (which we abbreviate as "FHS"). We combine the FHS primitives, which describe lower-income consumers, with corresponding estimates for higher-income Massachusetts households purchasing coverage on the unsubsidized side of the individual market, known as "the Connector." The latter estimates come from Hackmann, Kolstad and Kowalski (2015) (which we abbreviate as "HKK"). Both sets of demand and cost curves are well-identified using exogenous variation in net consumer prices. FHS use a regression discontinuity design based on three household income cutoffs that generate discrete changes in consumer subsidies. HKK use a difference-in-differences design leveraging the introduction of an uninsurance penalty in Massachusetts. Additional details about the estimation of the FHS and HKK curves can be found in Appendix C as well as in the respective papers.

We make two key modifications to the baseline FHS and HKK estimates. First, to allow for the widest possible set of policy counterfactuals, we must extrapolate the curves to generate estimates over the full range of s-types. For example, the W_L curve from FHS spans $s = 0.36$ to $s = 0.94$. We extend the range down to $s = 0$ and up to $s = 1$. Second, we combine the two sets of estimates to form one set of aggregated demand and cost curves. The aggregated curves reflect a single market

that combines the subsidized (low-income) and unsubsidized (high-income) enrollees, mirroring the mixing of these populations in the ACA Marketplaces. To do this, we assume that the low-income group makes up 60% of the market, reflecting national enrollment patterns in ACA Marketplaces. Details regarding these modifications and the construction of our demand and cost curves, as well as figures showing the final demand and cost curves, are found in Appendix C. Our procedure yields estimates of $W_H(s)$, $W_L(s)$, $C_H(s)$, and $C_L(s)$ for the full range of consumer types $s \in [0, 1]$.

Given these demand and cost curves, it is straightforward to estimate equilibrium prices and allocations of consumers across H, L, and U under a given set of policies. Our method for finding equilibrium is based on the approach described in Figure 3. We start by considering price vectors resulting in positive enrollment in both H and L. For each potential P_L we find the P_H such that $P_H = AC_H$ and for each potential P_H we find the P_L such that $P_L = AC_L$. We then find where these two "reaction functions" intersect. The intersection is the price vector at which both H and L break even. We then also consider price vectors where there is zero enrollment in H, zero enrollment in L, or zero enrollment in both H and L. We then use a Riley equilibrium concept to choose which breakeven price vector is the equilibrium price vector.[23] This method results in a unique equilibrium for each policy environment we consider.

We use these demand and cost curves to find equilibrium prices and allocations of consumers across H, L, and U under different specifications of a mandate penalty ($0 to $60 per month) and risk adjustment (no risk adjustment to transfers 3 times the size of ACA transfers). We separately consider welfare in Section 6. We study the effects of these policies in a 2×2 matrix of market environments. The first dimension we vary is the subsidy regime. We consider two subsidy regimes: (1) an ACA-like regime where low-income consumers receive a subsidy linked to the price of the lowest-price plan available in the market,[24] and (2) a fixed subsidy regime where low-income consumers receive subsidies that are set exogenously. For (1) we follow the ACA rules by setting the subsidy such that the net-of-subsidy price of the index plan is equal to 4% of income for someone with an income equal to 150% of the federal poverty line (FPL) in 2011, the year on which our estimated demand and cost curves are based.[25] In our setting, this is $55. In both subsidy cases, low-income consumers receive

---

[23]See Appendix D for additional details. The version of the Riley equilibrium concept we use says that a breakeven price vector is a Riley equilibrium if there is no weakly profitable deviation resulting in positive enrollment for the deviating plan that survives all possible weakly profitable responses to that deviation. We describe how we empirically implement this equilibrium concept in the appendix.

[24]The lowest priced plan is L if there is positive enrollment in L and H if the equilibrium results in L unravelling.

[25]The ACA subsidy rules actually set the subsidy according to the price of the second-lowest cost silver plan. Our

subsidies only if they purchase H or L, and the subsidy they receive is identical no matter which plan they choose. High-income consumers never receive subsidies.

The second dimension we vary is whether L is a perfect cream-skimmer (i.e. $C_L(s) = C_H(s)$ for all s) or has a cost advantage (i.e. $C_L(s) < C_H(s)$ for all s). FHS find no evidence that L has lower costs than H in CommCare, motivating our perfect cream-skimmer case. To illustrate possibilities under other market primitives, we additionally analyze the case where L is assumed to have a 15% cost advantage (i.e. $C_L(s) = 0.85C_H(s)$). Of particular interest is how the welfare consequences of risk adjustment and the uninsurance penalty vary across these two cases. We explore these in Section 6.

5   Results: Prices and Enrollment

In this section, we present results on how prices and market shares change under (1) stronger mandate penalties and (2) stronger risk adjustment. In Appendix E we also present results on how prices and market shares change under benefit regulation, where we implement benefit regulation by eliminating L from the consumers' choice set.

5.1   Mandate/Uninsurance Penalties

We first present equilibrium market shares for each option, H, L, and U, under different levels of a mandate penalty for remaining uninsured ($P_U \equiv M$). We consider penalties in increments from \$0 to \$60.[26] In all cases we include ACA-style risk adjustment (described in detail in Section 5.2 below). Equilibrium market shares are presented in the top row of Figure A2 for the case where L is a pure cream-skimmer and the top row of Figure A3 for the case where L has a 15% cost advantage. In each figure, the case with ACA-like price-linked subsidies is shown in the top-left panel, and the case with a fixed subsidy is shown in the top-right panel.[27] All results are also reported in Tables 1 and 2.

For the two ACA-like subsidy cases (top-left), the patterns are qualitatively similar regardless of modeling L as a cream skimmer (Figure A2) or as having a cost advantage (Figure A3). When there

---

subsidy rule mimics this rule in spirit (in a way that is compatible with our CommCare setting) by linking the subsidy to the price of L.

[26]We find that in all cases, $P_U = 60$ is sufficient to drive the uninsurance rate to 0.

[27]Fixed subsidies are equal to \$275 in the case where L is a pure cream-skimmer and \$250 in the case where L has a 15% cost advantage. These values were chosen in order to ensure that risk adjustment and the uninsurance penalty have some effect on market shares. With subsidies that are "too large" no consumers opt to be uninsured and with subsidies that are "too small" no consumers opt to purchase insurance, making the simulated policy modifications uninformative.

Figure 9: Cream-skimming L Plan Market Shares

(a) ACA-like subsidy, increase mandate

(b) Fixed $ 275 subsidy, increase mandate

(c) ACA-like subsidy, increase α

(d) Fixed $ 275 subsidy, increase α
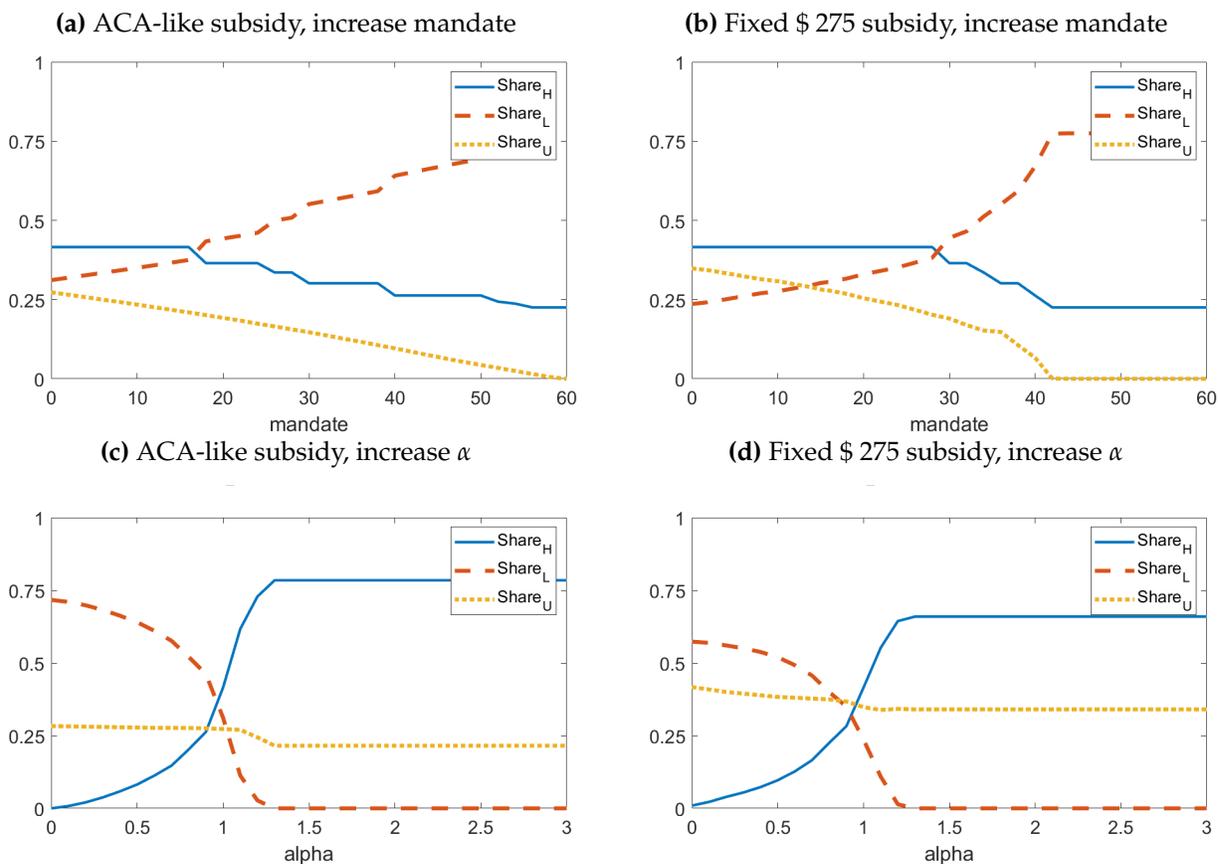
is no mandate penalty, some consumers choose each of the three options, H, L, and U, though the share in H is extremely low in the cost advantage case. As the penalty increases, the uninsurance rate decreases, with no consumers remaining uninsured at a penalty of $60/month. However, there are also intensive margin consequences: As the penalty increases there is a shift of consumers from H to L. In the case where L is a pure cream-skimmer, H's market share decreases from 42% with no penalty to 23% with a penalty of $60/month. This represents a significant decline in H's market share and a significant deterioration of the average generosity of coverage among the insured. In the case where L has a 15% cost advantage, the patterns are similar, though H's initial market share with no penalty is much lower ($\approx 2\%$), so the intensive margin consequences are less stark.

The two fixed subsidy cases are in the top-right panels of the figures, When L is a pure cream-skimmer, in the absence of a penalty consumers are split relatively evenly across H, L, and U. As the penalty increases, consumers move from U to L, the intended effect of the policy. However, at a penalty of just under $30/month the influx of relatively inexpensive consumers into L causes P_L to

Figure 10: 15% L-plan Cost Advantage Market Shares

(a) ACA-like subsidy, increase mandate

(b) Fixed $ 250 subsidy, increase mandate

(c) ACA-like subsidy, increase α

(d) Fixed $ 250 subsidy, increase α
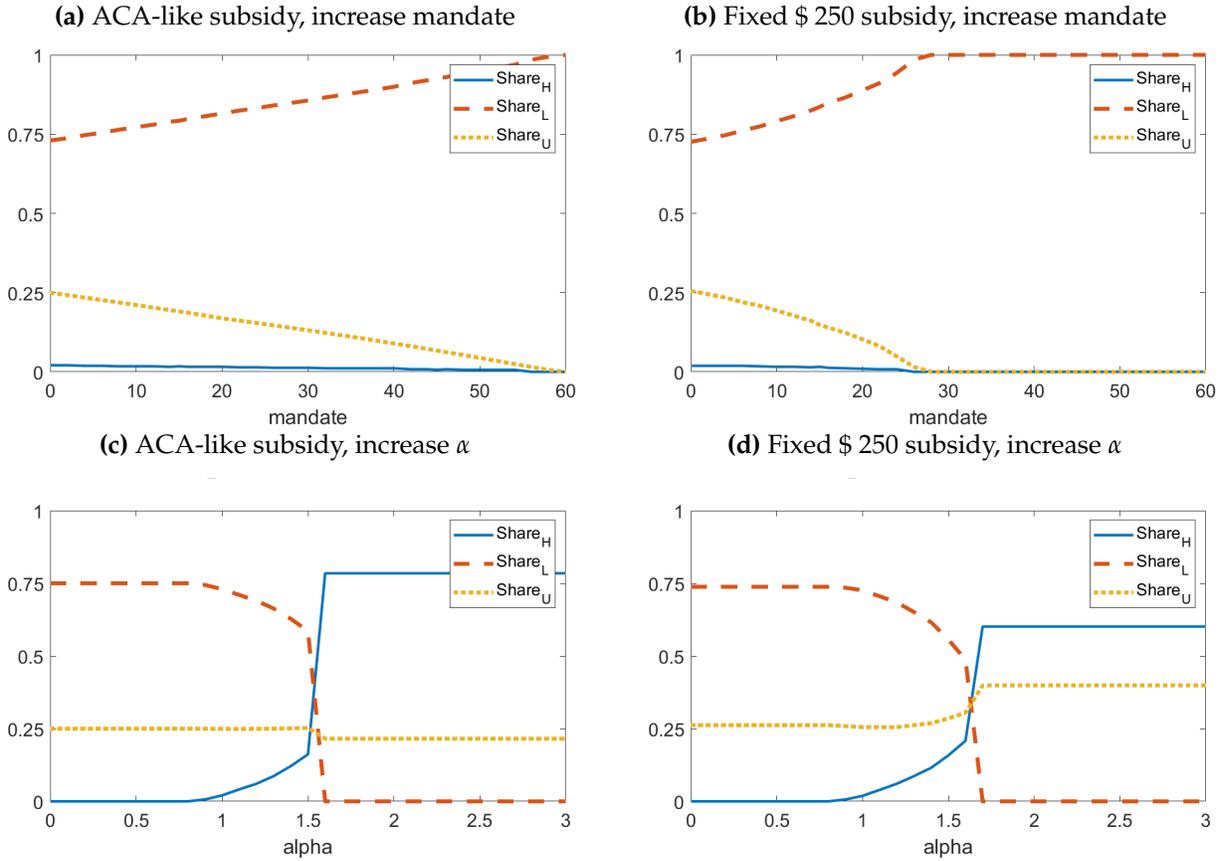
get low enough relative to P_H that some consumers previously in H now start to opt for L. As the penalty continues to increase, consumers move into L from both U and H until the mandate reaches just over $40/month and all consumers are enrolled in insurance. At this point 23% of the market is enrolled in H and 77% of the market is enrolled in L. This represents an intended decline in the uninsurance rate from 35% to 0% but also an unintended decline in H's market share from 42% to 23%. In the case where L has a 15% cost advantage, the penalty again decreases both the uninsurance rate (intended) and H's market share (unintended), but H's market share with a $0 penalty is so low (around 3.5%) that the decline in H's market share (to zero) is relatively insignificant.

In each of the empirical cases we consider in Figures A2 and A3, a larger insurance mandate penalty has the intended consequence of decreasing the portion of consumers opting to remain unin-sured *and* the unintended consequence of shifting consumers from H to L. This is consistent with implications of our graphical model as well as the comparative statics we outline in Sections 2 and 3. The unintended intensive margin effect is most stark in the case where L is a perfect cream-skimmer,

highlighting how the market primitives can amplify the cross-margin impacts of policy changes.[28]

5.2 Risk Adjustment

We now consider the effects of risk adjustment. We start with risk adjustment transfers implied by the ACA risk adjustment transfer formula first presented in Eq. (9). We first calculate risk scores for each individual using the HHS-HCC risk adjustment model used in the ACA Marketplaces. (This is a straightforward mechanical application of the regulator's algorithm to our individual-level claims data.) We then use those scores plus the FHS regression discontinuity design to estimate a "risk score curve" $RA(s)$ describing the average risk score across consumers of a given s-type. Because this curve is novel to this paper and not estimated by FHS, we describe the estimation of this curve in Appendix F. We plot this curve alongside the cost curve in Appendix Figure **??**. It is apparent that while risk scores explain part of the correlation between willingness-to-pay and costs, they do so only imperfectly. Specifically, we find that risk scores account for about half of the correlation between willingness-to-pay and costs, implying substantial selection on costs net of the ACA's imperfect risk adjustment policy.

We use the risk score curve to determine the average risk scores for H and L for any given allocation of consumers across H, L, and U. This is similar to constructing average cost curves from marginal costs. We then plug these average risk scores into the risk adjustment transfer formula (Equation 9 to determine the transfer from L to H for a given price vector $T(P)$. Finally, we find the equilibrium prices. These satisfy $P_H = AC_H(P) - T(P)$ and $P_L = AC_L(P) + T(P)$ when L and H have non-zero enrollment.

To vary the strength of risk adjustment transfers we maintain the original risk scores and structure of the transfer formula, but we multiply transfers by a scalar α (as in the comparative statics in Appendix B) so that transfers from L to H are some multiple of the transfers implied by the ACA formula. We allow α to vary from 0 (no risk adjustment) to 3 (risk adjustment transfers 3 times the size of ACA transfers. The case of ACA transfers occurs where $\alpha = 1$. This approach to evaluating strengthening or weakening risk adjustment reflects real-world policy experimentation: The federal government recently reduced α from 1 to 0.85 in the ACA Marketplaces and gave states the ability

---

[28]To see why the effect would be larger for the cream-skimmer case, note that for fixed consumer preferences, it is relatively more difficult to achieve high levels of enrollment in H when L has an actual cost advantage versus when L has similar costs to H. This leads to lower enrollment in H even at low levels of the mandate penalty, and less opportunity for a reduction in H's market share.

to further reduce α.[29] Our approach thus maps to feasible policy interventions, rather than assuming that the regulator can increase the predictive power of risk scores.

Equilibrium market shares for different levels of α in the cases without and with a cost advantage for L are found in the bottom row of Figures A2 and A3, respectively. Market shares under ACA-like subsidies are presented in the bottom-left panels of each figure, and market shares under fixed subsidies are found in the bottom-right panels. With ACA-like subsidies, patterns are qualitatively similar when L is a pure cream-skimmer and when L has a 15% cost advantage. In both cases, when there is no risk adjustment ($\alpha = 0$), the market unravels to L: No consumers choose H, and the market is split between L and uninsurance. As the strength of risk adjustment transfers increases, consumers shift from L to H. This is the intended consequence of risk adjustment. When L is a pure cream-skimmer, transfers about 1.25 times the size of ACA transfers are sufficient to cause the market to "upravel" to H. When L has a 15% cost advantage transfers need to be 1.6 times the size of ACA transfers to generate the same outcome. In both cases, there is no extensive margin effect except at the level of α where the market initially upravels to H where there is a small reduction in the uninsurance rate. This reduction is due to the fact that at this point the subsidy becomes linked to the (higher) price of H instead of the (lower) price of L due to the exit of L from the market. With the larger subsidy, more consumers purchase insurance.[30]

The bottom right panels of Figures A2 and A3 present market shares under fixed subsidies with different levels of α. Here, we again see that stronger risk adjustment transfers have the intended effect: Higher levels of α result in more consumers choosing H instead of L. In the case where L is a pure cream-skimmer, we see only a small extensive margin effect, with a small decrease in the uninsurance rate as α increases. This is consistent with our comparative statics from Section 3: The direct effect of increasing the transfer from L to H is more than fully offset by the indirect effect of the costliest (even conditional on imperfect risk adjustment) L enrollees leaving L and joining H, resulting in a decrease in P_L and a corresponding decrease in the uninsurance rate (see Section 3 and Appendix B for a fuller discussion of this result).

---

[29] The reduction of α from 1 to 0.85 occurred when the federal government decided to "remove administrative costs" from the benchmark premium that multiplies insurer risk scores to determine transfers in the transfer formula described by Eq. 9.

[30] This reduction seemingly goes against the intuition we present in Section 3 where we showed that in many cases risk adjustment may *increase* the uninsurance rate rather than decrease it as we see here. Recall, however, that in these cases the subsidy is linked to the extensive margin price. This results in risk adjustment having no effect on the net-of-subsidy extensive margin price faced by the consumer (except where L exits the market) and no unintended extensive margin consequence.

On the other hand, in the case where L has a 15% cost advantage we see a different unintended extensive margin consequence of stronger risk adjustment transfers: More consumers opt to remain uninsured. In this case, with no risk adjustment ($\alpha = 0$) all insured consumers opt for L, with no consumers choosing H and the market split between L and U. ACA risk adjustment transfers ($\alpha = 1$) barely alter these market shares. As transfers are strengthened above ACA levels, consumers begin to opt for H instead of L. At the higher levels of α, extensive margin consequences also start to appear with some consumers exiting the market and opting for uninsurance. When transfers are strengthened to two times the size of ACA transfers, the market upravels to H with all insured consumers opting for H instead of L. But at this level of α the uninsurance rate reaches almost 50%, an increase of 15 percentage points (a 60% increase above the uninsurance rate with no risk adjustment), indicating that this shift of consumers to more generous coverage on the intensive margin had a substantial extensive margin impact.

These results provide important lessons for where the unintended extensive margin effects of risk adjustment will matter most. First, ACA-like price-linked subsidies protect against the unintended extensive margin effects of risk adjustment (though there may be important effects on the size of the subsidies themselves, and thus the cost to the government). Second, the unintended extensive margin effects are more likely to occur when L has a larger cost advantage over H. In cases where L and H have similar costs, extensive margin effects are likely to be small. But when L has a large cost advantage, stronger risk adjustment can have significant effects on the portion of consumers in the market who opt to be uninsured.

6   Results: Welfare

We now show how our graphical model can be used to assess the welfare consequences of extensive and intensive margin policies, following our graphical analysis in Section 2. We first characterize welfare at a baseline equilibrium, then trace the gains and losses associated with illustrative policy changes, and finally determine optimal policy. Importantly, we show that the optimal size of the mandate is dependent on the parameter determining the strength of risk adjustment and vice versa. One straightforward implication is that if mandate penalties were altered by legislative action or court outcomes, a constrained optimal response from a regulator would likely be to adjust risk adjustment strength in concert. (Unlike altering a mandate penalty, a regulator would typically have authority

35

to tune risk adjustment without further changes to law.)

6.1 Baseline

We begin by noting the possibility that in many settings, social surplus may not be increased by policies that increase insurance take-up or that move consumers from less generous coverage to more generous coverage. This is because some consumers may not value insurance more than the cost of providing it to them and may not value the incremental coverage provided by more generous plans more than the incremental cost of providing that coverage. Further, we have shown above that policies may have opposing effects on the intensive and extensive margins, increasing enrollment in more generous coverage while simultaneously decreasing overall insurance take-up, or vice versa. For these reasons, it is important to understand the effects of policies not just on market allocations but also on social welfare.

As discussed in Section 2, it is straightforward to estimate overall social surplus associated with some equilibrium market outcome (enrollment shares), given the $W_H^{Net} = W_H - (C_H - C_L)$; W_L; and $C_L^{Net} = C_L - C_U$ curves. From Section 4, we have all necessary primitives except C_U. From Section 5, we have equilibrium market shares under a variety of policy environments, which we can contrast to the social optimum defined by the primitives. Therefore, the only missing piece for estimating welfare is the social cost of uninsurance. In Section 2 we assumed $C_U = 0$ for simplicity. However, this assumption ignores uncompensated care, care paid for by other state programs, or more difficult-to-measure parameters like a social preference against others being uninsured. Because we do not have any way to directly measure the social cost of uninsurance, we specify it as linked to the observed type-specific cost of enrolling in H. We write the social cost of uninsurance for type s as:

$$C_U(s) = \frac{(1-d)C_H(s)}{1+\phi} + \omega \tag{10}$$

where d is the share of total uninsured healthcare costs that the uninsured pay out of pocket, ϕ is the assumed moral hazard from insurance, and ω is some fixed cost of un-insurance. For d and ϕ, we use the values as derived from Finkelstein, Hendren and Shepard (2017) and assume that $d = 0.2$ and $\phi = 0.25$.[31] We set the fixed cost $\omega = 97$, which is the ω value consistent with 90% of the population

---

[31]We note that without this assumption (i.e. if we assume $C_U = 0$), it is inefficient for any consumer to purchase insurance, as no consumer values either H or L more than the cost of enrolling them in H or L. This fact plus a full discussion of the derivation of the assumed values of d and ϕ can be found in Finkelstein, Hendren and Shepard (2017).

being optimally insured in the cream-skimming case.

Before showing how to use our graphical model to estimate welfare, we provide an important caution: As is standard in the literature, our welfare estimation depends critically on inferring consumer valuation of H and L from estimates of demand-response to exogenous variation in the prices of these products. Our welfare estimates are accurate only to the extent that demand curves accurately reflect true valuations. Behavioral frictions might cause consumer demand to deviate from valuations (Handel, Kolstad and Spinnewijn, 2019). Liquidity constraints could also cause valuation and demand to diverge (Casaburi and Willis, 2018).[32] Although these considerations do not threaten the use of our model for the positive analysis of Section 5 (prediction of prices and market shares), these considerations do suggest caution in interpreting welfare estimates.

We now show how to estimate welfare with our graphical model. Figure 11 plots the empirical analogs to our welfare figures from Section 2. Panel (a) shows the case where L is a pure cream-skimmer, and Panel (b) shows the case where L has a 15% cost advantage. In both cases, instead of plotting C_L, we plot $C_L^{Net} = C_L - C_U$, as in Eq. (8) to account for the fact that $C_U \neq 0$. In both cases 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_{HL}^e and the extensive margin cutoff is s_{LU}^e. Thus, consumers with $s < s_{HL}^e$ enroll in H, consumers with $s_{HL}^e < s < s_{LU}^e$ enroll in L, and consumers with $s > s_{LU}^e$ remain uninsured.

In the case where L is a pure cream-skimmer (Panel (a)), 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 $s > s_{HU}^*$) 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

---

[32]A separate issue is that our specification of C_U is *ad hoc* and may not reflect the actual social costs of uninsurance. This latter issue is amenable to cycling through a range of assumptions on C_U and tracing the impacts.

Figure 11: Empirical Estimates of Foregone Surplus

(a) No cost advantage for L **(b)** 15% Cost Advantage for L
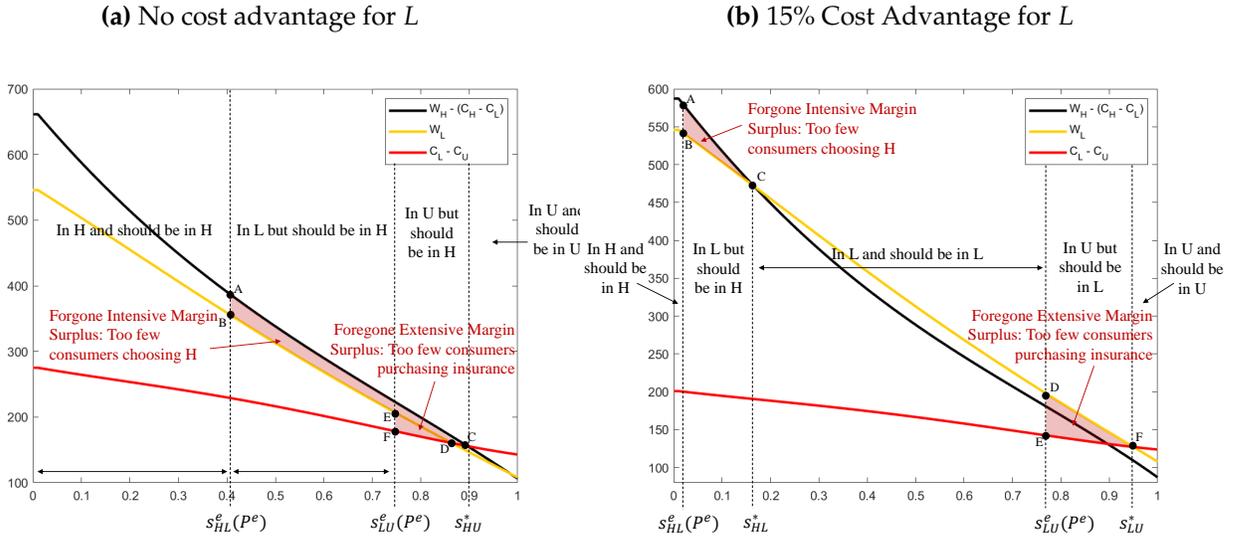
Notes: Panels (a) and (b) show welfare losses under ACA-like subsidies relative to efficient sorting, when L is a cream-skimmer and when L has a 15% cost advantage over H, respectively. In both settings, 60% of the population is low-income and 40% of the population is high-income, so WTP curves are weighted sums of both types. Efficient cutoffs are indicated with a * while equilibrium outcomes are denoted with an e superscript.

area represents a welfare loss of $16.58. The total foregone surplus in the baseline ACA setting in panel a of Figure 11 is $58.50.

In the case where L has a 15% cost advantage in panel (b) of Figure 11, the efficient sorting of consumers across H, L, and U involves some consumers in each of the three options. Consumers with $s < s^*_{HL}$ should be in H, consumers with $s^*_{HL} < s < s^*_{LU}$ should be in L, and the few consumers with $s > s^*_{LU}$ should be uninsured to maximize social surplus. Again, we indicate the foregone surplus via shading. Intensive margin foregone surplus, represented by the welfare triangle ABC is much smaller than in the case where L is a pure cream-skimmer, representing a welfare loss of $19.08. Extensive margin foregone surplus is represented by the welfare triangle DEF. Welfare loss on this margin amounts to $8.30. The total foregone surplus in the baseline ACA setting in panel (b) of Figure 11 where L has a 15% cost advantage is thus $27.38.

6.2   Welfare Consequences of Penalties and Risk Adjustment

We next estimate the welfare consequences of adding mandate penalties and tuning risk adjustment strength, relative to the baseline shown above in Figure 11. Panel (a) of Figure 12 adds a $60 uninsurance penalty to the baseline case in which L is a pure cream-skimmer. Recall from the top-left panel

of Figure A2 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.

Figure 12: Welfare Effects of Uninsurance Penalty and Stronger Risk Adjustment Transfers

(a) Impose \$60 Penalty
(ACA Subsidies, 0% Cost Advantage for L)

(b) Strengthen Risk Adjustment
(Fixed Subsidies, 15% Cost Advantage for L)
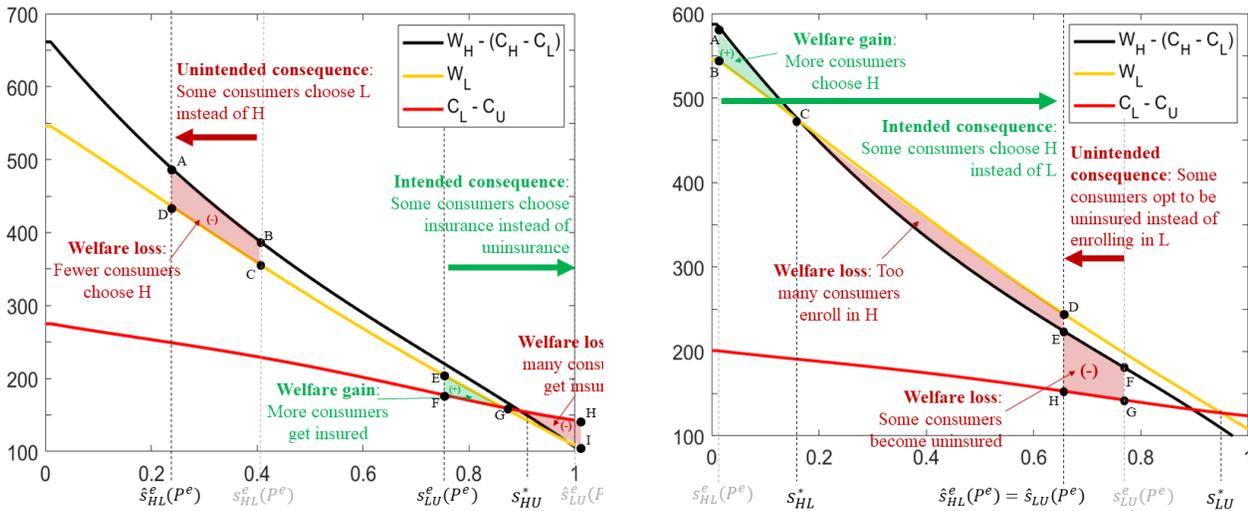
Notes: Panels (a) and (b) show welfare changes under different policies, relative to baseline policies. Panel (a) shows the welfare impact of imposing a \$60 mandate under a price-linked subsidy when L is a cream-skimmer. Panel (b) shows welfare changes from strengthening risk adjustment from $\alpha = 1$ to $\alpha = 2$ when the subsidy is fixed at \$250 and L has a 15% cost advantage over H.

We report welfare impacts of a mandate in other market settings in Appendix G. Those results, which correspond to the cases in Figures A2 and A3, 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.

Panel (b) of Figure 12 shows the welfare consequences of strengthening risk adjustment in a market where, similar to our previous setting, 60% of the population is receiving a subsidy, but unlike the previous-setting, this subsidy is set at a fixed amount unrelated to the costs of either plan. Specifically, we show how welfare changes when going from a setting where consumers receive a fixed subsidy equal to \$250/month and $\alpha = 1$ to a similar setting where $\alpha = 2$, so that risk adjustment transfers are increased to two-times the ACA transfers. We show this for the case where L has a 15% cost advantage. We assume there is no uninsurance penalty in either setting. Recall from the bottom-right panel of Figure A3 that moving from $\alpha = 1$ to $\alpha = 2$ in this setting (1) shifts 65% of consumers in the market from L to H but also (2) shifts 30% of consumers in the market from L to U. Overall, no consumers remain in L when $\alpha = 2$. Effect (1) is the intended consequence of risk adjustment, and here it implies both welfare gains and losses. Welfare gains occur when consumers whose incremental valuation for H vs. L exceeds the incremental cost of H vs. L (i.e. those with $W_H^{Net}(s) > W_L(s)$) enroll in H instead of L. These gains are represented by the green welfare triangle ABC, and they amount to \$19.08. Welfare losses occur when consumers whose incremental valuation for H vs. L is less than the incremental cost of H vs. L (i.e. those with $W_H^{Net}(s) < W_L(s)$) enroll in H instead of L. These offsetting welfare losses occur when "too many" consumers enroll in H, and they are represented by the red welfare triangle CDE and amount to \$19.24. In other settings, where it is always more efficient for consumers to be enrolled in H instead of L (such as the pure cream-skimming case), there will only be welfare gains on this margin. In the case of Figure 12 (b), the two effects nearly cancel each other out so that the net welfare loss due to the intended consequence of shifting consumers from L to H amounts to \$0.15.

Effect (2) is the unintended consequence of risk adjustment, and here it implies welfare losses. Because risk adjustment leads to a higher price of L, some consumers exit the market, increasing the uninsurance rate. In this case, all consumers exiting the market value insurance more than the (net) cost of insuring them, $C_L^{Net} - C_L - C_U$, causing the welfare consequences of this shift of consumers out of the market to be unambiguously negative. The size of the welfare loss is represented by the area of $DFGH$, which we estimate to be \$52.97. Combining the intended and unintended consequences of risk adjustment, we estimate that doubling risk adjustment transfers by shifting from $\alpha = 1$ to $\alpha = 2$

40

would decrease welfare by \$53.12.

Welfare results for all settings studied in Figures A2 and A3 and the full range of levels of α are found in Appendix G. These results indicate that with ACA-like subsidies, increasing the strength of risk adjustment transfers always improves welfare when L is a pure cream-skimmer. In this case, there is no effect of risk adjustment on the extensive margin due to the linkage of the subsidy to the price, leaving only intensive margin consequences. The intensive margin effects of moving consumers from L to H are also unambiguously positive, as it is inefficient for any consumer to be enrolled in L vs. H. When L has a cost advantage, increasing the strength of risk adjustment transfers improves welfare given low initial levels of α but decreases welfare given higher initial levels of α, with the welfare-maximizing risk adjustment policy having an α around 1.5, or 1.5 times the strength of ACA risk adjustment transfers. This non-monotonic result is due to the fact that increases in α from low initial levels of α induce only those consumers who value H highest relative to L to enroll in H, with consumers whose incremental WTP does not exceed incremental cost remaining enrolled in L.

With fixed subsidies, the welfare consequences again depend on whether L has a cost advantage. Recall that when L is a pure cream-skimmer, extensive margin consequences of risk adjustment are limited. It is inefficient for any consumers to be enrolled in L vs. H in the cream-skimmer case, implying that the intensive margin effects of moving consumers from L to H are unambiguously positive. When L has a cost advantage, patterns in the fixed subsidy case are similar to the ACA-like subsidy case, with welfare increasing with the strength of risk adjustment at low initial levels of α and decreasing at higher levels. Here, in addition to moving some consumers who should not be in H into H, stronger risk adjustment also pushes consumers out of the market, further worsening the negative effects of risk adjustment. Overall, risk adjustment is most likely to improve welfare in a setting with ACA-like subsidies and when L plans do not have a cost advantage. However, policymakers should be cautious when strengthening risk adjustment in settings where subsidies are fixed and/or plans are heterogeneous in their cost structures.

6.3 Welfare under Interacting Policies

We have shown that in some cases, policies targeted at the extensive margin have unintended effects on the intensive margin and vice versa. This implies the necessity of a second-best approach to policy:

optimal extensive margin policy (penalties and subsidies) will often depend on the intensive margin policies (risk adjustment and benefit regulation) currently in use in a market.

We now show how our model can be used to assess optimal policy, allowing for the interaction of simultaneous policies targeting selection on each margin. We again consider uninsurance penalties and risk adjustment. We compute social welfare for a grid of uninsurance penalties and levels of α. We do this for the case where L is a pure cream-skimmer and consumers receive a fixed subsidy equal to \$250 when purchasing insurance. We "cherry-pick" this case because the two policies interact in interesting ways. For completeness, we perform similar analyses for all other settings studied in Figures A2 and A3. Results are reported in Appendix G.

Figure 13: Welfare under Interacting Extensive and Intensive Margin Policies
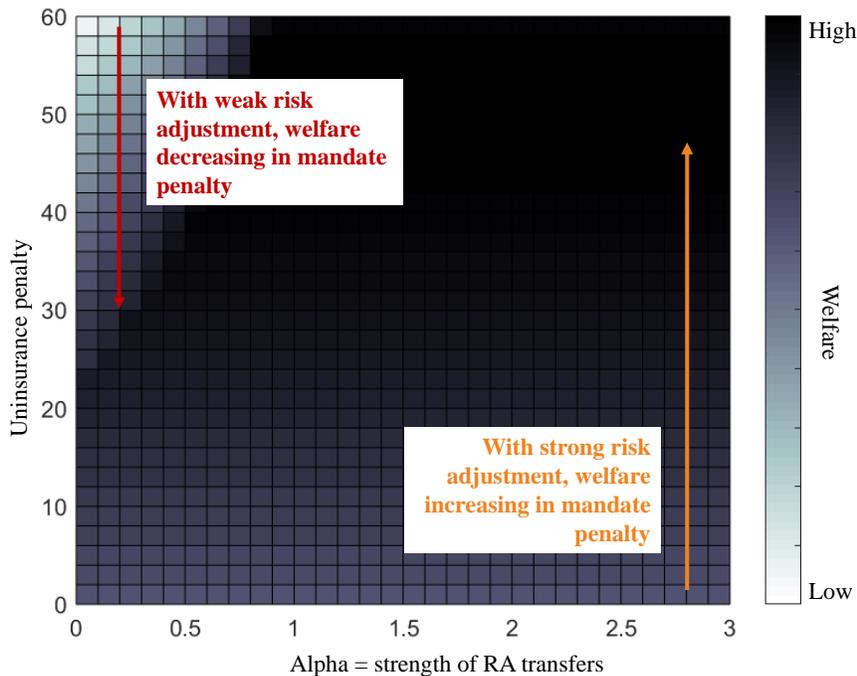
Figure 13 presents the welfare estimates graphically as a heat map, where darker areas represent higher values of social surplus. The figure shows that in this setting, when risk adjustment is strong (α is high), welfare is increasing in the penalty. When risk adjustment is weak, however, welfare is sometimes *decreasing* in the penalty, especially at particularly high penalty levels. Recall that in this setting, the socially efficient allocation of consumers across H, L, and U is to have almost all consumers enrolled in H. Thus, when risk adjustment is strong enough to eliminate L, a policy like

an uninsurance penalty that shifts consumers out of uninsurance and into the market unambiguously improves welfare because risk adjustment prevents the unintended intensive margin consequence of the penalty (shifting consumers from H to L) from occurring. When risk adjustment is weak, on the other hand, the potential harm of the mandate penalty is high, resulting in welfare-decreasing shifts of consumers from H to L that more than offset the welfare gains from decreasing the uninsurance rate.

We can also use Figure 13 to consider the optimal level of α for each level of the uninsurance penalty. In the case we consider, there is no ambiguity: Welfare is always increasing in the strength of the risk adjustment transfer. Here, welfare losses from the unintended extensive margin effects of risk adjustment never offset the welfare gains from the intended intensive margin effects. This is likely due to the fact that, as illustrated in Figure A2, when L is a pure cream-skimmer, the extensive margin consequences of risk adjustment are weak.[33]

Figure 13 can also be used to determine optimal policy in this setting. The figure reveals that welfare is highest when the uninsurance penalty is large and risk adjustment transfers are strong (high α). This is the combination of policies that induces all consumers in the market to enroll in H, which is close to the socially efficient outcome in this particular setting. In Appendix G we show that other settings have different optimal policies. In the case where L is a pure cream-skimmer and subsidies are linked to prices (ACA-like subsidies), optimal policy is to have strong risk adjustment (high α) and a weak mandate. In the case where L has a cost advantage, a weak mandate with moderate to weak risk adjustment is the optimal policy. In all cases, it is clear that these two policies interact with each other, implying that evaluating one policy in isolation from the other can be misleading. Ours is the first paper to show this.

7 Conclusion

Adverse selection in insurance markets can occur on either the extensive (insurance vs. uninsurance) or intensive (more vs. less generous coverage) margin. While this possibility has been recognized for a long time, most prior treatments of adverse selection focus on only one margin or the other. This myopic focus has caused important trade-offs inherent to policies often used to combat selection on

---

[33]Recall that the direct effect of the increased transfers away from L are offset by the substitution of the sickest L enrollees from L to H, resulting in little or no change in P_L (see Appendix B).

one margin or the other to be missed.

In this paper, we developed a new simple theoretical and graphical framework that allows for selection on both margins. We use this framework to build intuition for the unintended intensive margin consequences of extensive margin policies and the unintended extensive margin consequences of intensive margin policies. We show that policies that target selection on one margin will often exacerbate selection on the other. The extent to which this occurs depends on the primitives of the market. We build intuition for this trade-off with a simple graphical framework that generalizes the framework of Einav, Finkelstein and Cullen (2010) by adding the option to remain uninsured. We see this generalized graphical framework as a key contribution of the paper.

We also show that it is straightforward to take the graphical framework to data: With only demand and cost curves from the H and L plans, equilibrium prices and market shares can be found, *even in the setting where uninsurance is available as an option to consumers*. We do this with data from the Massachusetts Connector and show that the extensive/intensive margin trade-off is empirically relevant for evaluating the consequences of various policies. Specifically, we show that: (1) strengthening uninsurance penalties can help some consumers by getting them into the market while hurting other consumers by inducing them to enroll in lower-quality coverage and (2) strengthening risk adjustment transfers can help some consumers by inducing them to enroll in higher-quality coverage while hurting other consumers by forcing them out of the market. Additionally, we show that price-linked subsidies can weaken some of these trade-offs (i.e. effects of risk adjustment and benefit regulation) but not others (i.e. mandates/uninsurance penalties). Finally, we show that trade-offs related to risk adjustment are often more pronounced when L has a cost advantage.

We also show how our graphical model, like the model of Einav, Finkelstein and Cullen (2010), can be used to estimate the welfare consequences of policies. Because many policies lead to coverage gains on one margin and coverage losses on the other, such welfare analysis is critical for assessing the normative consequences of policies. We show that in some cases the unintended effects of policies are first order with respect to welfare, with the welfare losses from coverage losses on the unintended margin exceeding welfare gains from coverage gains on the intended margin. This happens most often with a penalty for choosing to be uninsured.

The simplicity of our approach is not without its costs. Specifically, our assumption of a vertical model of insurance demand is restrictive. Many of our insights apply to more general settings,

though in less-transparent ways. However, some of our insights may differ in more complex markets, and these complexities are an important area for future research.

These issues are highly relevant for future reform of the individual health insurance market in the U.S. In this market, many have observed that the overall quality of coverage available to consumers is low, with most plans characterized by tight provider networks, high deductibles, and strict controls on utilization. Additionally, others have observed that take-up is far from complete, with many young, healthy consumers opting out of the market altogether and choosing to remain uninsured (Domurat, Menashe and Yin, 2018). These two observations are consistent with adverse selection on the intensive and extensive margins, respectively. Our framework highlights the unfortunate but important conceptual point that budget-neutral policies that target one of these two problems are likely to exacerbate the other due to the inherent trade-off between extensive and intensive margin selection. This point is often absent from discussions of potential reforms by policymakers and economists, and our intention is to correct this potentially costly omission.

There are ways to address selection on both the intensive and extensive margins simultaneously, however. They just require additional resources to be injected into the market. For example, intensive margin selection problems can be addressed without exacerbating extensive margin selection via an incremental subsidy to H plans (or a larger penalty for uninsurance). In this case, the key trade-off is the welfare gain of higher quality coverage vs. the welfare cost of raising the funds to pay for the incremental subsidy. Additionally, any policy that severs the link between selection and prices on one of the two margins (for example, a strong mandate that induces complete take-up in *all* states of the world or price-linked subsidies available to *all* consumers) frees up policymakers to be aggressive as they feel necessary on the other margin without any unintended consequences. Though, again, such policies come with their own trade-offs.

In summary, common policies targeting the problems caused by adverse selection do not provide a "free lunch". Instead, they involve complex trade-offs. In this paper, we make an important step toward understanding one of the most important of these trade-offs.

References

Adams, William, Liran Einav, and Jonathan Levin. 2009. "Liquidity Constrains and Imperfect Information in Subprime Lending." *The American economic review*, 99(1): 49–84.

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

Brown, Jason, Mark Duggan, Ilyana Kuziemko, and William Woolston. 2014. "How does risk selection respond to risk adjustment? New evidence from the Medicare Advantage Program." *American Economic Review*, 104(10): 3335–64.

Brown, Jeffrey, Amy Finkelstein, and Norma Coe. 2007. "Medicaid Crowd-Out of Private Long Term Care Insurance Demand: Evidence from the Health and Retirement Survey." *Tax Policy and the Economy*, 21: 1–34.

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

Carey, Colleen. 2017*a*. "Technological Change and Risk Adjustment: Benefit Design Incentives in Medicare Part D." *American Economic Journal: Economic Policy*, 9(1): 38–73.

Carey, Colleen. 2017*b*. "A Time to Harvest: Evidence on Consumer Choice Frictions from a Payment Revision in Medicare Part D." Working Paper.

Casaburi, Lorenzo, and Jack Willis. 2018. "Time versus State in Insurance: Experimental Evidence from Contract Farming in Kenya." *American Economic Review*, 108(12): 3778–3813.

Domurat, Richard. 2018. "How Do Supply-Side Regulations in the ACA Impact Market Outcomes? Evidence from California." UCLA Working Paper.

Domurat, Richard, Isaac Menashe, and Wesley Yin. 2018. "Frictions in Health Insurance Take-up Decisions: Evidence from a Covered California Open Enrollment Field Experiment." UCLA Working Paper.

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

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

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

Einav, Liran, Mark Jenkins, and Jonathan Levin. 2012. "Contract Pricing in Consumer Credit Markets." *Econometrica*, 80(4): 1387–1432.

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

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

Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard. 2017. "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts." National Bureau of Economic Research Working Paper 23668.

Geruso, Michael. 2017. "Demand heterogeneity in insurance markets: Implications for equity and efficiency." *Quantitative Economics*, 8(3): 929–975.

Geruso, Michael, and Thomas G. McGuire. 2016. "Tradeoffs in the Design of Health Plan Payment Systems: Fit, Power and Balance." *Journal of Health Economics*, 47(1): 1–19.

Geruso, Michael, and Timothy Layton. 2015. "Upcoding: Evidence from Medicare on Squishy Risk Adjustment." National Bureau of Economic Research Working Paper 21222.

Geruso, Michael, and Timothy Layton. 2018. "Upcoding: Evidence from Medicare on Squishy Risk Adjustment." *"Journal of Political Economy"*, forthcoming.

Geruso, Michael, Timothy J Layton, and Daniel Prinz. 2018. "Screening in contract design: evidence from the ACA health insurance exchanges." forthcoming, American Economic Journal: Applied Economics.

Glazer, Jacob, and Thomas G. McGuire. 2000. "Optimal Risk Adjustment in Markets with Adverse Selection: An Application to Managed Care." *American Economic Review*, 90(4): 1055–1071.

Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski. 2015. "Adverse Selection and an Individual Mandate: When Theory Meets Practice." *American Economic Review*, 105(3): 1030–1066.

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

Handel, Benjamin R, and Jonathan T Kolstad. 2015. "Health insurance for" humans": Information frictions, plan choice, and consumer welfare." *American Economic Review*, 105(8): 2449–2500.

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

Handel, Benjamin R., Jonathan Kolstad, and Johannes Spinnewijn. 2019. "Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets." forthcoming, Review of Economics and Statistics.

Harrington, Scott E. 2017. "Stabilizing Individual Health Insurance Markets With Subsidized Reinsurance." *Penn LDI Issue Brief*, 21(7).

Ketcham, Jonathan D, Nicolai V Kuminoff, and Christopher A Powers. 2016. "Estimating the Heterogeneous Welfare Effects of Choice Architecture: An Application to the Medicare Prescription Drug Insurance Market." National Bureau of Economic Research Working Paper 22732.

Kling, Jeffrey R, Sendhil Mullainathan, Eldar Shafir, Lee C Vermeulen, and Marian V Wrobel. 2012. "Comparison friction: Experimental evidence from Medicare drug plans." *The Quarterly Journal of Economics*, 127(1): 199–235.

Lavetti, Kurt, and Kosali Simon. 2016. "Strategic Formulary Design in Medicare Part D Plans." National Bureau of Economic Research Working Paper 22338.

Layton, Timothy J. 2017. "Imperfect risk adjustment, risk preferences, and sorting in competitive health insurance markets." *Journal of Health Economics*, 56: 259–280.

Layton, Timothy J., Randall P. Ellis, Thomas G. McGuire, and Richard Van Kleef. 2017. "Measuring Efficiency of Health Plan Payment Systems in Managed Competition Health Insurance Markets." *Journal of Health Economics*, Forthcoming.

Layton, Timothy J, Thomas G McGuire, and Anna D Sinaiko. 2016. "Risk corridors and reinsurance in health insurance marketplaces: insurance for insurers." *American journal of health economics*, 2(1): 66–95.

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

Newhouse, Joseph P. 2017. "Risk adjustment with an outside option." *Journal of health economics*, 56: 256–258.

Newhouse, Joseph P, Mary Price, John Hsu, J Michael McWilliams, and Thomas G McGuire. 2015. "How much favorable selection is left in Medicare Advantage?" *American journal of health economics*, 1(1): 1–26.

Polyakova, Maria. 2016. "Regulation of insurance with adverse selection and switching costs: Evidence from Medicare Part D." *American Economic Journal: Applied Economics*, 8(3): 165–95.

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

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

Shepard, Mark. 2016. "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange." National Bureau of Economic Research Working Paper 22600.

Tebaldi, Pietro. 2017. "Estimating equilibrium in health insurance exchanges: Price competition and subsidy design under the aca."

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

Table 1: Cream-skimming L plan, some subsidized

(a) ACA-like subsidy

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

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

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

(d) Fixed $275

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

Table 2: 15 % Cost Advantage L plan, some subsidized

(a) ACA-like subsidy

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

(b) Fixed $275

mandate	0	15	30	45	60
price H	394	.	.	.	.
price L	278	273	273	273	273
share H	.0065	0	0	0	0
share L	.96	1	1	1	1
share U	.035	0	0	0	0
subsidy	275	275	275	275	275
welfare	.14	0	0	0	0

(c) ACA-like subsidy

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 $275

α	0	.5	1	1.5	2
price H	.	.	394	345	377
price L	280	280	278	299	.
share H	0	0	.0065	.17	.66
share L	.95	.95	.96	.69	0
share U	.046	.046	.035	.14	.34
subsidy	275	275	275	275	275
welfare	.17	.17	.14	.55	1

<div align="center">

Online Appendix for:
The Two Margin Problem in Insurance Markets

</div>

A   Appendix Section 1

A.1   Riley Equilibrium

We follow Handel, Hendel and Whinston (2015) and consider equilibria that meet the requirements of the Riley Equilibrium (RE) notion. In words, a price vector P 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.[34] 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 details). 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, and we denote the set of breakeven price vectors, $\mathcal{P}^{BE} = \{P : P_H = AC_H, P_L = AC_L\}$. This set consists of the following potential breakeven vectors:

1. **No 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 in $\mathcal{P}_4^{BE}$, or the set of breakeven vectors where there is positive enrollment in both H and L. However, we note note that under certain conditions there will not be an RE in $\mathcal{P}_4^{BE}$ and the competitive equilibrium will instead consist of positive enrollment in only one or neither one of the two plan options. We allow for these possibilities in the empirical portion of the paper.[35] Given the assumption that in equilibrium there is positive enrollment in H and L, we have the familiar equilibrium condition that prices are set equal to average costs:

$$P_H = AC_H \left(P^{cons} \right)$$
$$P_L = AC_L \left(P^{cons} \right) \tag{11}$$

We use this expression to define equilibrium throughout Section 2.

---

[34]Formally, a "Riley Deviation" (i.e. a deviation that would cause a price vector to not be a Riley Equilibrium) is a price offer P' that is strictly profitable when $P \cup P'$ is offered and for which there is no "safe" (i.e. weakly profitable) reaction P'' that makes the firm offering P' incur losses when $P \cup P' \cup P''$ is offered.

[35]Handel, Hendel and Whinston, 2015 show that there is a unique RE in the setting where there is no outside option. With an outside option, their definition of a Riley Equilibrium requires a slight modification in order to achieve uniqueness. Specifically, instead of requiring the deviation to be strictly profitable, we require the deviation to be weakly profitable but also to achieve positive enrollment for the deviating plan. In the empirical exercise below, we use this definition to find the competitive equilibrium for each counterfactual simulation.

B Appendix Section 2: Comparative Statics for Effects of Policies on Prices and Enrollment

In this appendix, we derive comparative statics describing the effects of increasing the size of the uninsurance penalty and increasing the strength of risk adjustment transfers. The setup is identical to what we introduced in Section 2, with $P = P_H, P_L$ describing insurer prices and $G = S_H, S_L, M$ describing the vector of 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 $P_U^{cons} = M$. Demand follows the vertical model and is defined, along with the cutoff s-types s_{HL} and s_{LU}, as in Section 2. Note that demand depends only on the relative consumer price of H vs. L ($\Delta P_{HL}^{cons} = P_H - P_L$) and on the relative price of L vs. U ($\Delta P_{LU}^{cons} = P_L - S - M$).

We also have average cost functions:

$$AC_H(P; G) = \frac{1}{D_H(P; G)} \int_0^{s_{HL}(\Delta P_{HL}^{cons})} C_H(s) ds$$

$$AC_L(P; G) = \frac{1}{D_L(P; G)} \int_{s_{HL}(\Delta P_{HL}^{cons})}^{s_{LU}(\Delta P_{LU}^{cons})} C_H(s) ds \tag{12}$$

Similarly, we can define average risk score functions:

$$\overline{R}_H(P; G) = \frac{1}{D_H(P; G)} \int_0^{s_{HL}(\Delta P_{HL}^{cons})} R(s) ds$$

$$\overline{R}_L(P; G) = \frac{1}{D_L(P; G)} \int_{s_{HL}(\Delta P_{HL}^{cons})}^{s_{LU}(\Delta P_{LU}^{cons})} R(s) ds \tag{13}$$

where $R(s)$ is the average risk score among type-s consumers. The baseline 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 ($\equiv \overline{P}(P; G)$):

$$T(P; G) = \left(\frac{\overline{R}_H(P; G) - \overline{R}_L(P; G)}{\overline{R}(P; G)} \right) \overline{P}(P; G) \tag{14}$$

Finally, we introduce a parameter $\alpha \in (0, 1)$ that multiplies the transfer, $\alpha T(P; G)$, 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. This mimics a policy option recently given to states.

As in Section 2 we define equilibrium prices as the set of prices that result in firms earning zero profits. Here, however, the equilibrium condition is not that prices equal average cost, but instead that prices equal average costs net of risk adjustment transfers:

$$P_H = AC_H(P; G) - \alpha T(P; G) \equiv AC_H^{RA}(P; G, \alpha)$$

$$P_L = AC_L(P; G) - \alpha T(P; G) \equiv AC_L^{RA}(P; G, \alpha) \tag{15}$$

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

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.

B.1   Increase in Uninsurance Penalty

In Section 3 we use the graphical model to show that, given the primitives in the figure, an increase in the uninsurance penalty results in some consumers moving from U to L and some consumers moving from H to L. We now derive the more general effects of an increase in the uninsurance penalty mathematically.

First, note that $D_H(P; G) = s_{HL}(P_H - P_L)$, so, given the assumptions of the vertical model, there is no direct effect of a slightly larger M on D_H. This is a consequence of the vertical model, where all marginal uninsured consumers are on the margin between L and U, not H. We explore the consequences of relaxing this assumption in Appendix XX. Given this assumption, the penalty affects D_H only to the extent that it affects P_H and P_L. Mathematically:

$$\frac{dD_H}{dM} = \underbrace{\left(-\frac{\partial s_{HL}}{\partial \Delta P_{HL}}\right)}_{(+)} \cdot \left(\frac{dP_L}{dM} - \frac{dP_H}{dM}\right) \tag{16}$$

where the first term is positive by the law of demand, so the sign of the effect depends on the sign of $\frac{dP_H}{dM}$. Differentiating P_H and P_L with respect to M gives

$$\frac{dP_H}{dM} = \frac{\partial AC_H^{RA}}{\partial M} + \frac{\partial AC_H^{RA}}{\partial P_H}\frac{dP_H}{dM} + \frac{\partial AC_H^{RA}}{\partial P_L}\frac{dP_L}{dM}$$

$$\frac{dP_L}{dM} = \frac{\partial AC_L^{RA}}{\partial M} + \frac{\partial AC_L^{RA}}{\partial P_H}\frac{dP_H}{dM} + \frac{\partial AC_L^{RA}}{\partial P_L}\frac{dP_L}{dM} \tag{17}$$

We note that under the vertical model $\frac{\partial AC_H^{RA}}{\partial M} = -\alpha\frac{\partial T}{\partial M}$, since there is no direct effect of M on enrollment in H and $\frac{\partial AC_H}{\partial M} = 0$. Solving this linear system for $\frac{dP_H}{dM}$ gives

$$\frac{dP_H}{dM} = \left[\underbrace{-\alpha\frac{\partial T}{\partial M}}_{\substack{\text{Ext. Margin Selection}(-)}} + \underbrace{\underbrace{\frac{\partial AC_L^{RA}}{\partial M}}_{(-)} \cdot \underbrace{\frac{\partial AC_H^{RA}}{\partial P_L}}_{(-)} \underbrace{\left(1 - \frac{\partial AC_L^{RA}}{\partial P_L}\right)^{-1}}_{(+)}}_{\text{Substitution to } L(+)} \right] \times \Phi_H^{-1} \tag{18}$$

where $\Phi_H \equiv 1 - \frac{\partial AC_H^{RA}}{\partial P_H} - \frac{\partial AC_H^{RA}}{\partial P_L}\frac{\partial AC_L^{RA}}{\partial P_H} \cdot \left(1 - \frac{\partial AC_L^{RA}}{\partial P_L}\right)^{-1} > 0$ is a term that must be positive for any stable equilibrium.[36] The effect of the penalty on P_H is thus made up of two components. The first component, $-\alpha\frac{\partial T}{\partial M}$ captures the effect of extensive margin selection on the price of H. Risk adjustment transfers cause some of the change in risk in L to be passed on to H. Specifically, a larger penalty brings healthier consumers into L, which increases $T(.)$ – implying that $\partial T/\partial M > 0$ and the first term in brackets is negative. The second term captures the indirect effect of the penalty on H's risk pool via the effect of the penalty on P_L leading to substitution between H and L. If there is extensive margin selection even after risk adjustment (due to imperfect risk adjustment), then $\frac{\partial AC_L^{RA}}{\partial M} < 0$ – i.e., costs in L will fall. Further, adverse selection implies that $\frac{\partial AC_H^{RA}}{\partial P_L} < 0$ and stability requires that

---

[36]Specifically, stability implies $1 - \frac{\partial AC_j^{RA}}{\partial P_j} > 0$, which, given that $\frac{AC_H^{RA}}{\partial P_L}$ and $\frac{AC_L^{RA}}{P_H}$ must have opposite signs, implies that $\Phi_H > 0$.

$1 - \frac{\partial AC_L^{RA}}{\partial P_L} > 0$. Thus, the second term must be positive. The intuition here is that the penalty induces healthier consumers to enroll in L, lowering the price of L, and at the new lower price of L, the healthiest H enrollees opt instead to enroll in L, driving up the average cost in H and thus the price.

We can also do the same for P_L, giving:

$$
\frac{dP_L}{dM} = \left[\underbrace{\frac{\partial AC_L^{RA}}{\partial M}}_{\text{Ext. Margin Selection}(-)} + \underbrace{\underbrace{\left(\alpha \frac{\partial T}{\partial M} \right)}_{(+)} \cdot \underbrace{\frac{\partial AC_L^{RA}}{dP_H}}_{(+)} \underbrace{\left(1 - \frac{\partial AC_H^{RA}}{dP_H} \right)^{-1}}_{(+)}}_{\text{Substitution to } L(-)} \right] \times \Phi_L^{-1} \tag{19}
$$

where $\Phi_L \equiv 1 - \frac{\partial AC_L^{RA}}{\partial P_L} - \frac{\partial AC_L^{RA}}{\partial P_H} \frac{\partial AC_H^{RA}}{\partial P_L} \cdot \left(1 - \frac{\partial AC_H^{RA}}{\partial P_H} \right)^{-1} > 0$ is a term that must be positive for any stable equilibrium.[37] Here, the direct effect of a larger penalty is captured by the first term in the brackets. If there is extensive margin adverse selection after risk adjustment, the marginal enrollees in L will have lower risk-adjusted costs, pushing down the price of L. The substitution effect is positive. This is because here the substitution effect captures changes to risk adjustment transfers: As L's risk pool gets healthier, it pays larger transfers to H, driving up the price of L. Thus, these two effects compete with each other to determine the overall effect of the mandate on P_L.

These comparative statics get notably simpler when there is no risk adjustment in the market. When this is the case $\alpha \frac{\partial T}{\partial M} = 0$ eliminating the extensive margin effect from $\frac{dP_H}{dM}$ and the substitution effect from $\frac{dP_L}{dM}$. Without risk adjustment, we thus have the following:

$$
\frac{dP_H}{dM} = \left[\underbrace{\frac{\partial AC_L^{RA}}{\partial M}}_{(-)} \cdot \underbrace{\frac{\partial AC_H^{RA}}{dP_L}}_{(-)} \underbrace{\left(1 - \frac{\partial AC_L^{RA}}{dP_L} \right)^{-1}}_{(+)} \right] \times \Phi_H^{-1} \tag{20}
$$
$$
\underbrace{\hspace{5cm}}_{\text{Substitution to } L(+)}
$$

$$
\frac{dP_L}{dM} = \left[\underbrace{\frac{\partial AC_L^{RA}}{\partial M}}_{\text{Ext. Margin Selection}(-)} \right] \times \Phi_L^{-1} \tag{21}
$$

Here, the effect of a larger penalty on P_H and the effect of a larger penalty on P_L are both unambiguous. For P_H, a larger penalty will always increase the price due to the substitution of the relatively healthy consumers on the margin between H and L leaving H. For P_L, a larger penalty will always decrease the price due to the enrollment of the relatively healthy consumers on the L vs. U margin.

---

[37]Specifically, stability implies $1 - \frac{\partial AC_j^{RA}}{\partial P_j} > 0$, which, given that $\frac{AC_H^{RA}}{\partial P_L}$ and $\frac{AC_L^{RA}}{P_H}$ must have opposite signs, implies that $\Phi_L > 0$.

B.2 Increase in the Strength of Risk Adjustment (i.e. α)

In Section 3 we use the graphical model to show that, given the primitives in the figure, perfect risk adjustment results in some consumers choosing H instead of L and other consumers choosing U instead of L. We now consider the effects of increasing the strength of imperfect risk adjustment transfers.

We consider the effects of a small increase in α the parameter that determines the size of the risk adjustment transfers. First, again note that stronger risk adjustment affects the share uninsured (D_U) only via its effect on the relative price of L:

$$\frac{\partial D_U}{\partial \alpha} = \underbrace{\left(-\frac{\partial s_{LU}}{\partial \Delta P_{LU}^{cons}} \right)}_{(+)} \cdot \frac{d\Delta P_{LU}^{cons}}{d\alpha} \tag{22}$$

Again, the first term is positive by the law of demand. Thus, the sign of the effect depends on the sign of $\frac{d\Delta P_L U}{d\alpha}$. This, in turn, depends on the nature of subsidies. With price-linked subsidies, $\Delta P_{LU}^{cons} = P_L - S - M$ is fixed by construction. Therefore, $\frac{dD_U}{d\alpha} = 0$.

The more interesting case is the case with fixed subsidies. In this case, $\frac{d\Delta P_L U}{d\alpha} = \frac{dP_L}{d\alpha}$. Differentiating Equation 13 with respect to P_H gives

$$\frac{dP_H}{d\alpha} = T(.) \times \left[\underbrace{-1}_{\text{Direct (-)}} + \underbrace{\frac{\partial AC_H^{RA}}{\partial P_L} \left(1 - \frac{\partial AC_L^{RA}}{\partial P_L} \right)^{-1}}_{\text{Substitution to } H \text{ (-)}} \right] \times \Phi_H^{-1} < 0 \tag{23}$$

where Φ_H is a term that must be positive under any stable equilibrium and is defined the same as in the case of the uninsurance penalty. 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, we know unambiguously that $\frac{dP_H}{d\alpha} < 0$ because both the direct and indirect effects push P_H down.

Doing the same for P_L gives

$$\frac{dP_H}{d\alpha} = T(.) \times \left[\underbrace{1}_{\text{Direct (+)}} + \underbrace{\frac{\partial AC_L^{RA}}{\partial P_H} \left(1 - \frac{\partial AC_H^{RA}}{\partial P_H} \right)^{-1}}_{\text{Substitution to } H \text{ (-)}} \right] \times \Phi_L^{-1} < 0 \tag{24}$$

where Φ_L is a term that must be positive under any stable equilibrium and is defined the same as in the case of the uninsurance penalty. 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 and $1 - \frac{\partial AC_H^{RA}}{\partial P_H} > 0$ by stability. Intuitively, stronger risk adjustment transfers increase the price of L, causing consumers on the H vs. 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. It is thus 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, thent he 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 adverse selection is very strong 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 case where L has no cost advantage may differ from the case where L has a cost advantage. In the no cost advantage case, the only reason L gets any demand is intensive margin adverse selection. When this adverse selection is weak enough, everyone who buys insurance purchases H. Here, when adverse selection is strong, L exists but the substitution effect is also large, muting the direct effect of risk adjustment. But when adverse selection is weak, L fails to exist, which we clearly see happen as α increases. Thus, it is more likely that increasing α will have little or no effect on P_L in the case where L has no cost advantage than in the case where L has a cost advantage.

C   Appendix Section 3: 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 (2017), 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   Low-Income Demand and Costs: FHS (2018)

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, 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.[38] 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 ideas.

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 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 $W_{HL}^{lin}(s) + 3 \times \frac{(0.31-s)}{0.31} \cdot W_{HL}^{lin}(0)$ 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,

$$W_{HL}^{enh}(s) = \begin{cases} W_{HL}^{lin}(s) & \text{for } s \in [0.31, 1] \\ W_{HL}^{lin}(s) + 3 \times \frac{(0.31-s)}{0.31} \cdot W_{HL}^{lin}(0) & \text{for } s \in [0, 0.31) \end{cases} \tag{25}$$

  and

$$W_H^{enh}(s) = W_L^{lin}(s) + W_{HL}^{enh}(s) \tag{26}$$

  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

---

[38]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.

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, for our baseline setting, we assume that L has a 15% cost advantage so that $C_L(s) = 0.85 C_H(s)$. We also a consider an alternative 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.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). Their demand curve takes the following form:

$$W_{HKK}(s) = -\$9,276.81 * s + \$12,498.68 \tag{27}$$

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 $W_H^{HIenh}(s)$: HKK only estimate demand for a single L plan. Similar to FHS, we estimate start by estimating a linearly extrapolated WTP for H, $W_H^{HIlin}(s)$, and then "enhance" demand for H among the highest WTP types, $W_H^{HIenh}(s)$, 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: low-income + enhanced demand for H and high-income + enhanced demand for H. 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.

D Appendix Section 4: Description of Reaction Function Approach to Finding Equilibrium

Evaluating demand, profits: For each un-insurance penalty, risk adjustment strength, L-plan cost advantage, and subsidy type setting, we find the equilibrium P_L and P_H pair using the following grid-search method. We construct a grid of P_L, P_H 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. 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 other, per the ACA risk adjustment formula (see equation 9). The transfer is multiplied by some α for later optimal policy simulations. 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 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 exists and is making zero profits (2) only L exists and is making zero profits (3) both H and L exist and are both making zero profits. Given the coarseness of the grid, there are usually multiple potential equibria of each type. We use the following process to refine this set down to the final equilibrium point.

- **Single plan equilibria:** First, we refine our $L-$only and $H-$only equilibria. For the remainder of this paragraph, I will refer to the potential $L-$only equilibria, but the methodology also applies to potential $H-$only equilibria. Given the curved nature of the primitives, for some settings, especially those where L has a large cost advantage, there are multiple $L-$only prices P_L^{L-only} that are potential $L-$only equilibria. For each of these P_L^{only}, we evaluate a single point (P_L^{L-only}, P_H). For a given $L-$only equilibrium price P_L^{L-only}, there are typically a set of $P_H > P_H^{min}$ that satisfy the conditions (1) L makes zero profit and (2) H has zero enrollment. To cut down on the number of potential equilibria we must evaluate, for each P_L^{L-only}, we evaluate only the pair $(P_H^{min}, P_L^{L-only})$. For each potential P_L^{L-only}, we need only to evaluate this minimum price since any potential H deviations from (P_H^{min}, P_L) would also be deviations from (P_H, P_L), $P_H > P^{min}$.

  Once a set of potential $L-$only equibria prices have been refined to unique (P_L^{only}, P_H) pairs, we then evaluate each P_L^{L-only} to determine if it is a Riley Equilibrium. We begin with the minimum P_L^{L-only}. The Riley Equilibrium code involves three nested loop. First, the outer loop evaluates each each grid point $\Pi^H(P_L^{L-only}, P_H')$, $P_H' < P_H$ to identify potential H-deviations where $\Pi_H(P_L^{L-only}, P_H') > T$. If no such potential H-deviations are found, (P_L^{L-only}, P_H) is considered a RE. If a potential H-deviation is found, the second loop is called. This loop evaluates each grid point (P_L', P_H'), $P_L' < P_L$ to identify potential L-retaliations where $\Pi^L(P_L', P_H') > -T, \Pi^H(P_L', P_H') < -T$. If no such potential retaliations are found for a given potential $H-$deviation, then (P_L^{only}, P_H) is not a riley equilibrium. If a potential retaliation is found, a third loop is activated to evaluate if there is any point (P_L', P_H''), $P_H'' < P_H'$ that makes a given retaliation "unsafe" where unsafe is defined as $\Pi_L < -T$. If no such "unsafe" point exists, then the retaliation point is safe and the potential deviation would not succeed. The next potential deviation point for this P_L^{L-only} is then evaluated. If no retaliation-proof deviation exists for a given (P_L^{only}, P_H), then the point is a RE. If a deviation does exist, the next larger P_L^{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 of any $H-$deviations or L-deviations exist, using the same set of RE loops described in the previous paragraph. If either a H deviation or an L deviation is found, the HL equilibrium is not a riley equilibrium.

E   Appendix Section 5: Benefit Regulation Results

Tables A1 and A2 characterizes 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. Consistent with our other

analyses, in all settings 60% of the population is subsidized and has primitive curves estimated from Finkelstein, Hendren and Shepard (2017) while the other 40% is unsubsidized with primitives derived from Hackmann, Kolstad and Kowalski (2015). 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 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 A1: Benefit Regulation : L-plan Cream Skimmer

	ACA-like		Avg. Costs		$ 300		$ 275		$ 250	
	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

Table A2: Benefit Regulation : L-plan 15% cost advantage

|	ACA-like		Avg. Costs		$ 300		$ 275		$ 250	
	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

F   Appendix Section 6: Estimation of Risk Score Curve

Like WTP and costs, we use FHS's regression discontinuity approach to estimate a continuum of risk adjustment values $r(s)$ which is indexed by the share of the population s and characterizes the expected cost of a given enrollee relative to the population average. All risk-adjustment scores are estimated using the Hierarchical Condition Categories (HCC), a risk-adjustment model used by the Centers for Medicare and Medicaid Services[39] Using the HCC condition categories, we estimated a

---

[39]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

risk score based on the claims of each individual. We then leveraged the same subsidy thresholds used in Finkelstein, Hendren and Shepard (2017) to estimate changes in average risk score across the discontinuities. We then connected these regression discontinuity coefficients in a similar fashion to our construction of the cost and WTP curves to estimate a full curve of risk scores for each share of the population enrolled.

**Figure A1:** Raw costs versus risk adjustment scores
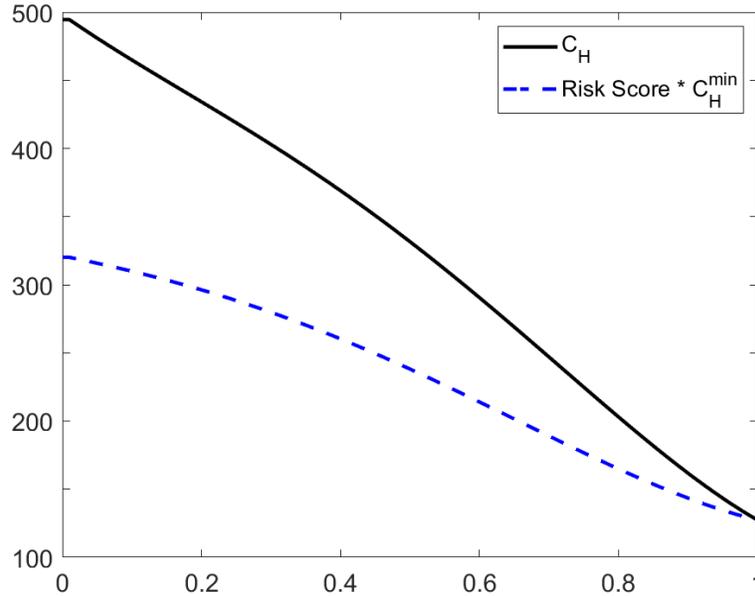
G   Appendix Section 7: Additional Welfare Results

Welfare figures below correspond to the market shares in figures A2 and A3. For a given parameter setting i, we report here welfare normalized as follows: $W_i = \frac{welfare - min(welfare)}{max(welfare) - min(welfare)}$

Figure A2: Cream-skimming L Plan Welfare

(a) ACA-like subsidy, increase mandate

(b) Fixed $ 275 subsidy, increase mandate

(c) ACA-like subsidy, increase α

(d) Fixed $ 275 subsidy, increase α
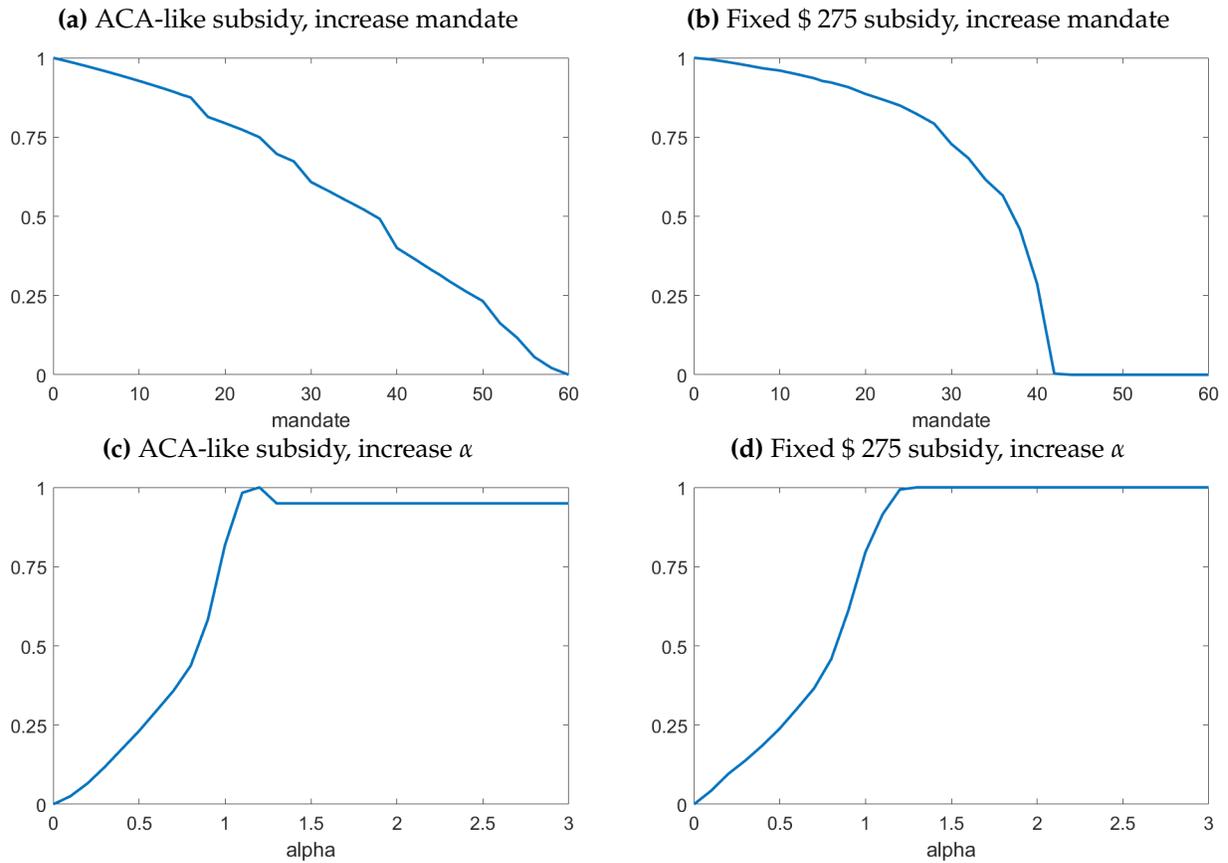
Figure A3: 15% L-plan Cost Advantage Welfare

(a) ACA-like subsidy, increase mandate

(b) Fixed $ 250 subsidy, increase mandate

(c) ACA-like subsidy, increase α

(d) Fixed $ 250 subsidy, increase α
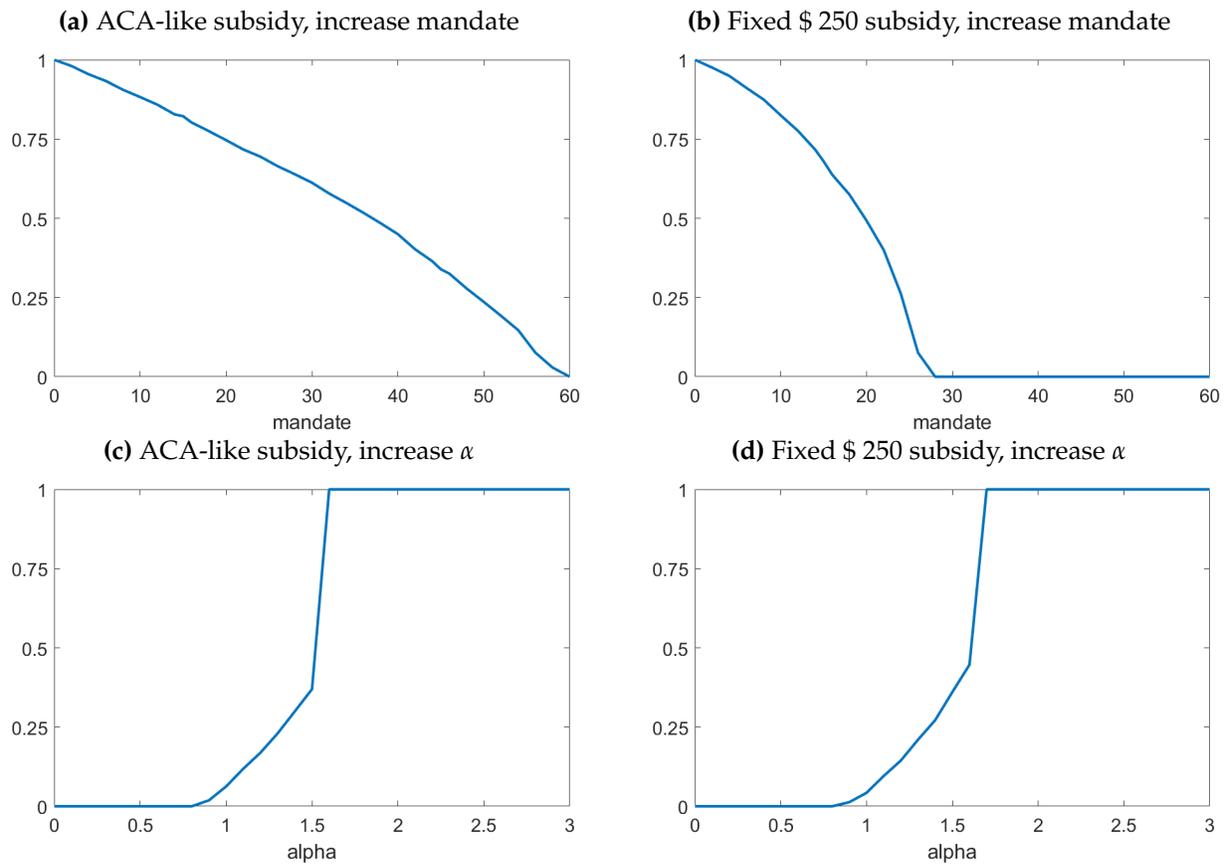
Figure A4: WTP Curves for H and L

(a) Low-Income
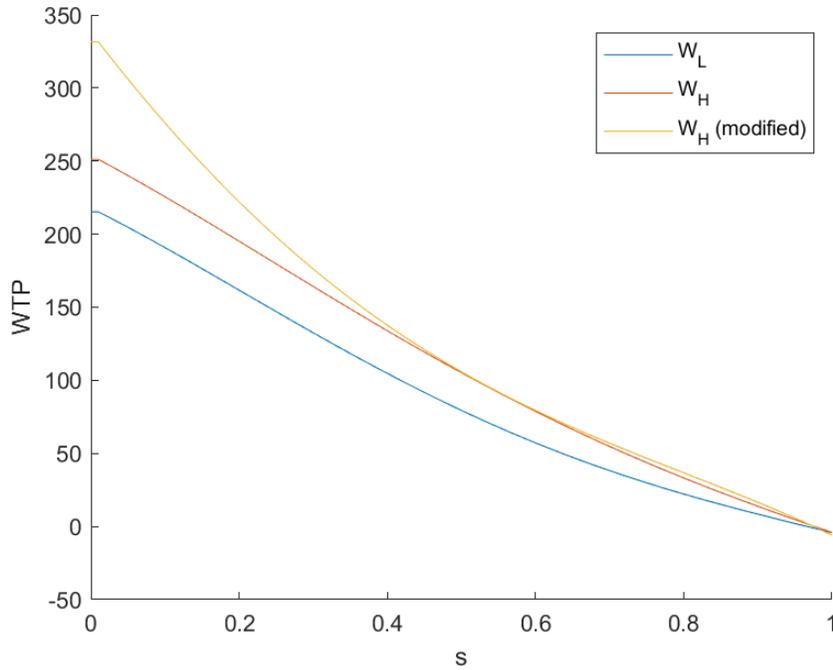
(b) High-Income
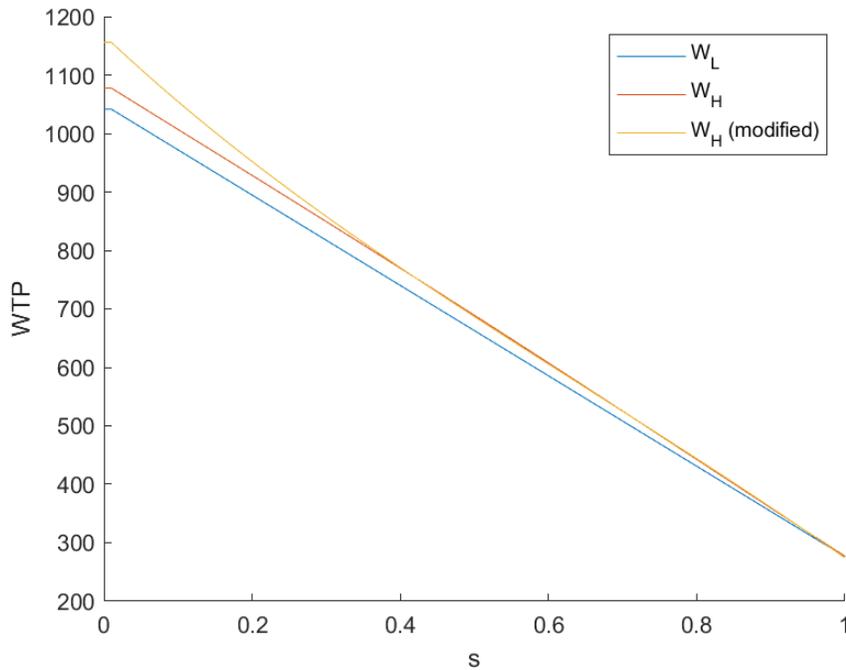
Note: Figure shows WTP Curves for H and L, $W_H(s)$ and $W_L(s)$. Left panel shows curves for low-income group which come from (Finkelstein, Hendren and Shepard, 2017). Right panel shows curves for high-income group which come from (Hackmann, Kolstad and Kowalski, 2015). Linear curves extrapolate linearly 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. The exact formula for the enhanced curves can be found in the appendix.