The Two Margin Problem in Insurance Markets*

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

February 1, 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. 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.

*We thank Ashley Swanson for serving as a discussant for this paper. We also thank Kate Bundorf, Marika Cabral, Vilsa Curto, Keith Ericson, Amy Finkelstein, Jon Gruber, Tom McGuire, Joe Newhouse, and participants at the 2019 American Economic Association meetings, 2018 American Society of Health Economists meeting, the 2018 Annual Health Economics Conference, the Chicago Booth Junior Health Economics Summit, the NBER Health Care meeting, and the University of Wisconsin for useful feedback. We gratefully acknowledge financial support for this project from the Laura and John Arnold Foundation and the Agency for Healthcare Research and Quality (K01-HS25786-01). No party had the right to review this paper prior to its circulation.

†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 (1) cause some consumers to inefficiently remain uninsured and (2) cause others to select inefficient levels of coverage. The prior literature aimed at understanding equilibria and optimal 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 (Rothschild and Stiglitz, 1976; Einav, Finkelstein and Cullen, 2010; Handel, 2013; 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 (Cawley and Philipson, 1999; Finkelstein and McGarry, 2006; Hendren, 2013; 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. But this simplification does not for allow for potential interactions between these two margins of selection (Azevedo and Gottlieb, 2017). 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. The intuition for this interaction is that the marginal consumers induced to enroll by the mandate may be the healthiest consumers in the market. When these healthy consumers enter the market, they are likely to enroll in the lowest-price (and lowest-quality) plan available. The healthier risk pool causes the plan's costs to drop, leading to lower prices. But the lower prices in the low-quality plan implies an increase in the incremental price between it and the higher-quality plans. This will induce some consumers on the margin between the high-quality and low-quality plans to opt for a low-quality option in the presence of the mandate.

Our first goal in this paper is to show how interventions that counteract selection across one margin can affect 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 managed competition settings, including: 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 addressing 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 affect net costs and bring down the price of the more generous plans in equilibrium, leading to greater takeup. This additional takeup of generous coverage is the intended effect. 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 provide a path forward for empirical studies to incorporate the cross-margin interaction without sacrificing the simplicity or transparency of a reduced form framework based on demand and cost curve estimation. Recent complementary work has pointed to the 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 imperfect competition, horizontal differentiation, selection on moral hazard, and irrational or ill-informed consumers.

With the intuition and price theory in place, we corroborate the model's insights empirically using data from the Massachusetts Connector. The Connector was an individual insurance market that was a precursor to the state Health Insurance Marketplaces established by the Affordable Care Act.

[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.

It was introduced by Massachusetts 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. Following the same transparent regression discontinuity design that exploits discontinuities in the income-based premium subsidy scheme, we construct demand and cost curves for the lower and higher quality plans. We use these demand and cost curves in a number of illustrative counterfactual 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 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 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 findings. The size of the unintended cross margin effects can be large enough to imply first-order impacts on contract allocations. We find that a strong mandate sufficient to move all consumers into insurance—increasing enrollment by 40 percentage points in our setting—can cause the market for more generous insurance to unravel completely.[2] In the other direction, enforcing minimum coverage generosity, via requirements on networks, actuarial values, or essential health benefits, can substantially reduce market-level consumer participation—in our setting by as much as 79 percent. 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

[2] We note that this result occurs under ACA-like subsidies that fixed across plans. During the period from which our data are drawn, Massachusetts used a unique subsidy scheme that provided larger subsidies to consumers choosing more expensive plans. Because of these incremental subsidies, in practice the Massachusetts market does not unravel to low-quality coverage in the presence of a strong mandate.

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.[3] 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 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).

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.[4] Our work also relates to several other parts of the insurance market literature. Ericson and Starc (2015), Handel, Hendel and Whinston (2015), Tebaldi (2017), Domurat (2018), and other empirical, policy-oriented studies have analyzed equilibria and optimal policy in an insurance exchange. A related, but distinct 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

[3] 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.

[4] 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 in 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.

empirical evidence that this occurs in practice in the context of health insurance (Carey, 2017*a*,*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 this ubiquitous regulatory tool. 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 decisions 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 in contexts where 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

Consider an insurance market with two fixed contracts, $j = \{H, L\}$, where H is more generous than L on some metric (e.g., cost sharing, provider network). There is also an outside option, U, which in the focal application of our model to the ACA's individual markets 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 receive a subsidy, S, and choose among options based on post-subsidy prices, $P_j^{cons} = P_j - S$, and based on the price of the outside option, $P_U^{cons} = M$. In our focal example, M is a mandate penalty. The distinguishing feature of U is that its price is exogenous and set by policymakers; 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. $P^{cons} = \{P_H^{cons}, P_L^{cons}, P_U^{cons}\}$ is the vector of consumer prices in the market.

Our goal is to graphically depict insurance market equilibrium and welfare, in the spirit of Einav, Finkelstein and Cullen (2010) (or "EFC"). The ability of this framework to represent selection markets in pictures has been quite useful for teaching and conceptual thinking about risk selection. But the framework is limited to two options – typically one plan plus an outside option. Extending to our richer setting generates additional insights but creates additional complexity for graphing. In the most general formulation, demand curves for two-product markets involve three dimensional curves, which are challenging to work with and difficult to visualize.

To make progress, we assume a *vertical model* of demand, which assumes contracts can be clearly ranked in terms of quality. Although the vertical assumptions are not necessary for many of our key insights to hold, they capture the key features of our two selection margins in a simple way that makes our conceptual contributions clearer and allows for two dimensional graphs. In the next subsections, we present the vertical model and its graphical representation, then add the cost curves, and show how to find equilibrium and welfare.

2.1 Vertical Demand Model: Math and Graphs

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 (so $W_{i,U} \equiv 0$). 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

[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 $U_{i,j} = W_{i,j} - P_j = \beta_i Q_j - P_j$.

are horizontally differentiated, 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.

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 (WTP for L relative to U) *and* lower WTP for H relative to L (smaller gap between $W_H(s)$ and $W_L(s)$).[6] Intuitively, s represents a person's percentile rank (note that it is uniformly distributed) in terms of "taste for generosity." 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 based on WTP (s), 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 (the intensive margin) and between L and U (the extensive margin), but *not* between H and U. If the price of U (the mandate penalty) increases modestly, the newly insured all buy L (the cheaper plan), not H. Similarly, after a small increase in P_H, all the consumers who leave H shift to L, not uninsurance. 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, as long as this margin is quantitatively less important.[7] Thus, Assumption 2 should be thought of as a useful simplification, rather than an assumption that generates knife-edge results.

[6] By construction, we can define s to order consumers by WTP for L vs. U. The assumption is that this same index indexes WTP for H relative to L.

[7] We plan to work out this generalization in an appendix in a future draft.

Figure 1: Consumer Sorting under Vertical Model

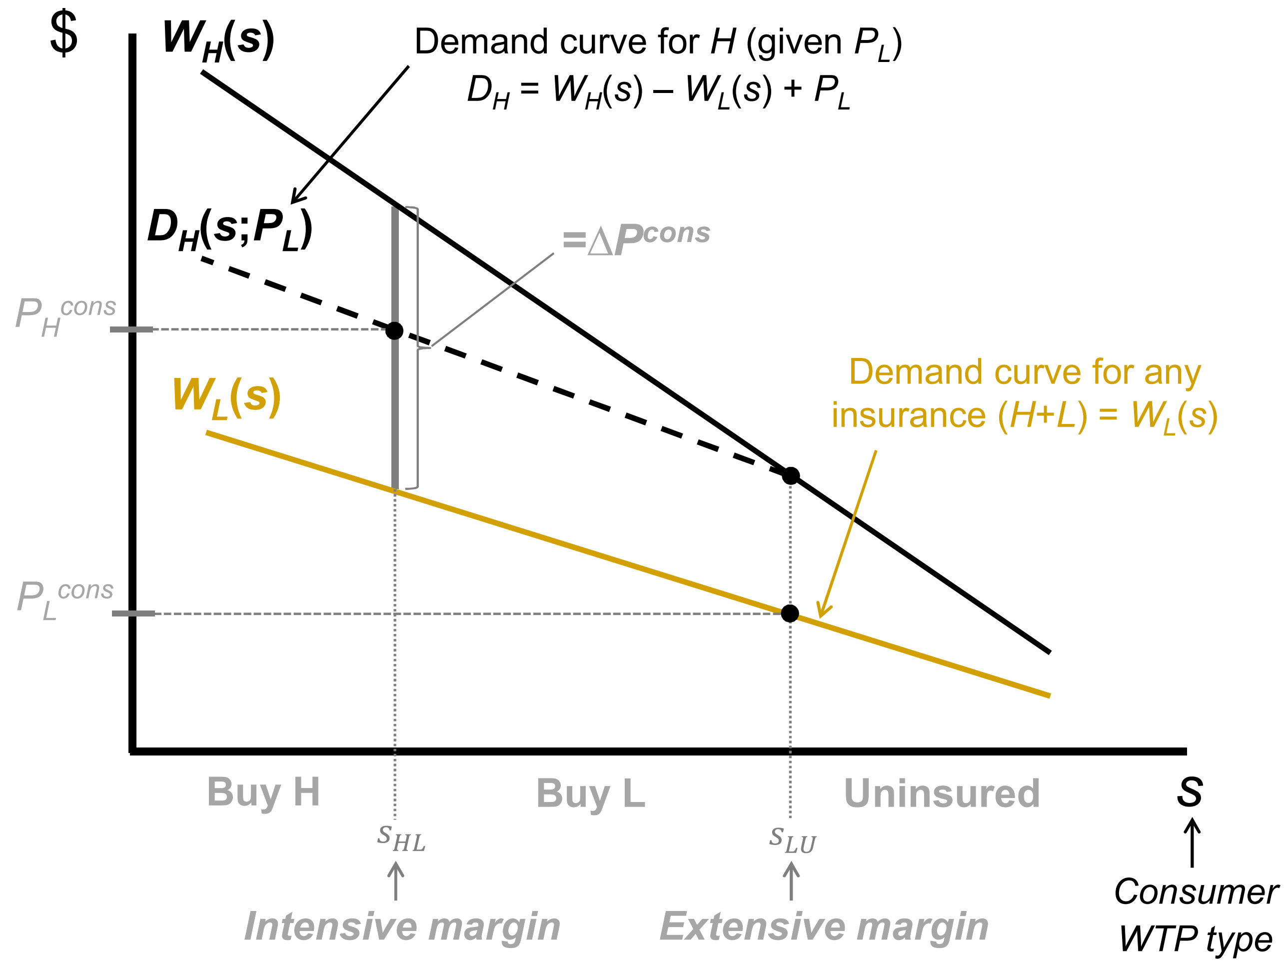

Figure 1 plots a simple linear example of $W_H(s)$ and $W_L(s)$ curves that satisfy these assumptions and shows how consumer sorting plays out. The x-axis is the WTP index s, so WTP declines from left to right as usual. As noted, consumers sort based on WTP, with two marginal types indifferent between options. Let $s_{LU}(P^{cons})$ be the marginal type indifferent between L and U – the extensive margin – at a given set of prices P^{cons}, which is defined by:

$$W_L(s_{LU}) = P_L^{cons} - M \tag{1}$$

Consumers to the right of s_{LU} go uninsured, while those to its left buy insurance. Therefore, $W_L(s)$ represents the *demand curve for any formal insurance* (H or L), with a relevant net price of insurance of $P_L^{cons} - M$ (which recall, equals $P_L - S - M$). This demand curve and net price determine the extensive margin sorting between insured and uninsured. Any policy that affects this net price will affect the insured rate.

Let $s_{HL}(P^{cons})$ be the marginal type indifferent between H and L, defined by:

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

Consumers to the left of s_{HL} all buy H, since their incremental WTP for H over L – which we label ΔW_{HL} – exceeds the incremental price. Therefore, the incremental WTP curve $\Delta W_{HL}(s)$ determines demand for H, with a relevant price of $\Delta P^{cons} \equiv P_H^{cons} - P_L^{cons}$. Together, these determine the intensive margin sorting. With demand for H and for $H + L$ thus determined, demand for L equals the difference between the two.[8]

It will be convenient to define a demand curve for H whose value is P_H^{cons} for the intensive margin type, s_{HL}. To do so, we rearrange equation (2) to define:

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

Figure 1 shows $D_H(s; P_L)$ with a dashed line. Notice several features of D_H. First, it slopes down (as guaranteed by assumption 2), but is flatter than WTP for H – since 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). Second, it is not a pure primitive but instead depends on both primitives ($W_H - W_L$) and, critically, on P_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. Finally, we can draw D_H by noting that it intersects the W_H curve at the cutoff type s_{LU} (since $W_L(s_{LU}) = P_L^{cons}$).[9] It then proceeds leftward at a slope equal to that of ΔW_{HL}, and its intersection with P_H^{cons} determines s_{HL}.

[8] Formally, the demand functions are defined by the following equations:

$$\begin{aligned} D_H(P^{cons}) &= s_{HL}(\Delta P^{cons}) \\ D_L(P^{cons}) &= s_{LU}(P_L^{cons} - M) - s_{HL}(\Delta P^{cons}) \\ D_U(P^{cons}) &= 1 - s_{LU}(P_L^{cons} - M) \end{aligned}$$

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

To preview one of the insights of this paper, we note here that the price of L could be a function of the composition of types choosing L—for example in a competitive equilibrium in which plan prices equal average plan costs. Because demand for H depends not only on the price of H but also on the price of L, this implies that 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. For example, adjusting a mandate penalty that only directly induces some consumers from uninsurance into the L plan could impact the take-up and even existence of the H plan. We next turn to describing cost functions and defining a competitive equilibrium to make such insights precise.

2.2 Costs

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

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

$C_j(s)$ is analogous to "marginal costs" 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 with two margins of adjustment, we refer to $C_j(s)$ as type-specific costs, or simply costs. In addition, we define $C_U(s)$ as the expected costs of uncompensated care – incurred by third parties, not insurers – of type-s consumers if uninsured. Along with adverse selection, uncompensated care motivates subsidies and mandates.

[10] The reason is that with community rated pricing, consumers sort into plans based only on WTP. There is no way to segregate consumers within WTP, 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.

Figure 2: Cost Curves under Vertical Model

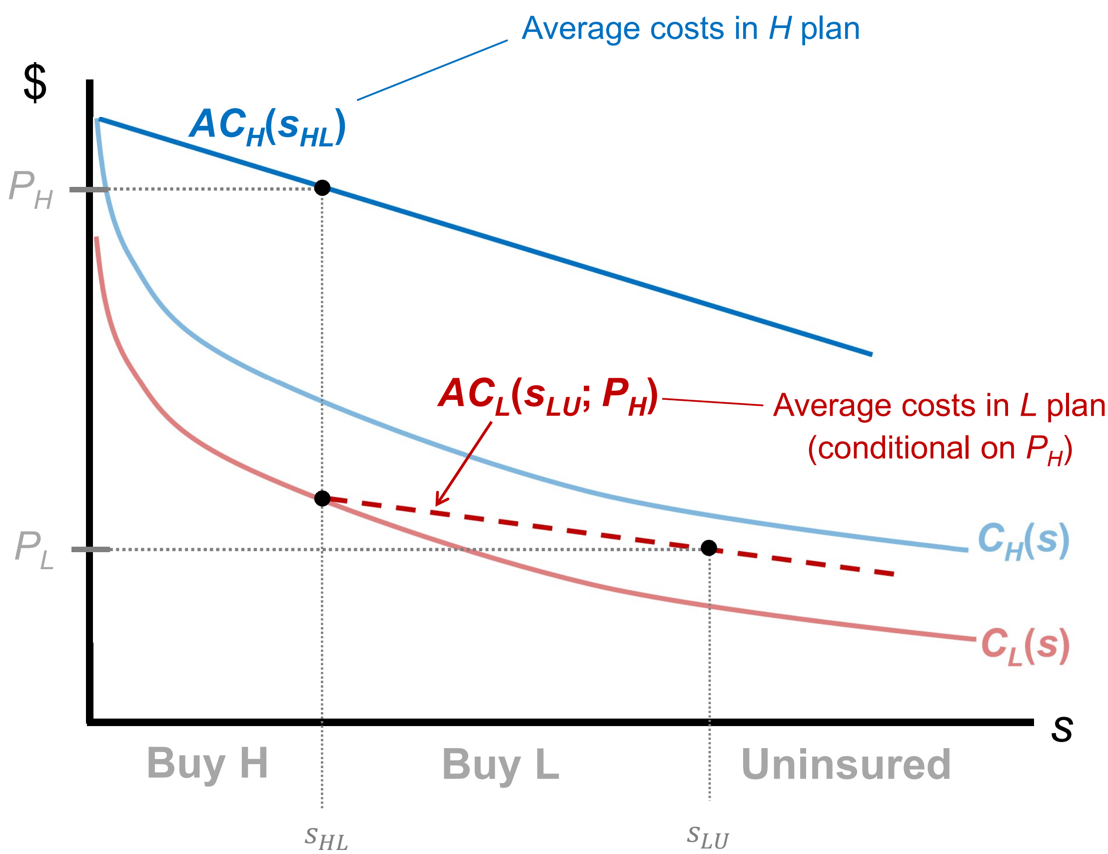

In addition to type-specific costs which is sufficient for welfare, EFC derive an average cost curve, AC_j, whose intersection with demand determines competitive equilibrium. We apply the analogous concept, though the multi-plan setting makes its construction more complex. Average costs are defined as 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 price 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. $C_H(s)$ describes the expected insurer cost in H across all consumers of type s, and $C_L(s)$ is the analogous cost for L. 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 $AC_H(P)$, average cost for the H plan. It is helpful for graphing to redefine the argument of AC_H as the marginal type, $s_{HL}(P)$, that buys H; we use this notation in Figure 2. To graph $AC_H(s_{HL})$, recognize that for $s_{HL} = 0$, the only consumers enrolled in H are the consumers for whom $s = 0$, implying that $AC_H(s = 0) = C_H(s = 0)$. Then, as s_{HL} increases moving right along the horizontal axis, H includes more and more relatively healthy consumers, resulting in a downward sloping average cost curve. Eventually, when $s_{HL} = 1$ and all consumers are enrolled in H, $AC_H(1)$ 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. Importantly, this curve can be calculated directly from a market primitive (by integrating over $C_H(s)$) and is not an equilibrium object conditional on knowing s_{HL}. In this sense, the AC_H curve is "fixed," which we denote by plotting it with a solid line in Figure 2. In other words, while the average cost in H depends on consumer sorting between H and L (and thus, indirectly, on the price of L), the average cost *curve*, or the *relationship* between average cost and consumer sorting between H and L, does not depend on the price of L, making the curve fixed.

The average cost curve for L is more complicated because it is an average over a range, $s \in [s_{HL}, s_{LU}]$ that has two endogenous margins. This makes it impossible to plot a single fixed AC_L curve as we did with AC_H. Nonetheless, it is possible to plot it if we first *condition* on P_H. We denote this curve $AC_L(s_{LU}; P_H)$ and plot it with a dashed red curve in Figure 2. The intersection of P_H with AC_H will in equilibrium imply a given $s_{HL}(P_H)$, the marginal type on the intensive margin who is also the highest-WTP type buying L. Just as AC_H equals C_H at $s = 0$, AC_L equals C_L at $s = s_{HL}$. Moving rightward, L includes more and more relatively healthy consumers, resulting in a downward sloping average cost curve. Eventually, when $s_{LU} = 1$ and all consumers not in H are enrolled in L (i.e. no consumers opt for U), $AC_L(1; P_H)$ is equal to the average cost in L across all types of consumers for whom $s > s_{HL}(P_H)$.

The critical difference between AC_H and AC_L is that while AC_H gives the average cost in H for

[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.

a given cutoff s_{HL}, AC_L gives the average cost in L for a given cutoff s_{LU}, *conditional on* P_H. While AC_H is a fixed curve and does not depend on the price of L, AC_L is an equilibrium object in that it changes as the price of H or (equivalently) the cutoff s_{HL} changes. This dependency represents the interaction between extensive and intensive margin selection that is the focus of this paper. For example, a subsidy targeted to H that results in a lower P_H would cause AC_L to shift downward and 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. We explore these equilibrium interactions in detail in the next section.

2.3 Competitive Equilibrium

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

$$\begin{aligned} P_H &= AC_H\left(P^{cons}\right) \\ P_L &= AC_L\left(P^{cons}\right) \end{aligned} \tag{6}$$

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 Einav, Finkelstein and Cullen (2010) and Handel, Hendel and Whinston (2015). In those settings, it is assumed that all consumers choose either H or L. This assumption conveniently simplifies the equilibrium condition from two expressions to one because without the outside option the price vector that results in $P_H - P_L = AC_H - AC_L$ also results in $P_H = AC_H\left(P^{cons}\right)$ and $P_L = AC_L\left(P^{cons}\right)$. Given one equilibrium condition, that the differential average cost must be set equal to the differential price, the equilibrium process can be plotted on a simple two-dimensional graph as shown in Handel, Hendel and Whinston (2015).

[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)

Figure 3: Competitive Equilibrium with H Only

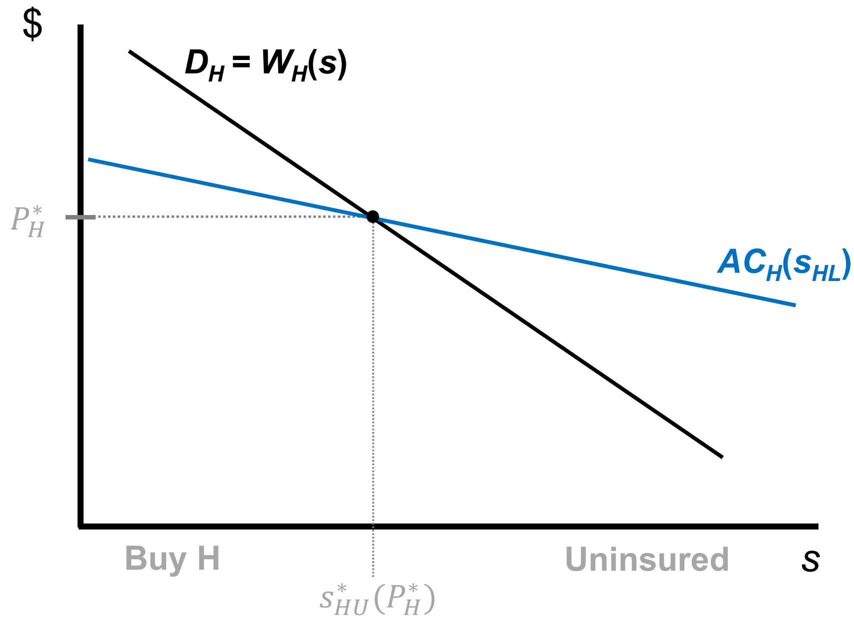

Once one considers the outside option of uninsurance, the equilibration process is more complex because the two equilibrium conditions described in equation (6) no longer simplify to one condition. To build intuition, we now describe this process graphically. We start by considering the simple case where there is only one plan, the H plan, and uninsurance. This case is analogous to the classic case considered by Einav, Finkelstein and Cullen (2010), if one defines their L plan as uninsurance. This case is presented in Figure 3, which includes two curves. The first curve is the demand curve for H which is equal to the WTP curve for H ($D_H = W_H$) because there is no L option, implying that 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. The second curve is the average cost curve for H which describes the average cost in H for the consumers who purchase H given an H vs. U cutoff s_{HU}, i.e. all consumers with $s < s_{HU}$. The equilibrium price, P_H^*, and cutoff s-type, $s_{HU}^*(P_H^*)$, are found where the demand and average cost curves intersect, as this is the cutoff s-type for which price is equal to average cost as specified in equation (6) above.

Figure 4: Equilibrating Process with H, L, and Outside Option

((A)) Panel A

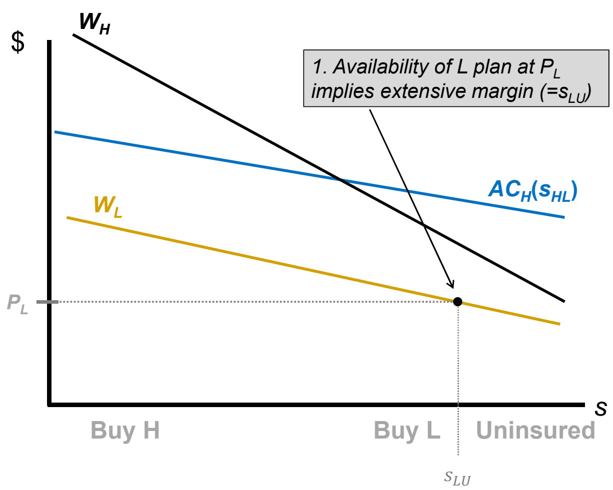

((B)) Panel B

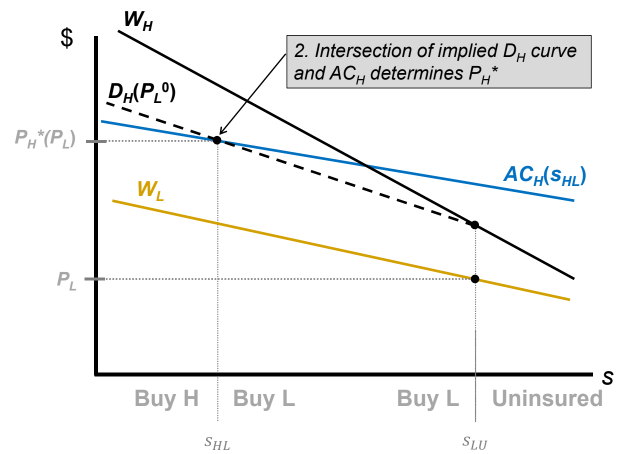

((C)) Panel C

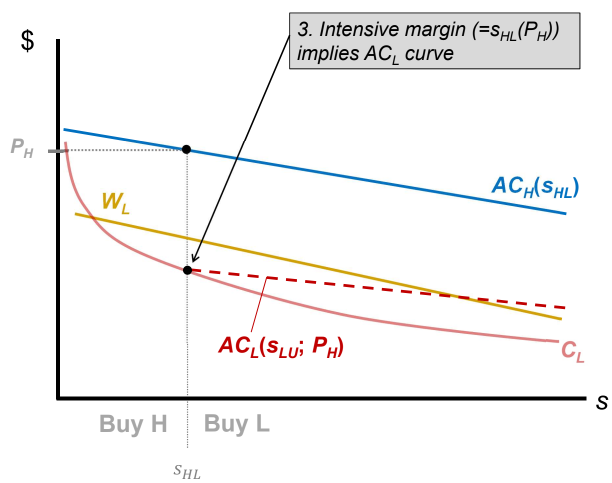

((D)) Panel D

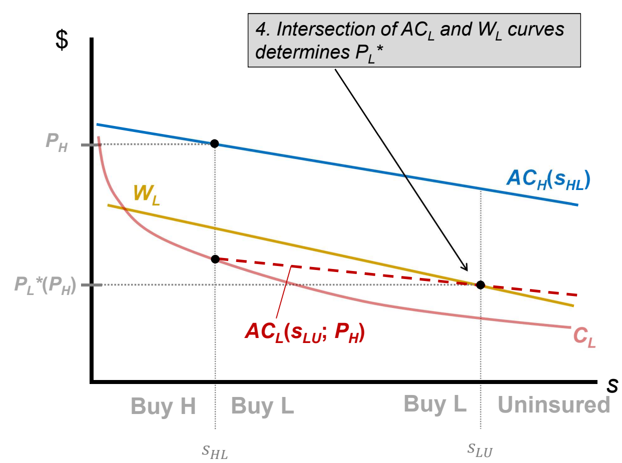

In Figure 4 we add L. Adding L causes a number of significant changes to the figure and the equilibration process. To provide intuition for these changes, we proceed in four steps, corresponding to the four panels in the figure. Panels A-B show how P_H is determined, given a fixed price of L. Panel A 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 fixed price of L.

Figure 5: Welfare

((A)) Equilibrium with H, L, and Outside Option

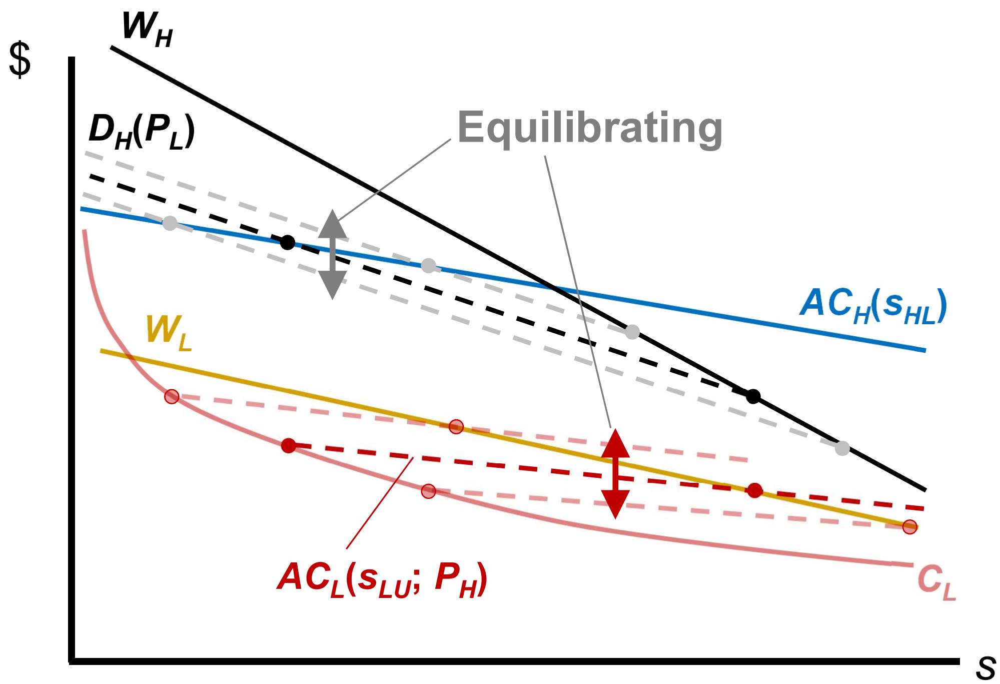

((B)) Final Equilibrium

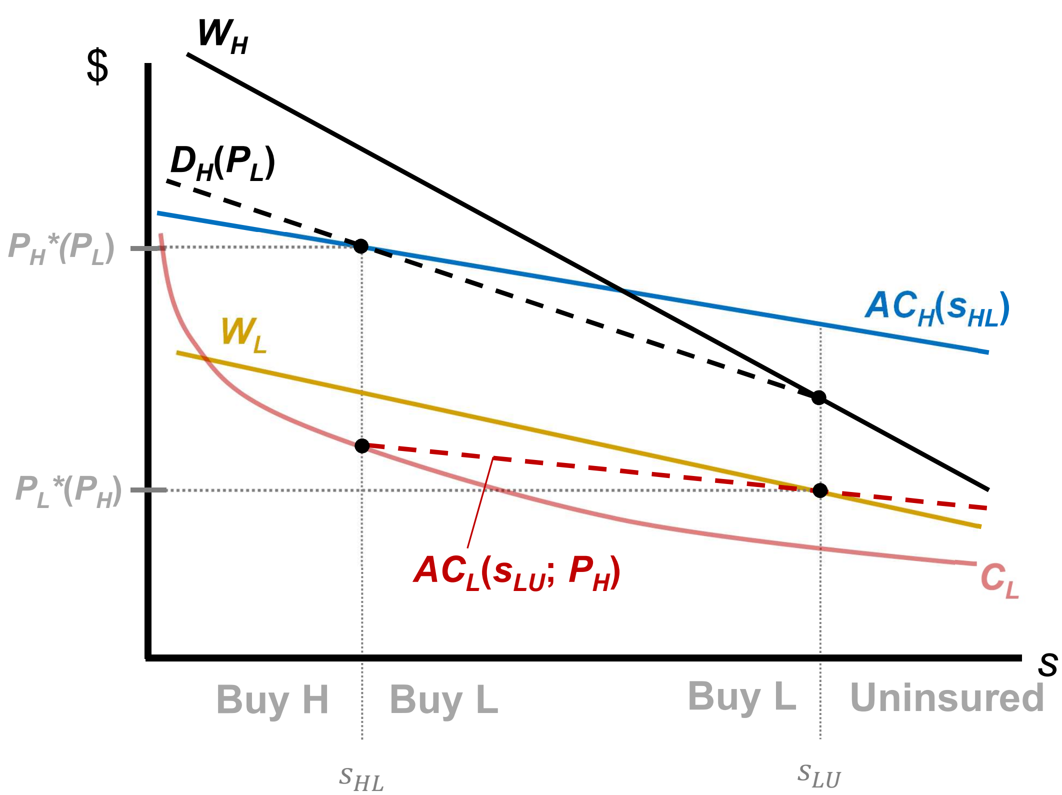

Panels C-D of Figure 4 show how P_L is determined, given a fixed P_H. Panel C shows that the fixed P_H implies a given intensive margin cutoff (s_{HL}), which in turn fixes the AC_L curve. Panel D shows how the intersection of AC_L with W_L determines P_L (and the extensive margin cutoff s_{LU}). This process determines the reaction function $P_L^*(P_H)$, which describes the breakeven price of L for a given fixed price of H.

In equilibrium, the reaction functions must equal each other: $P_H = P_H^*(P_L)$ and $P_L = P_L^*(P_H)$. The top panel of Figure 5 depicts this equilibrating process occurring. The key idea is that different values of P_H shift the average cost curve for L (thereby changing $P_L^*(P_H)$), while different values of P_L shift the demand curve for H (thereby changing $P_H^*(P_L)$). This process continues until a rest point is reached, as depicted in Panel B of Figure 5. This rest point occurs where $P_H = P_H^*(P_L^*)$ and $P_L = P_L^*(P_H^*)$.

To sum up, unlike in the case with only an H plan, equilibrium now consists of a vector of prices instead of a single price (or instead of a single price difference). The dashed lines in Figure 5—the H demand curve and the L average cost curve—are themselves equilibrium outcomes, even though we are holding fixed consumer preferences and costs. The equilibrium vector of prices are the prices at which the demand curve for L intersects the average cost curve for L *and* the demand curve for H simultaneously intersects the average cost curve for H.

2.4 Social Welfare

We now show how the framework can be used to assess the welfare consequences of different policies. We define social welfare as total social surplus, abstracting from any distributional concerns:

$$\widehat{SW}\left(P^{cons}\right) = \int_{0}^{s_{HL}\left(P^{cons}\right)} \left(W_H\left(s\right) - C_H\left(s\right)\right) ds + \int_{s_{HL}\left(P^{cons}\right)}^{s_{LU}\left(P^{cons}\right)} \left(W_L\left(s\right) - C_L\left(s\right)\right) ds - \int_{s_{LU}\left(P^{cons}\right)}^{1} C_U\left(s\right) ds \quad (7)$$

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

$SW = \widehat{SW} + \int_0^1 C_U(s)\,ds$. Rearranging and simplifying, this yields the following expression:

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

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 given $W_L(s)$, $\Delta W_{HL}(s)$, C_L^{net}, and $\Delta C_{HL}(s)$ as well as the equilibrium cutoff values s_{LU}^* and s_{HL}^*. Figure 6 illustrates this concept for the case where $C_U = 0$ for simplicity. In the figure we plot W_L, W_H, C_H, and C_L. As above, the consumers who buy H are in the region of the x-axis to the left of s_{HL}^*. Equation 8 shows that for these consumers, social surplus is equal to the area between the W_H curve and the C_H curve ($ABGH$), highlighted in green. The consumers who buy L are in the region of the x-axis between s_{HL}^* and s_{LU}^*. For these consumers, social surplus is equal to the area between the W_L curve and the C_L curve ($JKNO$), again highlighted in green. Surplus foregone due to consumers enrolling in lower quality coverage on the intensive margin (L) consists of three pieces: Surplus from L for consumers who choose H ($IJOP$), surplus from H for consumers who choose L ($BCFG$), and surplus from H for consumers who choose U, all highlighted in red. In some cases foregone surplus will be less than market surplus for the consumers choosing H. In these cases, it is optimal for these consumers to choose H. In other cases, foregone surplus will be more than market surplus for the consumers choosing H. In these cases, it is optimal for these consumers to choose L instead of H. The Figure illustrates that both of these possibilities can be captured by our model.

In addition to intensive margin foregone surplus, there is also foregone surplus on the extensive margin. For consumers choosing to be uninsured, those with $s > s_{LU}^*$, extensive margin foregone surplus is equal to the surplus from L ($KLMN$), again highlighted in red. Here, the foregone surplus can again be positive or negative, depending on whether consumers choosing L value L more than the cost of providing it to them, C_L.

Figure 6: Welfare with Two Margins of Adjustment

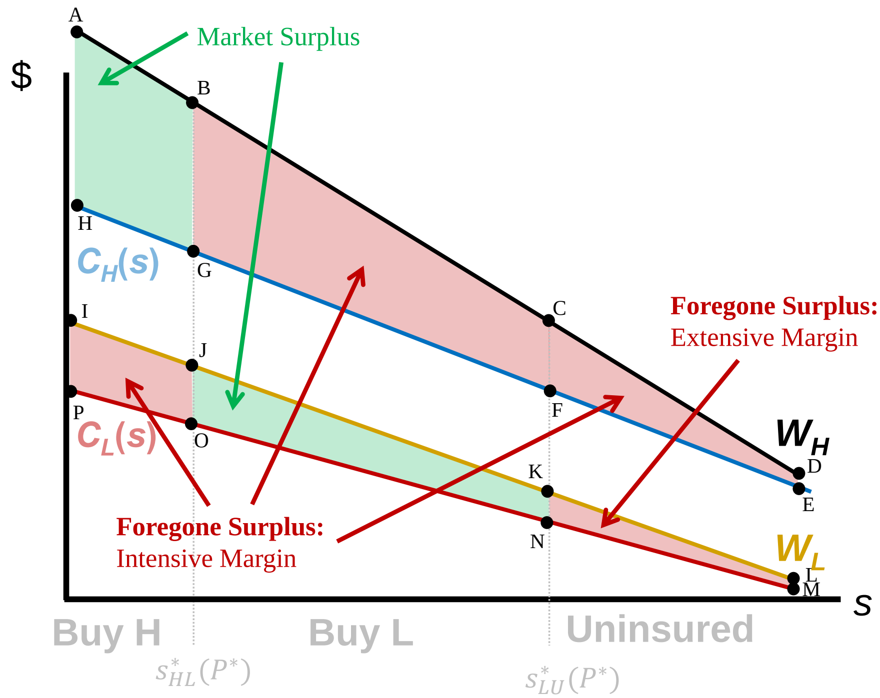

In the next section, we will use these insights and this graphical welfare model to illustrate the welfare consequences of various policies targeted and combating adverse selection problems on the extensive and intensive margins. The ability of our model to capture welfare as well as allocations of consumers across options in the market is critical, given that we will show that these policies typically have both intended consequences that often improve welfare and unintended consequences that often lead to lower levels of welfare. Thus, the welfare consequences are ambiguous and a strictly empirical question.

3 Equilibrium under Regulation

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, penalties for choosing uninsurance/subsidies for purchasing insurance, and risk adjustment transfers. Each of these policies is targeted at adverse selection problems on just one of the two margins, but we show affects the other.

We discuss each policy in turn and provide graphical illustrations for their intensive and extensive margin consequences. We also use the graphical model to illustrate the welfare consequences of each policy. We conclude with a discussion of how our framework can be used to understand the cross-margin effects of any market intervention that affect selection, including behavioral interventions for which selection effects may be a nuisance rather than an aim.

3.1 Benefit Regulation

We start by using our model to explore the consequences of benefit regulation. In Figure 7, we consider a regulation that eliminates L plans from the market. This thought experiment captures a variety of potential policies designed to limit the space of allowed insurance contracts, including network adequacy regulations, caps on maximum out-of-pocket costs under a plan, or the Essential Health Benefit regulations of the ACA. All of these policies can be thought of as eliminating low-quality options from the market.

The left panel of Figure 7 shows equilibrium prices and the allocation of consumers across plans in the case with no benefit regulation, i.e. where L is allowed to be offered in the market. The right panel illustrates equilibrium prices and allocations with benefit regulation, i.e. where L is eliminated from the market. When L is eliminated, the market can be described by the Hackmann, Kolstad and Kowalski (2015) adaptation of the classic model of Einav, Finkelstein and Cullen (2010). Demand for H is now equal to W_H, as *all* consumers who value H more than P_H will now purchase H instead of just the subset of those consumers whose incremental willingness-to-pay $W_H - W_L$ exceeds the incremental price $P_H - P_L$. The equilibrium price then occurs where the now fixed D_H crosses AC_H. This price implies an equilibrium H vs. U cutoff s-type consumer, $\hat{s}^*_{HU}$.

Comparing equilibrium sorting of consumers across plans for the cases with and without benefit regulation reveals that limiting the contract space has two competing effects: (1) some consumers who would have enrolled in L in the absence of benefit regulation now enroll in H (intended consequence) and (2) some consumers who would have enrolled in L in the absence of benefit regulation now opt to go without insurance (unintended consequence).

Figure 7: Equilibrium under Benefit Regulation

((A)) No benefit regulation

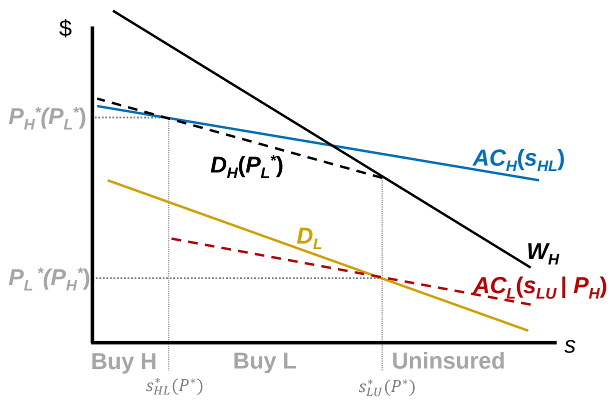

((B)) Plan L eliminated

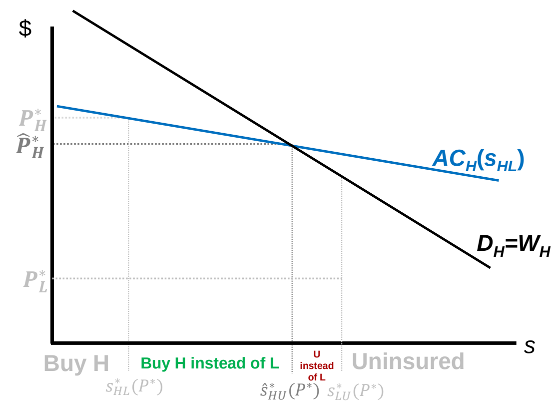

((C)) Net enrollment and welfare impacts (B-A)

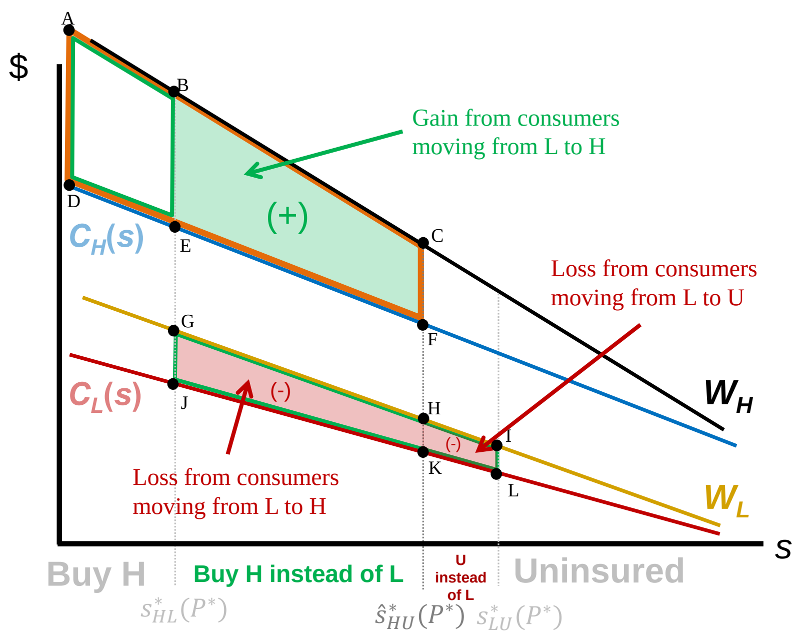

Notes:

These two effects are likely to have opposing welfare consequences, with the shift of consumers from L to H likely improving welfare and the shift of consumers from L to U likely reducing welfare. This implies that the welfare consequences of the policy are theoretically ambiguous. Fortunately, the model also allows for direct estimation of welfare effects of this type of policy. As discussed in Section 2, welfare can be estimated given W_H, W_L, C_H, C_L, and the equilibrium allocations of consumers across plans in the different policy environments.[13] These curves are plotted in Panel C of Figure 7 along with vertical lines indicating the cutoff consumer types without benefit regulation, s^*_{HL} and s^*_{LU}, and with benefit regulation, $\hat{s}^*_{HU}$. The social surplus without benefit regulation is outlined in green and consists of the surplus for consumers choosing H ($ABDE$) and the surplus for consumers choosing L ($GILJ$). The social surplus with benefit regulation is outlined in orange and consists of the surplus for consumers choosing H ($ACFD$).

The figure highlights the difference in welfare with and without benefit regulation, with welfare increases filled in green and welfare decreases filled in red. The welfare difference consists of two key pieces: (1) a shift in welfare due to the choice of some consumers to choose H instead of L and (2) a shift in welfare due to the choice of some consumers to choose U instead of L. The welfare impact of (2) is unambiguously negative as long as consumers value L more than its cost C_L, and the size of the loss is described by the area of $HILK$. The welfare impact of (1), on the other hand, is ambiguous, and depends on whether social surplus (valuation minus cost) of H is larger than that of L for the marginal consumers. For example, surplus from H might be smaller than surplus from L if the difference in moral hazard between H and L is large, such that the incremental value consumers place on H vs. L is less than the incremental cost of H vs. L. The size of the gain from (1) is equal to the area of $BCFE$ minus the area of $GHKJ$. The overall welfare effect of benefit regulation is found by subtracting the area of $GILJ$ from the area of $BCFE$.

This illustrates the interactions between the two margins of selection: Benefit regulation seeks to limit the effects of intensive margin selection by eliminating the intensive margin. It clearly does this. But this intensive margin policy has an extensive margin consequence: some consumers opt to go uninsured. The welfare welfare consequences, which are theoretically ambiguous and will vary across market settings, can be estimated directly in our model with two sources of price variation to identify the two sets of demand and cost curves.

3.2 Insurance Mandate/Uninsurance Penalty

Next we consider the consequences of a penalty assessed on consumers who opt not to purchase insurance, i.e. consumers who choose U. In a standard CARA utility framework, such a penalty is

[13]This is not quite true. one also needs an estimate of C_U, the social cost of someone choosing not to be insured, for a full welfare analysis. We abstract from C_U here by assuming that it is zero. In the empirical section of the paper we return to this issue and allow for non-zero C_U.

equivalent to a subsidy given to consumers who choose to purchase insurance, so all consequences of the penalty we consider here should also apply to an insurance subsidy.[14]

A penalty has a direct effect on insurance takeup, an indirect effect on insurance takeup, and a cross-margin effect on sorting between L and H. The direct effect of a penalty for choosing U is to cause demand for L to shift above the underlying valuation for L, W_L: With an uninsurance penalty, at a given P_L some consumers who value L less than P_L will purchase L due to the non-zero cost of choosing uninsurance. This is shown in Panel A of Figure 8. In the figure, equilibrium prices and cutoff s-types for the case without a mandate are shown in light gray and are labeled without hats: $P^*_H, P^*_L, s^*_{HL}, s^*_{LU}$. Without the penalty, equilibrium prices occur where W_L crosses AC_L and where the original D_L crosses AC_H. The introduction of an uninsurance penalty is illustrated as an upward shift of the demand curve for L such that D_L is now above W_L. The indirect effect is apparent in the figure illustrates as this upward shift in D_L leads to a new equilibrium price of L, $\hat{P}^*_L$, that is less than the equilibrium price of L without the penalty, P^*_L due to the downward-sloping AC_L curve. At this lower price of L, some consumers who would have chosen to remain uninsured without the penalty now choose to purchase L. The direct effect on demand and indirect effect on prices are the intended impacts of the penalty.

But there is also a cross-margin effect. Recall that demand for H is a function of the price of L: At the new lower price of L, H is relatively less attractive to consumers than previously. Thus, the penalty also leads to a downward shift in D_H. This effect of the penalty is also shown in Panel A of Figure 8. Lower demand for H leads to a higher equilibrium price of H, $\hat{P}^*_H$ due to the downward-sloping AC_H curve. At this higher price of H, some consumers who would have enrolled in H in the absence of the penalty now choose to enroll in L. This is the unintended effect of the policy.

Importantly, note that the equilibration process does not end with the shift in D_H. The new higher price of H leads to an upward shift in AC_L, causing the equilibrium price of L to shift back up, closer to the original price without the mandate, undoing some of the intended effect of the mandate. This leads to fewer consumers choosing L instead of U than would originally have been assumed when not considering the unintended intensive margin effects of the penalty. Thus our model shows how and why mandates may be less effective than the canonical models would predict.

[14] Non-CARA utility, and even non-neoclassical, non-rational consumer behavior pose no particular difficulty for our model, but the presence of such phenomenon do suggest that the particular price variation used in an empirical context—for example, changes in mandate penalties versus changes in subsidies—should match the counterfactual exercise to the extent possible.

It is straightforward to show, however, that the equilibrium price of L will *always* be decreasing in the penalty, while the equilibrium price of H will *always* be increasing in the penalty. We provide a proof in Appendix B.[15]

The competing effects—additional takeup for consumers and downgrading from H to L for others—again suggest that the welfare consequences of the policy are theoretically ambiguous. Panel B of Figure 8 provides graphical analysis of social surplus under each policy to assess the net affect. Again, we plot the sufficient statistics for welfare analysis: W_H, W_L, C_H, C_L. We also indicate the equilibrium cutoff s-types without the penalty, s^*_{HL} and s^*_{LU}, and with the penalty $\hat{s}^*_{HL}$ and $\hat{s}^*_{LU}$. In this case, the primitives chosen generate a full unravelling of the market for generous coverage (an empirically relevant case below), though the procedure is applicable to any case. The social surplus without the uninsurance penalty is outlined in green and consists of the surplus for consumers choosing H ($ABDE$) and the surplus for consumers choosing L ($GILJ$). The social surplus with the penalty is outlined in orange and consists only of the surplus for consumers choosing L ($ACFD$), as the penalty causes the market to unravel to L, leaving no consumers in H.

The figure highlights the difference in welfare with and without the uninsurance penalty, with welfare increases filled in green and welfare decreases filled in red. The welfare difference consists of two key pieces: (1) a shift in welfare due to the choice of some consumers to choose L instead of H and (2) a shift in welfare due to the choice of some consumers to choose L instead of U. The welfare impact of (2) is unambiguously positive as long as consumers value L more than its cost C_L, and the size of the loss is described by the area of $GHIJ$. This area describes the intended effect of the policy. The welfare impact of (1), on the other hand, is again ambiguous, and, as in the case of benefit regulation, depends on whether social surplus (valuation minus cost) is larger for H vs. L for the marginal consumers. The size of the gain from (1) is equal to the area of $ABDC$ minus the area of $EFLK$. The overall welfare effect of benefit regulation is found by summing the areas of $EFKL$ and $GHJI$ and subtracting the area of $ABDC$.

[15]This is true for any *stable* equilibrium. These claims need not hold for unstable equilibria. However, we deem unstable equilibria irrelevant by definition. We also note that these relationships need not hold in a market where there are risk adjustment transfers. In such a market, the cheaper consumers brought into the market due to the penalty have a direct negative effect on AC_H due to risk-pooling across the plans. If this negative effect dominates the substitution effect of consumers moving from L to H due to the lower relative price of L, the equilibrium price of H could be decreasing in the penalty.

Figure 8: Equilibrium under an Uninsurance Penalty

((A)) Effects of Uninsurance Penalty on Market Shares

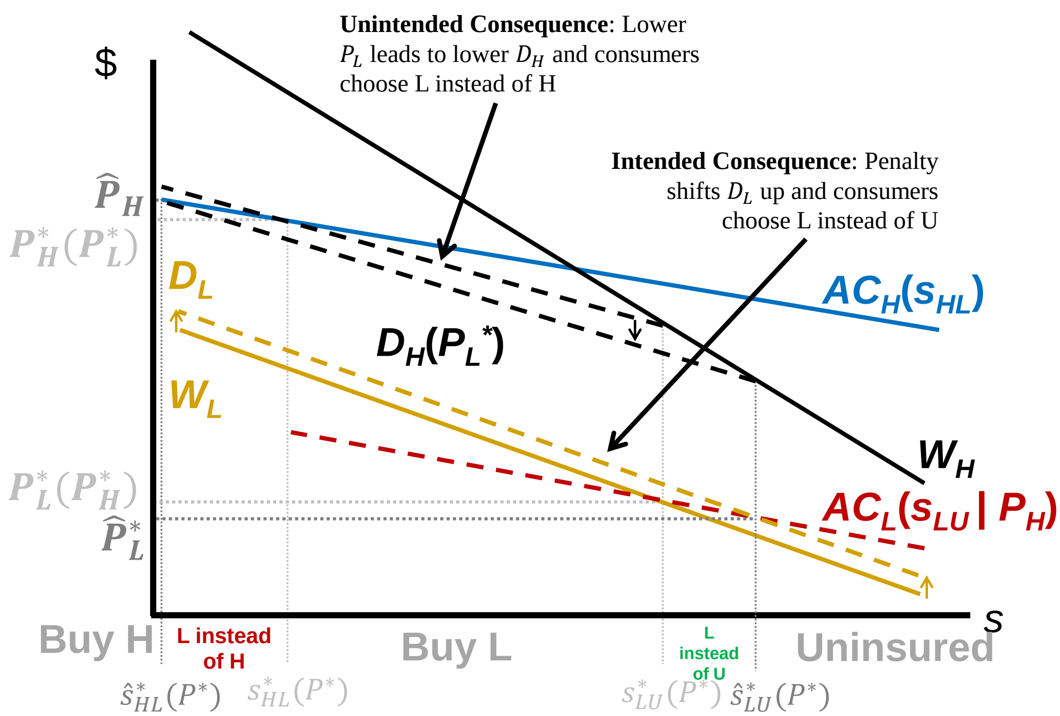

((B)) Welfare Effects of Uninsurance Penalty

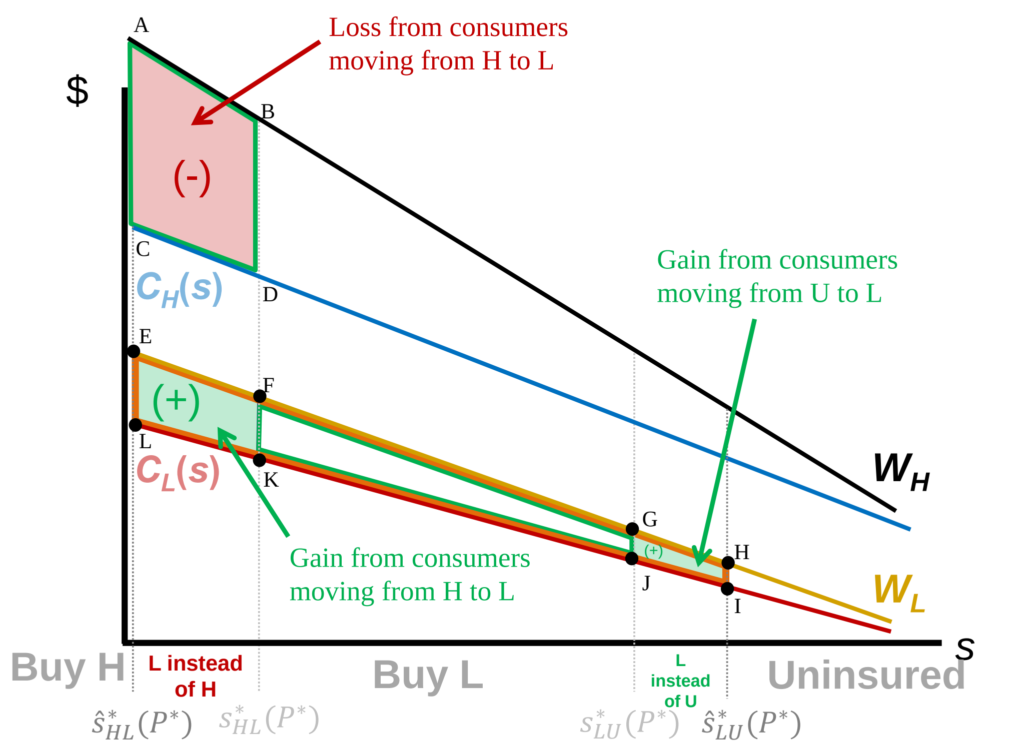

Again, this provides an illustration of the interaction between the two margins of selection. An uninsurance penalty seeks to limit the effects of adverse selection on the extensive margin by inducing healthy consumers to enroll in insurance. However, by adding healthy consumers to the market, the penalty may also lead some consumers previously enrolled in high-quality coverage to instead enroll in lower-quality coverage. In the extreme, a penalty could even lead to a market where no high-quality contracts are available 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 net cost curves that crowd the figure. Nonetheless, risk adjustment is an important policy lever used by regulators to combat intensive margin selection, so we show how our model can be used to describe the consequences of a simplified version of the policy. Specifically, we graphically explore how *perfect* risk adjustment, where transfers perfectly capture all variation in C_L across consumer types, affect equilibrium prices, allocations of consumers across plans, and welfare. We then describe the effects of magnifying *imperfect* risk adjustment transfers (a real-world policy option for states and the policy we study empirically in Section 5), relying on comparative statics derived in Appendix B.

To simplify exposition, we first make explicit the assumption that the causal difference in cost for a given consumer type enrolling in H vs. L is constant across all consumers and equal to δ. We then define perfect risk adjustment as a set of transfers such that the average cost in H net of risk adjustment is always equal to the average cost in L net of risk adjustment plus δ: $RAC_H(P^{cons}) = RAC_L(P^{cons}) + \delta$. 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 healthy consumers leave the market, then costs go up in both L *and* H. This occurs because the transfers cause the difference between the average cost in H and the average cost in L to always be equal to δ, so any policy that affects costs in L also affects costs in H.

We describe this setting in Panel A of Figure 9. In the figure, the equilibrium price of L, P_L^* occurs where $D_L(s)$ crosses $AC_L(s_{LU})$ (drawn as a semi-transparent red line) as before. The equilibrium price of H, P_H^* occurs where the original $D_H(s)$ (drawn as a dashed light gray line) crosses $AC_H(s_{HL})$

(drawn as a semi-transparent light blue line) as before. Under risk adjustment, equilibrium prices no longer occur where demand curves cross average cost curves. Instead, equilibrium prices now occur where demand curves cross *risk-adjusted* average cost curves, $RAC_L(s_{LU})$ and $RAC_H(s_{HL})$. With perfect risk adjustment, the risk-adjusted average cost in L no longer depends on the sorting of consumers between H and L. Instead, the risk-adjusted average cost in L now only depends on the sorting of consumers between L and U. We thus know that, unlike $AC_L(s_{LU})$ which was a function of P_H, $RAC_L(s_{LU})$ is a fixed object, independent of the price of H. Because under perfect risk adjustment $RAC_H(P^{cons}) = RAC_L(P^{cons}) + \delta$ we also know that $RAC_L(s_{LU})$ is parallel to AC_H, and the distance between RAC_L and AC_H is equal to δ. We can thus draw $RAC_L(s_{LU})$ as in Panel A of Figure 9 where $RAC_L(s_{LU})$ is represented by the solid red line which is parallel to $AC_H(s_{HL})$ and the gap between $AC_H(s_{HL})$ and $RAC_L(s_{LU})$ is equal to δ for all s.

It is thus straightforward to see that perfect risk adjustment induces an upward shift (as well as a slight rotation to a slightly steeper slope) in the relevant average cost curve for L from $AC_L(s_{LU})$ to $RAC_L(s_{LU})$. This shift/rotation results in the equilibrium price of L with risk adjustment $\hat{P}_L^*$ being higher than the equilibrium price of L without risk adjustment P_L^*. This results in some consumers choosing to be uninsured under perfect risk adjustment who would have chosen to enroll in L without risk adjustment. This is the unintended consequence of risk adjustment.

Now, the new higher price of L under risk adjustment also implies a new demand curve for H under risk adjustment, as demand for H remains a function of the price of L. This new demand curve for H is drawn in Panel A of Figure 9 as a black dashed line, which is above the demand curve for H without risk adjustment, drawn as a dashed gray line. Additionally, with perfect risk adjustment, there is a new average cost curve for H. As we previously highlighted, because risk adjustment is perfect, the risk-adjusted average cost of H does not depend on consumer sorting between H and L. This implies that $RAC_H(s_{HL})$ is flat, as the average risk-adjusted cost is not a function of $s_H L$. The curve is not fixed, however, as the average risk-adjusted cost in H clearly depends on who enrolls in the market. This is determined by $s_L U$ and the price of L. Thus, now, rather than $AC_L(s_{LU})$ being a function of the price of H, $RAC_H(s_{HL})$ is a function of the price of L. The figure shows that the level of $RAC_H(s_{HL})$ is equal to $AC_H(s_{HL} = \hat{s}_{LU}^*)$. In other words, the average cost in H of the consumers who choose to purchase insurance at a given P_L (consumers with $s < s_{LU}(P_L)$) is the average risk-adjusted cost of consumers who choose to purchase H given any price of H and any sorting of insured consumers between H and L.

Given the new demand curve for H and the risk-adjusted cost curve for H, the equilibrium price of H with risk adjustment occurs where these two curves cross, because it is at this price where firms offering H earn zero profits. Note that this new equilibrium price with perfect risk adjustment, $\hat{P}_H^*$, is lower than the equilibrium price without risk adjustment, P_H^*. This lower price results in some consumers who would have chosen L without risk adjustment to instead opt for H with risk adjustment. This is the intended consequence of risk adjustment transfers.

Our graphical model thus shows two competing effects of risk adjustment: (1) transfers lower the price of H and cause some consumers to choose H instead of L and (2) transfers raise the price of L and cause some consumers to opt for uninsurance instead of choosing L. This suggests that, like benefit regulation and uninsurance penalties, the welfare effects of risk adjustment are theoretically ambiguous. Again, however, our model provides a simple framework for estimating the net welfare effects. This is shown in Panel B of Figure 9. Again, the figure includes the sufficient statistics for welfare analysis: W_H, W_L, C_H, and C_L as well as vertical lines indicating the cutoff consumer types without risk adjustment, s_{HL}^* and s_{LU}^*, and with risk adjustment, $\hat{s}_{HL}^*$ and $\hat{s}_{LU}^*$. The social surplus without risk adjustment is outlined in green and consists of the surplus for consumers choosing H ($ABEF$) and the surplus for consumers choosing L ($GJKN$). The social surplus with risk adjustment is outlined in orange and consists of the surplus for consumers choosing H ($ACDF$) and the surplus for consumers choosing L ($HILM$).

Figure 9: Equilibrium under Perfect Risk Adjustment

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

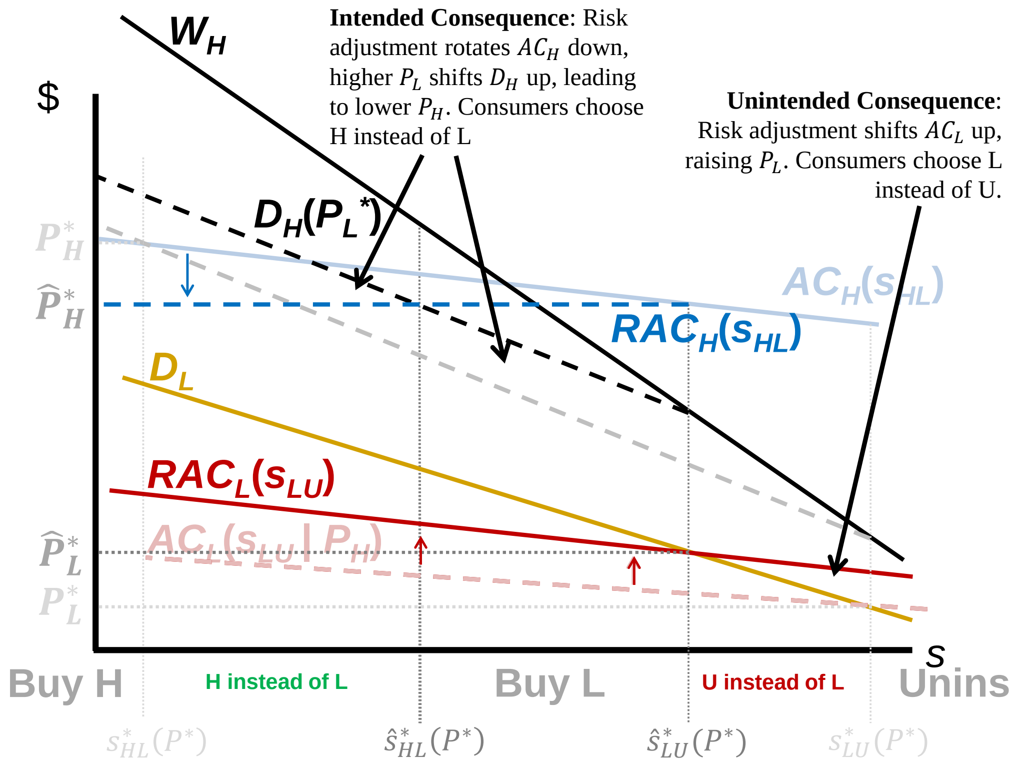

((B)) Welfare Effects of Perfect Risk Adjustment

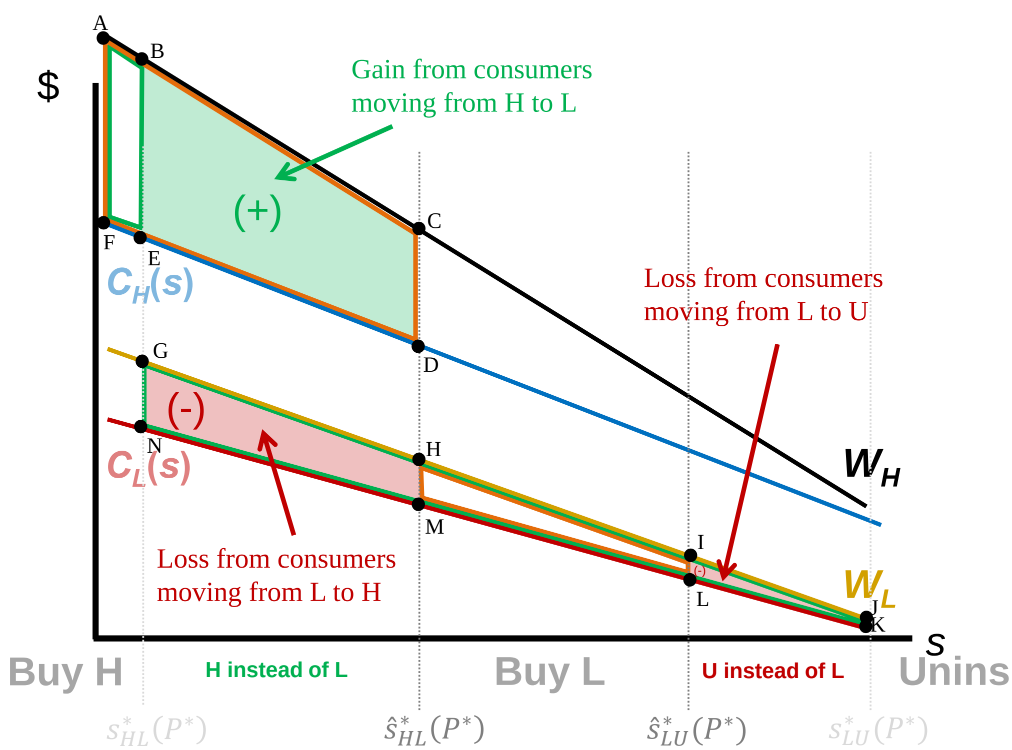

The figure highlights the difference in welfare with and without the risk adjustment, with welfare increases filled in green and welfare decreases filled in red. The welfare difference consists of two key pieces: (1) a shift in welfare due to the choice of some consumers to choose H instead of L (the intended effect of risk adjustment) and (2) a shift in welfare due to the choice of some consumers to choose U instead of L (the unintended effect of risk adjustment). The welfare impact of (2) is unambiguously negative as long as consumers value L more than its cost C_L, and the size of the loss is described by the area of $IJKL$. This area describes the unintended effect of the policy. The welfare impact of (1), on the other hand, is again ambiguous, and, as in the cases of benefit regulation and the uninsurance penalty, depends on whether social surplus (valuation minus cost) is larger for H vs. L for the marginal consumers. The size of the gain from (1) is equal to the area of $BCDE$ minus the area of $GHMN$. The overall welfare effect of benefit regulation is found by subtracting the areas of $GHMN$ and $IJKL$ and from the area of $BCDE$.

Perfect risk adjustment is not realistic, however. Instead, 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.) In the ACA Marketplaces, the transfer from L to H is determined according the following formula:[16]

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

where $\overline{R}_j(P^{cons})$ is the average risk score of the consumers enrolling in plan j given price vector P^{cons}, $\overline{R}(P^{cons})$ is the (share-weighted) average risk score among all consumers purchasing insurance given price vector P^{cons}, and $\overline{P}(P^{cons})$ 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 $\alpha \cdot T(P^{cons})$ 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 the price of H and the price of L. These comparative statics mimic the simulations we perform in the empirical section where we simulate equilibria under no risk ad-

[16] 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.

justment and with increasingly large risk adjustment transfers (i.e. increasingly large values for α). The comparative statics reveal that with imperfect risk adjustment, larger values of α (i.e. stronger transfers) unambiguously lower the price of H, as shown in the perfect risk adjustment case above.

The effect of an increase in α on the price of L, however, is ambiguous. The intuition for this result is that there are two effects of stronger imperfect risk adjustment: (1) the direct effect of taking money from L, raising L's risk-adjusted costs and leading to an increase in P_L and (2) the indirect effect of a higher P_L causing a set of the sickest consumers enrolling in L to opt instead to enroll in H, lowering L's risk-adjusted costs (to the extent that the transfers do not adjust for this shift in the composition of L's enrollees) and leading to an increase in P_L. The overall effect of an increase in α on P_L depends on whether the direct or the indirect effect dominates. Intuitively, the direct effect will dominate when the substitution effect is small, i.e. when the marginal consumers are either few in number or have costs that do not differ greatly from inframarginal consumers. On the other hand, the indirect effect will dominate when the substitution effect is large. In some cases, these effects may offset one another, leading to no extensive margin effect of increasing the strength of risk adjustment. In such settings, the regulator may be able to achieve welfare improvements by strengthening risk adjustment and shifting market share from L to H without the unintended consequence of forcing people out of the market.

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, this trade-off may or may not exist, depending on the size 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 the other, 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 states since the federal program expired.[17] To the extent reinsurance is implemented as a system of budget-neutral enforced transfers based 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 has affects similar to that of a mandate penalty.[18]

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 nuisance. 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 impacts plan choice in a way that affects the sorting of consumer risks (expected costs) across plans.

4 Simulations: Methods

To illustrate the trade-off between extensive and intensive margin selection we describe in Section 2, we perform a series of simulations. Our overarching goals are to (1) illustrate how our graphical model translates to empirical analaysis and (2) study the quantitative relevance of our theoretical points, using realistic demand and cost primitives. To do so, we draw on estimates from past work on health insurance exchanges that aligns closely with our vertical model (Hackmann, Kolstad and Kowalski, 2015; Finkelstein, Hendren and Shepard, 2017). In this section, we begin by giving an overview of the construction of the demand and cost curves and then describe our method for find-

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

[18] 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.

ing competitive equilibrium prices and allocations of consumers across plans, based on a reaction function approach.

4.1 Demand and Cost Estimates

For estimates of demand and cost, we draw on two recent empirical papers that estimate these primitives for Massachusetts' individual health insurance market. We use these papers' estimates to construct demand and cost curves for two groups that participate in the post-ACA individual market: low-income subsidized and high-income unsubsidized consumers. We will draw on these estimates to consider simulations either with just low-income subsidized consumers or with a mixed market that includes both types of consumers (with their relevant subsidy policies applied) as in the ACA Marketplaces.

Low-Income Subsidized Consumers: FHS (2017) For low-income consumers, we draw on estimates from Finkelstein, Hendren and Shepard (2017), which we abbreviate as "FHS." FHS study insurance demand in Massachusetts' pre-ACA subsidized health insurance exchange, known as Commonwealth Care or "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. Importantly, program 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.

FHS estimate demand and cost for CommCare in 2011. They argue that at this time the market featured a convenient vertical structure among competing plans. All plans follow the same (state-mandated) cost-sharing rules, but plans differ in the breadth of their provider networks, with plans falling into one of two groups: a more generous (broad-network) "H" option and less generous (narrower-network) "L" option. This two-plan vertical demand structure maps neatly into our vertical model, making these estimates a convenient way to parameterize our model.

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. The result of FHS's estimates are WTP curves ($W_L(s)$ and $W_H(s)$) and cost curves for H ($AC_H(s)$, $C_H(s)$) over an "in-sample" range of consumers, spanning the 31st to 94th highest percentile of the WTP distribution (i.e., $s \in [0.31, 0.94]$). We discuss additional details about the CommCare market and FHS's estimates in the appendix.

We extend the FHS estimates in two ways. First, we extrapolate their curves to generate estimates over the full $s \in [0, 1]$ range needed for our simulations. We consider two versions of this extrapolation: (1) a simple linear extrapolation (of demand and average costs, with marginal cost computed accordingly), and (2) an "enhanced demand" extrapolation that assumes a much higher WTP for the highest-WTP types ($s \in [0, 0.31]$). The final WTP curves are presented in Appendix Figure A1, and the details of these extrapolations are presented in the appendix.

Second, we need to produce estimates of $C_L(s)$ to complete the model. FHS provide suggestive evidence that $C_L(s)$ is quite similar to $C_H(s)$ – i.e., that for a given enrollee, L does not save money relative to H. 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 $C_L(s) = 0.85C_H(s)$. We also a consider an alternative setting where, consistent with the empirical evidence, L is a pure cream-skimmer, i.e. $C_L(s) = C_H(s)$.

High-Income Unsubsidized Consumers: HKK (2015) In a subset of our simulations, we consider the impact of unsubsidized (higher income) consumers selecting plans and impacting costs and risk adjustment transfers in the same market as the subsidized group. To do so, we must construct WTP curves for high-income households (who are not part of FHS's analysis). Rather than assuming similarity between the lower and higher income groups, we draw on estimates for individual-market health insurance coverage in Massachusetts from Hackmann, Kolstad and Kowalski (2015) ("HKK"). HKK estimate demand in the state's unsubsidized pre-ACA individual health insurance market for individuals with incomes above 300% of poverty (too high to qualify for CommCare). They do so using the introduction of the state's individual mandate in 2007-08 as a source of exogenous price variation to identify the insurance demand and cost curves.

The results of HKK's exercise are estimates of WTP curves for a single representative plan in the market. We map these estimates into into our two-plan vertical model by assuming that these represent estimates for the L plan – which we label $W_L^{HI}(s)$. We construct WTP estimates for H by simply adding the estimates of $W_{HL}(s) = W_H(s) - W_L(s)$ drawn from FHS. We also extrapolate HKK's estimates linearly outside of their sample range, following a similar procedure as for low-income consumers. We assume that the cost curves for the high-income consumers match the cost curves for the low-income consumers.[19] 4Further details of this method are discussed in the appendix.

4.2 Policy Simulations

We simulate a number of policy counterfactuals. In our simulations, we vary three policies: benefit design regulation, penalties for remaining uninsured, and risk adjustment transfers. Benefit design regulation is targeted at combating intensive margin selection problems by eliminating L. Penalties for remaining uninsured are targeted at extensive margin selection problems by raising the price of U. Risk adjustment transfers are targeted at combating intensive margin selection problems by lowering the price difference between H and L via transfers from L to H when H is adversely selected. The objective of our simulations is to show the unintended consequences of each policy for the margin of selection to which it is *not* targeted. We describe each policy simulation in detail below.

Our method for finding equilibrium is based on a reaction function approach. 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 modified version of the Riley equilibrium concept to choose which breakeven price vector is the equilibrium price vector.[20]

[19] An alternative approach would be to use the cost curve estimated bby HKK. We chose not to take this approach because the high-income consumers in the HKK setting are purchasing a different (and more expensive) set of products than the low-income consumers in the FHS setting, suggesting that the HKK cost curve id unlikely to provide a good picture of the high-income consumers' costs in the plans studied by FHS.

[20] See the appendix 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.

4.3 Subsidy Regimes

We consider each of the three policy simulations under two classes of subsidy regimes. We refer to the first class of subsidy regimes as "ACA Subsidies". ACA subsidies are linked to the price of the lowest price plan (L if there is positive enrollment in L and H if the market "upravels" to H). We follow the ACA rules by assuming that the subsidy for the lowest price plan is set such that the net-of-subsidy price of that plan is equal to \$55, or 4% of income for someone with an income equal to 150% of the federal poverty line (FPL) in 2011.The ACA subsidy rules actually set the subsidy according to the price of the second-lowest cost silver plan. Our 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. Consumers receive this subsidy if they choose to enroll in either H or L but not if they opt for U.

We consider two ACA subsidy regimes. The first focuses only on the low-income subsidized population. For this subsidy regime, we assume that the entire market is defined by the low-income enhanced ($W_L^{enh}(s)$, $W_H^{enh}(s)$) WTP curves, leaving results from simulations using the linear demand curves for the appendix. We thus effectively assume that there are no high-income unsubsidized consumers in the market. The second allows for high-income participants who do not receive subsidies. In this regime, the insurance market serves a population that is comprised of both high- and low-income individuals. Low income individuals, who comprise 60 % of the population, are eligible for the ACA subsidies as described in the previous subsidy regime. The cost and demand curves of these low-income individuals also remain the same as in the previous regime. High-income individuals, who comprise the remaining 40% of the population, differ from low-income individuals in two ways: they are ineligible for subsidies and have universally higher demand for insurance than the low-income population. The high-income demand curve is constructed from estimates found in Hackmann, Kolstad and Kowalski (2015) as described in Section 4.1. The cost curve for the high- and low-income populations remain the same so that a high- and a low-income individual of type s should have identical expected costs.[21]

We refer to the second class of subsidy regimes as "Fixed Subsidies". Under these regimes a fixed dollar amount is given towards either the purchase of the H or L plan. We simulate a number of different fixed subsidy amounts including the average cost of the entire population (\$321.55), \$300,

[21] We do not use the cost curve from Hackmann, Kolstad and Kowalski (2015) because their estimates come from a market with different products.

\$275, and \$250. Unlike the ACA subsidies, this amount is not linked to the offered price of either plan. For these simulations, all individuals are low-income types who are eligible for the subsidy.

We consider both fixed and ACA (price-linked) subsidies to show how the linkage between subsidies and prices under the ACA affects the intensive and extensive margin consequences of each of the policies. Clearly, if consumer subsidies are linked to the price of L, policies such as risk adjustment that target the intensive margin and may unintentionally lead to increases in P_L will have limited extensive margin consequences. This is due to the fact that if P_L increases by one dollar, the subsidy also increases by one dollar, resulting in no change in the consumer's price of L. ACA subsidies may thus soften the unintended extensive margin consequences of policies targeted at the intensive margin.

4.4 Outcomes

For all simulations, we report the competitive equilibrium price vector $P = P_H, P_L$ found using the reaction function approach described above. We also report the equilibrium market share for each plan and the subsidies for H and L.

5 Results

Here we document the intensive/extensive margin trade-off of common policies in terms of impacts on rates of uninsurance, consumer sorting between H and L, and plan prices. We defer evaluation of welfare impacts until Section 6, as the social surplus calculations require more stringent assumptions than the exercise of predicting market allocations and prices.

5.1 Benefit Regulation

The first policy we examine is the use of benefit regulation. The impact of benefit regulation is to weaken intensive margin selection either by forcing L to look more like H (thus, decreasing the valuation differential $W_{HL}(s)$ for all s) or by eliminating L altogether. As shown in Section 2, this type of policy may have the unintended consequence of pushing some consumers who are on the margin between L and U out of the market. Thus, benefit regulation may increase the quality of insurance coverage either by raising the quality of coverage offered by L or by moving marginal consumers from the lower-quality L plan to the higher-quality H plan, but it may also decrease the quality of coverage in the market by causing some consumers who would have otherwise enrolled in L to be uninsured.

To simulate benefit regulation, we assume that the regulator has the ability to identify and eliminate the L plan. In principle, this could be done via minimum actuarial value regulations, network adequacy rules, or Essential Health Benefits-type regulations. In practice, we simulate these regulations by computing equilibrium prices and market shares with a choice set that includes only H and U versus a choice set that includes H, L, and U. We assume that all other policy parameters (uninsurance penalties, risk adjustment transfers, etc.) remain constant across the two policy simulations.[22]

Table 1 reports the results of the benefit regulation simulations where we assume that L has a 15% cost advantage. Columns labeled "L offered" present outcomes for the setting where consumers can choose between H, L, and U. Columns labeled "No L" present outcomes for the setting where the regulator eliminates L from the consumer's choice set. The first two columns represent the ACA subsidy regime where we assume there are no unsubsidized consumers and the subsidy is set such that the net-of-subsidy price of the lowest-priced plan is \$55, consistent with the affordability standard used to set subsidies under the ACA. The next two columns add high-income unsubsidized consumers. The next two columns assume that subsidies are fixed rather than a function of plan prices and equal to the average cost across all consumers in the market at baseline (\$323 per month). The bottom panel of the table presents outcomes for additional levels of fixed subsidies: \$300 per month, \$275 per month, and \$250 per month. All settings include baseline ACA risk adjustment transfers and no penalty for choosing U.

With ACA subsidies and no unsubsidized consumers (columns 1 and 2), when L is offered, 57% of the market chooses L, 4.5% of the market chooses H, and 39% of the market opts to be uninsured. When L is removed from the consumer's choice set, all of the consumers who chose L when it was offered opt instead for H. Additionally, 9% of consumers chose U when L was offered but opt for H when L is removed from the choice set. These shifts in market share imply unambiguous improvements in insurance coverage in the market, with more consumers choosing H and fewer consumers choosing U. Here, there is no apparent trade-off between the two margins. Indeed, benefit regulation improves outcomes on both margins. However, this result is due to the linkage of the subsidy to the

[22]Risk adjustment, which is an intensive-margin policy as it is used in insurance Exchanges/Marketplaces, is degenerate in the case of only one plan.

price of the lowest-price plan: When L is eliminated from the choice set, the subsidy increases because it is now linked to the (higher) price of H, inducing more consumers to enter the market. In other words, the fact that the subsidy varies with the binary benefit regulation (with no limit on government outlays) drives this result. Similar effects of benefit regulation are observed when high-income unsubsidized consumers are added to the market (columns 3 and 4).

We now turn to the results for our fixed subsidy regimes. When the subsidy is set equal to the average cost across all consumers in the market ($323) and with L and H both offered, no consumers opt to remain uninsured. But the market unravels to L—that is, no consumers enroll in H. When L is removed from the choice set, again no consumers opt for U, so that now all consumers choose H, indicating unambiguous improvements in coverage generosity in the market. These same results hold with a $300 per month subsidy. When the fixed subsidy is lowered further to $275 per month, the market still fully unravels when L is offered, but now when L is removed from the choice set, only a relatively small portion of consumers (28%) opt to enroll in H. The rest of the consumers in the market (72%) opt to exit the market and remain uninsured. Similar effects (though smaller in magnitude) are observed with a subsidy equal to $250 per month. These settings clearly illustrate the trade-off between extensive and intensive margin selection, with the elimination of L having both its intended consequence (shifting enrollment from L to H) and the unintended consequence of increasing the uninsurance rate.

Results for the case where L has no cost advantage and is thus a perfect cream-skimmer are found in Table 2. Here, we find similar results for the ACA subsidies cases, where removing L from the choice set does not increase the uninsurance rate and sometimes decreases it. With fixed subsidies, we find that the market almost always "upravels" to H, with no consumers choosing L even when it is offered, leaving little scope for benefit regulation to affect the market equilibrium.

5.2 Mandate/Uninsurance Penalties

Next, we simulate the effects of an insurance mandate or penalty for choosing U. The economic efficiency purpose of a mandate/uninsurance penalty is to weaken extensive margin selection by raising the price of U, thus lowering the *relative* price of L and inducing marginal consumers who would otherwise choose to remain uninsured to enroll in L.[23] As shown in Section 3, because the

[23]If there are significant external costs of uninsurance, there is an economic efficiency rationale for subsidies even in the absence of too-high prices for L caused by adverse selection into the market.

consumers induced to enroll in L instead of U are healthier than the inframarginal L enrollees, this type of policy may also lower the price of L relative to H, thus causing a group of marginal consumers to enroll in L instead of H. Thus, uninsurance penalties may improve insurance coverage in the market by inducing some consumers to purchase insurance, but may simultaneously decrease the quality of insurance coverage in the market by inducing other consumers to move from the higher-quality H plan to the lower-quality L plan.

To simulate the effects of a mandate/uninsurance penalty, we change the price of uninsurance P_U. We find the equilibrium prices for L and H given a fixed price of uninsurance (i.e. penalty) of $P_U = 0, 5, 10, ..., 60$ per month.[24] The results for the uninsurance penalty simulations are presented in Figure 10 and Table 3. Each panel of the figure represents a different subsidy regime, and the figure plots 3 lines, each showing how the market share of one of the three plan options (H, L, or U) changes as the uninsurance penalty is increased from \$0 per month to \$60 per month. Prices, market shares, and the subsidy are reported in Table 3.

The top-left panel of Figure 10 shows how market shares are affected by an increasingly strong mandate under ACA subsidies with no unsubsidized consumers. As expected, as the penalty increases, the fraction of uninsured consumers drops. With no penalty, 39% of consumers choose to remain uninsured, but with a penalty of \$60 per month, all consumers in the market choose to purchase insurance. However, the \$60 per month penalty also causes the market to unravel to L: The 4.5% of consumers who would opt for H when there is no penalty opt for L when a \$60 penalty is assessed.

Although a shift of 4.5% of consumers from H to L is small and perhaps unlikely to have a first-order effect on welfare, the top-left panel of Figure 11 shows that the unintended extensive margin effects of a \$60 uninsurance penalty are quite significant in the setting where L has no cost advantage and is thus a pure cream skimmer. There, when there is no penalty, the market upravels to H and 70% of the market enrolls in H while 30% opts out of the market and remains uninsured. With a penalty of \$60, however, all consumers enter the market (i.e. U's share goes to zero), but H's market share decreases from 70% to 24%. Here, the downward shift in the quality of coverage in the market is significant with 46% of consumers choosing L instead of H. The offsetting improvements in coverage from consumers choosing L instead of U affect only 30% of consumers. Again, results are

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

similar when high-income unsubsidized consumers are added to the market. These results clearly illustrate the trade-off between extensive and intensive margin selection: a strong mandate causes some consumers to be better off by choosing L instead of U but also causes other consumers to be worse off by choosing L instead of H.

The bottom 4 panels of Figure 10 present outcomes for the fixed subsidy cases. For the first three fixed subsidies (average cost in the population, \$300, \$275), there is no effect of the uninsurance penalty because the subsidy is high enough that all consumers choose to enroll in insurance even without a penalty. With a subsidy of \$250 per month, when there is no penalty 66% of consumers opt to remain uninsured while 9.2% choose to enroll in H and 24% choose to enroll in L. As the penalty is increased, there are again two effects: Fewer consumers choose U and fewer consumers choose H. A penalty of \$35 is enough to cause all consumers to enroll in insurance, but it also causes the market to unravel to L with no consumers choosing H. Figure 11 shows that these competing effects are even more significant when L has no cost advantage. This once more illustrates the extensive/intensive margin trade-off.

5.3 Risk Adjustment

Finally, we simulate the effects of risk adjustment transfers. The intention of these transfers is to weaken intensive margin selection by transferring money from L to H and flattening the H and L risk adjusted average cost curves. As shown in Section 3, the transfers may also affect the sorting of consumers between L and U, as the transfer from L to H may raise the price of L, inducing consumers on the L versus U margin to exit the market.[25] Thus, risk adjustment transfers may increase quality of coverage in the market by moving some consumers from L to H, but they may simultaneously decrease quality of coverage in the market by inducing other consumers to exit the market and opt for U.

To simulate the effects of risk adjustment, we vary the strength of the risk adjustment transfers by introducing a strength parameter α to the ACA risk adjustment transfer formula from Eq. (9):

$$T\left(P^{cons}\right) = \left[\left(\frac{\overline{R}_H(P^{cons}) - \overline{R}_L(P^{cons})}{\overline{R}(P^{cons})}\right) \cdot \overline{P}(P^{cons})\right] \times \alpha \tag{10}$$

[25] Recall from Section 3 that imperfect risk adjustment need not always raise the price of L due to the presence of the offsetting "direct" and "substitution" effects of the transfers.

Values of α smaller than 1 represent risk adjustment policies that are weaker than the baseline ACA policy, while values of α larger than 1 represent policies that are stronger than the baseline policy. The introduction of $\alpha < 1$ by us is similar to recent guidance provided to states by CMS to adjust the size of risk adjustment transfers down according to the conditions present in the state's local market. With the exception of the α parameter, we hold the rest of the transfer formula constant across policy simulations. We find equilibrium prices and market shares for $\alpha = 0, 0.1, 0.2, ...3$.[26]

The results for the risk adjustment simulations are presented in Figure 12 and Table 5. As with the uninsurance penalty simulations, each panel of the figure represents a different subsidy regime. Again, we plot 3 lines, with each line showing how the market share of one of the three plan options changes as we strengthen the risk adjustment transfers by increasing α from 0 (no risk adjustment) to 3 (transfers three times the size of the ACA transfers), as described by Equation 10. Prices, market shares, and subsidies are reported in Table 5.

The top-left panel of Figure 12 shows how market shares are affected by increasingly strong risk adjustment under ACA subsidies with no unsubsidized consumers. Here, we see that with no risk adjustment H unravels, with no enrollment in equilibrium. Consumers split between L (61%) and uninsurance (39%). As risk adjustment is strengthened, eventually the transfers become large enough such that H can be offered at a price that induces positive enrollment. Despite inducing some enrollment in H (4.5%), ACA-strength risk adjustment ($\alpha = 1$) has little effect on the price of L, suggesting that the direct and substitution effects of strengthening the transfers offset one another. This offsetting continues until the transfers become large enough to cause the market to "upravel" to H: At risk adjustment transfers two times the size of the ACA transfers, there are no consumers enrolling in L. At the point where the market upravels, there is a small decrease in the uninsurance rate as well: A group of marginal uninsured consumers enter the market because once L is eliminated the price-linked subsidy becomes attached to the higher priced H plan, causing the subsidy to increase. Similar results are obtained when we add higher-income unsubsidized consumers to the market, as seen in Panel B of Figure 12. We also obtain similar results when we assume that L has no cost advantage, as presented in Figure 13 and Table 6.

Thus, with ACA price-linked subsidies, risk adjustment addresses the intensive margin selection problem while not making things worse on the extensive margin. In this case there is no apparent

[26] We find that in all cases the equilibrium price vector does not change with increases in α for $\alpha > 3$.

trade-off between extensive and intensive margin selection: Risk adjustment can improve the quality of coverage in the market without pushing anyone out of the market. The reason that this finding does not contradict the arguments in Section 3 is that under the mechanical rules of the ACA subsidy scheme, any increase in P_L automatically induces an identical increase in the subsidy, thus leaving the consumer's net price of L unchanged. There is, in fact, an implicit trade-off here: In order to hold the uninsurance rate steady, the regulator (automatically) ratchets up subsidies—with ambiguous effects on welfare, as we address in Section 6.

The bottom four panels of 12 present the fixed subsidy cases. When the fixed subsidy is equal to the average cost in the market, we see that the subsidy is large enough to drive the uninsurance rate to zero under all levels of α. Again we see that with no risk adjustment, H unravels, with no consumers enrolling in H in equilibrium. As risk adjustment transfers are strengthened, consumers begin to opt for H instead of L. At a level of α of 1.5, the market "upravels" to H.[27] With smaller fixed subsidies of \$300 per month and \$275 per month, we find that when there is no risk adjustment, again the market unravels to L and no consumers choose U. As risk adjustment transfers are strengthened, however, there are two consequences. First, some consumers start to choose H instead of L (33% of the market under the \$300 subsidy), as in the previous subsidy regimes. Second, some consumers exit the market and opt to be uninsured (52% of the market under the \$300 subsidy). The source of this second, extensive margin effect is clear when observing the equilibrium price of L in Table 5: With no risk adjustment $P_L = \$273$ for the \$300 and \$275 subsidy cases, but with risk adjustment two times the strength of ACA risk adjustment (i.e. $\alpha = 2.0$) the P_L increases to \$383 and \$395, respectively. Here, the direct effect of stronger transfers dominates the indirect effect, resulting in large increases in P_L and large increases in the portion of the market choosing U. With a smaller subsidy of \$250, we observe a similar shift from L to H accompanied by a similar increase in the portion of consumers opting to remain uninsured, though to a lesser degree than with the larger subsidies. These results clearly illustrate the two effects of risk adjustment transfers: (1) the intended consequence of shifting consumers from L to H and (2) the unintended consequence of causing some consumers to exit the market and opt for U.

Figure 13 and Table 6 present results for the case where L has no cost advantage. Here, we find that stronger risk adjustment shifts consumers from L to H as before, but we find no effect of risk

[27] For $\alpha > 1.5$ and with fixed subsidies equal to average costs, there is again no extensive margin effect of risk adjustment. Here, the lack of an extensive margin effect is due to the subsidy being so high that no consumer ever wants to be uninsured.

adjustment on the extensive margin, i.e. stronger risk adjustment does not cause consumers to exit the market. Thus, the substitution effect (the sickest L consumers shifting to H) offsets the direct effect (transfer of money from L to H), such that increasingly strong risk adjustment does not cause the price of L to increase. Instead, stronger risk adjustment only causes the price of H to decrease. This is a novel finding consistent with the comparative statics we present in Appendix B: When L is a cream skimmer, the unintended extensive margin effects of stronger risk adjustment are minimal, implying that consumers can be shifted from the cream-skimming L plan to the higher-quality H plan without forcing other consumers out of the market.

6 Welfare

We have shown that in a variety of settings, there appears to be a trade-off between selection on the intensive (H versus L) and extensive (insurance versus uninsurance) margins in terms of prices and enrollment allocations. As described in Section 3, our graphical model also presents a clear way forward to assessing welfare consequences of this trade-off. Overall social surplus can be determined by summing the area between $W_H(s)$ and $C_H(s)$ for people choosing H; the area between $W_L(s)$ and $C_L(s)$ for people choosing L; and the (negative) area between $W_U(s) = 0$ (a normalization without substantive implications) and $C_U(s)$ for people choosing U. Fortunately, we have all of the pieces of this problem, with the exception of $C_U(s)$. This missing cost term includes items like uncompensated care, care paid for by other state programs, or more difficult-to-measure parameters like a social preference against others being uninsured. It is necessary for considering the full welfare consequences of the policies we consider.

6.1 Cost of Uninsurance

To characterize welfare under a range of possible values of $C_U(s)$, we assume that $C_U(s)$ is made up of two components, a fixed component and a component that is proportional to a consumer's costs in H. Formally, we assume that $C_U(s) = \sigma + \phi C_H(s)$, where $\sigma = \$0, \$50, \$100$ and $\phi = 0, 0.25, 0.5, 0.75, 1$. This implies that s-types with higher insured costs also generate a higher social cost when uninsured. Thus, we estimate total welfare according to Equation 7 for a variety of specifications of the social cost of uninsurance, including zero social cost of uninsurance ($\phi = 0, \sigma = 0$) and a social cost of uninsurance equal to the cost of insuring the consumer in Plan H ($\phi = 1, \sigma = 0$). Higher values of ϕ and σ correspond to more weight being placed on the welfare of the uninsured. The range of costs explored covers any plausible unobserved value of $C_U(s)$. This exercise therefore allows the reader to assess the welfare consequences of a given policy given her particular priors on ϕ and σ.

For ease of interpretation, rather than reporting the dollar denominated estimate of total social surplus, we report for each subsidy structure (B) and L cost advantage (c) the percent of the gap between minimum welfare and maximum welfare achieved at the policy parameter values (α or the uninsurance penalty) value p:

$$w_{pcB} = \frac{W_{pcB} - W_{cs}^{min}}{W_{cs}^{max} - W_{cs}^{min}} \tag{11}$$

where W_{pcB} is total social welfare and w_{pcB} is our relative welfare measure.

We note that under the primitives from the setting we study, all consumers value L less than the cost of enrolling them in L. Finkelstein, Hendren and Shepard (2017) argue that this is likely due to the presence of uncompensated care and other state programs that subsidize health care for the uninsured. The practical implication of this for our simulations is that when the social cost of uninsurance is low (i.e. low ϕ and σ) welfare will be *increasing* in the uninsurance rate. When the social cost of uninsurance is high, on the other hand, welfare will be *decreasing* in the uninsurance rate.

The welfare consequences of consumers shifting from L to H do not depend on the social cost of uninsurance. They do, however, depend on whether L has a cost advantage or is a pure cream-skimmer. If L is a pure cream-skimmer, then welfare is *increasing* in H's share of insured consumers. If L has a 15% cost advantage, however, most consumers' incremental valuation of H vs. L is less than the incremental cost of enrolling them in H vs. L. In the welfare figures in Section 3, this would appear as the welfare gains from consumers enrolling in H being smaller than the welfare losses from those same consumers leaving L. This implies that when L has a 15% cost advantage, welfare is *decreasing* in H's share of insured consumers. These relationships are critical to keep in mind when interpreting our welfare estimates and applying the insights to other market contexts.

6.2 Welfare Effects of Policies in Isolation

6.2.1 Benefit Regulation

We first present welfare estimates that correspond to the market share allocations described in Section 5, beginning with benefit regulation. We then consider interactions between the uninsurance penalty and the strength of risk adjustment and consider optimal policy combinations.

Tables 1 and 2 contain welfare estimates of benefit regulation for the case when L has a 15% cost advantage and for the case when L has no cost advantage, respectively. Table 2 shows that when L does not have a cost advantage, under ACA subsidies welfare is weakly higher when L is not in the choice set than when L is available. Recall that with ACA subsidies, the subsidy is linked to the price of the low priced plan, so when L is removed from the market the subsidy becomes linked to the higher price of H, resulting in a higher subsidy and a lower uninsurance rate. This result does not hold when L has a 15% cost advantage. There, welfare is only higher under benefit regulation when the social cost of uninsurance is large. This is due to the fact that when L has a 15% cost advantage, welfare is decreasing in the share of consumers choosing H, so only when C_U is large do the welfare gains from the decrease in the uninsurance rate offset the welfare losses from the shift of consumers from L to H.

Tables 1 and 2 also present results for the fixed subsidy cases. When L does not have a cost advantage (Table 2), benefit regulation only binds in the case where the fixed subsidy is equal to the average cost in the population. In that case, welfare is higher when L is removed from the market, as the effect of the removal of L is to move consumers from lower-surplus L to higher-surplus H.

6.2.2 Uninsurance Penalty

Figures 14 and 15 present welfare estimates for different levels of the uninsurance penalty where L has a 15% cost advantage and where L is a pure cream-skimmer, respectively. Each panel of the figure represents a different subsidy regime, and each line represents a different value of ϕ. All figures assume $\sigma = \$50$. Welfare estimates for additional values of σ are found in Tables 3 and 4.

The results presented in Figure 15 reveal that when L has a 15% cost advantage, welfare is often decreasing in the size of the uninsurance penalty under the ACA subsidy regimes. Only when ϕ is large does a larger mandate lead to welfare improvements, and those welfare improvements are small. This suggests that the welfare gains from the intended effect of the penalty (moving consumers from U to L) are often more than fully offset by welfare losses from the unintended effect of the policy (moving consumers from H to L). This is likely due to the fact that in this setting the consumers on the margin of U vs. L place very low valuation on insurance coverage, while the consumers on the margin of H vs. L place high incremental valuation on H vs. L.

This result does not hold, however, when L is assumed to have a 15% cost advantage. Figure 14 shows that in that setting under ACA subsidies welfare is improving in the size of the mandate penalty except for the lowest values of ϕ. This is because wiith a 15% cost advantage, the unintended effect of the mandate (moving consumers from H to L) is actually welfare improving, as it is inefficient for any consumer to enroll in H. Thus, only in the case where it is socially efficient for most consumers to be uninsured (small ϕ) does the penalty lead to lower levels of welfare.

Moving to the fixed subsidy regimes we see that when L has a 15% cost advantage, the mandate is binding for all levels of the subsidy except for the highest level. Here, welfare is increasing in the size of the penalty for high levels of ϕ (when insurance is socially efficient) and decreasing in the size of the penalty for low levels of ϕ (when insurance is socially inefficient). When L does not have a cost advantage, the penalty only binds for the lowest subsidy level. At that subsidy level, the results are similar, with welfare increasing in the size of the penalty for high levels of ϕ and decreasing in the size of the penalty for low levels of ϕ.

6.2.3 Risk Adjustment Transfers

Figures 16 and 17 present welfare estimates for different strengths (i.e. levels of α) of the risk adjustment transfers where L has a 15% cost advantage and where L is a pure cream-skimmer, respectively. $\alpha = 0$ corresponds to no risk adjustment, $\alpha = 1$ corresponds to ACA risk adjustment, $\alpha = 2$ corresponds to transfers twice as large as ACA risk adjustment transfers, and so on. Each panel of the figure represents a different subsidy regime, and each line represents a different value of ϕ. All figures assume $\sigma = \$50$. Welfare estimates for additional values of σ are found in Tables 5 and 6.

The results in Figure 17 indicate that with ACA subsidies, welfare is almost always increasing in the strength of the risk adjustment transfers when L has no cost advantage. This is because the link between the subsidy and the price of L essentially eliminates any unintended extensive margin effects of strengthened risk adjustment transfers, meaning the only effect of stronger risk adjustment is a shift of consumers from L to H which is welfare improving when L is a pure cream-skimmer. The opposite is true when L has a cost advantage, as indicated by the results in Figure 16. Here, stronger risk adjustment still shifts consumers from L to H but in this case this shift in market share is socially inefficient because the social surplus from L exceeds the social surplus from H, causing welfare to be decreasing in H's market share among the insured.

With fixed subsidies, welfare is again always weakly increasing in the strength of the risk adjustment transfers when L has no cost advantage. This is because in this setting, the direct and substitution effects perfectly offset each other, causing strong risk adjustment to have no effect on the price of L. This result does not hold for the case where L has a 15% cost advantage, however. Figure 14 reveals that when L has a cost advantage, welfare is increasing in α for low levels of ϕ and decreasing in α for high levels of ϕ. Here, there are two effects of increasing α: (1) some consumers shift from L to H and (2) other consumers shift from L to U. When L has a 15% cost advantage, effect (1) always leads to lower welfare because the social surplus from L exceeds the social surplus from H. Effect (2), on the other hand, increases welfare when ϕ is small (and insurance is inefficient) and decreases welfare when ϕ is large (and insurance is efficient). Thus, when ϕ is small, the welfare gains from moving consumers out of inefficient insurance sometimes offset the welfare losses from moving consumers out of inefficiently generous coverage (H). When ϕ is large, on the other hand, stronger risk adjustment unambiguously decreases welfare because the two effects of risk adjustment (moving consumers from H to L and moving consumers from L to U) are both welfare-decreasing.

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: optimal extensive margin policy will often depend on the intensive margin policies 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. We consider the most commonly used extensive margin policy, an uninsurance penalty, and the most commonly used intensive margin policy, risk adjustment transfers. We compute social welfare for a grid of uninsurance penalties and levels of α. We do this separately for each subsidy regime and each (ϕ, σ) pair. We repeat the procedure with and without assuming a cost advantage for L. For parsimony we "cherry-pick" for discussion one case that clearly illustrates the interaction between extensive and intensive margin policies, and we relegate results for all other cases to the appendix.

The case we highlight is the one in which L has no cost advantage, there is a fixed subsidy equal to \$275 per month, $\phi = 0.5$, and $\sigma = 100$. There is thus a moderate social cost of uninsurance, making it socially efficient for most consumers to be enrolled in insurance. With these parameters there is more total surplus generated by consumers enrolling in H than in L. Figure 18 presents the welfare estimates graphically as a heat map, where darker areas represent higher levels of welfare. The figure shows that in this setting, when risk adjustment is strong (α is high), welfare is increasing in the mandate. When risk adjustment is weak, however, welfare is *decreasing* in the mandate. Recall that in this particular setting, the socially efficient allocation of consumers across H, L, and U is to have all consumers enrolled in H. The intuition for this result is thus that when risk adjustment is strong, the market "upravels" to H and no consumers enroll in L. 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 unintended intensive margin effect of risk adjustment 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 18 to consider the optimal level of α for each level of the uninsurance penalty. Here, there is no ambiguity: Welfare is always increasing in the strength of the risk adjustment transfer. Here, it appears that the welfare losses from the unintended extensive margin effects of risk adjustment never offset the welfare gains from the intended intensive margin effects. This may be due to the incremental H vs. L social surplus for the H vs. L marginal consumers being large relative to the social surplus from L for the L vs. U marginal consumers. It may also be due to the direct effect of the increased transfers away from L being offset by the substitution effect of the sickest L enrollees shifting from L to H, resulting in little or no change in P_L and little or no extensive margin effect of increasing α.

Figure 18 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 the socially efficient outcome. Given alternative levels of ϕ and σ and alternative assumptions about L's cost advantage, the socially efficient outcome will differ from all consumers being enrolled in H. In these settings, the optimal combination of policies will also differ. This can be seen in the heat maps presented in the appendix, where in some cases optimal policy consists of strong risk adjustment and a weak mandate and in other cases optimal policy consists of weak risk adjustment and a strong mandate.

Our goal in this section is not to make general statements about optimal policy. Instead, our goal is to illustrate how our model can be used to determine optimal policy given market primitives and to show how extensive and intensive margin policies interact. Policymakers need to be aware of these interactions and consider them when determining regulation for a given market.

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 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) eliminating cream-skimming plans can help some consumers by increasing the quality of their coverage while hurting other consumers by forcing them out of the market; (2) 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 (3) 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 these trade-offs 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, though for most (but not all) policies the welfare consequences depend critically on the social cost of uninsurance, which is unobserved.

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.

Cawley, John, and Tomas Philipson. 1999. "An empirical examination of information barriers to trade in insurance." *American Economic Review*, 89(4): 827–846.

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, and Kathleen McGarry. 2006. "Multiple dimensions of private information: evidence from the long-term care insurance market." *American Economic Review*, 96(4): 938–958.

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.

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

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

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.

Figure 10: Plan Market Shares under Increasingly Larger Uninsurance Penalty, 15% L Cost Advantage

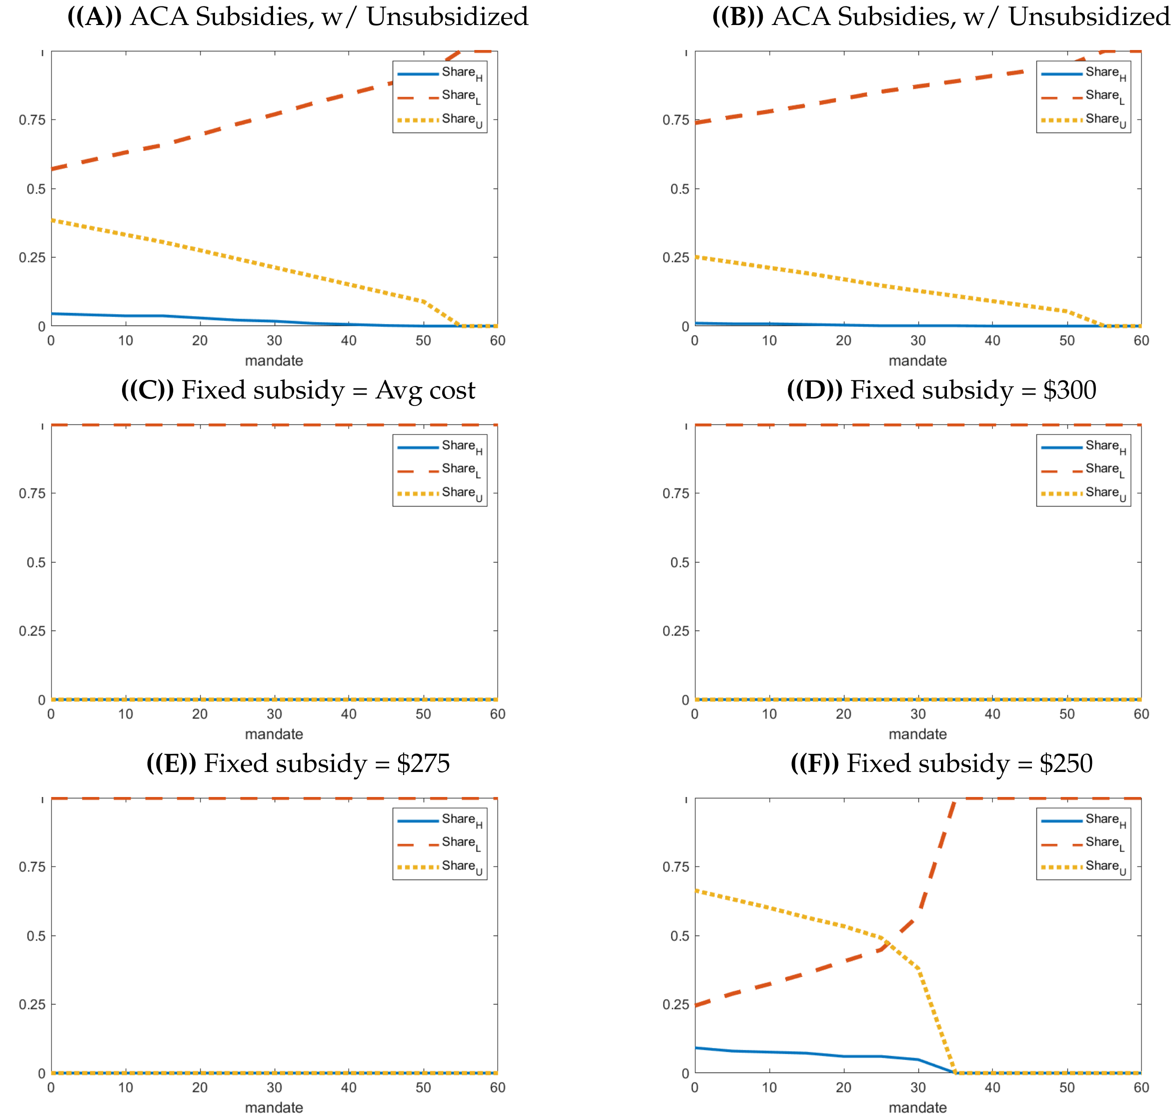

Note: Figures show equilibrium market shares of H, L, and U under different penalties for chooding U. Equilibrium prices are found using the reaction function approach described in the text where we simulate costs and market shares for a large grid of H and L prices and then find the pair of prices where both H and L break even and for which there is no Riley Deviation. Each panel describes a different subsidy regime. The "ACA Subsidies, No Unsubsidized" considers only demand of low-income individuals and assumes that all individuals receive a subsidy linked to the price of the lowest-price plan such that the net-of-subsidy price is $55. The "ACA Subsidies, w/ Unsubsidized" case adds high-income individuals (who have different demand) who do not receive subsidies. The fixed subsidy cases assume that all individuals receive the same fixed subsidy that is not a function of plan prices, with the "avg cost" case setting the subsidy equal to the average cost across all consumers in the population.

Figure 11: Plan Market Shares under Increasingly Larger Uninsurance Penalty, No L Cost Advantage

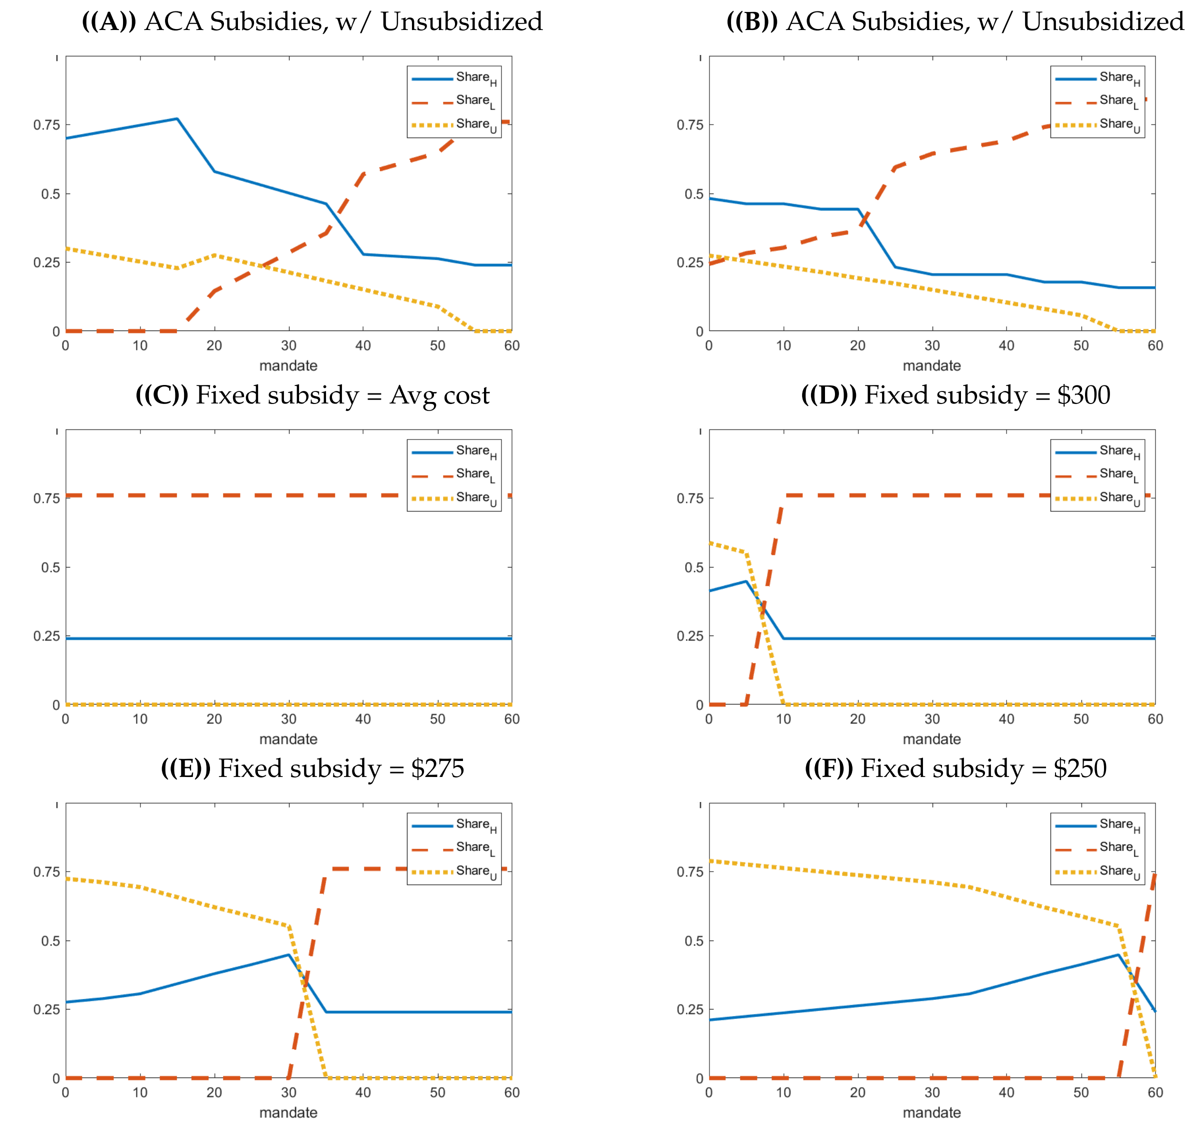

Note: Figures show equilibrium market shares of H, L, and U under different penalties for choosing U. Plan L is assumed to have a 15% cost advantage for these simulations. Equilibrium prices are found using the reaction function approach described in the text where we simulate costs and market shares for a large grid of H and L prices and then find the pair of prices where both H and L break even and for which there is no Riley Deviation. Each panel describes a different subsidy regime. The "ACA Subsidies, No Unsubsidized" considers only demand of low-income individuals and assumes that all individuals receive a subsidy linked to the price of the lowest-price plan such that the net-of-subsidy price is $55. The "ACA Subsidies, w/ Unsubsidized" case adds high-income individuals (who have different demand) who do not receive subsidies. The fixed subsidy cases assume that all individuals receive the same fixed subsidy that is not a function of plan prices, with the "avg cost" case setting the subsidy equal to the average cost across all consumers in the population.

Figure 12: Plan Market Shares under Increasingly Strong Risk Adjustment Transfers (i.e. Higher α), 15% L Cost Advantage

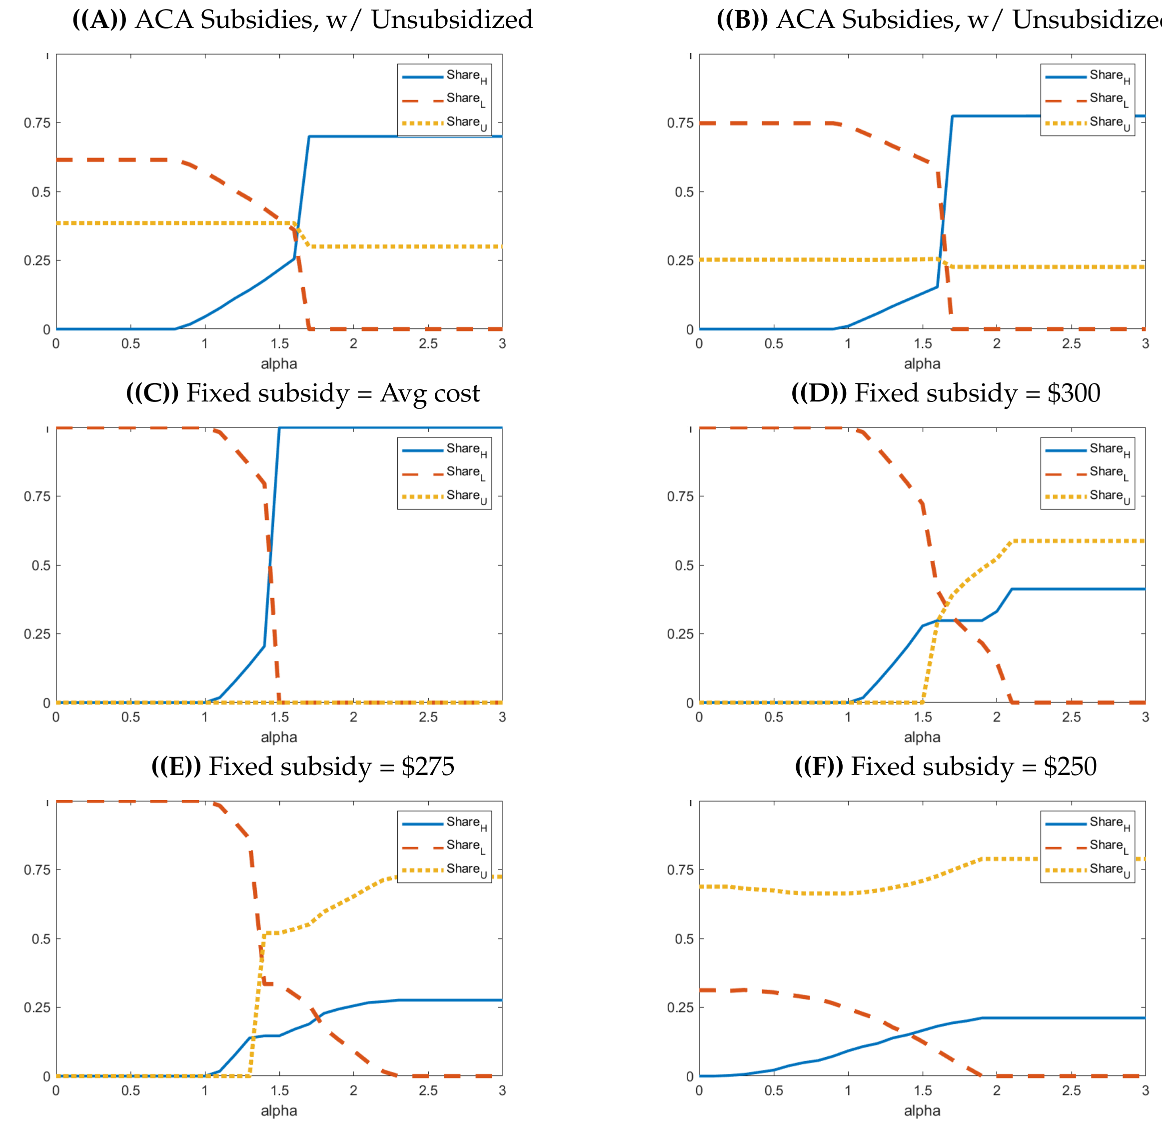

Note: Figures show equilibrium market shares of H, L, and U under different levels of α which multiplies the standard ACA risk adjustment transfers, as described in Equation 10 in the text. Equilibrium prices are found using the reaction function approach described in the text where we simulate costs and market shares for a large grid of H and L prices and then find the pair of prices where both H and L break even and for which there is no Riley Deviation. Each panel describes a different subsidy regime. The "ACA Subsidies, No Unsubsidized" considers only demand of low-income individuals and assumes that all individuals receive a subsidy linked to the price of the lowest-price plan such that the net-of-subsidy price is $55. The "ACA Subsidies, w/ Unsubsidized" case adds high-income individuals (who have different demand) who do not receive subsidies. The fixed subsidy cases assume that all individuals receive the same fixed subsidy that is not a function of plan prices, with the "avg cost" case setting the subsidy equal to the average cost across all consumers in the population.

Figure 13: Plan Market Shares under Increasingly Strong Risk Adjustment Transfers (i.e. Higher α), No L Cost Advantage

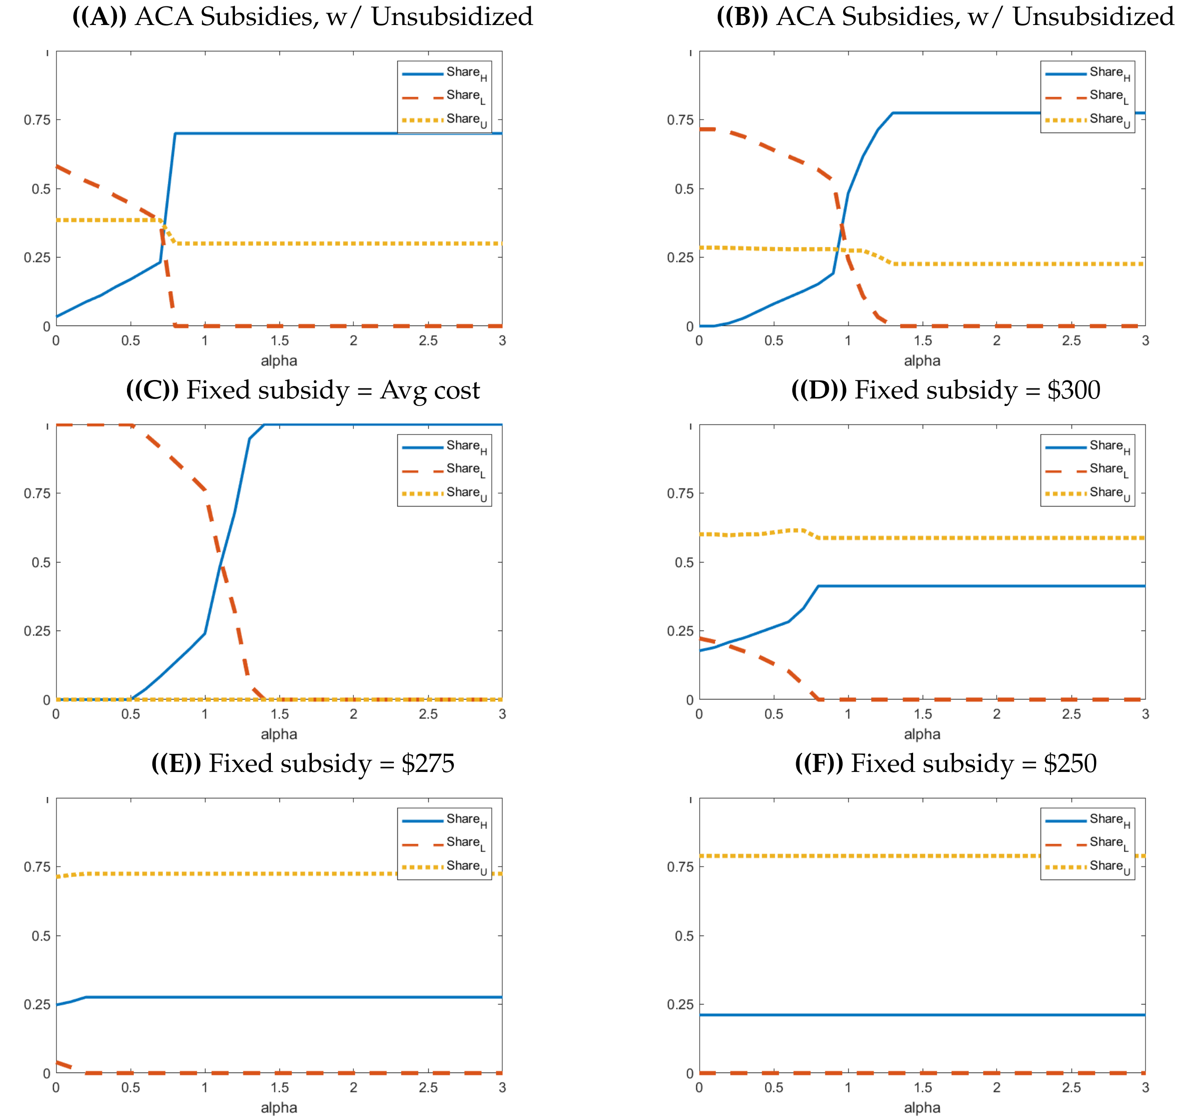

Note: Figures show equilibrium market shares of H, L, and U under different levels of α which multiplies the standard ACA risk adjustment transfers, as described in Equation 10 in the text. Plan L is assumed to have a 15% cost advantage for these simulations. Equilibrium prices are found using the reaction function approach described in the text where we simulate costs and market shares for a large grid of H and L prices and then find the pair of prices where both H and L break even and for which there is no Riley Deviation. Each panel describes a different subsidy regime. The "ACA Subsidies, No Unsubsidized" considers only demand of low-income individuals and assumes that all individuals receive a subsidy linked to the price of the lowest-price plan such that the net-of-subsidy price is \$55. The "ACA Subsidies, w/ Unsubsidized" case adds high-income individuals (who have different demand) who do not receive subsidies. The fixed subsidy cases assume that all individuals receive the same fixed subsidy that is not a function of plan prices, with the "avg cost" case setting the subsidy equal to the average cost across all consumers in the population.

Figure 14: Social Welfare under Increasingly Larger Uninsurance Penalty, 15% L Cost Advantage, $\sigma = \$50$

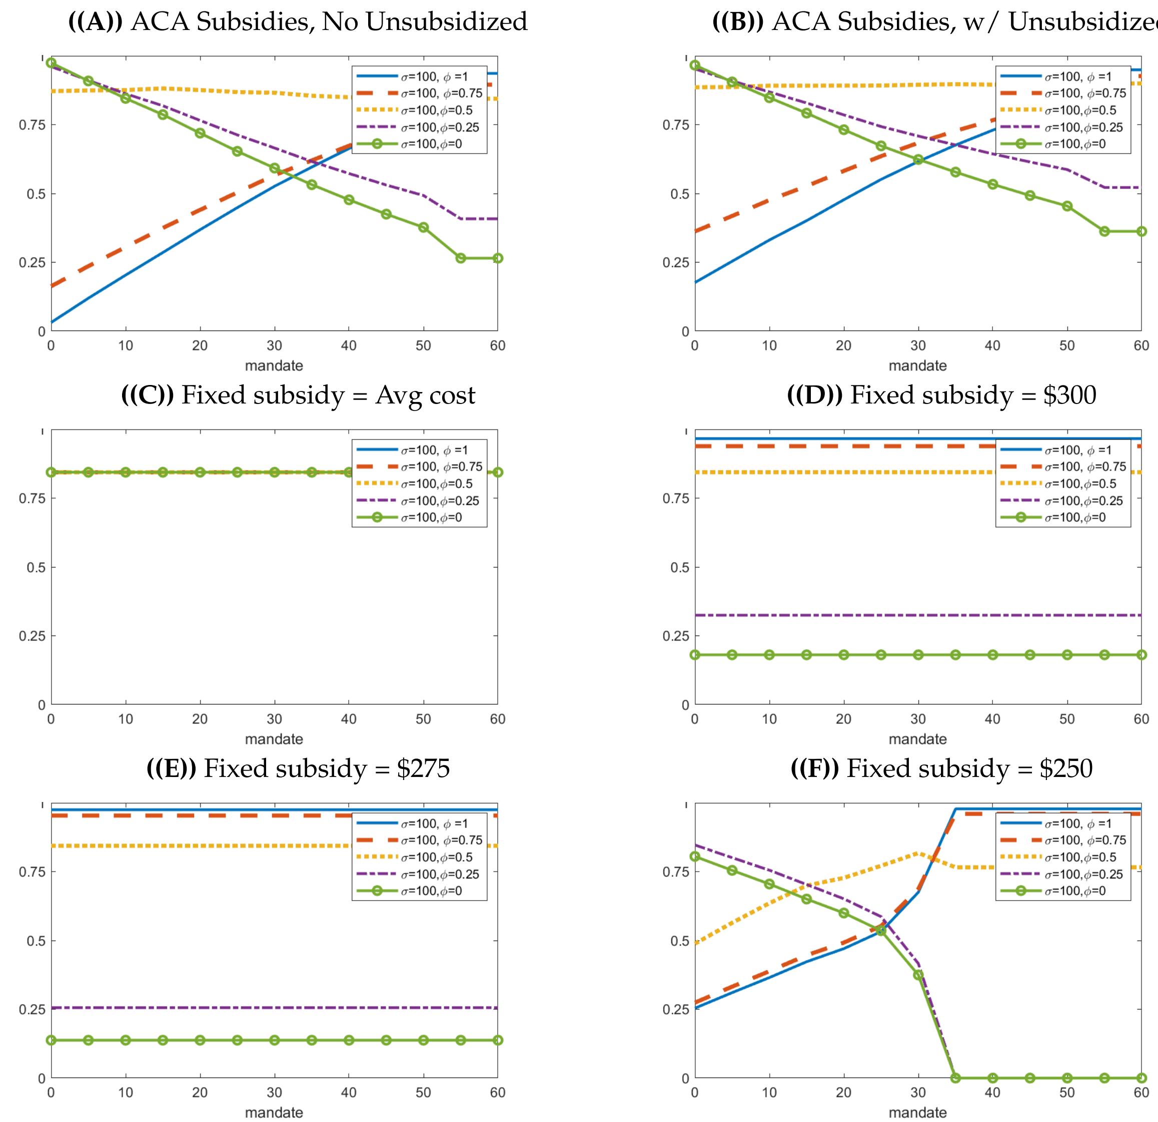

Note: Figures show estimates of overall social welfare as described by Equation 7 under increasingly larger uninsurance penalties. Each line represents a different value of ϕ, where $C_U = \phi C_H$ and describes the social cost of uninsurance. Each panel describes a different subsidy regime. The "ACA Subsidies, No Unsubsidized" considers only demand of low-income individuals and assumes that all individuals receive a subsidy linked to the price of the lowest-price plan such that the net-of-subsidy price is $55. The "ACA Subsidies, w/ Unsubsidized" case adds high-income individuals (who have different demand) who do not receive subsidies. The fixed subsidy cases assume that all individuals receive the same fixed subsidy that is not a function of plan prices, with the "avg cost" case setting the subsidy equal to the average cost across all consumers in the population.

Figure 15: Social Welfare under Increasingly Larger Uninsurance Penalty, No L Cost Advantage, $\sigma = \$50$]

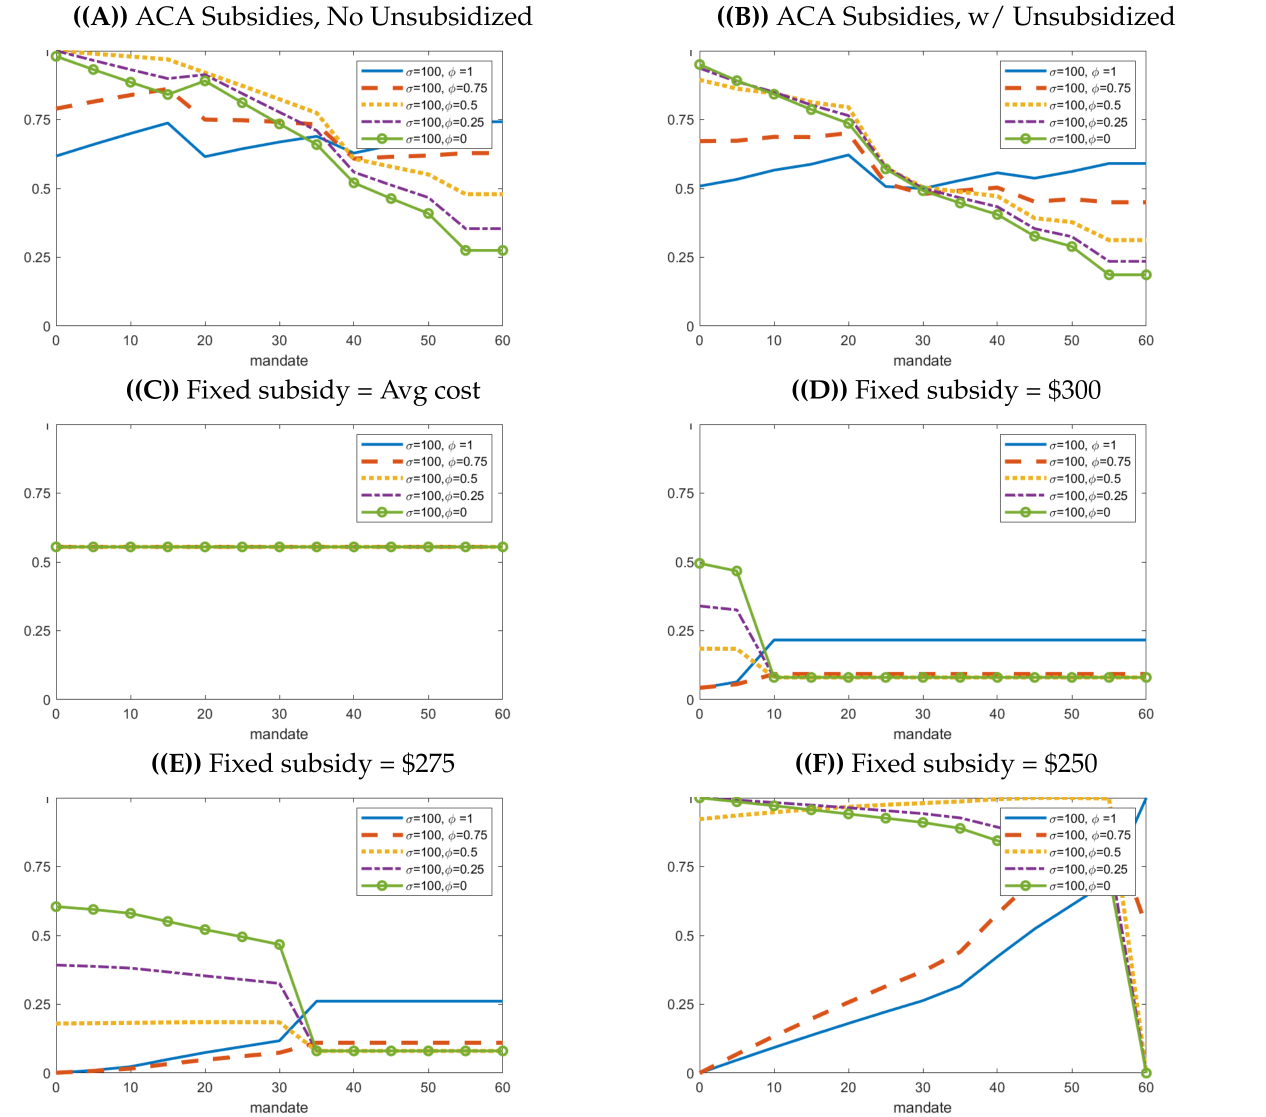

Note: Figures show estimates of overall social welfare as described by Equation 7 under increasingly larger uninsurance penalties. Each line represents a different value of ϕ, where $C_U = \phi C_H$ and describes the social cost of uninsurance. Each panel describes a different subsidy regime. The "ACA Subsidies, No Unsubsidized" considers only demand of low-income individuals and assumes that all individuals receive a subsidy linked to the price of the lowest-price plan such that the net-of-subsidy price is $55. The "ACA Subsidies, w/ Unsubsidized" case adds high-income individuals (who have different demand) who do not receive subsidies. The fixed subsidy cases assume that all individuals receive the same fixed subsidy that is not a function of plan prices, with the "avg cost" case setting the subsidy equal to the average cost across all consumers in the population.

Figure 16: Social Welfare under Increasingly Strong Risk Adjustment Transfers (higher α), 15% L Cost Advantage, σ = \$50

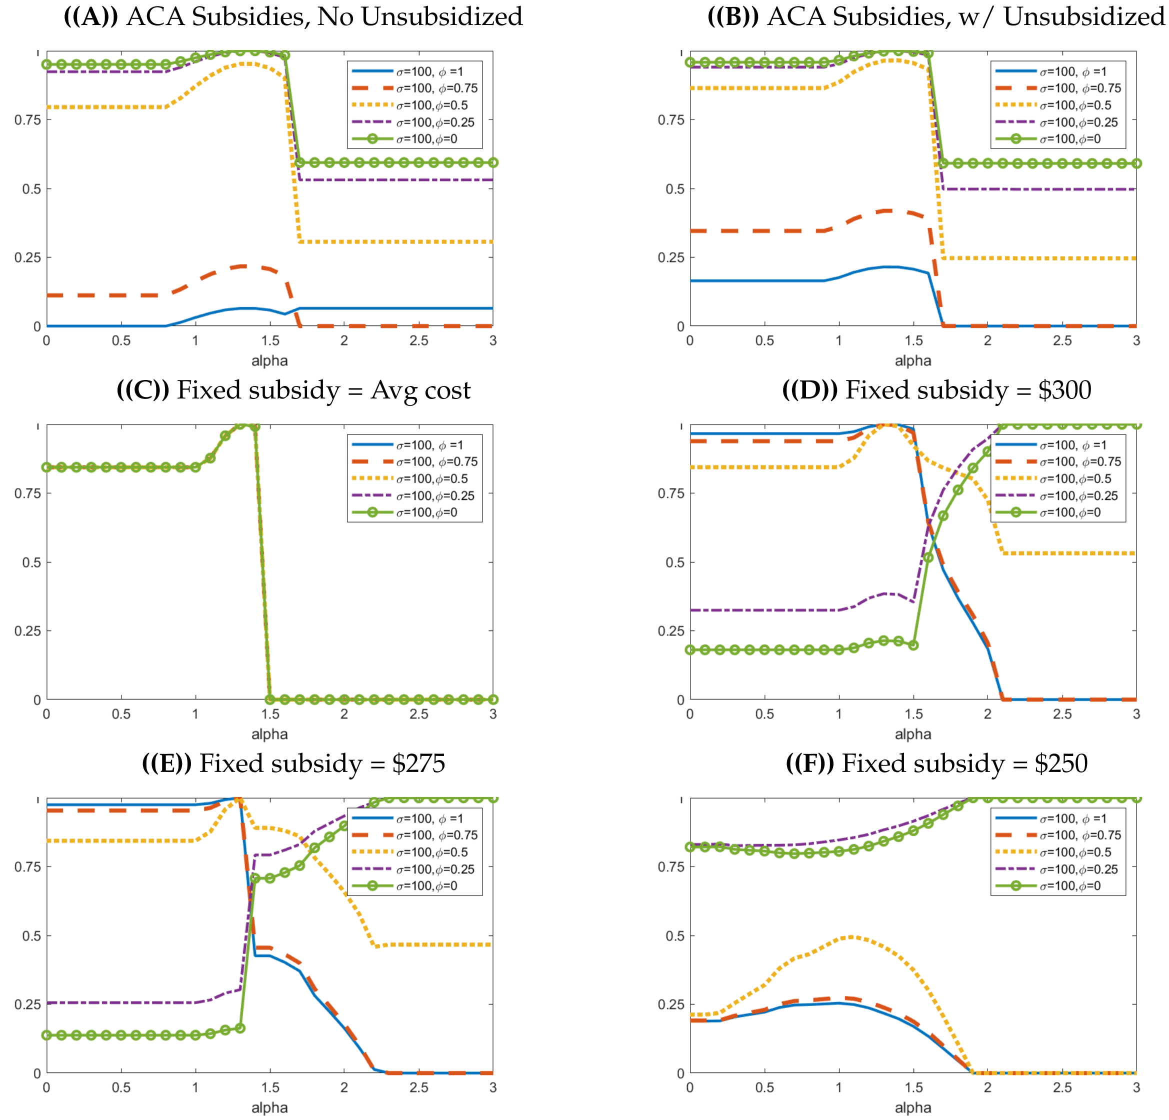

Note: Figures show estimates of overall social welfare as described by Equation 7 under increasingly strong risk adjustment transfers (larger α). Each line represents a different value of ϕ, where $C_U = \phi C_H$ and describes the social cost of uninsurance. Each panel describes a different subsidy regime. The "ACA Subsidies, No Unsubsidized" considers only demand of low-income individuals and assumes that all individuals receive a subsidy linked to the price of the lowest-price plan such that the net-of-subsidy price is \$55. The "ACA Subsidies, w/ Unsubsidized" case adds high-income individuals (who have different demand) who do not receive subsidies. The fixed subsidy cases assume that all individuals receive the same fixed subsidy that is not a function of plan prices, with the "avg cost" case setting the subsidy equal to the average cost across all consumers in the population.

Figure 17: Social Welfare under Increasingly Strong Risk Adjustment Transfers (higher α) [No L Cost Advantage, $\sigma = \$50$]

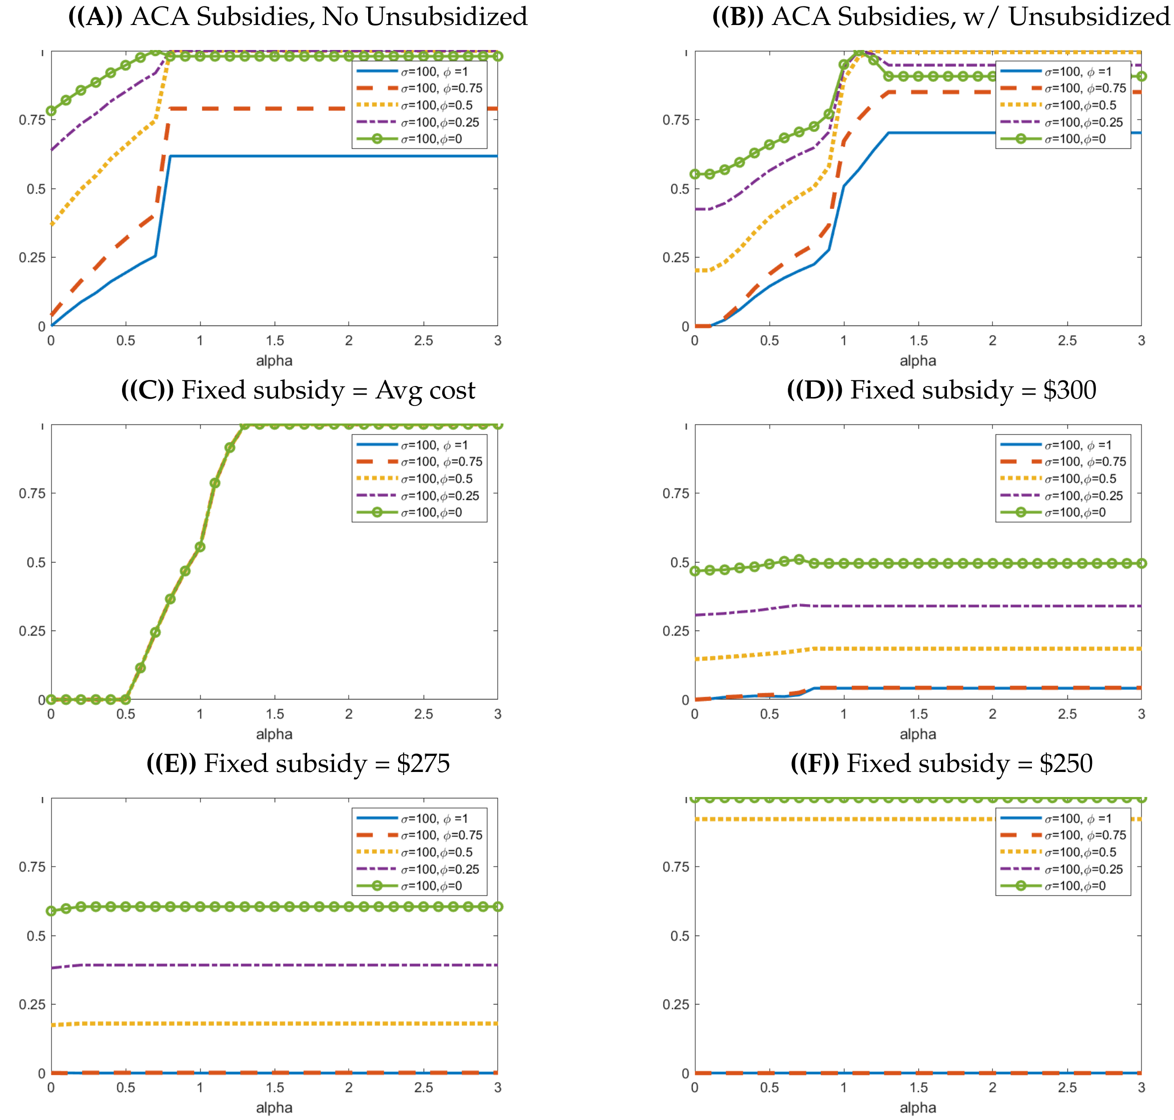

Note: Figures show estimates of overall social welfare as described by Equation 7 under increasingly strong risk adjustment transfers (larger α). Each line represents a different value of ϕ, where $C_U = \phi C_H$ and describes the social cost of uninsurance. Each panel describes a different subsidy regime. The "ACA Subsidies, No Unsubsidized" considers only demand of low-income individuals and assumes that all individuals receive a subsidy linked to the price of the lowest-price plan such that the net-of-subsidy price is $55. The "ACA Subsidies, w/ Unsubsidized" case adds high-income individuals (who have different demand) who do not receive subsidies. The fixed subsidy cases assume that all individuals receive the same fixed subsidy that is not a function of plan prices, with the "avg cost" case setting the subsidy equal to the average cost across all consumers in the population.

Figure 18: Welfare under Different Uninsurance Penalties and Levels of α with a Fixed Subsidy of $275 [No L Cost Advantage, $\phi = 0.5$, $\sigma =$ $100]

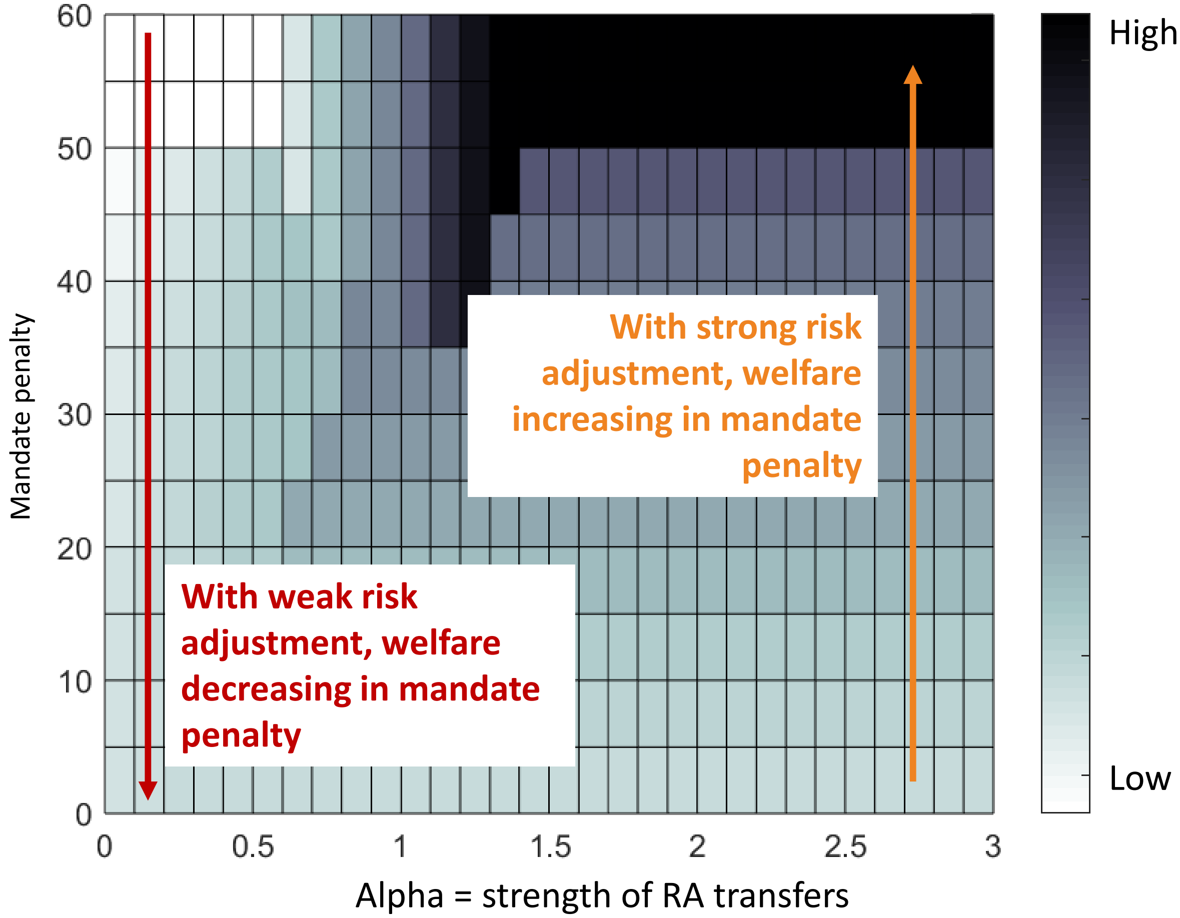

Table 1: Benefit Regulation : L-plan 15% cost advantage

	ACA-like all sub		ACA-like some sub		Fixed sub = avg. cost		Fixed sub = $300		Fixed sub = $275		Fixed sub = $250	
	L offered	No L	L offered	No L	L offered	No L	L offered	No L	L offered	No L	L offered	No L
price H	432	380	412	364	.	321	.	.	.	453	458	463
price L	335	.	308	.	273	.	273	273	273	.	373	.
share H	.045	.7	.011	.77	0	1	0	0	0	.28	.092	.21
share L	.57	0	.74	0	1	0	1	1	1	0	.24	0
share U	.39	.3	.25	.23	0	0	0	0	0	.72	.66	.79
subsidy H	280	325	253	309	0	322	0	0	0	275	250	250
subsidy L	280	0	253	0	322	0	300	300	275	0	250	0
welfare	.	.	.	.	.	.	.	.	.	.	.	.
phi1 fix100	.64	.68	.77	.74	1	.91	1	1	1	.12	.23	0
phi.75 fix100	.63	.63	.76	.69	1	.86	1	1	1	.12	.24	0
phi.5 fix100	.61	.5	.73	.55	1	.7	1	1	1	.11	.25	0
phi.25 fix100	.68	.18	.63	.12	.66	0	.66	.66	.66	.94	.96	1
phi0 fix100	.44	.23	.34	.16	.19	0	.19	.19	.19	.9	.82	1
phi1 fix50	.68	.69	.79	.75	1	.88	1	1	1	.13	.26	0
phi.75 fix50	.71	.64	.81	.68	1	.77	1	1	1	.14	.29	0
phi.5 fix50	1	.35	.95	.26	.97	0	.97	.97	.97	.54	.76	.38
phi.25 fix50	.56	.31	.43	.23	.24	0	.24	.24	.24	.93	.88	1
phi0 fix50	.46	.28	.33	.2	.12	0	.12	.12	.12	.9	.83	1
phi1 fix0	.77	.73	.85	.76	1	.82	1	1	1	.16	.31	0
phi.75 fix0	1	.67	.99	.63	.99	.42	.99	.99	.99	.23	.48	0
phi.5 fix0	.76	.45	.6	.34	.32	0	.32	.32	.32	.98	.99	1
phi.25 fix0	.53	.34	.39	.25	.14	0	.14	.14	.14	.92	.87	1
phi0 fix0	.47	.3	.33	.22	.093	0	.093	.093	.093	.91	.83	1
phi0 fix0	.	.	.	.	.	.	.	.	.	.	.	.

Table 2: Benefit Regulation : No L-plan cost advantage

	ACA-like all sub		ACA-like some sub		Fixed sub = avg. cost		Fixed sub = $300		Fixed sub = $275		Fixed sub = $250	
	L offered	No L	L offered	No L	L offered	No L	L offered	No L	L offered	No L	L offered	No L
price H	380	380	379	364	357	321	431	431	453	453	463	463
price L	.	.	352	.	310	.	.	.	.	.	.	.
share H	.7	.7	.48	.77	.24	1	.41	.41	.28	.28	.21	.21
share L	0	0	.24	0	.76	0	0	0	0	0	0	0
share U	.3	.3	.27	.23	0	0	.59	.59	.72	.72	.79	.79
subsidy H	325	325	297	309	322	322	300	300	275	275	250	250
subsidy L	0	0	297	0	322	0	0	0	0	0	0	0
welfare	.	.	.	.	.	.	.	.	.	.	.	.
phi1 fix100	.74	.74	.73	.81	.9	1	.37	.37	.13	.13	0	0
phi.75 fix100	.74	.74	.7	.81	.83	1	.37	.37	.14	.14	0	0
phi.5 fix100	.71	.71	.6	.78	.56	1	.4	.4	.16	.16	0	0
phi.25 fix100	.51	.51	.39	.47	0	.4	.82	.82	.97	.97	1	1
phi0 fix100	.35	.35	.3	.3	0	.16	.71	.71	.91	.91	1	1
phi1 fix50	.79	.79	.75	.85	.86	1	.41	.41	.15	.15	0	0
phi.75 fix50	.83	.83	.74	.88	.7	1	.47	.47	.18	.18	0	0
phi.5 fix50	.84	.84	.63	.78	0	.62	1	1	.95	.95	.86	.86
phi.25 fix50	.44	.44	.37	.38	0	.2	.78	.78	.94	.94	1	1
phi0 fix50	.36	.36	.31	.29	0	.11	.72	.72	.91	.91	1	1
phi1 fix0	.89	.89	.81	.93	.78	1	.5	.5	.19	.19	0	0
phi.75 fix0	1	1	.73	.94	0	.7	.82	.82	.47	.47	.2	.2
phi.5 fix0	.59	.59	.48	.5	0	.25	.89	.89	.98	.98	1	1
phi.25 fix0	.42	.42	.36	.35	0	.13	.77	.77	.93	.93	1	1
phi0 fix0	.36	.36	.32	.29	0	.087	.73	.73	.92	.92	1	1
phi0 fix0	.	.	.	.	.	.	.	.	.	.	.	.

Table 3: Mandate Simulations with 15%L Cost Advantage

(a) ACA-like, all subsidized

α	0	15	30	45	60
price H	432	422	412	401	.
price L	335	323	308	293	273
share H	.045	.037	.018	.0022	0
share L	.57	.66	.77	.88	1
share U	.39	.31	.21	.12	0
subsidy	280	268	253	238	218
welfare	.	.	.	.	.
phi1 fix100	0	0	0	0	0
phi.75 fix100	0	0	0	0	0
phi.5 fix100	.16	.16	.16	.16	.16
phi.25 fix100	.8	.8	.8	.8	.8
phi0 fix100	.92	.92	.92	.92	.92
phi1 fix50	0	0	0	0	0
phi.75 fix50	.11	.11	.11	.11	.11
phi.5 fix50	.8	.8	.8	.8	.8
phi.25 fix50	.92	.92	.92	.92	.92
phi0 fix50	.95	.95	.95	.95	.95
phi1 fix0	.062	.062	.062	.062	.062
phi.75 fix0	.72	.72	.72	.72	.72
phi.5 fix0	.92	.92	.92	.92	.92
phi.25 fix0	.95	.95	.95	.95	.95
phi0 fix0	.96	.96	.96	.96	.96

(b) ACA-like, some subsidized

0	15	30	45	60
412	407	400	.	.
308	301	292	284	273
.011	.006	.0013	0	0
.74	.8	.87	.93	1
.25	.19	.13	.072	0
253	246	237	229	218
.	.	.	.	.
.092	.092	.092	.092	.092
.19	.19	.19	.19	.19
.37	.37	.37	.37	.37
.85	.85	.85	.85	.85
.94	.94	.94	.94	.94
.16	.16	.16	.16	.16
.35	.35	.35	.35	.35
.86	.86	.86	.86	.86
.94	.94	.94	.94	.94
.96	.96	.96	.96	.96
.32	.32	.32	.32	.32
.83	.83	.83	.83	.83
.94	.94	.94	.94	.94
.96	.96	.96	.96	.96
.97	.97	.97	.97	.97

(c) Fixed sub = avg. monthly cost

0	15	30	45	60
.	.	.	.	.
273	273	273	273	273
0	0	0	0	0
1	1	1	1	1
0	0	0	0	0
300	300	300	300	300
.	.	.	.	.
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84
.84	.84	.84	.84	.84

(d) Fixed sub = $300

α	0	15	30	45	60
price H	.	.	.	.	.
price L	273	273	273	273	273
share H	0	0	0	0	0
share L	1	1	1	1	1
share U	0	0	0	0	0
subsidy	300	300	300	300	300
welfare	.	.	.	.	.
phi1 fix100	.98	.98	.98	.98	.98
phi.75 fix100	.96	.96	.96	.96	.96
phi.5 fix100	.93	.93	.93	.93	.93
phi.25 fix100	.84	.84	.84	.84	.84
phi0 fix100	.28	.28	.28	.28	.28
phi1 fix50	.97	.97	.97	.97	.97
phi.75 fix50	.94	.94	.94	.94	.94
phi.5 fix50	.84	.84	.84	.84	.84
phi.25 fix50	.32	.32	.32	.32	.32
phi0 fix50	.18	.18	.18	.18	.18
phi1 fix0	.95	.95	.95	.95	.95
phi.75 fix0	.84	.84	.84	.84	.84
phi.5 fix0	.38	.38	.38	.38	.38
phi.25 fix0	.2	.2	.2	.2	.2
phi0 fix0	.13	.13	.13	.13	.13

(e) Fixed sub = $275

0	15	30	45	60
.	.	.	.	.
273	273	273	273	273
0	0	0	0	0
1	1	1	1	1
0	0	0	0	0
275	275	275	275	275
.	.	.	.	.
.98	.98	.98	.98	.98
.97	.97	.97	.97	.97
.94	.94	.94	.94	.94
.7	.7	.7	.7	.7
.21	.21	.21	.21	.21
.97	.97	.97	.97	.97
.95	.95	.95	.95	.95
.84	.84	.84	.84	.84
.26	.26	.26	.26	.26
.14	.14	.14	.14	.14
.96	.96	.96	.96	.96
.88	.88	.88	.88	.88
.33	.33	.33	.33	.33
.16	.16	.16	.16	.16
.1	.1	.1	.1	.1

(f) Fixed sub = $250

0	15	30	45	60
458	450	430	.	.
373	360	334	273	273
.092	.072	.049	0	0
.24	.36	.57	1	1
.66	.57	.38	0	0
250	250	250	250	250
.	.	.	.	.
.17	.17	.17	.17	.17
.17	.17	.17	.17	.17
.15	.15	.15	.15	.15
.7	.7	.7	.7	.7
.79	.79	.79	.79	.79
.19	.19	.19	.19	.19
.19	.19	.19	.19	.19
.21	.21	.21	.21	.21
.83	.83	.83	.83	.83
.82	.82	.82	.82	.82
.22	.22	.22	.22	.22
.26	.26	.26	.26	.26
.9	.9	.9	.9	.9
.85	.85	.85	.85	.85
.84	.84	.84	.84	.84

Table 4: Mandate Simulations with No *L* Cost Advantage

(a) ACA-like, all subsidized

α	0	15	30	45	60
price H	380	366	373	372	357
price L	.	.	347	333	310
share H	.7	.77	.5	.27	.24
share L	0	0	.29	.61	.76
share U	.3	.23	.21	.12	0
subsidy	325	311	292	278	255
welfare	.	.	.	.	.
phi1 fix100	0	0	0	0	0
phi.75 fix100	0	0	0	0	0
phi.5 fix100	.06	.06	.06	.06	.06
phi.25 fix100	.38	.38	.38	.38	.38
phi0 fix100	.66	.66	.66	.66	.66
phi1 fix50	0	0	0	0	0
phi.75 fix50	.038	.038	.038	.038	.038
phi.5 fix50	.37	.37	.37	.37	.37
phi.25 fix50	.64	.64	.64	.64	.64
phi0 fix50	.78	.78	.78	.78	.78
phi1 fix0	.015	.015	.015	.015	.015
phi.75 fix0	.35	.35	.35	.35	.35
phi.5 fix0	.62	.62	.62	.62	.62
phi.25 fix0	.78	.78	.78	.78	.78
phi0 fix0	.83	.83	.83	.83	.83

(b) ACA-like, some subsidized

α	0	15	30	45	60
price H	379	373	376	366	356
price L	352	344	341	330	315
share H	.48	.44	.21	.18	.16
share L	.24	.34	.64	.74	.84
share U	.27	.21	.15	.081	-1.1e-16
subsidy	297	289	286	275	260
welfare	.	.	.	.	.
phi1 fix100	0	0	0	0	0
phi.75 fix100	0	0	0	0	0
phi.5 fix100	0	0	0	0	0
phi.25 fix100	.19	.19	.19	.19	.19
phi0 fix100	.42	.42	.42	.42	.42
phi1 fix50	0	0	0	0	0
phi.75 fix50	0	0	0	0	0
phi.5 fix50	.2	.2	.2	.2	.2
phi.25 fix50	.42	.42	.42	.42	.42
phi0 fix50	.55	.55	.55	.55	.55
phi1 fix0	0	0	0	0	0
phi.75 fix0	.21	.21	.21	.21	.21
phi.5 fix0	.43	.43	.43	.43	.43
phi.25 fix0	.56	.56	.56	.56	.56
phi0 fix0	.64	.64	.64	.64	.64

(c) Fixed sub = avg. monthly cost

α	0	15	30	45	60
price H	357	357	357	357	357
price L	310	310	310	310	310
share H	.24	.24	.24	.24	.24
share L	.76	.76	.76	.76	.76
share U	0	0	0	0	0
subsidy	322	322	322	322	322
welfare	.	.	.	.	.
phi1 fix100	0	0	0	0	0
phi.75 fix100	0	0	0	0	0
phi.5 fix100	0	0	0	0	0
phi.25 fix100	0	0	0	0	0
phi0 fix100	0	0	0	0	0
phi1 fix50	0	0	0	0	0
phi.75 fix50	0	0	0	0	0
phi.5 fix50	0	0	0	0	0
phi.25 fix50	0	0	0	0	0
phi0 fix50	0	0	0	0	0
phi1 fix0	0	0	0	0	0
phi.75 fix0	0	0	0	0	0
phi.5 fix0	0	0	0	0	0
phi.25 fix0	0	0	0	0	0
phi0 fix0	0	0	0	0	0

(d) Fixed sub = \$300

α	0	15	30	45	60
price H	431	413	357	357	357
price L	.	.	310	310	310
share H	.41	.52	.24	.24	.24
share L	0	0	.76	.76	.76
share U	.59	.48	0	0	0
subsidy	300	300	300	300	300
welfare	.	.	.	.	.
phi1 fix100	0	0	0	0	0
phi.75 fix100	0	0	0	0	0
phi.5 fix100	.12	.12	.12	.12	.12
phi.25 fix100	.83	.83	.83	.83	.83
phi0 fix100	.91	.91	.91	.91	.91
phi1 fix50	0	0	0	0	0
phi.75 fix50	0	0	0	0	0
phi.5 fix50	.79	.79	.79	.79	.79
phi.25 fix50	.9	.9	.9	.9	.9
phi0 fix50	.93	.93	.93	.93	.93
phi1 fix0	0	0	0	0	0
phi.75 fix0	.67	.67	.67	.67	.67
phi.5 fix0	.89	.89	.89	.89	.89
phi.25 fix0	.93	.93	.93	.93	.93
phi0 fix0	.94	.94	.94	.94	.94

(e) Fixed sub = \$275

α	0	15	30	45	60
price H	453	442	425	357	357
price L	.	.	.	310	310
share H	.28	.34	.45	.24	.24
share L	0	0	0	.76	.76
share U	.72	.66	.55	0	0
subsidy	275	275	275	275	275
welfare	.	.	.	.	.
phi1 fix100	.012	.012	.012	.012	.012
phi.75 fix100	.0058	.0058	.0058	.0058	.0058
phi.5 fix100	.14	.14	.14	.14	.14
phi.25 fix100	.96	.96	.96	.96	.96
phi0 fix100	.97	.97	.97	.97	.97
phi1 fix50	.011	.011	.011	.011	.011
phi.75 fix50	.0024	.0024	.0024	.0024	.0024
phi.5 fix50	.94	.94	.94	.94	.94
phi.25 fix50	.97	.97	.97	.97	.97
phi0 fix50	.97	.97	.97	.97	.97
phi1 fix0	.0085	.0085	.0085	.0085	.0085
phi.75 fix0	.7	.7	.7	.7	.7
phi.5 fix0	.98	.98	.98	.98	.98
phi.25 fix0	.98	.98	.98	.98	.98
phi0 fix0	.98	.98	.98	.98	.98

(f) Fixed sub = \$250

α	0	15	30	45	60
price H	463	457	451	436	420
price L	.	.	.	.	.
share H	.21	.25	.29	.38	.48
share L	0	0	0	0	0
share U	.79	.75	.71	.62	.52
subsidy	250	250	250	250	250
welfare	.	.	.	.	.
phi1 fix100	0	0	0	0	0
phi.75 fix100	0	0	0	0	0
phi.5 fix100	0	0	0	0	0
phi.25 fix100	1	1	1	1	1
phi0 fix100	1	1	1	1	1
phi1 fix50	0	0	0	0	0
phi.75 fix50	0	0	0	0	0
phi.5 fix50	.72	.72	.72	.72	.72
phi.25 fix50	1	1	1	1	1
phi0 fix50	1	1	1	1	1
phi1 fix0	0	0	0	0	0
phi.75 fix0	0	0	0	0	0
phi.5 fix0	1	1	1	1	1
phi.25 fix0	1	1	1	1	1
phi0 fix0	1	1	1	1	1

Table 5: Risk Adjustment Simulations with 15%L Cost Advantage

(a) ACA-like, all subsidized

α	0	.5	1	1.5	2
price H	.	.	432	396	380
price L	337	337	335	343	.
share H	0	0	.045	.22	.7
share L	.61	.61	.57	.4	0
share U	.39	.39	.39	.39	.3
subsidy	282	282	280	288	325
welfare	.	.	.	.	.
phi1 fix100	0	0	.021	.04	.11
phi.75 fix100	0	0	.031	.058	.033
phi.5 fix100	.16	.16	.21	.25	0
phi.25 fix100	.8	.8	.88	.94	.23
phi0 fix100	.92	.92	.96	.99	.5
phi1 fix50	0	0	.031	.058	.065
phi.75 fix50	.11	.11	.16	.21	0
phi.5 fix50	.8	.8	.87	.94	.31
phi.25 fix50	.92	.92	.96	.99	.53
phi0 fix50	.95	.95	.97	1	.59
phi1 fix0	.062	.062	.12	.16	0
phi.75 fix0	.72	.72	.79	.85	.34
phi.5 fix0	.92	.92	.96	.99	.57
phi.25 fix0	.95	.95	.97	1	.62
phi0 fix0	.96	.96	.98	1	.64

(b) ACA-like, some subsidized

	0	.5	1	1.5	2
price H	.	.	412	365	364
price L	309	309	308	312	.
share H	0	0	.011	.13	.77
share L	.75	.75	.74	.62	0
share U	.25	.25	.25	.25	.23
subsidy	254	254	253	257	309
welfare	.	.	.	.	.
phi1 fix100	.092	.092	.1	.12	0
phi.75 fix100	.19	.19	.2	.23	0
phi.5 fix100	.37	.37	.39	.43	.00017
phi.25 fix100	.85	.85	.87	.94	.16
phi0 fix100	.94	.94	.95	.99	.45
phi1 fix50	.16	.16	.18	.21	0
phi.75 fix50	.35	.35	.36	.41	.00023
phi.5 fix50	.86	.86	.89	.96	.25
phi.25 fix50	.94	.94	.95	1	.5
phi0 fix50	.96	.96	.97	1	.59
phi1 fix0	.32	.32	.33	.39	.00029
phi.75 fix0	.83	.83	.85	.92	.31
phi.5 fix0	.94	.94	.95	1	.54
phi.25 fix0	.96	.96	.97	1	.62
phi0 fix0	.97	.97	.97	1	.66

(c) Fixed sub = avg. monthly cost

	0	.5	1	1.5	2
price H	.	.	.	321	321
price L	273	273	273	.	.
share H	0	0	0	1	1
share L	1	1	1	0	0
share U	0	0	0	0	0
subsidy	300	300	300	322	322
welfare	.	.	.	.	.
phi1 fix100	.84	.84	.84	0	0
phi.75 fix100	.84	.84	.84	0	0
phi.5 fix100	.84	.84	.84	0	0
phi.25 fix100	.84	.84	.84	0	0
phi0 fix100	.84	.84	.84	0	0
phi1 fix50	.84	.84	.84	0	0
phi.75 fix50	.84	.84	.84	0	0
phi.5 fix50	.84	.84	.84	0	0
phi.25 fix50	.84	.84	.84	0	0
phi0 fix50	.84	.84	.84	0	0
phi1 fix0	.84	.84	.84	0	0
phi.75 fix0	.84	.84	.84	0	0
phi.5 fix0	.84	.84	.84	0	0
phi.25 fix0	.84	.84	.84	0	0
phi0 fix0	.84	.84	.84	0	0

(d) Fixed sub = $300

α	0	.5	1	1.5	2
price H	.	.	.	319	414
price L	273	273	273	282	383
share H	0	0	0	.28	.33
share L	1	1	1	.72	.15
share U	0	0	0	0	.52
subsidy	300	300	300	300	300
welfare	.	.	.	.	.
phi1 fix100	.98	.98	.98	.99	.16
phi.75 fix100	.96	.96	.96	.98	.17
phi.5 fix100	.93	.93	.93	.96	.17
phi.25 fix100	.84	.84	.84	.92	.96
phi0 fix100	.28	.28	.28	.31	.91
phi1 fix50	.97	.97	.97	.98	.18
phi.75 fix50	.94	.94	.94	.97	.2
phi.5 fix50	.84	.84	.84	.92	.72
phi.25 fix50	.32	.32	.32	.35	.95
phi0 fix50	.18	.18	.18	.2	.9
phi1 fix0	.95	.95	.95	.97	.23
phi.75 fix0	.84	.84	.84	.92	.49
phi.5 fix0	.38	.38	.38	.41	1
phi.25 fix0	.2	.2	.2	.21	.93
phi0 fix0	.13	.13	.13	.14	.9

(e) Fixed sub = $275

	0	.5	1	1.5	2
price H	.	.	.	428	438
price L	273	273	273	357	395
share H	0	0	0	.15	.26
share L	1	1	1	.33	.092
share U	0	0	0	.52	.65
subsidy	275	275	275	275	275
welfare	.	.	.	.	.
phi1 fix100	.98	.98	.98	.39	.15
phi.75 fix100	.97	.97	.97	.39	.15
phi.5 fix100	.94	.94	.94	.39	.15
phi.25 fix100	.7	.7	.7	.92	.99
phi0 fix100	.21	.21	.21	.7	.9
phi1 fix50	.97	.97	.97	.43	.16
phi.75 fix50	.95	.95	.95	.46	.18
phi.5 fix50	.84	.84	.84	.89	.66
phi.25 fix50	.26	.26	.26	.79	.93
phi0 fix50	.14	.14	.14	.71	.9
phi1 fix0	.96	.96	.96	.5	.2
phi.75 fix0	.88	.88	.88	.74	.31
phi.5 fix0	.33	.33	.33	.94	.99
phi.25 fix0	.16	.16	.16	.76	.92
phi0 fix0	.1	.1	.1	.71	.9

(f) Fixed sub = $250

	0	.5	1	1.5	2
price H	.	479	458	452	463
price L	380	376	373	386	.
share H	0	.022	.092	.17	.21
share L	.31	.3	.24	.13	0
share U	.69	.67	.66	.71	.79
subsidy	250	250	250	250	250
welfare	.	.	.	.	.
phi1 fix100	.17	.2	.23	.15	0
phi.75 fix100	.17	.2	.23	.16	0
phi.5 fix100	.15	.18	.23	.16	0
phi.25 fix100	.7	.75	.9	.99	.99
phi0 fix100	.79	.78	.78	.87	1
phi1 fix50	.19	.22	.25	.17	0
phi.75 fix50	.19	.23	.27	.19	0
phi.5 fix50	.21	.32	.49	.37	0
phi.25 fix50	.83	.83	.85	.92	1
phi0 fix50	.82	.81	.81	.88	1
phi1 fix0	.22	.26	.3	.2	0
phi.75 fix0	.26	.33	.41	.29	0
phi.5 fix0	.9	.92	.97	1	.99
phi.25 fix0	.85	.84	.84	.91	1
phi0 fix0	.84	.82	.82	.89	1

Table 6: Risk Adjustment Simulations with No *L* Cost Advantage

(a) ACA-like, all subsidized

α	0	.5	1	1.5	2
price H	491	444	380	380	380
price L	391	379	.	.	.
share H	.033	.17	.7	.7	.7
share L	.58	.45	0	0	0
share U	.39	.39	.3	.3	.3
subsidy	336	324	325	325	325
welfare	.	.	.	.	.
phi1 fix100	0	.15	.52	.52	.52
phi.75 fix100	0	.19	.6	.6	.6
phi.5 fix100	.06	.33	.76	.76	.76
phi.25 fix100	.38	.68	1	1	1
phi0 fix100	.66	.87	1	1	1
phi1 fix50	0	.19	.62	.62	.62
phi.75 fix50	.038	.32	.79	.79	.79
phi.5 fix50	.37	.66	1	1	1
phi.25 fix50	.64	.85	1	1	1
phi0 fix50	.78	.95	.98	.98	.98
phi1 fix0	.015	.3	.82	.82	.82
phi.75 fix0	.35	.63	1	1	1
phi.5 fix0	.62	.83	1	1	1
phi.25 fix0	.78	.95	1	1	1
phi0 fix0	.83	.96	.94	.94	.94

(b) ACA-like, some subsidized

α	0	.5	1	1.5	2
price H	.	437	379	364	364
price L	373	363	352	.	.
share H	0	.081	.48	.77	.77
share L	.71	.64	.24	0	0
share U	.29	.28	.27	.23	.23
subsidy	318	308	297	309	309
welfare	.	.	.	.	.
phi1 fix100	0	.12	.41	.6	.6
phi.75 fix100	0	.14	.5	.68	.68
phi.5 fix100	0	.19	.66	.82	.82
phi.25 fix100	.19	.39	.9	.98	.98
phi0 fix100	.42	.56	.94	.93	.93
phi1 fix50	0	.14	.51	.7	.7
phi.75 fix50	0	.19	.67	.85	.85
phi.5 fix50	.2	.39	.89	1	1
phi.25 fix50	.42	.56	.94	.95	.95
phi0 fix50	.55	.66	.95	.91	.91
phi1 fix0	0	.19	.68	.88	.88
phi.75 fix0	.21	.4	.88	1	1
phi.5 fix0	.43	.57	.94	.96	.96
phi.25 fix0	.56	.66	.95	.92	.92
phi0 fix0	.64	.72	.96	.89	.89

(c) Fixed sub = avg. monthly cost

α	0	.5	1	1.5	2
price H	.	.	357	321	321
price L	321	321	310	.	.
share H	0	0	.24	1	1
share L	1	1	.76	0	0
share U	0	0	0	0	0
subsidy	322	322	322	322	322
welfare	.	.	.	.	.
phi1 fix100	0	0	.55	1	1
phi.75 fix100	0	0	.55	1	1
phi.5 fix100	0	0	.55	1	1
phi.25 fix100	0	0	.55	1	1
phi0 fix100	0	0	.55	1	1
phi1 fix50	0	0	.55	1	1
phi.75 fix50	0	0	.55	1	1
phi.5 fix50	0	0	.55	1	1
phi.25 fix50	0	0	.55	1	1
phi0 fix50	0	0	.55	1	1
phi1 fix0	0	0	.55	1	1
phi.75 fix0	0	0	.55	1	1
phi.5 fix0	0	0	.55	1	1
phi.25 fix0	0	0	.55	1	1
phi0 fix0	0	0	.55	1	1

(d) Fixed sub = $300

α	0	.5	1	1.5	2
price H	468	448	431	431	431
price L	405	407	.	.	.
share H	.18	.26	.41	.41	.41
share L	.22	.13	0	0	0
share U	.6	.61	.59	.59	.59
subsidy	300	300	300	300	300
welfare	.	.	.	.	.
phi1 fix100	0	.029	.11	.11	.11
phi.75 fix100	0	.055	.16	.16	.16
phi.5 fix100	.12	.25	.4	.4	.4
phi.25 fix100	.83	.93	1	1	1
phi0 fix100	.91	.98	1	1	1
phi1 fix50	0	.046	.15	.15	.15
phi.75 fix50	0	.11	.27	.27	.27
phi.5 fix50	.79	.9	1	1	1
phi.25 fix50	.9	.97	1	1	1
phi0 fix50	.93	.98	.98	.98	.98
phi1 fix0	0	.085	.24	.24	.24
phi.75 fix0	.67	.78	.9	.9	.9
phi.5 fix0	.89	.96	1	1	1
phi.25 fix0	.93	.98	.99	.99	.99
phi0 fix0	.94	.98	.98	.98	.98

(e) Fixed sub = $275

α	0	.5	1	1.5	2
price H	457	453	453	453	453
price L	412	.	.	.	.
share H	.25	.28	.28	.28	.28
share L	.04	0	0	0	0
share U	.71	.72	.72	.72	.72
subsidy	275	275	275	275	275
welfare	.	.	.	.	.
phi1 fix100	.012	0	0	0	0
phi.75 fix100	.0058	.00067	.00067	.00067	.00067
phi.5 fix100	.14	.16	.16	.16	.16
phi.25 fix100	.96	1	1	1	1
phi0 fix100	.97	1	1	1	1
phi1 fix50	.011	0	0	0	0
phi.75 fix50	.0024	.0066	.0066	.0066	.0066
phi.5 fix50	.94	.97	.97	.97	.97
phi.25 fix50	.97	1	1	1	1
phi0 fix50	.97	1	1	1	1
phi1 fix0	.0085	.00077	.00077	.00077	.00077
phi.75 fix0	.7	.72	.72	.72	.72
phi.5 fix0	.98	1	1	1	1
phi.25 fix0	.98	1	1	1	1
phi0 fix0	.98	1	1	1	1

(f) Fixed sub = $250

α	0	.5	1	1.5	2
price H	463	463	463	463	463
price L	.	.	.	.	.
share H	.21	.21	.21	.21	.21
share L	0	0	0	0	0
share U	.79	.79	.79	.79	.79
subsidy	250	250	250	250	250
welfare	.	.	.	.	.
phi1 fix100	0	0	0	0	0
phi.75 fix100	0	0	0	0	0
phi.5 fix100	0	0	0	0	0
phi.25 fix100	1	1	1	1	1
phi0 fix100	1	1	1	1	1
phi1 fix50	0	0	0	0	0
phi.75 fix50	0	0	0	0	0
phi.5 fix50	.72	.72	.72	.72	.72
phi.25 fix50	1	1	1	1	1
phi0 fix50	1	1	1	1	1
phi1 fix0	0	0	0	0	0
phi.75 fix0	0	0	0	0	0
phi.5 fix0	1	1	1	1	1
phi.25 fix0	1	1	1	1	1
phi0 fix0	1	1	1	1	1

Online Appendix for:
The Two Margin Problem in Insurance Markets

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.[28] 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.[29] 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:

$$\begin{aligned} P_H &= AC_H\left(P^{cons}\right) \\ P_L &= AC_L\left(P^{cons}\right) \end{aligned} \tag{12}$$

We use this expression to define equilibrium throughout Section 2.

[28] 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.

[29] 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:

$$\begin{aligned} AC_H(P;G) &= \frac{1}{D_H(P;G)} \int_0^{s_{HL}(\Delta P_{HL}^{cons})} C_H(s)ds \\ AC_L(P;G) &= \frac{1}{D_L(P;G)} \int_{s_{HL}(\Delta P_{HL}^{cons})}^{s_{LU}(\Delta P_{LU}^{cons})} C_H(s)ds \end{aligned} \tag{13}$$

Similarly, we can define average risk score functions:

$$\begin{aligned} \overline{R}_H(P;G) &= \frac{1}{D_H(P;G)} \int_0^{s_{HL}(\Delta P_{HL}^{cons})} R(s)ds \\ \overline{R}_L(P;G) &= \frac{1}{D_L(P;G)} \int_{s_{HL}(\Delta P_{HL}^{cons})}^{s_{LU}(\Delta P_{LU}^{cons})} R(s)ds \end{aligned} \tag{14}$$

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 ($\equiv \overline{P}(P;G)$):

$$T(P;G) = \left(\frac{\overline{R}_H(P;G) - \overline{R}_L(P;G)}{\overline{R}(P;G)} \right) \overline{P}(P;G) \tag{15}$$

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:

$$\begin{aligned} P_H &= AC_H(P;G) - \alpha T(P;G) \equiv AC_H^{RA}(P;G,\alpha) \\ P_L &= AC_L(P;G) - \alpha T(P;G) \equiv AC_L^{RA}(P;G,\alpha) \end{aligned} \tag{16}$$

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:

$$\frac{dD_H}{dM} = \underbrace{\left(-\frac{\partial s_{HL}}{\partial \Delta P_{HL}}\right)}_{(+)} \cdot \left(\frac{dP_L}{dM} - \frac{dP_H}{dM}\right) \tag{17}$$

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

$$\begin{aligned} \frac{dP_H}{dM} &= \frac{\partial AC_H^{RA}}{\partial M} + \frac{\partial AC_H^{RA}}{dP_H}\frac{dP_H}{dM} + \frac{\partial AC_H^{RA}}{dP_L}\frac{dP_L}{dM} \\ \frac{dP_L}{dM} &= \frac{\partial AC_L^{RA}}{\partial M} + \frac{\partial AC_L^{RA}}{dP_H}\frac{dP_H}{dM} + \frac{\partial AC_L^{RA}}{dP_L}\frac{dP_L}{dM} \end{aligned} \tag{18}$$

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

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

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.[30] 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

[30] 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 intution 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:

$$\frac{dP_L}{dM} = \left[\underbrace{\frac{\partial AC_L^{RA}}{\partial M}}_{\text{Ext. Margin Selection}(-)} + \underbrace{\underbrace{\left(\alpha \frac{\partial T}{\partial M}\right)}_{(+)} \cdot \underbrace{\frac{\partial AC_L^{RA}}{dP_H}}_{(+)} \underbrace{\left(1 - \frac{\partial AC_H^{RA}}{dP_H}\right)^{-1}}_{(+)}}_{\text{Substitution to } L(-)} \right] \times \Phi_L^{-1} \tag{20}$$

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.[31] 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:

$$\frac{dP_H}{dM} = \left[\underbrace{\underbrace{\frac{\partial AC_L^{RA}}{\partial M}}_{(-)} \cdot \underbrace{\frac{\partial AC_H^{RA}}{dP_L}}_{(-)} \underbrace{\left(1 - \frac{\partial AC_L^{RA}}{dP_L}\right)^{-1}}_{(+)}}_{\text{Substitution to } L(+)} \right] \times \Phi_H^{-1} \tag{21}$$

$$\frac{dP_L}{dM} = \left[\underbrace{\frac{\partial AC_L^{RA}}{\partial M}}_{\text{Ext. Margin Selection}(-)} \right] \times \Phi_L^{-1} \tag{22}$$

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.

[31] 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:

$$\frac{\partial D_U}{\partial \alpha} = \underbrace{\left(-\frac{\partial s_{LU}}{\partial \Delta P_{LU}^{cons}}\right)}_{(+)} \cdot \frac{d\Delta P_{LU}^{cons}}{d\alpha} \tag{23}$$

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 14 with respect to P_H gives

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

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

$$\frac{dP_H}{d\alpha} = T(.) \times \left[\underbrace{1}_{\text{Direct (+)}} \underbrace{+ \frac{\partial AC_L^{RA}}{\partial P_H} \left(1 - \frac{\partial AC_H^{RA}}{\partial P_H}\right)^{-1}}_{\text{Substitution to } H \text{ (-)}} \right] \times \Phi_L^{-1} < 0 \tag{25}$$

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

C.1 Low-Income Demand and Costs: FHS (2018)

As discussed in Section 4.1, we draw on demand and cost estimates for low-income subsidized consumers from Finkelstein, Hendren and Shepard (2017), which we abbreviate as "FHS." As described in section 4.1, FHS estimate insurance demand in Massachusetts' pre-ACA subsidized health insurance exchange, known as "CommCare." Here we describe some additional details about FHS's estimates and our construction of the demand and cost curves.

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.

To estimate demand and costs, 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.[32] 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.

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 generate two sets of modified WTP curves: (1) linear demand and (2) "enhanced" demand. We focus throughout the paper on

[32] 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.

"enhanced" demand, as we view this as more realistic. Results using the lienar demand curves are found in the appendix. We discuss each set of demand curves in turn.

(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,

$$W_{HL}^{enh}(s) = \begin{cases} W_{HL}^{lin}(s) & \text{for } s \in [0.31, 1] \\ W_{HL}^{lin}(s) + 3 \times \frac{(0.31-s)}{0.31} \cdot W_{HL}^{lin}(0) & \text{for } s \in [0, 0.31) \end{cases} \tag{26}$$

and

$$W_H^{enh}(s) = W_L^{lin}(s) + W_{HL}^{enh}(s) \tag{27}$$

Both the linear and the enhanced WTP curves are shown in the top panel of Figure A1.

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 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). To do so, they 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.

We construct both linear and enhanced versions of these curves and we denote the linear curves $W_L^{HI,lin}(s)$ and $W_H^{HI,lin}(s)$ and the enhanced curves $W_L^{HI,enh}(s)$ and $W_H^{HI,enh}(s)$. We assume that the cost curves for this group are equivalent to the cost curves of the subsidized population, $C_L(s)$ and $C_H(s)$.[33]

We start by constructing $W_L^{HI,lin}(s)$, based on the estimates from Hackmann, Kolstad and Kowalski (2015). Their demand curve takes the following form:

$$W_{HKK}(s) = -\$9,276.81 * s + \$12,498.68 \tag{28}$$

This demand curve is "in-sample" in the range of $0.70 < s < 0.97$. As with the low-income, subsidized

[33] We note that this assumption implies 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.

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. Similar to the low-income, subsidized consumers, we also specify $W_H^{HI,lin}(s)$ as $\hat{W}_H^{HI,lin}(s) = W_L^{HI,lin}(s) + W_{HL}(s)$. For this demand system, the relevant curves are thus $W_L^{HI,lin}(s)$, $W_H^{HI,lin}(s)$, $C_L(s)$, and $C_H(s)$.

Similar to the low-income, subsidized consumers, we also construct a version of the high-income demand system with enhanced demand for H. Under this demand system, we leave $W_L^{HI,lin}(s)$, $C_L(s)$, and $C_H(s)$ unchanged. We construct a modified version of $W_H^{HI,lin}(s)$ which we call $\hat{W}_H^{HI,enh}(s)$. As above, we set $\hat{W}_H^{HI,enh}(s) = W_L^{HI,lin}(s) + \hat{W}_{HL}^{enh}(s)$. We thus have four demand systems: low-income + normal demand for H, low-income + enhanced demand for H, high-income + normal demand for H, high-income + enhanced demand for H.

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

[COMING SOON]

E Appendix Section 5: Results with Linear Demand Curves

[COMING SOON]

F Appendix Section 6: Results from Simulations with CommCare Subsidies

[COMING SOON]

G Additional Figures and Tables

Figure A1: WTP Curves for H and L

((A)) Low-Income

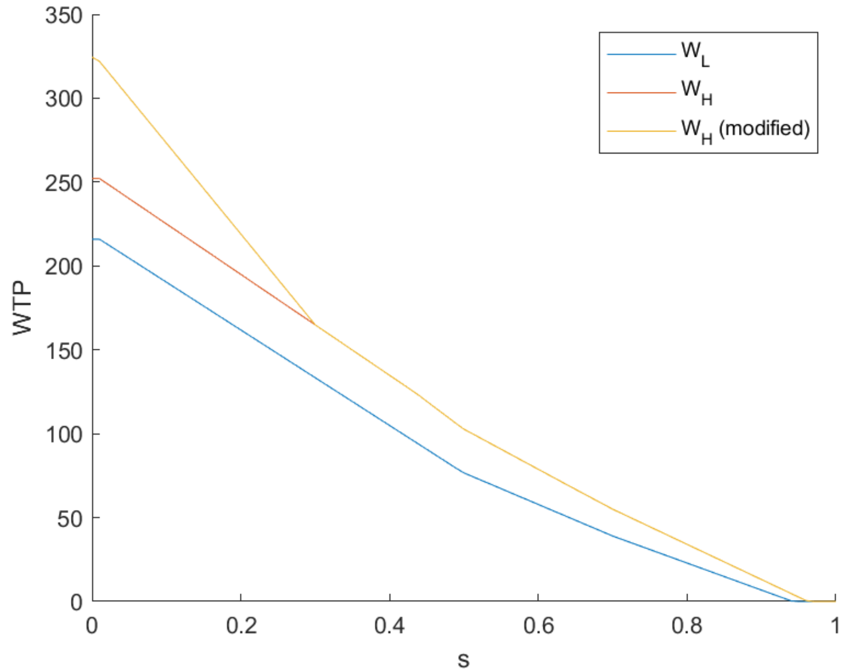

((B)) High-Income

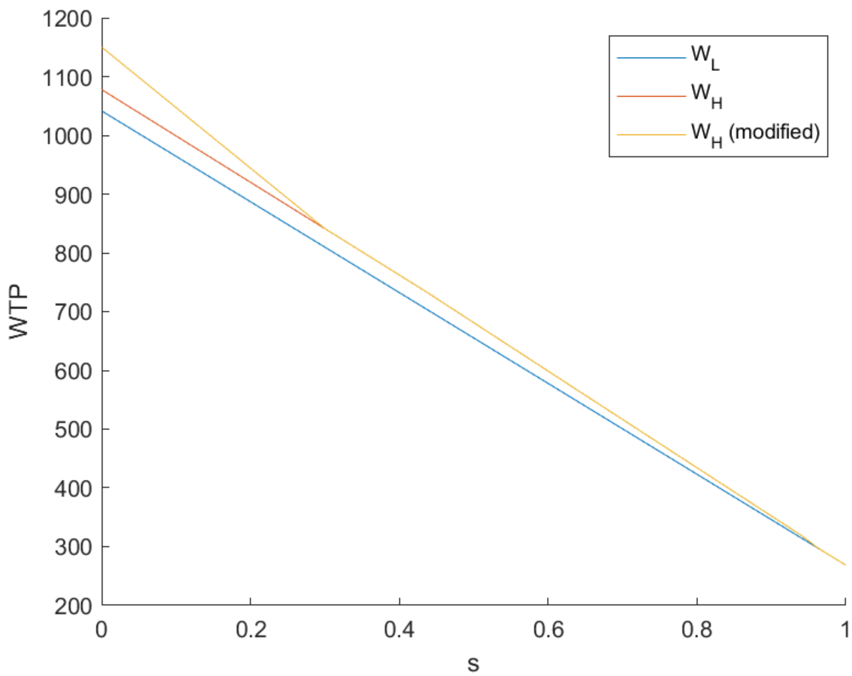

Note: Figure shows WTP Curves for H and L, $W_H(s)$ and $W_L(s)$. Left panel shows curves for low-income group which come from (Finkelstein, Hendren and Shepard, 2017). Right panel shows curves for high-income group which come from (Hackmann, Kolstad and Kowalski, 2015). Linear curves extrapolate linearly 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.