Trade-offs between Extensive and Intensive Margin Selection in Competitive Insurance Markets*

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

October 12, 2018

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 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 tradeoff 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. We illustrate these trade-offs empirically using a sufficient statistics approach incorporating demand and cost curves for more- and less-generous plans from Massachusetts, showing 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 Tom McGuire, Marika Cabral, and participants at ASHEcon, and the Chicago Booth Junior Health Economics Summit 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 a persistent problem in insurance markets. The tendency for higher-risk consumers to purchase insurance can lead to distorted prices and can cause consumers to sort inefficiently between insured and uninsured states (Akerlof, 1970). Additionally, the tendency for higher-risk consumers to demand generous insurance contracts can lead to equilibria in which insurers offer contracts that are inefficiently distorted away from socially efficient contracts (Rothschild and Stiglitz, 1976). These "extensive" (insurance vs. uninsurance) and "intensive" (high vs. low generosity) margin selection problems are especially salient in health insurance markets (Geruso and Layton, 2017).

In this paper, we study the interaction between these extensive and intensive margin selection problems. The previous economics literature has largely considered these two forms of selection in isolation, either assuming that all consumers always choose a plan (Einav, Finkelstein and Cullen, 2010) or that all plans in the market are effectively identical (Hackmann, Kolstad and Kowalski, 2015), reducing selection to either the in vs. out or high vs. low margin.[1] For some settings, one of these two assumptions may be a valid approximation. In other settings, however, simultaneous selection on both the extensive and intensive margins may be empirically relevant. For these settings, a more flexible framework is necessary for understanding the consequences of policies meant to combat selection.

We construct a stylized framework suitable to illustrating the two margin problem, and then use data from an individual insurance market to corroborate the model's insights empirically. We build on the models of Einav, Finkelstein and Cullen (2010) and Finkelstein, Hendren and Shepard (2017) to develop a simple, tractable model of equilibrium in a health insurance market featuring three options: a high generosity plan (H), a low generosity plan (L), and an outside option (uninsurance, U). The key insight extracted from our model is that there is indeed an interaction between extensive and intensive margin selection, and that that interaction often takes the form of an inherent trade-off between selection on one margin vs. the other.

The dual problems described by these margins—rates of uninsurance and the prices and availability of generous plans—are *the* most pressing concerns in individual and small group markets for many state policy makers and regulators. The framework we develop sheds light on a little un-

[1] A notable exception is Azevedo and Gottlieb (2017) where the authors consider a large contract space that includes uninsurance as well as many other potential contracts. Recent work by Saltzman (2017) and Domurat (2018) also allow for contract spaces that include multiple insurance options as well as the outside option of uninsurance.

derstood fact: Policies such as risk adjustment or regulatory constraints on benefit design that are aimed at combating intensive margin selection problems can exacerbate selection on the extensive margin. Likewise, policies that are aimed at combating extensive margin selection problems—such as mandates, penalties for remaining uninsured, or subsidies to encourage insurance take-up—can exacerbate selection on the intensive margin.

Why? The intuition is surprisingly straightforward: Intensive margin-focused policies, such as risk adjustment transfers, tend to cause the prices of the high and low generosity plans to converge, leading to lower prices for H but higher prices for L. As L prices rise, some marginal consumers are forced out of the market and into U. These policies thus improve the quality of coverage for some consumers who move from H to L (the intended effect) while simultaneously worsening the quality of coverage for other consumers who move from L to U (the unintended effect). Extensive margin-focused policies like penalties for uninsurance, on the other hand, tend to lower the effective price of L, causing healthy marginal consumers to move out of U and into L, lowering the price of L and increasing the gap between the prices of H and L. This larger price gap then induces the healthiest marginal H enrollees to move out of H and into L. These policies thus improve the quality of coverage for some consumers who move from U to L (the intended effect) while simultaneously worsening the quality of coverage for other consumers who move from H to L (the unintended effect).

We provide intuition for this interaction using a series of figures that illustrate equilibrium prices, sorting, and welfare in a market where consumers choose between H, L, and U. We use these figures to show the potential effects of various policies, including penalties for remaining uninsured, benefit regulation like the Essential Health Benefits that make up part of the ACA, and risk adjustment transfers present in most individual health insurance markets around the world. We also show that even with the addition of the outside option of uninsurance, social welfare under a given set of prices can be recovered using a set of sufficient statistics similar to those described by Einav, Finkelstein and Cullen (2010) that can be described in a simple diagram consisting of demand and cost curves. Finally, we also show that equilibrium prices and sorting can still be recovered using a similar set of sufficient statistics, again consisting only of the demand and cost curves for H and L, even with the more complex contract environment. These insights imply that the standard for reduced form empirical research on health insurance markets need not be restricted to the 2 option setting but instead standard methods can easily be applied to more complex market settings without adopting

structural approaches.

We illustrate our methodological insights as well as the empirical trade-off between extensive and intensive margin selection in the context of the Massachusetts Connector. The Connector was a precursor to the state Health Insurance Marketplaces established by the Affordable Care Act. It was introduced by Massachusetts to provide subsidized health insurance coverage to low-income Massachusetts residents who did not qualify for Medicaid. Finkelstein, Hendren and Shepard (2017) document significant adverse selection both into the market and within the market across a narrow-network, lower-quality option and a set of wider-network, higher-quality plans. Following the same regression discontinuity design, we construct modified versions of the demand and cost curves for H and L recovered by Finkelstein, Hendren and Shepard (2017). We use these demand and cost curves in a number of illustrative counterfactual exercises that examine enrollment and welfare as we vary benefit design rules, mandates and penalties, and risk adjustment strength. These counterfactuals are closely tied to relevant policy alternatives, including the flexibility recently granted to states by CMS in the regulation of individual and small group Marketplaces.[2]

First, we analyze the consequences of benefit regulation by comparing equilibrium prices and sorting across the three options (H, L, and U) in a setting where L is allowed and a setting where H and U are the only options. We show that in many plausible settings allowing for a low quality option has two effects: it causes some consumers to move out of uninsurance and to purchase low-quality coverage, and it simultaneously causes some consumers to move out of H and into L.

We then analyze the consequences of an insurance mandate or penalty for remaining uninsured. We show that strengthening the penalty again has two consequences. First, it causes some relatively low cost consumers who would have remained uninsured to opt to purchase L. Second, because the penalty leads to a lower price for L relative to H, it also causes some consumers who would have chosen H to move to L. The key implication here is that even if all consumers value L more than the social cost of providing it, it may not be optimal to impose a penalty large enough to induce all consumers to enroll in L due to the unintended consequence of worsening adverse selection on the intensive margin.

We then turn to the consequences of risk adjustment. We show that strong risk adjustment always leads to more enrollment in H. In fact, we show that risk adjustment is often necessary to

[2]Beginning in 2019, 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.

generate nonzero enrollment in H in many settings. However, in some (but not all) settings stronger risk adjustment can also cause the price of L to increase, inducing some consumers to move to U and remain uninsured. This result is consistent with recent structural work on the consequences of risk adjustment in other markets (Saltzman, 2017; Domurat, 2018). In other settings, however, the monetary transfers from L to H that are induced by risk adjustment are offset by selection of relatively sick consumers out of L, causing risk adjustment policies to improve intensive margin selection with minimal extensive margin consequences.

Finally, we explore the welfare implications of the trade-off between extensive and intensive margin selection. With demand and cost curves for L and H, we are missing only one sufficient statistic for welfare analysis: the social cost of uninsurance for each consumer type. We follow Finkelstein, Hendren and Shepard (2017) and assume that the social cost of insurance for a given consumer type is that type's expected spending when insured (i.e. more expensive types have a higher social cost of being uninsured) multiplied by some constant, ϕ. We then present the welfare consequences of each of the policy options for a variety of values of ϕ. The choice of ϕ can be conceptually linked to a state regulator's assessment of which problem is more pressing: the uninsurance rate, or selection on the intensive margin leading to an inadequate offering of generous plans. We show that in our setting, under a variety of counterfactual subsidy regimes a high penalty for uninsurance almost always *decreases* welfare relative to the case of no mandate/penalty no matter the value of ϕ, implying that the welfare losses from moving consumers from H to L (intensive margin selection) more than offset the welfare gains reducing the uninsurance rate. This is a critical result that would be missed by more typical frameworks that focus only on one margin of selection. Risk adjustment, on the other hand, appears to increase welfare in almost all settings. The welfare consequences of benefit regulation (i.e. elimination of L) ultimately depend on ϕ.

We see this paper as providing an important link between two strands of the economics literature on adverse selection that are largely disconnected. The first strand focuses on the consequences of adverse selection for the sorting of consumers between insurance and uninsurance. This literature has focused largely on the individual health insurance market (Hackmann, Kolstad and Kowalski, 2015; Tebaldi, 2017) as well as markets for other types of insurance such as life insurance (Hendren, 2013) and long-term care insurance (Finkelstein and McGarry, 2006). The second strand focuses on the consequences of adverse selection for the sorting of consumers between more and less generous

insurance options. This literature has focused largely on the employer-provided health insurance market (Einav, Finkelstein and Cullen, 2010; Bundorf, Levin and Mahoney, 2012; Geruso, 2017).[3] Our paper shows that in many markets, the ideas from these two literatures are intimately connected, and policies intended to intervene on one margin may have unintended consequences for the other.

This idea builds on important insights in Azevedo and Gottlieb (2017) and Saltzman (2017). Our paper complements this small body of prior work by providing an intuitive graphical framework and by showing that the intensive-extensive margin trade-off is general across many policy interventions. This carries the practical implication that regulators may need to decide which type of problem is most pressing in their market: the uninsurance rate, or steep prices for plans with generous networks and coverage.[4] Further, our framework describes the conditions under which risk adjustment on the intensive margin will or will not have the unintended consequence of exacerbating selection out of the market into uninsurance. Understanding such conditions is key to the continuing reform of health insurance market regulation.

Today, states are being given increasing regulatory flexibility over selection-related policies, and our framework addresses the expected and unintended consequences of tuning such regulations.[5] For example, our framework makes clear why regulation that eliminates low-quality, cream-skimming plan options could have the important negative side effect of driving up prices and forcing some consumers out of the market. The previous literature, which has focused on either the intensive *or* the extensive margin, offers no guidance on this tradeoff.

We also see this paper as contributing to a related, but separate literature on the effects of adverse selection on contract design (Geruso and Layton, 2017). This literature shows that adverse selection can lead insurers to neglect to offer socially efficient contracts in equilibrium (Rothschild and Stiglitz, 1976; Glazer and McGuire, 2000; Veiga and Weyl, 2016). There is empirical evidence that this occurs in practice (Carey, 2017*a*,*b*; Lavetti and Simon, 2016; Shepard, 2016; Geruso, Layton and Prinz, 2018).

[3] Though see Cohen and Einav (2007) for a treatment in the auto insurance market.

[4] Saltzman (2017) is closely related in showing that risk adjustment can exacerbate the problem of uninsurance; we show that this type of phenomenon is not unique to risk adjustment. It applies to any policy targeting intensive margin sorting through explicit or implicit cross-plan transfers (including the ACA's transitional mandatory reinsurance program) or that targets plan quality by directly regulating acceptable benefits (such as the ACA's Essential Health Benefits rules). Likewise, we show that policies aimed at moving consumers into the market, such as mandates or tax penalties, can have undesirable effects on sorting *within* the market.

[5] For example, states were recently given guidance by CMS that they could choose to decrease the size of risk adjustment transfers according to local market conditions. However, to our knowledge, prior to this paper there has been no research exploring the trade-offs/consequences involved with strengthening/weakening these transfers. In a companion white paper, we describe how states can use the insights from this paper to make actual simple policy modifications that can allow states to optimally thread the needle of the extensive/intensive margin trade-off (Geruso and Layton 2018).

However, this literature has not previously incorporated the outside option of uninsurance. The lesson here is that neglecting 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.

Finally, we contribute to a literature that focuses on the design of the ACA Health Insurance Marketplaces. 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 framework shows that many of these proposals may have previously unrecognized unintended consequences for adverse selection on the intensive margin, potentially lowering the quality of coverage available in this market while simultaneously decreasing the number of uninsured.

2 Model

Consider an insurance market with two fixed insurance contracts, H and L, where H is more generous than L. There is also an outside option of uninsurance, U, and a mandate penalty M for uninsured consumers.

2.1 Demand

Let $W_{i,H}$ be willingness-to-pay (WTP) of consumer i for plan H, and $W_{i,L}$ be WTP for L, both defined as WTP relative to U (so $W_{i,U} = 0$). We adopt the vertical model of Finkelstein, Hendren and Shepard (2017). While the vertical assumptions are not necessary for many of our insights to hold, they make the presentation of our key conceptual contributions much clearer. They are also likely to be valid in our empirical setting where all consumer types appear to prefer H to L in the absence of prices. The vertical model applies well to settings where plan rankings are clear – e.g., a low-deductible vs. high-deductible plan, or a narrow vs. complete provider network – but would work less well when plans are significantly horizontally differentiated.

The assumptions of the vertical model are:

1. **Vertical preferences:** $W_{iH} > W_{iL}$ for all i

2. **Single index of WTP heterogeneity:** Consumers are indexed by a WTP type $1 - s \in [0, 1]$, with $W_L'(1-s) > 0$ and $W_H'(1-s) - W_L'(1-s) > 0$ globally.

The first assumption implies that all consumers would choose H over L if the prices of H and L were equal. The second assumption implies that consumer demand heterogeneity can be summarized along a single index $1-s$, where higher $1-s$ types have higher WTP for L relative to U *and* for H relative to L.[6] If (as we will assume below) costs are decreasing in s, this implies that L is adversely selected relative to U and H is adversely selected relative to L – i.e. there is adverse selection on both the extensive and intensive margins. In this way, we naturally extend the assumptions typical of the prior literature to the case of simultaneous two-margin selection.

Let P_j be the price of plan $j \in \{L, H\}$ set by the insurer and P_j^{cons} be the post-subsidy consumer premium. In the simplest case, $P_j^{cons} = P_j - Subs$ for some subsidy $Subs$, though one can imagine alternative subsidy rules as well.[7] Under consumer prices $P^{cons} = \{P_H^{cons}, P_L^{cons}, M\}$ that ensure positive demand for each option (see Finkelstein, Hendren and Shepard, 2017 for the conditions), it is simple to show that types sort into plans based on WTP rank $(1-s)$, with the highest-WTP types choosing H, intermediate-WTP types choosing L, and lowest-WTP types choosing U. Let $s_{HL}(P^{cons})$ be the cutoff type who is indifferent between H and L, and $s_{LU}(P^{cons})$ be the analogous cutoff type indifferent between L and U. Formally, these are defined by:

$$\Delta W_{HL}\left(s_{HL}\left(P^{cons}\right)\right) = P_H^{cons} - P_L^{cons} \tag{1}$$

where $\Delta W_{HL}(s) = W_H(s) - W_L(s)$, and

$$W_L\left(s_{LU}\left(P\right)\right) = P_L^{cons} - M \tag{2}$$

Note that s_{HL} is a function only of $\Delta P^{cons} \equiv P_H^{cons} - P_L^{cons}$, while s_{LU} is a function only of $P_L^{cons} - M$. Without loss of generality, let the indexing variable s be uniformly distributed on [0,1]. Then demand

[6] By construction, we can define $1-s$ to index increasing WTP for L vs. U. The substantive assumption is that this same index perfectly captures increasing WTP for H relative to L.

[7] In the empirical section, we simulate a number of subsidy rules including a single fixed subsidy for purchasing insurance, a single subsidy linked to the price of the lowest price plan, and plan-specific subsidies linked to the each plan's own price.

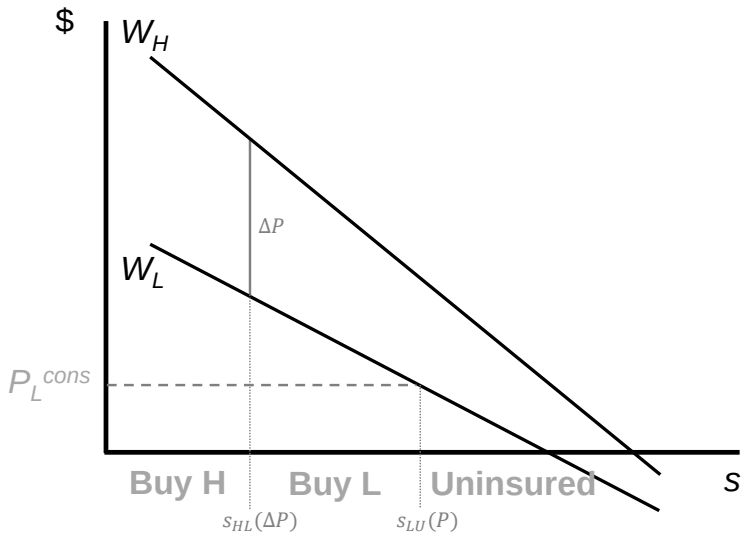
is defined by these cutoff types:

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

It will also be convenient to refer to the demand for formal insurance, $D_{Ins}\left(P^{cons}\right) = 1 - D_U\left(P^{cons}\right)$.

This setting is illustrated in Figure 1, which comes from Finkelstein, Hendren and Shepard (2017). The figure shows the $W_L(P^{cons})$ and $W_H(P^{cons})$ curves, describing WTP for H and L for each s-type. The curves give the price of L at which each s-type consumer is indifferent between purchasing L vs. U, and similarly the price difference between H and L, ΔP^{cons}, at which each s-type is indifferent between H vs. L. The figure illustrates that the $W_L(P^{cons})$ curve indicates the cutoff type, s_{LU}, for a given price of L, P_L^{cons}. Similarly, the $W_{HL}(P^{cons})$ curve indicates the cutoff type, s_{HL}, for a given incremental price, ΔP^{cons}, of H vs. L. In the figure, the consumers in the region of the x-axis labeled "Buy H" purchase H, the consumers in the region labeled "Buy L" purchase L, and the consumers in the region labeled "Uninsured" take up neither.

Figure 1: Consumer Sorting under Vertical Model



To preview one of the insights of this paper, we note here that both the price of L and the incremental price difference between L and H 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. This implies that what happens on the margin of insurance/uninsurance can affect the sorting of consumers across the H/L margin and vice versa. For example, adjusting a mandate penalty that only induces some consumers from uninsurance into the L plan could impact the take-up and 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

Let C_{ij} be the expected insurer cost of consumer i in plan $j \in \{L, H\}$. Type-s-specific insurer costs are defined as:

$$C_j(s) = E\left[C_{ij} \mid s_i = s\right] \tag{3}$$

$C_j(s)$ is analogous to "marginal costs" in the Einav, Finkelstein and Cullen (2010) setting. However, the term "marginal" is less useful in this setting (especially for L) since there are two margins of adjustment. In addition we define $C_U(s)$ as the expected costs of uncompensated care (incurred by third parties) of type-s consumers if uninsured. In addition to adverse selection, uncompensated care costs motivate the mandate penalty.

Given the demand curves defined above, we can also define average costs $AC_j(P^{cons})$ of enrollees in each plan j. For H this is:

$$AC_H(P^{cons}) = \frac{1}{D_H(P^{cons})} \int_0^{s_{HL}(P^{cons})} C_H(s)\, ds \tag{4}$$

For L, average costs are:

$$AC_L(P) = \frac{1}{D_L(P^{cons})} \int_{s_{HL}(P^{cons})}^{s_{LU}(P^{cons})} C_L(s)\, ds \tag{5}$$

Recall here that $s = 0$ corresponds to the highest willingness-to-pay types.

2.3 Competitive Equilibrium

We consider competitive equilibria where at the equilibrium price vector, P^*, the average cost of each plan's enrollees is equal to the plan's price.[8]

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

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 the appendix.

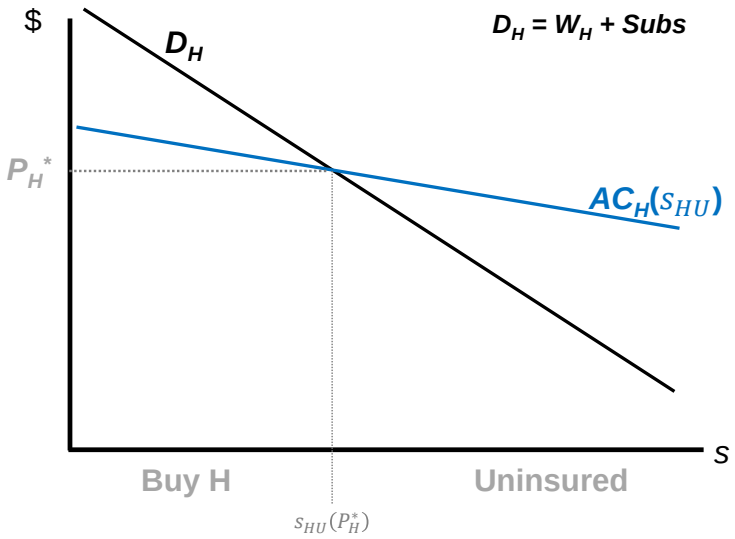
With the outside option of uninsurance, the equilibration process for the prices of H and L differs somewhat from the more familiar settings explored by Einav, Finkelstein and Cullen (2010) and Handel, Hendel and Whinston (2015). In those settings, it is assumed that all consumers choose either H or 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).

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. This case is analogous to the classic case considered by Einav, Finkelstein and Cullen (2010) and it is illustrated in Figure 2. As in Figure 1, the single index of WTP heterogeneity, s, is plotted on the x-axis, and the price P_H is plotted on the y-axis. The demand curve for H, D_H gives the (pre-subsidy) price at which each s-type is indifferent between enrolling in H and remaining uninsured, equal to $W_H(s) + Subs$. The average cost curve for H gives

[8] 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.

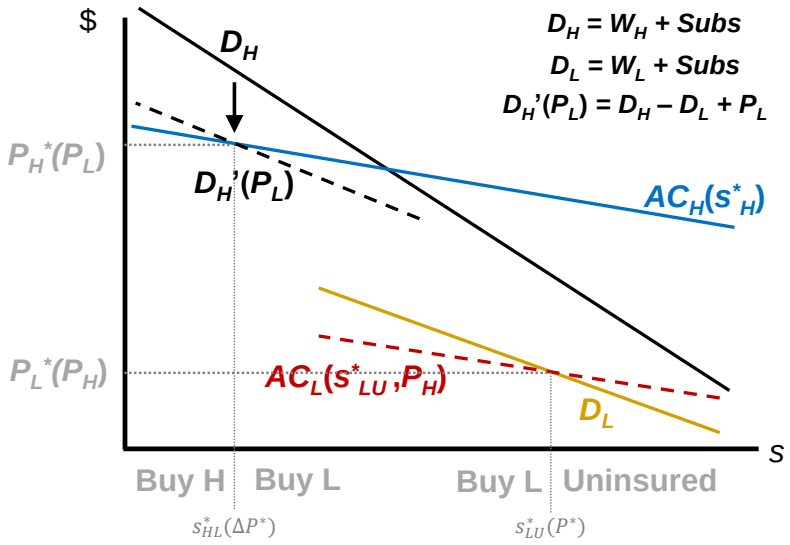
the average cost of the consumers enrolling in H for a given cutoff s-type, $s_{HU}(P_H)$. The equilibrium price, P_H^* and cutoff s-type, s_{HU}^* are found where the average cost curve crosses the demand curve as illustrated in the figure.

Figure 2: Competitive Equilibrium with H Only



In Figure 3 we add L. Adding L causes a number of significant changes to the figure and the equilibration process. First, the relevant demand curve for H is no longer based on the price at which a given s-type is indifferent between H and U. Instead, what matters for sorting is the price at which a given s-type is indifferent between H and L. Instead of being given by $D_H(s) = W_H(s) + Subs$ this price is given by $D_H^{'}(s, P_L) = W_H(s) - W_L(s) + P_L$. It is critical to note that this new demand curve is not fixed but is instead an equilibrium object because it is a function of P_L which is determined in equilibrium. We indicate this in Figure 2 by drawing $D_H^{'}$ (below D_H) as a dashed line. Note, however, that the average cost curve for H, $AC_H(s_H)$, is fixed and does not depend on L because for a given s_H all consumers to the left of s_H enroll in H.

Figure 3: Equilibrium with H, L, and the Outside Option



The second change is that it is now necessary to include demand and average cost curves for L. The demand curve for L gives the price at which a given s-type is indifferent between L and U, given by $D_L(s) = W_L(s) + Subs$. It is fixed in that it does not depend on P_H.[9] The average cost curve for L gives the average cost of the consumers enrolling in L for a given cutoff s-type *at a given* P_H. This curve is not fixed but is instead an equilibrium condition. This is because as P_H increases, consumers move from H to L, shifting the average cost curve for L upwards because the marginal consumers are sicker than the average L enrollee.[10]

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 single price difference). The dashed lines in Figure 2—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.

[9] The demand curve for L is undefined for some s types whose incremental WTP for H vs. L is large enough that there is no P_L that would induce them to enroll in L given P_H. The undefined region is a function of P_H, making the curve technically an equilibrium object, but because the slope and intercept of the curve do not depend on P_H we simplify by calling this curve fixed.

[10] Again, the average cost curve for L is only defined over the region for which the demand curve for L is defined.

2.3.1 Equilibrium under Regulation

We now use our model to determine the consequences of two policies commonly used to combat adverse selection in insurance markets: a mandate/penalty for remaining uninsured and risk adjustment transfers from plans attracting healthier than average consumers to plans attracting sicker than average consumers. The goal of uninsurance mandates/penalties is to weaken selection on the extensive margin, while the goal of risk adjustment is to weaken selection on the intensive margin. In each case we focus on the consequences of the policies for P_H and P_L and, thus, how consumers sort across plans.

Mandates/Uninsurance Penalties We start by considering the effects of an increase in the penalty for being uninsured, M. We start by differentiating the equilibrium price equations from Equation 6 with respect to M, which gives (proof in appendix)

$$\begin{aligned}\frac{dP_L}{dM} &= \left[\frac{\partial AC_L}{\partial M}\right] \times \Phi_L^{-1} \\ \frac{dP_H}{dM} &= \left[\frac{\partial AC_L}{\partial M} \cdot \frac{\partial AC_H}{\partial P_L}\left(1 - \frac{\partial AC_L}{\partial P_L}\right)^{-1}\right] \times \Phi_H^{-1}\end{aligned} \tag{7}$$

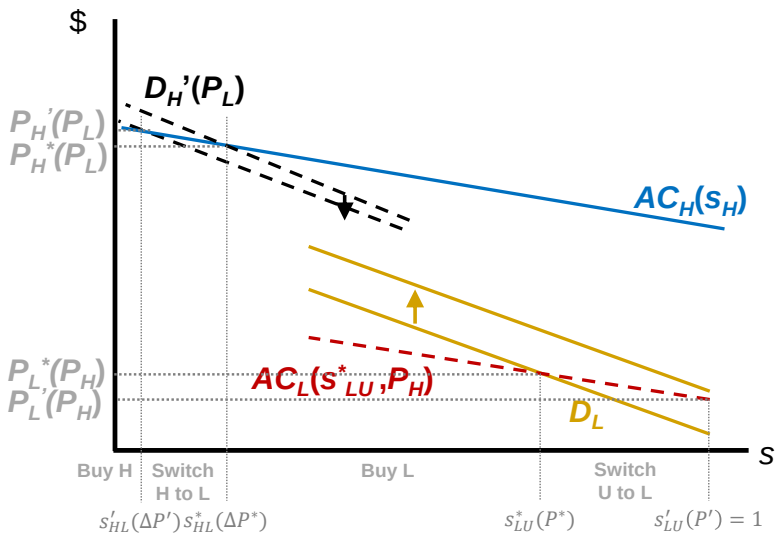
where $\Phi_H = 1 - \frac{\partial AC_H}{\partial P_H} - \frac{\partial AC_H}{\partial P_L}\frac{\partial AC_L}{\partial P_H} \cdot \left(1 - \frac{\partial AC_L}{\partial P_L}\right)^{-1}$ and $\Phi_L = 1 - \frac{\partial AC_L}{\partial P_L} - \frac{\partial AC_L}{\partial P_H}\frac{\partial AC_H}{\partial P_L} \cdot \left(1 - \frac{\partial AC_H}{\partial P_H}\right)^{-1}$. The first expression shows the effect of an increase in M on P_L. If there is adverse selection on the extensive margin the portion in the brackets, $\frac{\partial AC_L}{\partial M}$, must be negative. It is also straightforward to show that Φ_L must be positive for any stable equilibrium, implying that an increase in M unambiguously decreases the price of L. This is the (intuitive) direct effect of the penalty: a larger M results in a group of relatively healthy consumers joining L, driving down the average cost and thus the price of L.

The second expression shows the effect of an increase in M on P_H. Again, if there is adverse selection on the extensive margin, we know that $\frac{\partial AC_L}{\partial M}$ must be negative. Likewise, if there is adverse selection on the intensive margin, we know that $\frac{\partial AC_H}{\partial P_L}$ must also be negative. Finally, if the equilibrium is stable, $\frac{\partial AC_L}{dP_L}$ must be less than 1, implying that $1 - \frac{\partial AC_L}{dP_L}$ must be positive and that Φ must be positive. Taken together, these imply that an increase in M will unambiguously increase the price of H. Here, there is no direct effect of the penalty, as the penalty only directly applies to people on the margin of L vs. U. The positive effect of M on P_H is an indirect effect, where as an increase in M leads to a decrease in P_L, consumers on the margin between H and L move to L. Because these marginal

consumers are healthier than the inframarginal H enrollees, this leads to a higher average cost in H and thus a higher price.

Figure 4 provides graphical intuition for these mathematical results. An increase in the mandate penalty is equivalent to an upward shift in the demand curve for L. We consider the imposition of a penalty that is large enough such that D_L is always everywhere above AC_L at the equilibrium P_H. The direct effect of the penalty can be seen as the intersection of D_L and AC_L shifts to the right, leading to a lower price of L, as described by Equation 7. This lower price of L will lead to a downward shift in the demand curve for H, as illustrated in the figure. This downward shift in D_H causes the intersection of D_H and AC_H to shift left, leading to a higher price of H, as described by Equation 7. This is the indirect effect of the mandate.[11]

Figure 4: Equilibrium with an Insurance Mandate



This shows the trade-off between intensive and extensive margin selection. While the goal of a large uninsurance penalty is to weaken selection on the extensive margin by pushing some consumers out of uninsurance and into L, it has the unambiguous unintended consequence of worsening adverse selection on the intensive margin by causing other consumers to move out of H and into

[11] We note that there are additional equilibration effects that may weaken the direct and indirect effects. Specifically, as the price of H increases, this will cause the AC_L curve to shift up, potentially causing the price of L to move back toward P_L^*. However, as Equation 7 shows, P_L will never increase past P_L^* as such a price increase would require that $\frac{\partial AC_L}{dP_L}$ which would violate our requirement that the equilibrium be stable.

L, making the overall welfare consequences of such a policy unclear.[12]

Risk adjustment We now consider the effects of risk adjustment transfers. Risk adjustment transfers are based on individual risk scores computed from diagnoses appearing in health insurance claims (Geruso and Layton, 2015). In the ACA Marketplaces, the transfer from L to H is determined according the following formula:[13]

$$T\left(P^{cons}\right) = \left(\frac{\overline{R}_H(P^{cons}) - \overline{R}_L(P^{cons})}{\overline{R}(P^{cons})}\right) \cdot \overline{P}(P^{cons}) \tag{8}$$

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. 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.

With risk adjustment, prices are set equal to average cost *net of transfers* instead of to pure average cost:

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

To determine how prices are affected by risk adjustment transfers, we consider the effects of a small

[12] In the appendix, we show that the effects of the mandate on the intensive margin are not unambiguous in the presence of risk adjustment. With risk adjustment, pushing healthier people into L has a spillover effect on the price of H by increasing the gap between the risk scores of the H and L enrollees thus increasing the transfer from L to H. If the transfer effect dominates the substituton effect (people moving from H to L) then the price of H could fall.

[13] 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.

increase in α by differentiating Equation 9 with respect to the prices of H and L:

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

$$\frac{dP_L}{d\alpha} = T(\cdot) \times \left[\underbrace{1}_{\text{Direct (+)}} + \underbrace{\left(-\frac{\partial AC_L^{RA}}{\partial P_H}\right)\left(1 - \frac{\partial AC_H^{RA}}{\partial P_H}\right)^{-1}}_{\text{Substitution to } H \text{ (-)}} \right] \times \Phi_L^{-1} \tag{10}$$

where $\Phi_H = 1 - \frac{\partial AC_H^{RA}}{\partial P_H} - \frac{\partial AC_H^{RA}}{\partial P_L}\frac{\partial AC_L^{RA}}{\partial P_H} \cdot \left(1 - \frac{\partial AC_L^{RA}}{\partial P_L}\right)^{-1}$ and $\Phi_L = 1 - \frac{\partial AC_L^{RA}}{\partial P_L} - \frac{\partial AC_L^{RA}}{\partial P_H}\frac{\partial AC_H^{RA}}{\partial P_L} \cdot \left(1 - \frac{\partial AC_H^{RA}}{\partial P_H}\right)^{-1}$. The first expression shows the effect of stronger risk adjustment on the price of H. The two components inside the brackets represent the direct and indirect effects of strengthening risk adjustment. The first component represents the fact that if the transfer (to H) is positive a larger transfer implies a lower price, and if the transfer is negative a larger transfer implies a higher price. The second component represents the indirect effect. With adverse selection on *costs net of risk adjustment* on the intensive margin, $\frac{\partial AC_H^{RA}}{\partial P_L}$ must be negative, and with stability $\left(1 - \frac{\partial AC_L^{RA}}{\partial P_L}\right)$ must be positive, implying that the overall indirect effect must be negative.[14] Intuitively, as the transfer to H gets larger and pushes down the price of H, consumers who are marginal to H vs. L move into H. Because these marginal consumers are healthier than the inframarginal H enrollees, they drive down the average cost of H and thus the price of H decreases. As in the case of the mandate, Φ_H must be positive at any stable equilibrium, implying that the entire expression must be negative if the transfer to H is negative and positive if the transfer to H is positive. Given that the transfer to H will be positive as long as there is adverse selection into H on risk scores which is highly likely to be the case, the expression indicates that increasing the strength of risk adjustment will lower P_H. This is the intended consequences of risk adjustment.

The second expression describes the effect of stronger risk adjustment on the price of L. Again, the two components inside the brackets represent the direct and indirect effects of strengthening risk adjustment. However, this time the effects have the opposite signs. Again, the first component is the direct effect, representing the fact that if the transfer (to H) is positive a larger transfer implies

[14] Adverse selection on costs net of risk adjustment differs from adverse selection on gross costs. Here, it must be the case that there is adverse selection on the "residual" costs not explained by risk adjustment. Because risk adjustment is imperfect, it is likely often the case that when there is adverse selection on gross costs there is also adverse selection on residual costs (Layton, 2017). Indeed, we find this to be the case in our empirical application below.

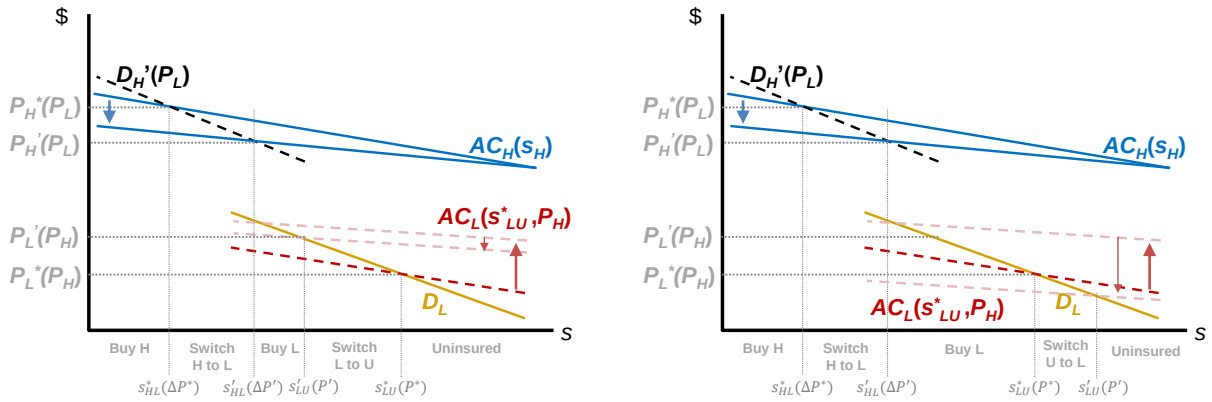
a higher price, and if the transfer is negative a larger transfer implies a lower price. The second component represents the indirect effect. With adverse selection on *costs net of risk adjustment* on the intensive margin, $\frac{\partial AC_L^{RA}}{\partial P_H}$ must be negative, and with stability $\left(1 - \frac{\partial AC_H^{RA}}{\partial P_H}\right)$ must be positive, implying that the overall indirect effect must be negative. Intuitively, as the transfer away from L gets larger and pushes up the price of L, consumers who are marginal to H vs. L move into H. Because these marginal consumers are sicker than the inframarginal L enrollees, they drive down the average cost of L and thus the price of L decreases. Note that both $T(\cdot)$ and Φ_L remain positive, so this "substitution effect" can partially or fully offset the direct effect and potentially lead to settings where strengthening risk adjustment lowers the price of L, despite leading to larger transfers from L to H. Indeed, we find this to be the case in some of the counterfactual simulations we explore empirically below. Thus, the effect of risk adjustment on P_L is ambiguous.[15]

Figure 5 provides graphical intuition for these results. The left panel illustrates the setting where the direct effect dominates, while the right panel illustrates the setting where the indirect substitution effect dominates. In both settings, risk adjustment flattens AC_H and AC_L. Thus, the price of H drops in both cases. In both cases, there is also a larger transfer away from L, initially causing AC_L to shift up (fat arrow). However, when AC_H rotates and the price of H falls, a group of relatively sick marginal consumers move from L to H, shifting AC_L back down (skinny arrow). The left panel depicts the case where the downward shift is not enough to fully offset the transfer, resulting in a higher price of L under strengthened risk adjustment. The right panel depicts the case where the downward shift is enough to fully offset the transfer, leading to a lower level of AC_L and a lower price of L under strengthened risk adjustment.[16]

[15] Again, the extensive margin effect appears in Φ_L. Increasing adverse selection on the extensive margin will weaken the overall effect of increasing α, but it can never flip the sign of Φ_L because doing so would violate our stability condition

[16] Note that there will also be equilibration effects via shifts in D_H but as we show in Equation 10 these shifts will never be enough to cause P_H to increase.

Figure 5: Equilibrium under Risk Adjustment



Notes:

Note that in both cases, risk adjustment increases the portion of people choosing H but the implications for the portion choosing L vs. U differ. When the direct effect dominates (left panel), the price of L increases under risk adjustment, driving some consumers out of the market and into uninsurance. When the substitution effect dominates (right panel), the price of L decreases under risk adjustment, leading some uninsured consumers to choose to become insured by purchasing L. Thus, the trade-off between intensive and extensive margin selection we describe above may not always hold: In some settings, increasing the strength of risk adjustment transfers may lead to more consumers choosing H and fewer consumers choosing to be uninsured.

To summarize: whereas a mandate has the tendency to increase the fraction of the population insured at the cost of reducing the fraction of the population with higher quality coverage, risk adjustment has the tendency to increase the fraction of the population with higher quality coverage, potentially (but not necessarily) at the cost of increasing the uninsured population. The innovation of modeling an uninsurance option alongside a high and low plan is that it allows us to capture these tradeoffs. Such tradeoffs are by construction assumed away when applying the models from the prior litearture (e.g., Einav, Finkelstein and Cullen, 2010; Handel, Hendel and Whinston, 2015) to a market like the ACA Exchanges.

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:

$$\widehat{SW}\left(P^{cons}\right) = \int_{0}^{s_{HL}(P^{cons})} \left(W_H\left(s\right) - C_H\left(s\right)\right) ds + \int_{s_{HL}(P^{cons})}^{s_{LU}(P^{cons})} \left(W_L\left(s\right) - C_L\left(s\right)\right) ds - \int_{s_{LU}(P^{cons})}^{1} C_U\left(s\right) ds \quad (11)$$

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

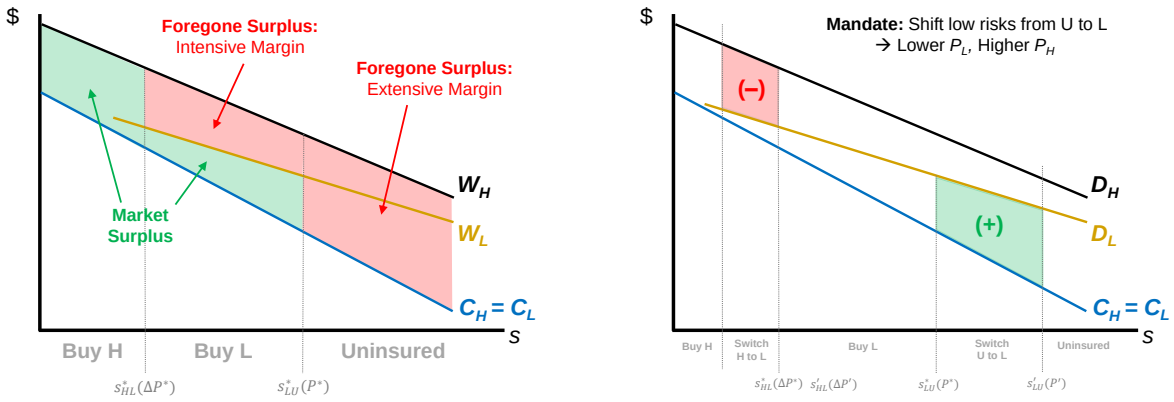
$$SW = \underbrace{\int_{0}^{s_{LU}(P^{cons})} \left(W_L\left(s\right) - C_L^{net}\left(s\right)\right) ds}_{\text{Net Surplus from Insurance in } L} + \underbrace{\int_{0}^{s_{HL}(P^{cons})} \left(\Delta W_{HL}\left(s\right) - \Delta C_{HL}\left(s\right)\right) ds}_{\text{Extra Surplus from } H} \quad (12)$$

where $\Delta C_{HL}\left(s\right) \equiv C_H\left(s\right) - C_L\left(s\right)$ and $C_L^{net}\left(s\right) \equiv C_L\left(s\right) - C_U\left(s\right)$. 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 12 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}^*. The left panel of Figure 6 illustrates this concept for the simple case where $C_L^{net} = C_L$ and L is a pure cream-skimming plan such that $C_L = C_H$. In the figure we plot W_L, W_H, and $C_H = C_L$. As above, the consumers who buy H are in the region of the x-axis to the left of s_{HL}^*. Equation 12 shows that for these consumers, social surplus is equal to the area between the W_H curve and the C_H curve, 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, again highlighted in green. Surplus foregone due to consumers enrolling in lower quality coverage than is optimal for them (highlighted in red) occurs for two reasons: (1) some consumers who value H more than the social cost of enrolling them choose L and (2) other consumers who also value H more than the social cost of enrolling them choose U. Importantly, this figure shows that all that is necessary for estimating welfare is estimates of W_H, W_L, C_H, and C_L, as well as the equilibrium cutoff s-types, s_{HL}^* and s_{LU}^*. Below, we show that all of these

curves are straightforward to recover given exogenous variation in the premiums of H and L.

Figure 6: Welfare



Notes:

In the right panel of Figure 6 we show the welfare consequences of risk adjustment. As shown above, risk adjustment causes P_L and P_H to converge. Sometimes (though not always) this leads to an increase in P_L. This increases s^*_{HL} (increasing enrollment in H) while simultaneously decreasing s^*_{LU} (increasing the uninsurance rate). These two effects have opposing welfare consequences in this setting, with the shift of consumers from L to H improving welfare while the shift of consumers from L to U worsens welfare. Clearly, this result is specific to the setting we illustrate here where it is optimal for all consumers to enroll in H. In other settings where some consumers value H or L below the cost of enrolling them, the welfare consequences of risk adjustment would differ. However, this example clearly illustrates the possibility of an efficiency trade-off, where policies (such as risk adjustment) aimed at weakening intensive margin selection lead to welfare increases by improving the quality of coverage for some consumers but also lead to welfare decreases by pushing some consumers out of the market. Similarly, this framework reveals that while policies (such as insurance mandates) aimed at weakening extensive margin selection do lead to welfare increases by getting more consumers to purchase insurance, they also may lead to welfare losses by causing some consumers to move to lower quality coverage. Without allowing for all three options (H, L, and U) this trade-off is missed, and it often appears like many policies meant to combat a particular kind of selection problem (i.e. extensive or intensive margin) are unambiguously welfare improving when in

fact the welfare consequences are ambiguous.

3 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 goal is to 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 give an overview of the construction of the demand and cost curves and describe our method for finding competitive equilibrium prices.

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, in L, or 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.[17]

3.1 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).

[17] 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.

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 in two groups: a more generous (broad-network) "H" option and less generous (narrower-network) "L" option. They use this vertical structure along with a regression discontinuity design leveraging discontinuous cutoffs in subsidy amounts based on household income to estimate demand and cost curves. This two-plan vertical demand structure maps neatly into our vertical model – indeed, their paper partly motivated our work – making these estimates a convenient way to parameterize our model. We discuss additional details about the CommCare market and FHS's estimates in the appendix.

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 extend their 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. For our baseline, we therefore set $C_L(s) = C_H(s)$, which implies that L is a pure cream-skimming plan without any efficiency advantage. In addition, we consider the case where L has a fixed percentage cost advantage for all enrollees: $C_L(s) = (1-\eta)C_H(s)$ for various values of η.

High-Income Unsubsidized Consumers: HKK (2015) We construct WTP curves for high-income households using 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 and cost 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 $W_L^{HI}(s)$ and $C_L^{HI}(s)$. 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. Further details of this method are discussed in the appendix.

3.2 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.

It is useful to define a baseline case to which all other simulations will be compared. For our baseline, the choice set includes all three options: H, L and U. We assume there is no mandate or

penalty for remaining uninsured. Finally, we assume that the regulator imposes risk adjustment transfers that follow a slightly modified version of the ACA formula where T_j represents the transfer payment paid to plan $j \in H, L$:

$$T\left(P^{cons}\right) = \left(\frac{\overline{R}_H(P^{cons}) - \overline{R}_L(P^{cons})}{\overline{R}(P^{cons})}\right) \cdot \overline{P}(P^{cons}) \tag{13}$$

In the reaction function approach, we compute these transfers for each price vector and define breakeven price vectors as those for which the plan prices are equal to the plan risk-adjusted average cost, where the risk adjusted average cost is equal to $RAC_j = AC_j + T_j$, as we describe in detail in Section 2.3.1.

3.2.1 Simulation 1: Benefit Regulation

The simplest policy we examine is the use of benefit regulation. The intention of benefit regulation is to weaken intensive margin selection either by forcing L to look more like H (thus, decreasing $W_{HL}(s)$ 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 on the margin between L and U out of the market. Thus, benefit regulation may increase consumer utility either by raising the utility of those enrolled in L (as L becomes more generous) or by moving marginal consumers from the lower-utility L plan to the higher-utility H plan, but it may also decrease consumer utility 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 vs. 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.

3.2.2 Simulation 2: Mandate/Uninsurance Penalties

Next, we simulate the effects of an insurance mandate or penalty for choosing U. The intention 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. As shown in Section 2, because the 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 consumer utility by inducing some consumers to purchase insurance, but they may simultaneously decrease consumer utility by inducing other consumers to move from the higher-utility H plan to the lower-utility 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 $P_U = 0, 5, 10, ..., 60$.[18]

3.2.3 Simulation 3: Risk Adjustment Transfers

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 2, 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 vs. U margin to exit the market.[19] Thus, risk adjustment transfers may increase consumer utility by moving some consumers from L to H, but they may simultaneously decrease consumer utility by inducing other consumers to exit the market.

To simulate the effects of risk adjustment, we vary the strength of the risk adjustment transfers by introducing a strength parameter α to the transfer formula as follows:

$$T\left(P^{cons}\right) = \left(\frac{\overline{R}_H(P^{cons}) - \overline{R}_L(P^{cons})}{\overline{R}(P^{cons})}\right) \cdot \overline{P}(P^{cons})\alpha \tag{14}$$

Values of α smaller than 1 represent risk adjustment policies that are weaker than the baseline 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 up or 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

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

[19] Recall from Section 2.3.1 that risk adjustment need not always raise the price of L due to the presence of the offsetting "direct" and "substitution" effects of the transfers.

policy simulations. We find equilibrium prices and market shares for $\alpha = 0, 0.1, 0.2, ...2$.[20]

3.3 Subsidy Regimes

We consider each of the three policy simulations under each of several subsidy regimes. Pre-2014, Massachusetts used a different subsidy policy than the one currently in place in for low-income individuals in the ACA Marketplaces. Additionally, high-income households face a different subsidy regime than low-income households. Finally, in recent years policymakers have called for reforms that would amend the new ACA subsidy regime. Because the overarching goal of our exercise is to shed new light on the context-general tradeoff between intensive and extensive margin selection (rather than to focus on precise quantitative results for the CommCare environment per se), we run each of our policy simulations separately under the following subsidy regimes and examine outcomes under a broad spectrum of relevant policies.[21]

3.3.1 Subsidy regime 1: ACA subsidies, no unsubsidized population

Our first subsidy regime mimics the subsidy policies from the ACA but focusing 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. 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.

[20] We find that in all cases the equilibrium price vector does not change with increases in α for $\alpha > 2$.

[21] In addition to the subsidy regimes described below, in the appendix we include results where we assume subsidies are set according to the rules in place in the Massachusetts Connector during our sample period. We do not report these results in the main text because this subsidy regime involved incremental subsidies (i.e. larger subsidies for H vs. L) in addition to the base subsidy, weakening the trade-off between intensive and extensive margin selection. Given that Massachusetts has abandoned this subsidy design and no other state has implemented a similar design, we determined that it was not as relevant as the regimes described below. We report results from the CommCare regime in the appendix.

3.3.2 Subsidy regime 2: ACA subsidies, some portion unsubsidized

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 section. The cost and demand curves of these low-income individuals also remain the same as in the previous section. 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 the appendix. 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.[22]

3.3.3 Subsidy regime 3: Fixed subsidies

For this setting, a fixed subsidy amount is given towards the purchase of both the H and L plan. We simulate a numbe of different fixed subsidy amounts including the average cost of the entire population ($321.55), $300, $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.

3.4 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.

Calculating welfare in this setting is conceptually straightforward but has the practical difficulty that the social cost of remaining uninsured is unobserved. All components of the expression for social welfare (demand and cost curves for H and L) described by Equation 11 are observed, with the exception of the social cost of remaining uninsured, $C_U(s)$. This cost includes items like uncompensated care, care paid for by other state programs, and a general distaste for others being uninsured,

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

and it is necessary for considering the full welfare consequences of the policies we consider.

To characterize welfare without observing $C_U(s)$, we assume that $C_U(s)$ is a multiple of $C_H(s)$: $C_U(s) = \phi C_H(s)$. 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 11 for a variety of different values of ϕ ranging from $\phi = 0$ (zero social cost of uninsurance) to $\phi = 1$ (social cost of uninsurance equals cost of insuring the consumer in Plan H). Higher values of ϕ correspond to more weight being placed on the welfare of the uninsured. This exercise allows the reader to observe the welfare consequences of a given policy given her particular prior for the true value of ϕ.

4 Results

4.1 Benefit Regulation

Table 1 reports the results of the benefit regulation simulations. Columns labeled "With L Plan" present outcomes for the baseline setting where consumers can choose between H, L, and U. Columns labeled "No L Plan" 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. The next two columns add high-income unsubsidied consumers. The next two columns assume that subsidies are fixed and equal to the average cost across all consumers in the market (\$322.55 per month). The second row 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), benefit regulation has no effect. This is because no consumers choose to enroll in L when it is offered. This is likely due to the relatively strong risk adjustment policy in place in this market.

When higher-income unsubsidized consumers are added to the market (columns 3 and 4), if L is included in the choice set it attracts positive market share. The equilibrium prices of L and H are \$352 and \$379, respectively. 48% of the market chooses H, 24% chooses L, and 27% opts to remain uninsured. When L is eliminated, all of the consumers previously choosing L opt to enroll in H. Additionally, some consumers who previously chose U now choose to enroll in H, implying

unambiguous improvements for consumers, with more consumers choosing H and fewer consumers choosing U. This may seem to go against our argument that a fundamental trade-off exists between intensive and extensive margin selection. However, this result is due to the linkage of the subsidy to the 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.

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, no consumers opt to remain uninsured when L is offered. When L is eliminated, however, while some of the consumers who would have enrolled in L were it available choose to enroll in H instead, other consumers (37% of the market) opt to remain uninsured. This same pattern holds, though to a smaller degree, for a fixed subsidy of $300 per month. This clearly illustrates the trade-off between extensive and intensive marin 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. For smaller fixed subsidies ($275, $250), we find that the elimination of L has no effect on equilibrium market shares and prices. This is because the subsidy is not high enough to get the highest s-type consumers (those with the smallest incremental WTP for H, W_{HL}) to enter the market, so no consumers choose L even when it is included in the choice set.

4.2 Mandate/Uninsurance Penalties

The results for the uninsurance penalty simulations are presented in Figure A5 and Table 2. 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 2.

The top-left panel of Figure A5 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 portion of uninsured consumers drops. With no penalty, 30% of consumers choose to remain uninsured, but with a penalty of $60 per month, all consumers in the market choose to purchase insurance. More interesting are the effects of the penalty on enrollment in H. Initially, strengthening the penalty causes consumers to move from uninsurance to H (the marginal consumers have

relatively high incremental WTP for H, W_{HL}). However, as the penalty gets stronger, the consumers that move out of uninsurance have lower incremental WTP for H and choose instead to move from U to L. Because these consumers are healthy, the price of L is low, causing some consumers to move from H to L. As the penalty continues to strengthen, the additional marginal consumers choosing to purchase insurance are even healthier, causing the price of L to drop further, resulting in more consumers switching from H to L. With a penalty of \$60 per month, all consumers purchase insurance, with 76% choosing L and only 24% choosing H. This clearly illustrates 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 top-right panel of Figure A5 adds high-income unsubsidized consumers to the market. Here, the pattern is similar to the case without unsubsidized consumers, except here there is no increase in enrollment in H at low levels of the penalty. Instead, as the penalty increases the shares in U and H decrease monotonically, while the share in L increases monotonically, again clearly illustrating the extensive/intensive margin trade-off.

The bottom 4 panels of Figure A5 present outcomes for the fixed subsidy cases. For the fixed subsidy equal to the average cost in the population, there is no effect of the uninsurnace penalty because the subsidy is high enough that all consumers choose to enroll in insurance even without a penalty. With a subsidy of \$300 per month, when there is no penalty 59% of consumers opt to remain uninsured while the remainder choose to enroll in H, leaving L with zero enrollment. However, a small penalty of \$15 is sufficient to get all consumers to enroll, with the uninsurance rate dropping to zero. Importantly, however, the marginal consumers that the penalty pushes into the market (1) have low incremental WTP for H and (2) are very healthy, leading to positive enrollment in L at a low equilibrium price. At this low price of L some consumers (17%) move from H to L while other consumers (59%) move from U to L. This once more illustrates the extensive/intensive margin trade-off. A similar pattern is observed for the case of a \$275 subsidy, though the trade-off occurs at the higher \$35 penalty instead of at the lower \$15 penalty. This is because with the lower subsidy, a higher penalty is necessary to get the low-cost/low W_{HL} consumers to enroll. At the lowest subsidy (\$250), the trade-off doesn't occur with a penalty of \$60 or less, so in the range of penalties we simulate, larger penalties only move consumers from U to H.

4.3 Risk Adjustment

The results for the risk adjustment simulations are presented in Figure 8 and Table 3. 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 2 (transfers twice the size of the ACA transfers), as described by Equation 14. Prices, market shares, and subsidies are reported in Table 3.

The top-left panel of Figure 8 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 effectively unravels, with almost no enrollment in equilibrium. Consumers split between L (58%) and uninsurance (39%). As risk adjustment is strengthened, both the price of L and the price of H fall, while the uninsurance rate remains relatively steady. This implies that the substitution effect of sick consumers moving from L to H from equation 10 dominates the direct effect of larger transfers from L to H, resulting in declining costs in L. Importantly, however, the price of H declines more rapidly than the price of L, leading to a decline in the differential price, $P_H - P_L$ and an increase in the market share of H. This continues to occur until the market "upravels" to H, eliminating L. At this point, 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 and decreasing the uninsurance rate. Here, we see that risk adjustment addresses the intensive margin selection problem while not making things any worse on the extensive margin. It appears that there is no 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 due to the offsetting substitution and direct effects of strengthening risk adjustment transfers. The top-right panel of Figure 8 shows similar results when high-income unsubsidized consumers are added to the market.

The middle-left panel of Figure 8 presents outcomes for the case of a fixed subsidy equal to the average monthly cost across the entire population. Here, we again see that with no risk adjustment, H almost fully unravels, with only 12% of consumers enrolling in H in equilibrium. Initially, as risk adjustment is strengthened, market shares remain relatively stable: Both H and L decrease some, with the price of H decreasing faster than the price of L, resulting in a small increase in H's mar-

ket share and a small decrease in the uninsurance rate. However, once enough relatively healthy consumers move out of uninsurance and into the market, P_L drops rapidly causing H to once again unravel back to almost no enrollment. Importantly, however, H holds on to just enough of the sickest consumers in the market to allow L's price to remain low enough that the uninsurance rate falls to zero. After this point, strengthening risk adjustment again shifts consumers from L to H as intended, with no extensive margin consequences because no consumers are uninsured. However, at some point risk adjustment gets strong enough to cause L to upravel out of existence, at which point the low-price plan switches from L (which no longer exists) to H (the only option left), resulting in the healthiest consumers exiting the market and returning to uninsurance. This final shift from a market with no uninsurance and positive enrollment in H and L to a market split between H and U illustrates the trade-off between extensive and intensive margin selection: With weaker risk adjustment more consumers are insured, but some are enrolled in lower-quality insurance; but with stronger risk adjustment, some consumers become uninsured, but all insured consumers are enrolled in higher-quality insurance. This simulation shows that increasing the strength of risk adjustment transfers can have very different consequences at different levels of α, sometimes decreasing the uninsurance rate and shifting consumers from H to L and other time increasing the uninsurance rate and shifting consumers from L to H, making it clear that the intended effects of risk adjustment (shifting consumers from L to H) are often far from guaranteed.

The middle-right panel of Figure 8 presents outcomes for the case of a fixed subsidy equal to \$300. Here, we see that stronger risk adjustment has the intended effect, shifting consumers from L to H, with limited extensive margin consequences. With smaller fixed subsidies (\$275 and \$250), strengthening risk adjustment transfers has little or no effect on market shares, as only the consumers with the highest incremental WTP are in the market, and they never choose L, even at very low levels of α.

4.4 Cost Advantage for L

In many settings, the L plan will have a cost advantage over the H plan, i.e. the same person will have lower healthcare costs when they enroll in L vs. H. This could be due to higher cost sharing or tighter controls on healthcare utilization. In such a case, the consequences of policies may differ. For example, in our current setting where L is a pure cream-skimming plan, if there were no adverse

selection no consumer would choose to enroll in L in equilibrium because no consumer values L more than H and without adverse selection the prices of H and L would always be identical in competition. This suggests that with adverse selection, small changes in the price difference between H and L will lead to large shifts in enrollment between the two plans. Thus, as risk adjustment is strengthened, more money flows from L to H, but the subsequent shift of L's sickest consumers to H is so strong that it fully offsets the cost effects of the larger risk adjustment transfers, with stronger risk adjustment having no effect on the price of L. If, however, L has a cost advantage, we might not expect the substitution effect to be so strong: With a cost advantage, consumers can save more money by enrolling in L, leading to fewer consumers responding to a change in the price difference between H and L and causing the direct effect of the larger risk adjustment transfers to dominate the substitution effect. Then, larger transfers might lead to a higher price of L.

To determine whether the effects of policies differ when L has a cost advantage, we perform several additional counterfactual simulations. In all simulations, we assume that $C_L(s) = 0.85C_H(s)$, or that L has a 15% cost advantage.

4.4.1 Benefit Regulation with a Cost Advantage

We start with simulations of benefit regulation. These simulations are identical to those in Section 4.1 but with the new L cost curve. Results are reported in Table 4. For the two cases with ACA price-linked subsidies, we see that without benefit regulation very few (if any) consumers choose H. A substantial number (39% or 25%) of consumers also remain uninsured. When L is removed from the choice set we see that the price of H drops and the vast majority (70% or 77%) choose H, with the new H enrollees coming from both L and U so that the uninsurance rates drop to 30% and 23%. This result is similar to the results from Section 4.1. Again, the decrease in the uninsurance rate is due to the price-linked nature of the subsidy: When L is removed, the subsidy becomes linked to the (higher) price of H, increasing the subsidy for enrolling in insurance, and inducing more consumers to leave U and enter the market.

In the first fixed subsidy case (subsidy = average cost), all consumers choose H in the case with no benefit regulation. When L is eliminated from the choice set, however, consumers split between H and uninsurance. With a subsidy of \$300 per month, the opposite occurs: without benefit regulation all consumers again choose L but When L is removed from the choice set consumers split between

H and U. Here, the subsidy is not high enough to get all consumers to enroll in the higher-cost H plan when L is not available. We see similar results with other fixed subsidies ($275 and $250). All of these cases illustrate the trade-off between extensive and intensive margin selection, with the removal of the L option inducing some consumers to purchase higher-quality coverage while simultaneously inducing others to opt out of insurance altogether and go uninsured. This differs from the case where L was a pure cream-skimmer, where benefit regulation had no effect on the uninsurance rate for the fixed subsidy cases. This difference is likely due to the ability of the L plan to price lower when it has a cost advantage vs. when it does not.

4.4.2 Mandate/Uninsurance Penalties with a Cost Advantage

The results for simulations of uninsurance penalties when L has a 15% cost advantage are found in Table 5 and Appendix Figure A2. For the ACA price-linked subsidy cases, a higher uninsurance penalty leads to a lower uninsurance rate and higher enrollment in L. However, it also leads to lower enrollment in H (though enrollment in H was already very low without the penalty). When we add high-income unsubsidized consumers to the market, this trade-off becomes more evident, with larger penalties clearly decreasing the uninsurance rate while also decreasing H's market share. This is similar to the results with a cream-skimming plan, where the penalty induces more consumers to enroll in insurance, but these consumers are relatively healthy, driving down the price of L relative to the price of H and exacerbating selection on the intensive margin.

For most of the fixed subsidy cases, the subsidy is high enough that no consumers ever choose U, making uninsurance penalties unnecessary (though in all cases the market unravels to include only L). For the case with a fixed subsidy equal to $250, however, we again observe the extensive/intensive margin trade-off: As the penalty is increased, more people move from U to L, driving down the price of L relative to the price of H, causing people to move from H to L. Eventually, intensive margin adverse selection is so severe that there is a death spiral in H. At this point, the equilibrium price of L is low enough that all consumers choose to enroll in L. Again, higher penalties lead some consumers to enroll in insurance while simultaneously inducing others to move from higher-quality to lower-quality insurance.

4.4.3 Risk Adjustment with a Cost Advantage

The results for simulations of increasingly strong risk adjustment transfers when L has a 15% cost advantage are found in Table 6 and Appendix Figure A3. The ACA price-linked subsidy cases are again similar to the results from the cream-skimming plan case: Stronger risk adjustment transfers have little effect on the price of L but cause the price of H to drop rapidly, resulting in a re-sorting of consumers from L to H with little effect on the uninsurance rate. Only when L disappears entirely does the uninsurance rate move, at which point it drops due to the subsidy shifting from being linked to the (lower-priced) L plan to being linked to the (higher-priced) H plan.

With the fixed subsidies equal to \$275 and \$250 we see clear evidence of the extensive/intensive margin trade-off. In these cases, stronger risk adjustment induces consumers to move from L to H. Unlike the cream-skimming case, however, we see that stronger risk adjustment causes the price of L to *increase* rather than decrease (i.e. the direct effect dominates the substitution effect in Equation 10). This causes some consumers to exit the market. Here, the extensive/intensive margin trade-off is clear: Stronger risk adjustment raises the average quality of coverage purchased in the market, but at the cost of a much higher uninsurance rate.

Welfare

We have shown that in a variety of settings, there appears to be a trade-off between selection on the intensive and extensive margins. Stronger uninsurance penalties decrease the uninsurance rate but also decrease H market share. Benefit regulation sometimes causes more consumers to enroll in H but also increases the uninsurance rate. Stronger risk adjustment transfers sometimes increase H market share with little effect on the uninsurance rate, but other times increase H market share at the cost of a higher uninsurance rate. Many of these settings imply that the welfare consequences of these policy options are ambiguous.

We now present estimates of social welfare for each counterfactual simulation so as to assess the extent to which the unintended negative effect of each policy offsets the intended positive effect. As discussed in Section 3, we present welfare estimates for a variety of specifications of the cost of uninsurance. Specifically, we assume that $C_H(s)$: $C_U(s) = \phi C_H(s)$ and we show welfare estimates for values of ϕ between 0 (no social cost of uninsurance) and 1 (social cost of uninsurance = cost of insuring the consumer).

Welfare estimates for the benefit regulation simulations are found in Tables 1 and 4. Recall that

for the baseline case (Table 1), benefit regulation only affected equilibrium market shares for ACA subsidies with some unsubsidized case and the fixed subsidy equal to the average cost in the population case.[23] For the fixed subsidy case, benefit regulation has an unambiguously positive effect on welfare. For all values of ϕ, the welfare gain from moving consumers from L to H totally offsets the welfare loss from the increase in the uninsurance rate. For the ACA case, the welfare consequences are less clear. When the social cost of uninsurance is high, the benefits seem to outweigh the costs. But when the social cost of uninsurance is low, the sign flips and benefit regulation seems lead to a worse social outcome. When L has a cost advantage (Table 4), under ACA subsidy policies benefit regulation is clearly welfare reducing, with lower welfare estimates when L is eliminated from the choice set for all values of ϕ. With fixed subsidies, benefit regulation seems to lead to higher overall welfare with high fixed subsidies and lower overall welfare with low fixed subsidies, with results varying by the level of ϕ.

Welfare estimates for the mandate penalty simulations are found in Tables 2 and 5 as well as Figure 9. In all panels of Figure 9 the y-axis provides the welfare estimate, the level of the penalty is on the x-axis, and each line in the figure represents a different value of ϕ. The top row of Figure 9 shows the welfare estimates for the two ACA subsidy cases: without unsubsidized and with unsubsidized. The figures clearly show that in both cases, welfare is *decreasing* in the size of the uninsurance penalty *for all values of* ϕ. This implies that the unintended negative effect of the penalty—causing consumers to choose L instead of H—fully offsets the intended positive effect—decreasing the uninsurance rate—even when the social cost of uninsurance is equal to the cost of insuring a given consumer, likely an over-estimate of the true social cost. This result also holds for all but the highest value of ϕ for all of the fixed subsidy cases: In this market, larger uninsurance penalties almost always lead to lower overall welfare. The results in Table 5 show that this result also holds for most values of ϕ when L has a cost advantage.

Welfare estimates for the risk adjustment simulations are found in Tables 3 and 6 as well as in Figure 10. Here, the results for the baseline case (no cost advantage for L) seem to suggest that stronger risk adjustment transfers are almost always welfare-improving. Only in the case of the fixed subsidy equal to the average cost in the population is welfare lower with the strongest risk

[23] Welfare estimates are always negative because consumers value both H and L less than cost (Finkelstein, Hendren and Shepard, 2017). Importantly, however, social welfare is also negative when everyone is uninsured in all cases except when $\phi = 0$. Thus, welfare implications should be drawn from comparisons of welfare estimates across simulations rather than the actual levels reported in the table.

adjustment transfers than with no risk adjustment. When L has a cost advantage, however, this result does not always hold. Indeed, the results for the two ACA subsidy cases indicate that risk adjustment can be "too strong," leading to lower welfare when α exceeds 1.75. Under fixed subsidies, however, the results are more ambiguous. With fixed subsidies of \$275 and \$250 , stronger risk adjustment increases welfare for high values of ϕ while decreasing welfare for low values of ϕ. This is due to the fact that there is a clear extensive/intensive margin trade-off in these cases and the cost of rising uninsurance rates depends on the value of ϕ.

To summarize, the potential extensive/intensive margin trade-off make the welfare implications of common policies highly ambiguous. In our setting, we find that for higher penalties for uninsurance, the costs outweight the benefits and stronger penalties are welfare-decreasing. The welfare consequences of stronger risk adjustment transfers and benefit regulation are less-clear and depend on the subsidy regime, the cost advantage of L, and the social cost of uninsurance.

5 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 inherent trade-off between selection on the intensive and extensive margins. 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 on 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. These trade-offs are critical for understanding the full welfare implications of these policies.

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. 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

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

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

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.

Cohen, Alma, and Liran Einav. 2007. "Estimating risk preferences from deductible choice." *American economic review*, 97(3): 745–788.

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.

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 Timothy J. Layton. 2017. "Selection in Insurance Markets: Policy through the Lenses of Fixed and Endogenous Contracts." *Journal of Economic Perspectives*, Forthcoming.

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

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

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

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.

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."

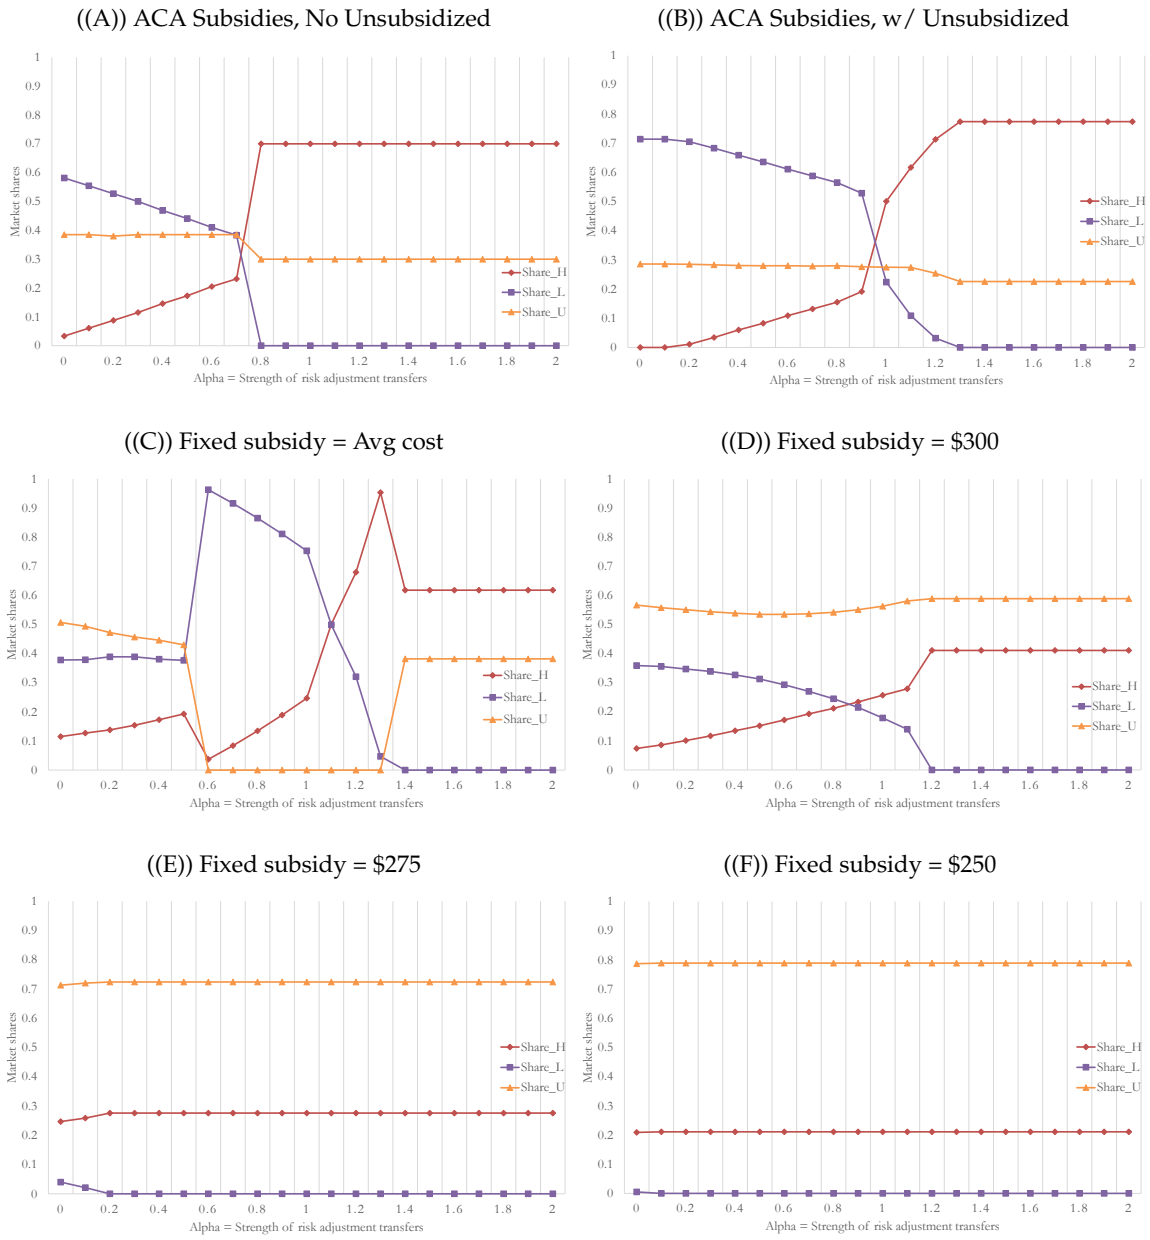
Veiga, André, and E. Glen Weyl. 2016. "Product Design in Selection Markets." *Quarterly Journal of Economics*, 131(2): 1007–1056.

Figure 7: Plan Market Shares under Increasingly Larger Uninsurance Penalty



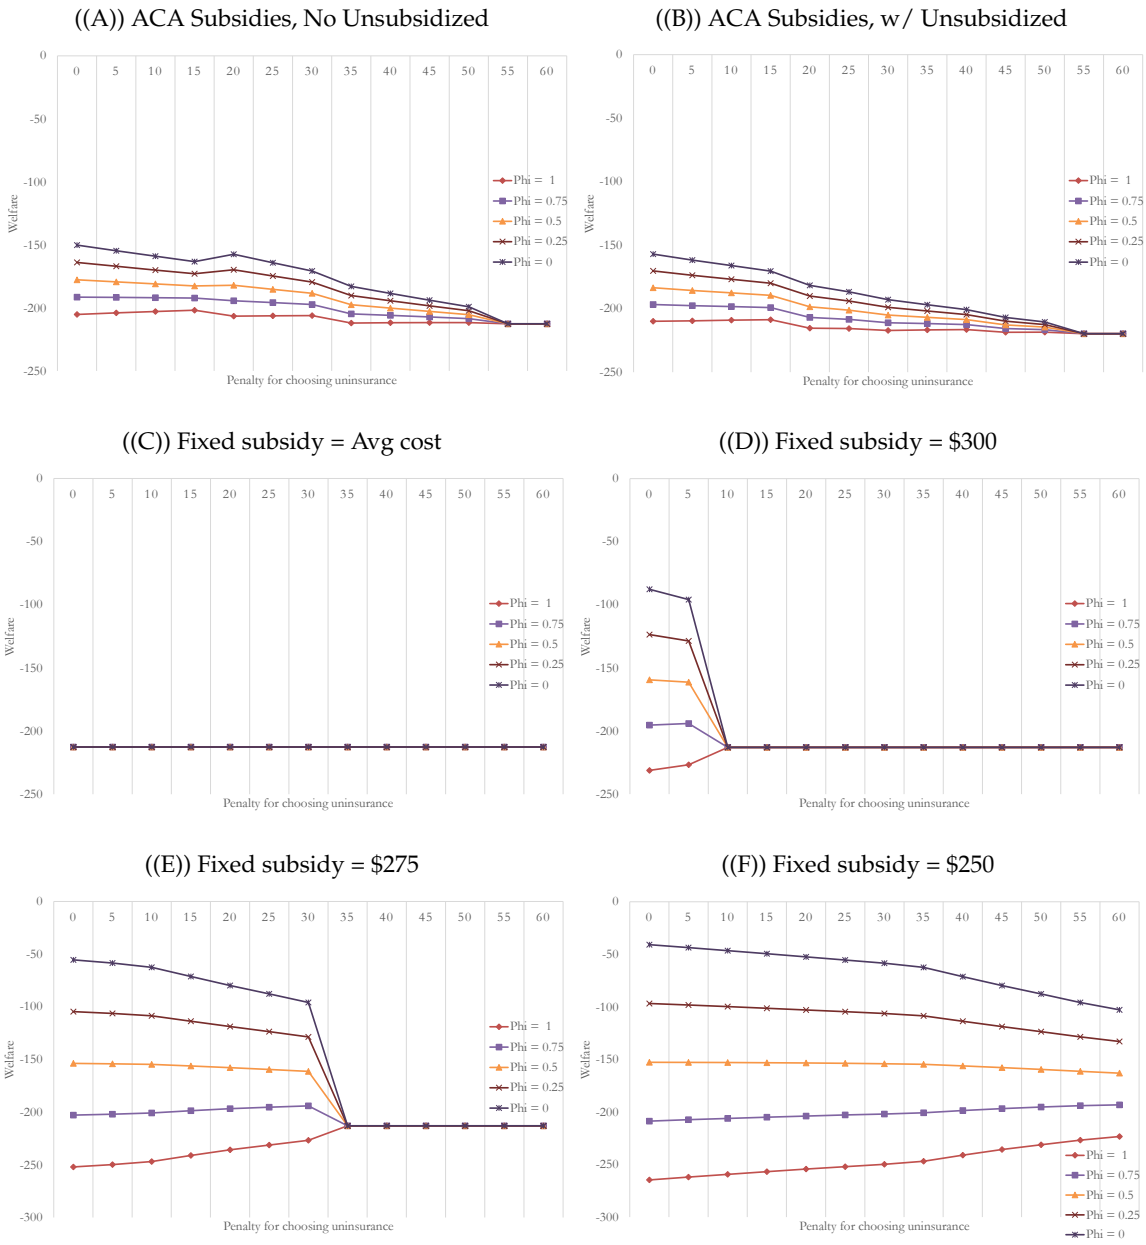
Note: Figures show equilibrium market shares of H, L, and U under different penalties for chooding U. Equilibrium prices are found using the reaction function approach described in the text where we simulate costs and market shares for a large grid of H and L prices and then find the pair of prices where 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 8: Plan Market Shares under Increasingly Strong Risk Adjustment Transfers (higher α)



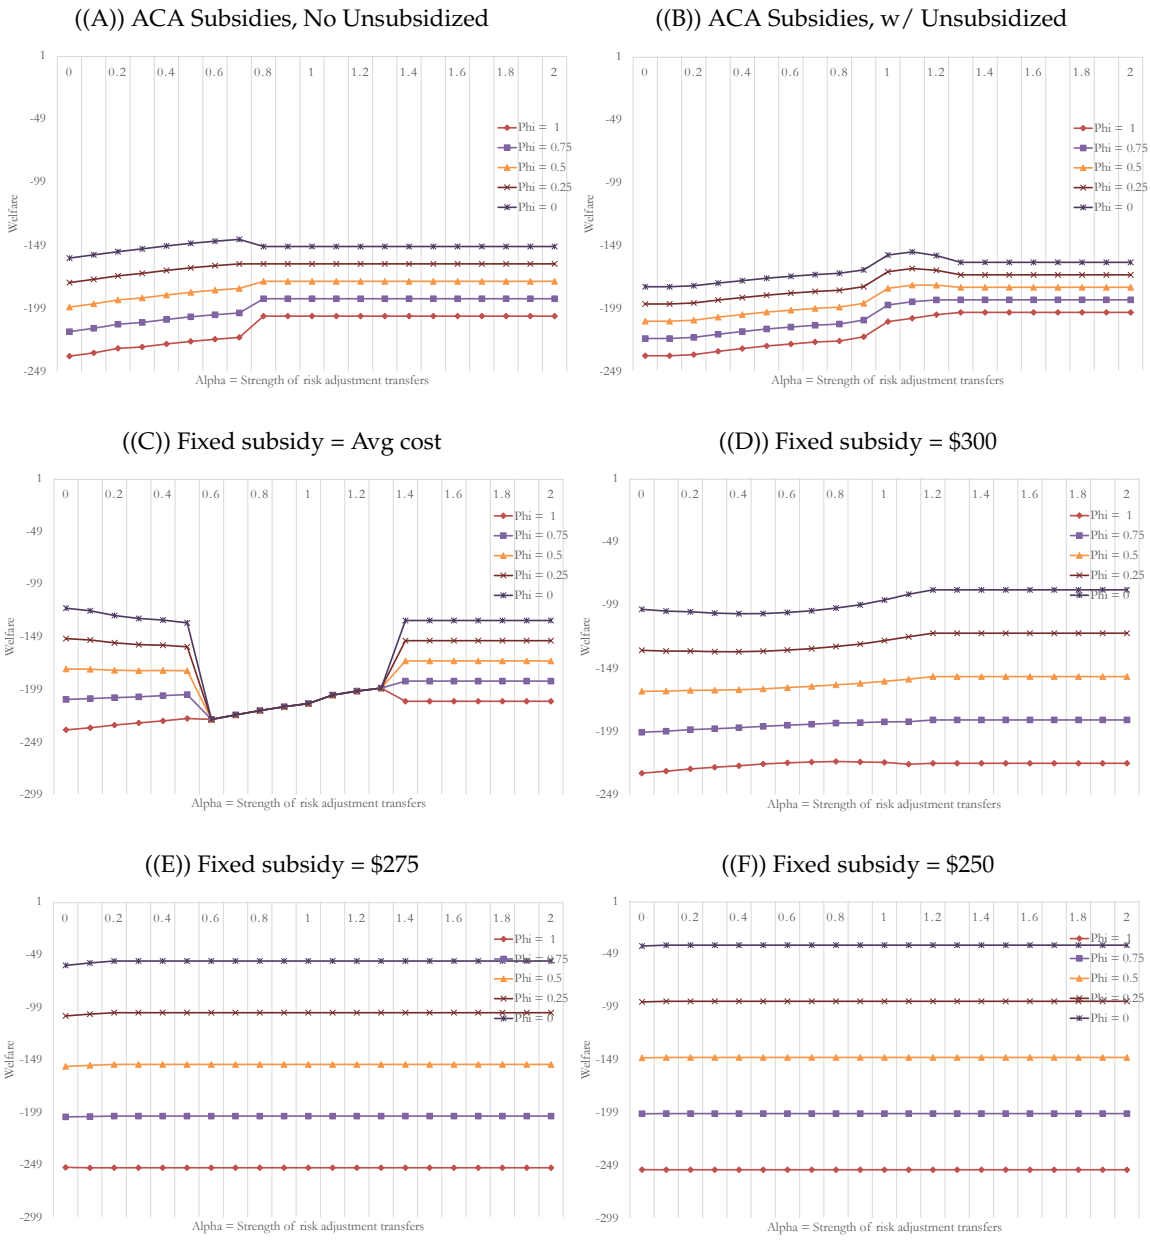
Note: Figures show equilibrium market shares of H, L, and U under different levels of α which multiplies the standard ACA risk adjustment transfers, as described in Equation 14 in the text. Equilibrium prices are found using the reaction function approach described in the text where we simulate costs 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 9: Social Welfare under Increasingly Larger Uninsurance Penalty



Note: Figures show estimates of overall social welfare as described by Equation 11 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 10: Social Welfare under Increasingly Strong Risk Adjustment Transfers (higher α)



Note: Figures show estimates of overall social welfare as described by Equation 11 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.

Table 1: Benefit Regulation Simulations

	ACA subsidies, no unsubsidized		ACA subsidies, w/ unsubsidized		Fixed Subsidy = Avg Monthly Cost	
	With L Plan	No L Plan	With L Plan	No L Plan	With L Plan	No L Plan
Price of L	N/A	N/A	352	N/A	310	N/A
Price of H	380	380	379	364	357	394
Share in H	0.7	0.7	0.48	0.77	0.24	0.63
Share in L	0	0	0.24	0	0.76	0
Share in U	0.3	0.3	0.27	0.23	0	0.37
Subsidy	325	325	297	309	321.55	321.55
Welfare Phi=1	-205.22	-205.22	-209.96	-202.57	-212.62	-209.91
Phi=0.75	-191.46	-191.46	-196.73	-192.63	-212.62	-191.33
Phi=0.5	-177.69	-177.69	-183.51	-182.68	-212.62	-172.76
Phi=0.25	-163.93	-163.93	-170.28	-172.74	-212.62	-154.19
Phi=0	-150.17	-150.17	-157.05	-162.79	-212.62	-135.62
	Fixed Subsidy = $300/month		Fixed Subsidy = $275/month		Fixed Subsidy = $250/month	
	With L Plan	No L Plan	With L Plan	No L Plan	With L Plan	No L Plan
Price of L	N/A	N/A	N/A	N/A	N/A	N/A
Price of H	431	431	453	453	463	463
Share in H	0.26	0.41	0.28	0.28	0.21	0.21
Share in L	0.18	0	0	0	0	0
Share in U	0.56	0.59	0.72	0.72	0.79	0.79
Subsidy	300	300	275	275	250	250
Welfare Phi=1	-223.45	-231.20	-252.00	-252.02	-253.47	-264.66
Phi=0.75	-191.42	-195.31	-202.85	-202.85	-200.31	-208.70
Phi=0.5	-159.39	-159.41	-153.70	-153.69	-147.15	-152.75
Phi=0.25	-127.37	-123.52	-104.56	-104.53	-93.98	-96.80
Phi=0	-95.34	-87.62	-55.41	-55.37	-40.82	-40.84

Note: The table reports equilibrium plan prices, market shares, and subsidies when L is included in the choice set ("with L plan") and when L is removed from the choice set ("no L plan). 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. We consider 6 subsidy regimes. 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.

Table 2: Mandate/Uninsurnace Penalty Simulations

	ACA subsidies, no unsubsidized					ACA subsidies, w/ unsubsidized					Fixed Subsidy = Avg Monthly Cost				
	$0	$15	$30	$45	$60	$0	$15	$30	$45	$60	$0	$15	$30	$45	$60
Price of L	N/A	N/A	347	334	311	353	345	341	330	315	311	311	311	311	311
Price of H	381	367	372	371	356	379	373	376	366	355	356	356	356	356	356
Share in H	0.7	0.77	0.52	0.28	0.25	0.5	0.46	0.21	0.21	0.16	0.24	0.24	0.24	0.24	0.24
Share in L	0	0	0.67	0.6	0.75	0.22	0.32	0.64	0.69	0.84	0.76	0.76	0.76	0.76	0.76
Share in U	0.3	0.23	0.21	0.12	0	0.28	0.22	0.15	0.1	0	0	0	0	0	0
Subsidy	326	312	272	279	256	298	290	286	275	260	322	322	322	322	322
Welfare Phi=1	-205	-202	-206	-212	-213	-210	-209	-217	-219	-220	-213	-213	-213	-213	-213
Phi=0.75	-192	-192	-197	-207	-213	-197	-199	-211	-216	-220	-213	-213	-213	-213	-213
Phi=0.5	-178	-183	-188	-203	-213	-184	-190	-205	-213	-220	-213	-213	-213	-213	-213
Phi=0.25	-164	-173	-180	-198	-213	-170	-180	-199	-210	-220	-213	-213	-213	-213	-213
Phi=0	-150	-163	-171	-194	-213	-157	-170	-193	-207	-220	-213	-213	-213	-213	-213

	Fixed Subsidy = $300/month					Fixed Subsidy = $275/month					Fixed Subsidy = $250/month				
	$0	$15	$30	$45	$60	$0	$15	$30	$45	$60	$0	$15	$30	$45	$60
Price of L	N/A	310	310	310	310	N/A	N/A	N/A	310	310	N/A	N/A	N/A	N/A	N/A
Price of H	431	357	357	357	357	453	442	425	357	357	463	457	451	436	420
Share in H	0.41	0.24	0.24	0.24	0.24	0.28	0.34	0.45	0.24	0.24	0.21	0.25	0.29	0.38	0.48
Share in L	0	0.76	0.76	0.76	0.76	0	0	0	0.76	0.76	0	0	0	0	0
Share in U	0.59	0	0	0	0	0.72	0.66	0.55	0	0	0.79	0.75	0.71	0.62	0.52
Subsidy	300	300	300	300	300	275	275	275	275	275	250	250	250	250	250
Welfare Phi=1	-231	-213	-213	-213	-213	-252	-241	-227	-213	-213	-265	-257	-250	-236	-223
Phi=0.75	-195	-213	-213	-213	-213	-203	-199	-194	-213	-213	-209	-205	-202	-197	-193
Phi=0.5	-159	-213	-213	-213	-213	-154	-156	-161	-213	-213	-153	-153	-154	-158	-163
Phi=0.25	-124	-213	-213	-213	-213	-105	-114	-129	-213	-213	-96.8	-101	-106	-119	-133
Phi=0	-87.7	-213	-213	-213	-213	-55.4	-71.2	-95.9	-213	-213	-40.8	-49.5	-58.4	-79.7	-103

Note: The table reports equilibrium plan prices, market shares, and subsidies for various levels of the penalty for choosing 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. We consider 6 subsidy regimes. 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.

Table 3: Risk Adjustment Simulations

	ACA subsidies, no unsubsidized				ACA subsidies, w/ unsubsidized				Fixed Subsidy = Avg Monthly Cost			
alpha	0	0.5	1	1.5	0	0.5	1	1.5	0	0.5	1	1.5
Price of L	391	379	N/A	N/A	374	364	353	N/A	400	385	311	N/A
Price of H	491	443	381	381	N/A	437	379	364	479	444	356	396
Share in H	0.03	0.17	0.7	0.7	0	0.08	0.5	0.77	0.12	0.19	0.25	0.62
Share in L	0.58	0.44	0	0	0.71	0.64	0.22	0	0.38	0.38	0.75	0
Share in U	0.39	0.39	0.3	0.3	0.29	0.28	0.28	0.23	0.51	0.43	0	0.38
Subsidy	336	324	326	326	319	309	298	309	322	322	322	322
Welfare Phi=1	-237	-225	-205	-205	-237	-229	-210	-203	-238	-227	-213	-211
Phi=0.75	-217	-206	-191	-191	-223	-216	-197	-193	-209	-204	-213	-191
Phi=0.5	-198	-186	-178	-178	-210	-202	-184	-183	-180	-181	-213	-172
Phi=0.25	-179	-167	-164	-164	-196	-189	-170	-173	-151	-159	-213	-153
Phi=0	-159	-148	-150	-150	-182	-175	-157	-163	-122	-136	-213	-134
	Fixed Subsidy = $300/month				Fixed Subsidy = $275/month				Fixed Subsidy = $250/month			
alpha	0	0.5	1	1.5	0	0.5	1	1.5	0	0.5	1	1.5
Price of L	396	387	395	N/A	412	N/A	N/A	N/A	408	N/A	N/A	N/A
Price of H	485	456	437	432	457	453	453	453	463	463	463	463
Share in H	0.07	0.15	0.26	0.41	0.25	0.28	0.28	0.28	0.21	0.21	0.21	0.21
Share in L	0.36	0.31	0.18	0	0.04	0	0	0	0	0	0	0
Share in U	0.57	0.54	0.56	0.59	0.71	0.72	0.72	0.72	0.79	0.79	0.79	0.79
Subsidy	300	300	300	300	275	275	275	275	250	250	250	250
Welfare Phi=1	-232	-225	-223	-224	-252	-252	-252	-252	-253	-253	-253	-253
Phi=0.75	-200	-195	-191	-190	-204	-203	-203	-203	-200	-200	-200	-200
Phi=0.5	-167	-165	-159	-156	-156	-154	-154	-154	-148	-148	-148	-148
Phi=0.25	-135	-136	-127	-121	-108	-105	-105	-105	-94.6	-94.6	-94.6	-94.6
Phi=0	-103	-106	-95.3	-87.2	-59.6	-55.4	-55.4	-55.4	-41.6	-41.6	-41.6	-41.6

Note: The table reports equilibrium plan prices, market shares, and subsidies for various strengths of the risk adjustment transfers (i.e. values of α from Equation 14. 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. We consider 6 subsidy regimes. 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.

Table 4: Benefit Regulation Simulations with 15% L Cost Advantage

	ACA subsidies, no unsubsidized		ACA subsidies, w/ unsubsidized		Fixed Subsidy = Avg Monthly Cost	
	With L Plan	No L Plan	With L Plan	No L Plan	With L Plan	No L Plan
Price of L	335	N/A	309	N/A	N/A	N/A
Price of H	432	380	N/A	364	321	394
Share in H	0.05	0.7	0	0.77	1	0.63
Share in L	0.57	0	0.75	0	0	0
Share in U	0.39	0.3	0.25	0.23	0	0.37
Subsidy	282	325	254	309	321.55	321.55
Welfare Phi=1	-202.38	-205.22	-195.99	-202.57	-198.00	-209.91
Phi=0.75	-183.01	-191.46	-183.66	-192.63	-198.00	-191.33
Phi=0.5	-163.63	-177.69	-171.33	-182.68	-198.00	-172.76
Phi=0.25	-144.26	-163.93	-159.00	-172.74	-198.00	-154.19
Phi=0	-124.89	-150.17	-146.67	-162.79	-198.00	-135.62
	Fixed Subsidy = $300/month		Fixed Subsidy = $275/month		Fixed Subsidy = $250/month	
	With L Plan	No L Plan	With L Plan	No L Plan	With L Plan	No L Plan
Price of L	300	N/A	350	N/A	373	N/A
Price of H	N/A	431	443	453	458	463
Share in H	0	0.41	0.06	0.28	0.09	0.21
Share in L	1	0	0.45	0	0.25	0
Share in U	0	0.59	0.49	0.72	0.66	0.79
Subsidy	300	300	275	275	250	250
Welfare Phi=1	-183.36	-231.20	-213.84	-252.02	-239.79	-264.66
Phi=0.75	-183.36	-195.31	-186.28	-202.85	-196.67	-208.70
Phi=0.5	-183.36	-159.41	-158.72	-153.69	-153.55	-152.75
Phi=0.25	-183.36	-123.52	-131.16	-104.53	-110.42	-96.80
Phi=0	-183.36	-87.62	-103.60	-55.37	-67.30	-40.84

Note: The table reports equilibrium plan prices, market shares, and subsidies when L is included in the choice set ("with L plan") and when L is removed from the choice set ("no L plan). Here, we assume that L has a 15% cost advantage, i.e. $C_L(s) = 0.85C_H(s)$. 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. We consider 6 subsidy regimes. 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.

Table 5: Mandate/Uninsurance Penalty Simulations with 15% L Cost Advantage

	ACA subsidies, no unsubsidized					ACA subsidies, w/ unsubsidized					Fixed Subsidy = Avg Monthly Cost				
	$0	$15	$30	$45	$60	$0	$15	$30	$45	$60	$0	$15	$30	$45	$60
Price of L	335	323	308	293	273	308	301	292	284	273	273	273	273	273	273
Price of H	432	422	412	401	N/A	412	407	400	N/A	N/A	N/A	N/A	N/A	N/A	N/A
Share in H	0.1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Share in L	0.6	0.7	0.8	0.9	1	0.7	0.8	0.9	0.9	1	1	1	1	1	1
Share in U	0.4	0.3	0.2	0.1	0	0.3	0.2	0.2	0.1	0	0	0	0	0	0
Subsidy	280	268	253	238	218	253	246	240	229	218	322	322	322	322	322
Welfare Phi=1	-202	-195	-190	-186	-183	-196	-191	-189	-185	-183	-183	-183	-183	-183	-183
Phi=0.75	-183	-181	-181	-182	-183	-183	-183	-182	-182	-183	-183	-183	-183	-183	-183
Phi=0.5	-164	-167	-172	-177	-183	-171	-174	-176	-180	-183	-183	-183	-183	-183	-183
Phi=0.25	-144	-153	-163	-173	-183	-159	-165	-170	-177	-183	-183	-183	-183	-183	-183
Phi=0	-125	-139	-155	-169	-183	-147	-157	-164	-175	-183	-183	-183	-183	-183	-183

	Fixed Subsidy = $300/month					Fixed Subsidy = $275/month					Fixed Subsidy = $250/month				
	$0	$15	$30	$45	$60	$0	$15	$30	$45	$60	$0	$15	$30	$45	$60
Price of L	273	273	273	273	273	273	273	273	273	273	373	360	273	273	273
Price of H	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	458	450	N/A	N/A	N/A
Share in H	0	0	0	0	0	0	0	0	0	0	0.1	0.1	0	0	0
Share in L	1	1	1	1	1	1	1	1	1	1	0.2	0.4	1	1	1
Share in U	0	0	0	0	0	0	0	0	0	0	0.7	0.6	0	0	0
Subsidy	300	300	300	300	300	275	275	275	275	275	250	250	250	250	250
Welfare Phi=1	-183	-183	-183	-183	-183	-183	-183	-183	-183	-183	-196	-190	-183	-183	-183
Phi=0.75	-183	-183	-183	-183	-183	-183	-183	-183	-183	-183	-178	-175	-183	-183	-183
Phi=0.5	-183	-183	-183	-183	-183	-183	-183	-183	-183	-183	-153	-156	-183	-183	-183
Phi=0.25	-183	-183	-183	-183	-183	-183	-183	-183	-183	-183	-110	-122	-183	-183	-183
Phi=0	-183	-183	-183	-183	-183	-183	-183	-183	-183	-183	-67	-88	-183	-183	-183

Note: The table reports equilibrium plan prices, market shares, and subsidies when L is included in the choice set ("with L plan") and when L is removed from the choice set ("no L plan"). Here, we assume that L has a 15% cost advantage, i.e. $C_L(s) = 0.85C_H(s)$. 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. We consider 6 subsidy regimes. 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.

Table 6: Risk Adjustment Simulations with 15% L Cost Advantage

	ACA subsidies, no unsubsidized					ACA subsidies, w/ unsubsidized					Fixed Subsidy = Avg Monthly Cost				
alpha	0	0.5	1	1.5	2	0	0.5	1	1.5	2	0	0.5	1	1.5	2
Price of L	337	337	335	343	N/A	309	309	309	312	N/A	N/A	N/A	N/A	N/A	N/A
Price of H	N/A	N/A	432	396	380	N/A	N/A	N/A	365	364	321	321	321	321	321
Share in H	0	0	0.05	0.22	0.7	0	0	0	0.13	0.77	1	1	1	1	1
Share in L	0.62	0.62	0.57	0.4	0	0.75	0.75	0.75	0.62	0	0	0	0	0	0
Share in U	0.39	0.39	0.39	0.39	0.3	0.25	0.25	0.25	0.25	0.23	0	0	0	0	0
Subsidy	282	282	280	288	325	254	254	254	257	309	322	322	322	322	322
Welfare Phi=1	-204	-204	-202	-201	-205	-196	-196	-196	-195	-203	-198	-198	-198	-198	-198
Phi=0.75	-184	-184	-183	-182	-191	-184	-184	-184	-182	-193	-198	-198	-198	-198	-198
Phi=0.5	-165	-165	-164	-162	-178	-171	-171	-171	-170	-183	-198	-198	-198	-198	-198
Phi=0.25	-146	-146	-144	-143	-164	-159	-159	-159	-158	-173	-198	-198	-198	-198	-198
Phi=0	-126	-126	-125	-124	-150	-147	-147	-147	-145	-163	-198	-198	-198	-198	-198

	Fixed Subsidy = $300/month					Fixed Subsidy = $275/month					Fixed Subsidy = $250/month				
alpha	0	0.5	1	1.5	2	0	0.5	1	1.5	2	0	0.5	1	1.5	2
Price of L	300	300	300	300	300	356	356	350	361	375	380	376	373	387	N/A
Price of H	N/A	N/A	N/A	N/A	N/A	N/A	N/A	443	426	427	N/A	479	458	452	463
Share in H	0	0	0	0	0	0	0	0.06	0.17	0.22	0	0.02	0.09	0.17	0.21
Share in L	1	1	1	1	1	0.48	0.48	0.45	0.3	0.2	0.31	0.3	0.25	0.13	0
Share in U	0	0	0	0	0	0.52	0.52	0.49	0.53	0.58	0.69	0.67	0.66	0.71	0.79
Subsidy	300	300	300	300	300	275	275	275	275	275	250	250	250	250	250
Welfare Phi=1	-183	-183	-183	-183	-183	-219	-219	-214	-218	-226	-246	-243	-240	-248	-265
Phi=0.75	-183	-183	-183	-183	-183	-189	-189	-186	-187	-190	-201	-199	-197	-200	-209
Phi=0.5	-183	-183	-183	-183	-183	-159	-159	-159	-156	-155	-155	-155	-154	-152	-153
Phi=0.25	-183	-183	-183	-183	-183	-130	-130	-131	-125	-119	-110	-111	-110	-105	-96.8
Phi=0	-183	-183	-183	-183	-183	-100	-100	-104	-93.8	-83.9	-64.3	-66.5	-67.3	-57.2	-40.8

Note: The table reports equilibrium plan prices, market shares, and subsidies when L is included in the choice set ("with L plan") and when L is removed from the choice set ("no L plan"). Here, we assume that L has a 15% cost advantage, i.e. $C_L(s) = 0.85C_H(s)$. 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. We consider 6 subsidy regimes. 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.

Online Appendix for:
Trade-offs between Extensive and Intensive Margin Selection in Competitive 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.[24] 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.[25] 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:

$$\begin{aligned} P_H &= AC_H\left(P^{cons}\right) \\ P_L &= AC_L\left(P^{cons}\right) \end{aligned} \tag{15}$$

We use this expression to define equilibrium throughout Section 2.

[24] 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.

[25] 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: Demand and Cost Curves

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

As discussed in Section 3.1, 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 3.1, 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.[26] 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 "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

[26] 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.

H relative to L, which seems likely to be true in the real world. Thus,

$$W_{HL}^{enh}(s) = \begin{cases} W_{HL}^{lin}(s) & \text{for } s \in [0.31, 1] \\ W_{HL}^{lin}(s) + 3 \times \frac{(0.31-s)}{0.31} \cdot W_{HL}^{lin}(0) & \text{for } s \in [0, 0.31) \end{cases} \tag{16}$$

and

$$W_{H}^{enh}(s) = W_{L}^{lin}(s) + W_{HL}^{enh}(s) \tag{17}$$

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

B.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 $C_H(s)$.[27]

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:

$$W_{HKK}(s) = -\$9,276.81 * s + \$12,498.68 \tag{18}$$

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

[27] 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.

C Appendix Section 3: Description of Reaction Function Approach to Finding Equilibrium

[COMING SOON]

D Appendix Section 4: Results with Linear Demand Curves

[COMING SOON]

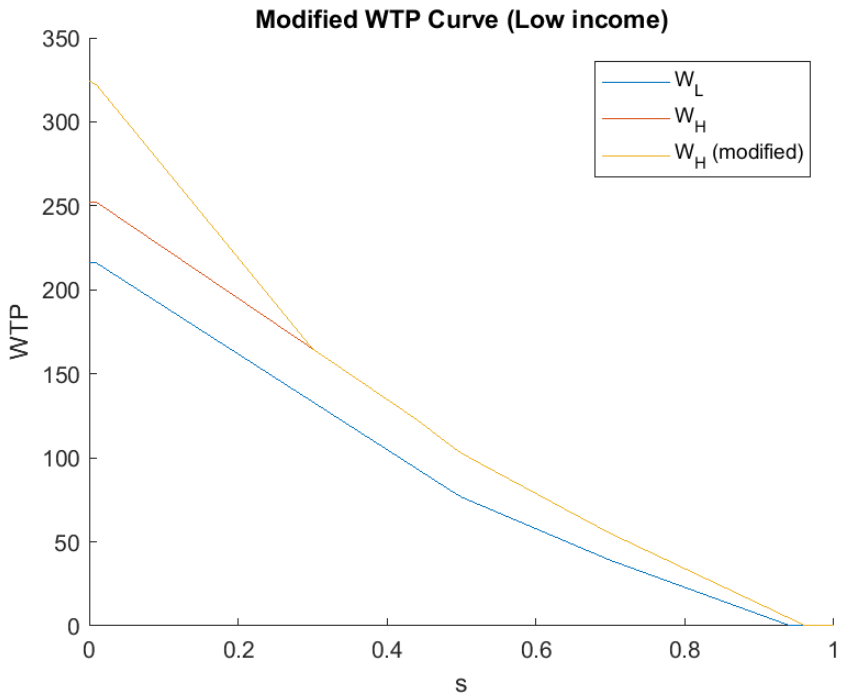
E Appendix Section 5: Results from Simulations with CommCare Subsidies

[COMING SOON]

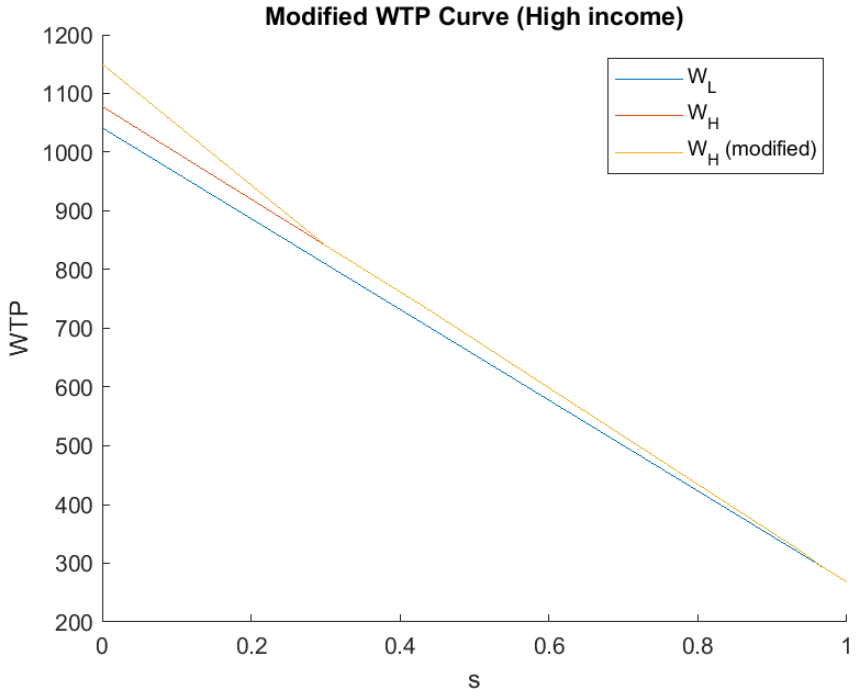
F Additional Figures and Tables

Figure A1: WTP Curves for H and L

((A)) Low-Income

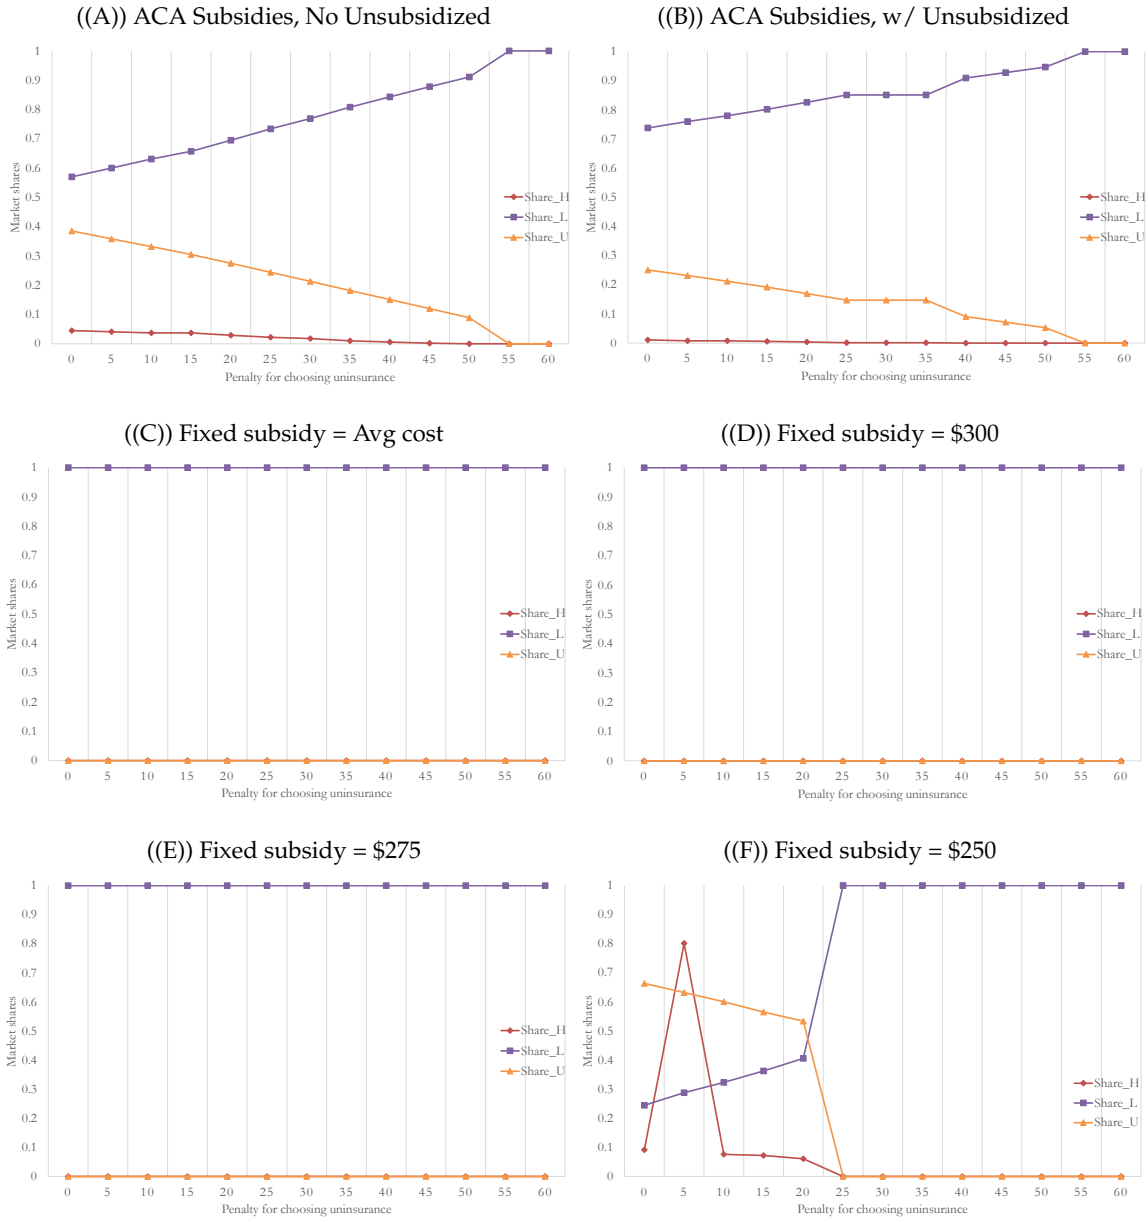


((B)) High-Income



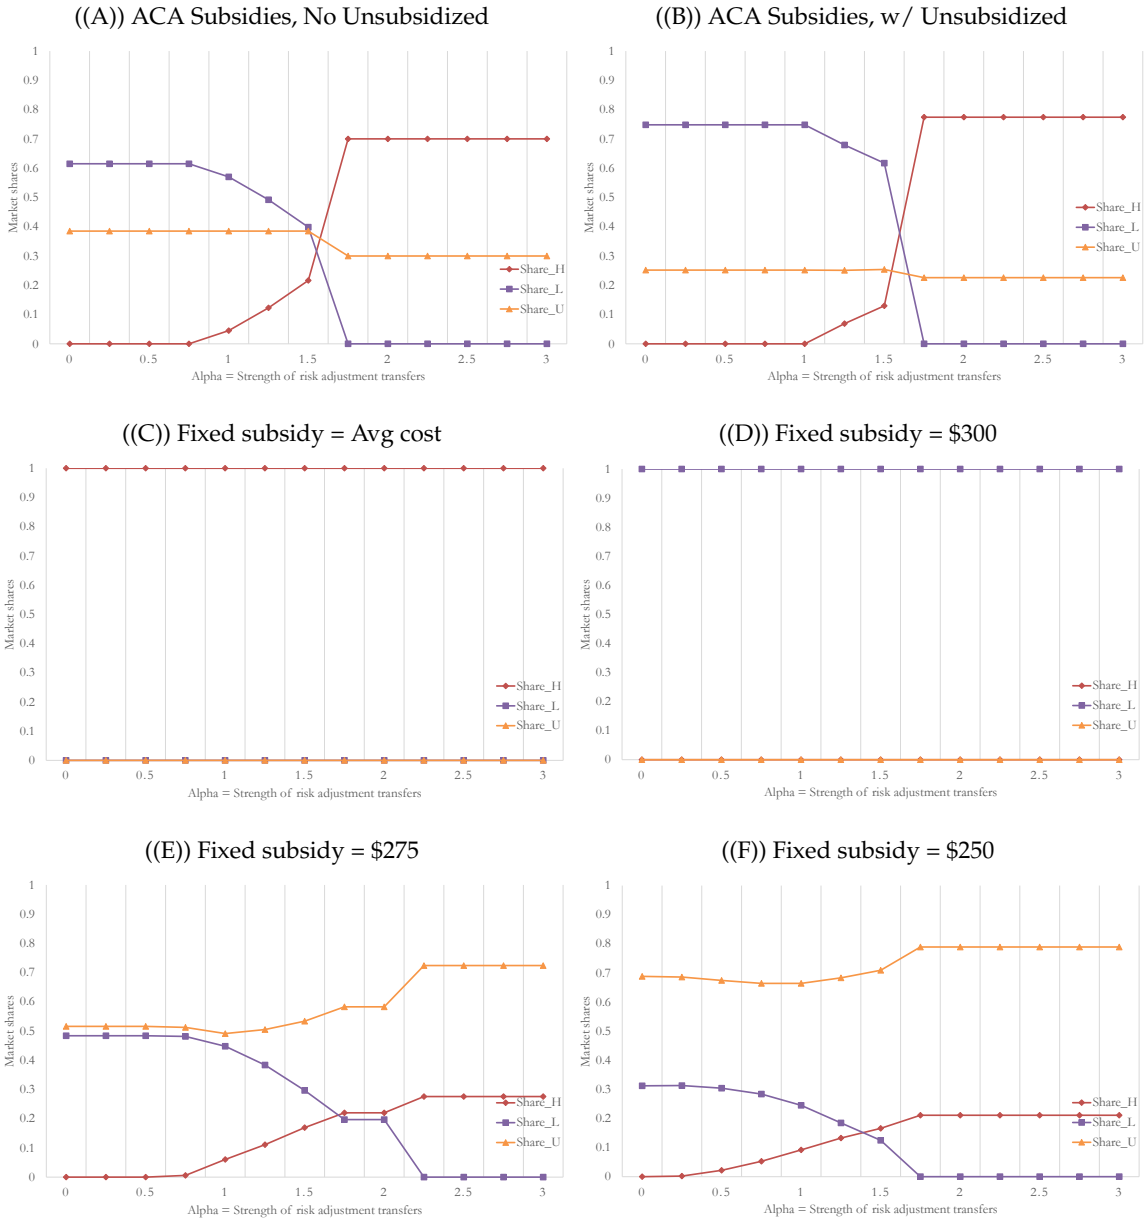
Note: Figure shows WTP Curves for H and L, $W_H(s)$ and $W_L(s)$. Left panel shows curves for low-income group which come from (Finkelstein, Hendren and Shepard, 2017). Right panel shows curves for high-income group which come from (Hackmann, Kolstad and Kowalski, 2015). Linear curves extrapolate linearly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

Figure A2: Plan Market Shares under Increasingly Strong Uninsurance Penalties (with 15% *L* cost advantage



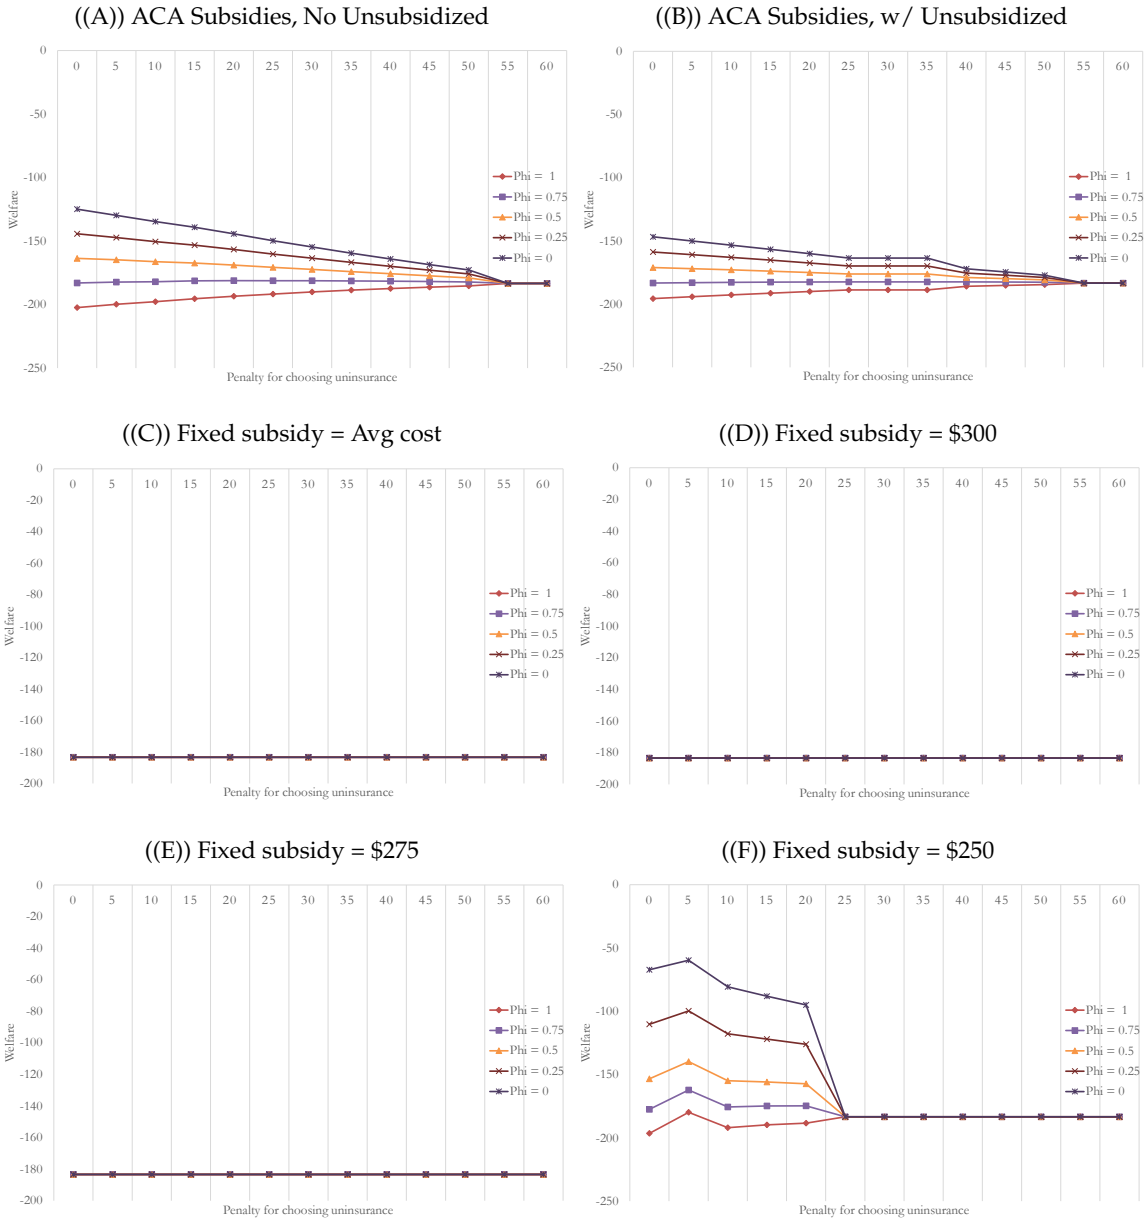
Note: Figures show equilibrium market shares of *H*, *L*, and *U* under different penalties for choosing *U*. Plan *L* is assumed to have a 15% cost advantage for these simulations. Equilibrium prices are found using the reaction function approach described in the text where we simulate costs and market 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 A3: Plan Market Shares under Increasingly Strong Risk Adjustment Transfers (higher α) (with 15% L cost advantage)



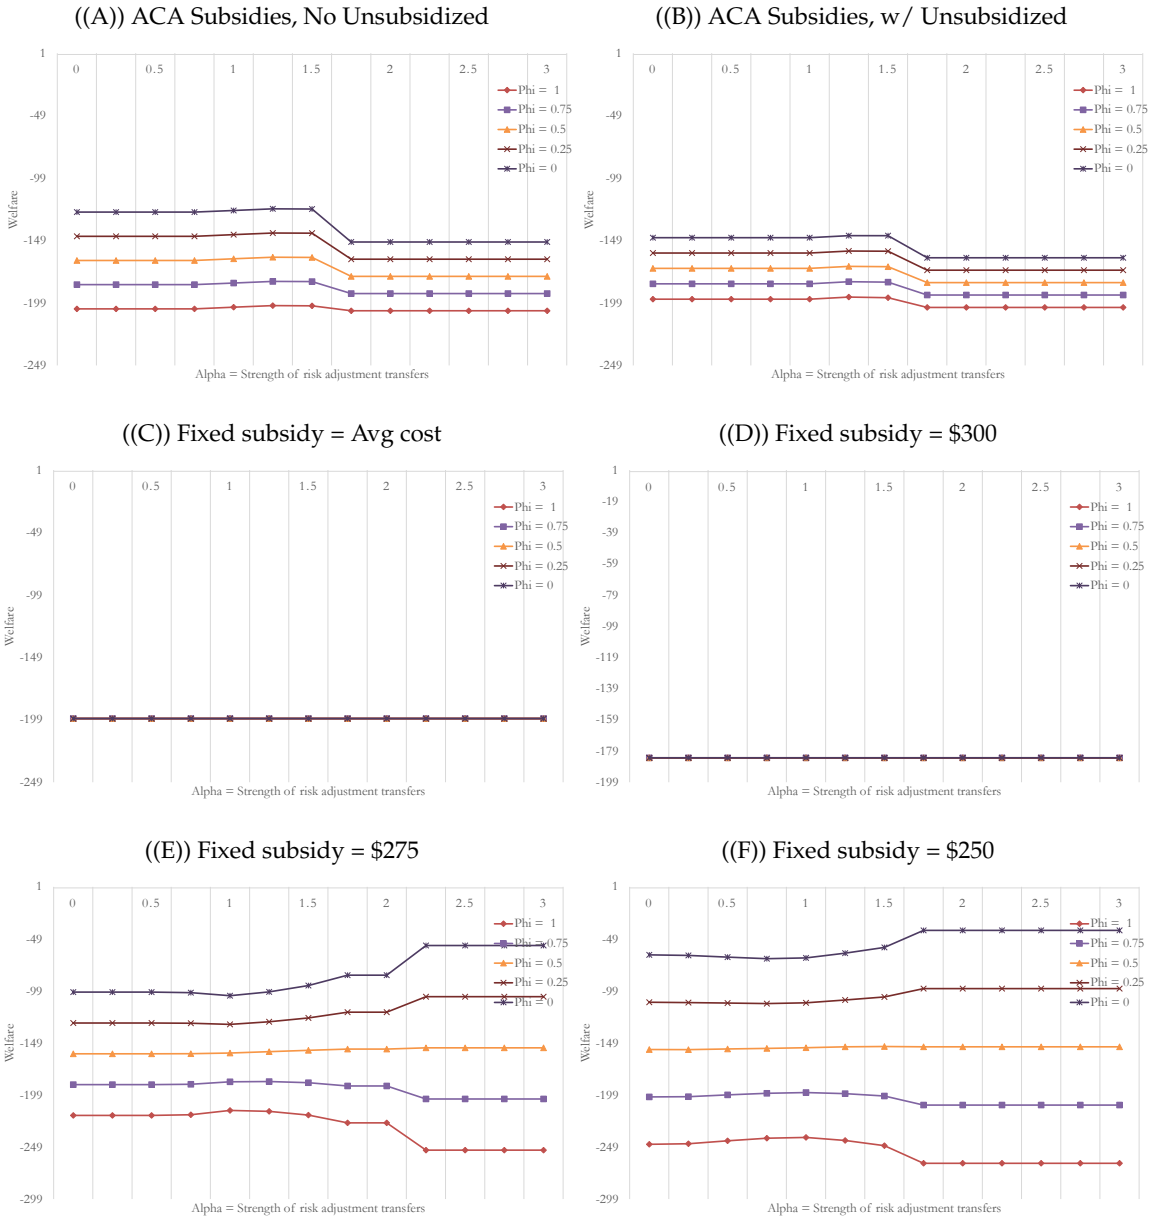
Note: Figures show equilibrium market shares of H, L, and U under different levels of α which multiplies the standard ACA risk adjustment transfers, as described in Equation 14 in the text. Plan L is assumed to have a 15% cost advantage for these simulations. Equilibrium prices are found using the reaction 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 A4: Social Welfare under Increasingly Strong Uninsurance Penalties (with 15% L cost advantage



Note: Figures show estimates of overall social welfare as described by Equation 11 under increasingly larger uninsurance penalties. Plan L is assumed to have a 15% cost advantage for these simulations. 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 A5: Social Welfare under Increasingly Strong Risk Adjustment Transfers (higher α) (with 15% L cost advantage)



Note: Figures show estimates of overall social welfare as described by Equation 11 under increasingly strong risk adjustment transfers (larger α). Plan L is assumed to have a 15% cost advantage for these simulations. 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.