

THE TWO-MARGIN PROBLEM IN INSURANCE MARKETS

Michael Geruso, Timothy J. Layton, Grace McCormack, and Mark Shepard*

Abstract—Insurance markets often feature consumer sorting along both an extensive margin (whether to buy) and an intensive margin (which plan to buy). We present a new graphical theoretical framework that extends a workhorse model to incorporate both selection margins simultaneously. A key insight from our framework is that policies aimed at addressing one margin of selection often involve an economically meaningful trade-off on the other margin in terms of prices, enrollment, and welfare. Using data from Massachusetts, we illustrate these trade-offs in an empirical sufficient statistics approach that is tightly linked to the graphical framework we develop.

I. Introduction

SOME of the most important problems in health insurance markets stem from adverse selection, or the tendency of sicker consumers to exhibit higher demand for insurance. Concerns about adverse selection have motivated a variety of regulatory interventions in the United States and around the rest of world, including insurance mandates, penalties for being uninsured, subsidies for purchasing insurance, risk adjustment transfers, benefit regulation, and reinsurance. Policy discussions about how to address adverse selection have become salient in the United States as many public programs have shifted toward providing health insurance via regulated markets (Gruber, 2017).

But a deeper look reveals that not all policies combating adverse selection are targeted at the same problem. On the

one hand, policies such as mandates and subsidies combat selection on the *extensive margin* (or “against the market”). This type of selection is characterized by sicker people being more likely to buy insurance. It leads to higher insurer costs and higher consumer prices and causes some healthy people to opt out. Policies such as risk adjustment and benefit regulation, on the other hand, combat selection on the *intensive margin* (or “within the market”). This type of selection is characterized by sicker people being more likely to purchase more generous plans within the market. Intensive margin selection drives up the price of generous plans relative to skimpy ones and results in too many consumers choosing skimpy plans. In some cases, selection within the market may be so strong that generous contracts cannot be sustained, and the market for them unravels entirely (Cutler & Reber, 1998).

Prior work has recognized these two problems and has studied policies targeted at each. However, this literature has largely considered these two forms of selection in isolation—either assuming all consumers buy insurance and focusing on the intensive margin (Handel, Hendel, & Whinston, 2015) or assuming all contracts within the market are identical and focusing on the extensive margin (Hackmann, Kolstad, & Kowalski, 2015). By ignoring one margin or the other, the selection problem is usefully simplified. In empirical work, it becomes amenable to a sufficient statistics approach based on demand and cost curves defined in reference to a single price—either the price of insurance or the price difference between a generous versus a skimpy plan (Einav, Finkelstein, & Cullen, 2010). However, this simplification does not allow for potential interactions between these two margins of selection.

In this paper, we generalize the canonical insurance market framework to address both margins simultaneously. The benefit of doing so is not merely a technical curiosity. It has first-order policy importance in settings like the ACA marketplaces where both the generosity of coverage and rates of uninsurance are serious concerns. To see why, consider an insurance mandate—a policy that aims to correct extensive margin selection by bringing healthy marginal consumers into the market. Our framework shows how a mandate that succeeds in increasing rates of insurance coverage will likely worsen selection on the intensive margin. Intuitively, the mandate brings more healthy and low-cost consumers into the market. Because these new consumers tend to select the

Received for publication June 8, 2020. Revision accepted for publication February 19, 2021. Editor: Benjamin R. Handel.

*Geruso: University of Texas at Austin and NBER; Layton: Harvard University and NBER; McCormack: Harvard University; Shepard: Harvard University and NBER.

We thank Sebastian Fleitas, Bentley MacLeod, Maria Polyakova, and Ashley Swanson for serving as discussants for this paper. We also thank Kate Bundorf, Marika Cabral, Amitabh Chandra, Vilsa Curto, Leemore Dafny, Keith Ericson, Amy Finkelstein, Jon Gruber, Tom McGuire, Neale Mahoney, Joe Newhouse, Evan Saltzman, Brad Shapiro, and Pietro Tebaldi; participants at NBER Health Care, NBER Insurance Working Group, CEPRA/NBER Workshop on Aging and Health, the 2019 Becker Friedman Institute Health Economics Initiative Annual Conference at the University of Chicago, the 2019 American Economic Association meetings, the 2018 American Society of Health Economists meeting, the 2018 Annual Health Economics Conference, the 2018 Chicago Booth Junior Health Economics Summit; and seminars at the Brookings Institution and the University of Wisconsin for useful feedback. We gratefully acknowledge financial support for this project from the Laura and John Arnold Foundation, the Eunice Kennedy Shriver National Institute of Child Health and Human Development Center grant P2CHD042849 awarded to the Population Research Center at UT-Austin, the Agency for Healthcare Research and Quality (K01-HS25786-01), and the National Institute on Aging, grant T32-AG000186. No party had the right to review this paper prior to its circulation.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01070.

lower-price (and lower-quality) plans, the risk pools of those plans will get even healthier. In equilibrium, these plans will further reduce prices, siphoning additional consumers away from higher-quality plans on the intensive margin, causing prices for high-quality coverage to spiral upward. These two offsetting effects (improving take-up and inducing within-market unraveling) represent a clear example of the intensive/extensive margin interactions that are the focus of our paper.¹

One of our main contributions is to provide a graphical demand-cost framework that lets economists visualize (and teach) the two-margin selection problem in a transparent way. To do so, we build on the influential work of Einav et al. (2010) and Einav and Finkelstein (2011), who show how to visualize selection markets in terms of demand, average cost, and marginal cost curves. We generalize their model to allow for two plans—a more generous H plan and a less generous L plan—plus an outside option of uninsurance (U). Although stylized, our vertical model captures the core intuition of the two selection margins: an intensive margin difference in generosity (H versus L) and an extensive margin option to exit the market (by choosing U). It also captures the key feature of adverse selection: that higher-risk consumers have greater willingness to pay for generous coverage—both for H relative to L , and for L relative to U . Our vertical model is the simplest framework that captures these features and is useful for developing intuition around a potentially multidimensional problem by allowing the market to be represented in standard two-dimensional graphs with familiar demand and cost curves. Equilibrium prices, market shares, and social surplus can all be easily visualized. We also show the extent to which the core intuitions hold as various assumptions on the model are relaxed, including, for example, allowing for horizontal differentiation across plans.

As in Einav et al. (2010), there is a tight link between our model and the estimation of sufficient statistics used to characterize equilibrium and welfare. Econometric identification is analogous, though exogenous price variation along two margins is required—for example, independent variation in the price of a skimpy plan and in the price of a generous plan.²

After developing the graphical framework, we use it to show how policies and regulatory actions that counteract selection on one margin can interact with the other. The relevance of these cross-margin interactions is the key conceptual message of our paper. We show that a mandate's impact on plan generosity is, in fact, an instance of a broader phenomenon that encapsulates many relevant policy interventions currently in place in insurance markets. These in-

clude plan benefits requirements, network adequacy rules, risk adjustment, reinsurance, subsidies, and behavioral interventions like plan choice architectures or autoenrollment. Each involves a potential trade-off: Policies that aim to address intensive margin selection tend to worsen extensive margin selection, and vice versa.

The graphical model helps show why these cross-margin interactions occur. The key insight is that for each plan, either its demand or average cost curve is not a price-invariant model primitive (as is true in a two-option model) but an equilibrium object that depends on the other plan's price. Policies that target one selection margin typically influence market prices (e.g., the mandate lowers P_L relative to P_H), which in turn shifts demand or cost curves that determine the other margin (e.g., the lower P_L reduces demand for H). This cross-plan dependence of demand and average costs is the key missing piece when the two margins are analyzed separately. We show how the geometry of the demand and cost curves generates this dependence. We also develop a more general nongraphical version of our model that allows for horizontal differentiation and use it to show that many of the key intuitions will hold with a modest amount of horizontal differentiation (i.e., consumers on the margin between H and U).

With the intuition and price theory in place, we analyze the model's insights empirically using demand and cost estimates from the Massachusetts Commonwealth Care program, a subsidized insurance exchange that was a precursor to the ACA health insurance marketplace. We draw on demand and cost estimates from Finkelstein et al. (2019) to simulate equilibrium in counterfactuals where we vary benefit design rules, mandate penalties, and risk adjustment strength.³ Beyond demonstrating how our framework can be used, the empirical exercise generates several policy insights. The size of the unintended cross-margin effects can be quite large. We find that a strong mandate sufficient to move all consumers into insurance—increasing enrollment by around 25 percentage points—can reduce the market share of generous plans by more than 15 percentage points, or 35% of baseline market share. In the other direction, strengthening risk-adjustment transfers until the market “unravels” to include only generous coverage can substantially reduce market-level consumer participation—in our setting, by as much as 15 percentage points or 60% of the baseline uninsurance rate. With the additional assumption that consumer choices reveal plan valuations, we find that the cross-margin welfare impacts can be similarly large (and often first order).

Further, we show that in some settings, cross-margin interactions are critical for determining optimal policy. When intensive margin policies (such as risk adjustment) are weak, it can be optimal to also have weak extensive margin policies

¹Recent theoretical insights from Azevedo and Gottlieb (2017) and empirical findings from Saltzman (2021) indicate that this is an important omission in contexts like the ACA marketplaces. We similarly find that these interactions are first-order for plan choices and welfare.

²Or alternatively, variation in a market-wide subsidy for selecting any plan and independent variation in the price difference between bare bones and generous plans.

³Finkelstein et al. (2019) use a regression discontinuity design to document significant adverse selection both into the market and within the market between a narrow-network, lower-quality option and a set of wider-network, higher-quality plans.

(such as an uninsurance penalty). But when intensive margin policies are strong, it can be optimal to also have strong extensive margin policies. These results show that in these markets, regulators are operating in a world of the second-best and must consider interactions between the two margins of selection in order to determine constrained optimal policy. This is true whether optimality is viewed from a formal social surplus perspective or reflects a political preference over rates of insurance coverage on the one hand and insurance quality on the other. While we stop short of prescribing *the* optimal policy in a given market, our results indicate that when extensive margin policies become stronger, intensive margin policies should often strengthen (and vice versa).

Our paper contributes to a growing literature on adverse selection in health insurance markets. Our main contribution is to provide a graphical model that unites two key strands of this literature. The first strand focuses on extensive margin selection and stems from the seminal work of Akerlof (1970).⁴ The second strand focuses on intensive margin selection, studying either consumer sorting across a fixed set of contracts within a market⁵ or how consumer selection is endogenously reflected in the characteristics of the contracts offered.⁶

The most directly connected work is a prior theoretical contribution by Azevedo and Gottlieb (2017) that points out the potential cross-margin effects of a mandate in a setting with vertically differentiated contracts that differ in their coinsurance rates. Our framework maintains the vertical assumption of Azevedo and Gottlieb (2017) while allowing differentiation to be more flexible (i.e., based on factors other than cost sharing) in a two-contract setting. Similar to Azevedo and Gottlieb (2017), our paper also takes a step toward bridging the gap between the Akerlof (1970) and Einav et al. (2010) fixed-contracts approach and the Rothschild and Stiglitz (1976) endogenous-contracts approach to modeling adverse selection in insurance markets by allowing some contracts to death-spiral out of existence in equilibrium while others remain available. This possibility that policies can affect which contracts are ultimately offered in equilibrium is a key feature of our model that was originally highlighted by Rothschild and Stiglitz (1976) but is generally overlooked by the Einav et al. (2010) workhorse model. Finally, Saltzman (2021) provides a complementary analysis (concurrent with ours) that investigates cross-margin effects using a structural model estimated with ACA data from California.

Our insights about cross-margin interactions are relevant for active policy debates in the ACA and other insurance settings. For example, within the last 5 years, the federal gov-

ernment has gone back-and-forth with respect to the level of flexibility it provides to states to weaken ACA Essential Health Benefits or risk adjustment transfers (intensive margin policies). The stated goal of more flexibility has been to lower plan prices and reduce uninsurance, and the stated goal of less flexibility has been to increase the quality of ACA insurance plans (cross-margin effects). However, state efforts to simplify enrollment (Domurat, Menashe, & Yin, 2021) or enact mandate penalties (all extensive margin policies) may create unintended consequences on the intensive margin. More broadly, our model is also relevant to other settings with two selection margins, including the Medicare program (with its Medicare Advantage option), employer programs with a plan choice decision and a participation decision (e.g., CalPERS), national health insurance systems with an opt-out (e.g., Germany), and other selection markets (outside health insurance) with both an extensive and intensive margin choice.

The rest of the paper is organized as follows. Section II presents the graphical vertical model. Section III applies the model to show two-margin impacts of various policies. Sections IV to VI apply the model with simulations: section IV discusses methods; section V shows price and enrollment results; and section VI shows welfare results. Section VII concludes.

II. Model

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

Each plan $j \in \{H, L\}$ sets a single community-rated price P_j that (along with any risk adjustment transfers; see below) must cover its costs. Consumers make choices based on these prices and on the price of the outside option, $P_U = M$.⁷ In our focal example, M is a mandate penalty. The distinguishing feature of U is that its price is exogenously determined; it does not adjust based on the consumers who select into it. This is natural for the case where U is uninsurance or a public plan like Traditional Medicare.⁸ $P = \{P_H, P_L, P_U\}$ is the vector of prices in the market.

In the most general formulation, demand in this market cannot be easily depicted in two-dimensional figures. To make the cross-margin effects of interest clearer, we impose a vertical model of demand, which assumes contracts are identically preference-ranked across consumers. Although the strict vertical assumption is not necessary for many of

⁴Recent theoretical advances in this strand include Hendren (2013) and Mahoney and Weyl (2017) and empirical applications by Bundorf, Levin, and Mahoney (2012), Hackmann et al. (2015), Tebaldi (2017), and others.

⁵See, e.g., Handel et al. (2015); Shepard (2022).

⁶See, e.g., Glazer and McGuire (2000); Veiga and Weyl (2016); Carey (2017); Lavetti and Simon (2018); and Geruso, Layton, and Prinz (2019). Geruso and Layton (2017) provide an overview comparing the fixed- and endogenous-contracts approaches to modeling intensive margin selection.

⁷Below, we allow that consumers may receive a subsidy, S , so that choices are based on post-subsidy prices, $P_j^{cons} = P_j - S$.

⁸We adapt the model to the case of Medicare in appendix B.2.

our main insights to hold, it captures the key features of the issues raised by simultaneous selection on two margins in a simple way that allows for graphical representation. We next present the vertical model, then add the cost curves, and finally show how to find equilibrium and welfare. Throughout the paper, we discuss the implications of relaxing the vertical demand assumption for our findings.

A. Demand

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

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

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

These assumptions, which are a slight generalization of the textbook vertical model,⁹ involve two substantive restrictions on the nature of demand. First, the products are vertically ranked: all consumers would choose H over L if their prices were equal and would similarly prefer L to U if their prices were equal.¹⁰ This is a statement about the type of setting to which our model applies. The vertical model applies best when plan rankings are clear—for example, a low- versus high-deductible plan, or a narrow versus complete provider network plan. Importantly, these are precisely the settings where intensive margin risk selection is most relevant. When plans are horizontally differentiated (such as in the Covered California market; see Tebaldi, 2017), it is less likely that high-risk consumers will heavily select into a single plan or type of plan. In such cases, the existing EFC framework can capture the main way risk selection matters: in versus out of the market (the extensive margin). Our model is designed to study the additional issues that arise when both intensive and extensive margins matter simultaneously.¹¹

Second, consumers' WTP for H and L —which in general could vary arbitrarily over two dimensions—are assumed to collapse to a single-dimensional index, $s \in [0, 1]$. Higher s types have both lower W_L and a smaller gap between W_H and

W_L . Lower s types care more about having insurance (L versus U) and more about the generosity of coverage (H versus L). This assumption is a natural approximation that captures the primary pattern of selection in many cases; indeed it holds exactly in a model where plans differ purely in their coinsurance rate (see Azevedo & Gottlieb, 2017). Substantively, assumption 2 restricts consumer sorting and substitution patterns among options when prices change. The primary consequence of this assumption is that consumers are only on the margin between adjacent-generosity options—between H and L or between L and U . No consumer is on the margin between H and U , so if the price of U (the mandate penalty) increases modestly, the newly insured all buy L (the cheaper plan), not H . This restriction captures in a strong way the general (and testable) idea that these are the main ways consumers substitute in response to price changes. With this restriction in place (and under a price vector at which all options are chosen), consumers sort into plans with the highest-WTP types choosing H , intermediate types choosing L , and low types choosing U . We show that weakening this assumption—allowing an H - U margin—does not change the key implications of the model as long as most consumers exhibit vertical preferences. We describe a more general (nongraphical) model in appendix A that allows for both horizontal and vertical differentiation. As we describe, horizontal differentiation tends to dampen the cross-margin effects we study. Throughout, we provide supplementary (theoretical and empirical) results that show the extent to which the relative degree of horizontal differentiation affects our main results.

Figure 1a plots a simple linear example of $W_H(s)$ and $W_L(s)$ curves that satisfy these assumptions. The x -axis is the WTP index s , so WTP declines from left to right as usual. Let $s_{LU}(P)$ be the extensive-marginal type who is indifferent between L and U at a given set of prices P . Assuming for now that $P_U \equiv M = 0$, this cutoff type is defined by the intersection of L 's WTP curve W_L and L 's price, where $W_L(s_{LU}) = P_L$. Consumers to the right of s_{LU} go uninsured. Those to the left buy insurance. Therefore, $W_L(s)$ represents the (inverse) demand curve for any formal insurance (H or L).¹²

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

$$\Delta W_{HL}(s_{HL}) \equiv W_H(s_{HL}) - W_L(s_{HL}) = P_H - P_L. \quad (1)$$

Consumers to the left of s_{HL} buy H because their incremental WTP for H over L —which we label ΔW_{HL} —exceeds the incremental price. With demand for H and for $H + L$ thus determined by these cutoffs, demand for L equals the

⁹Our vertical model follows the format of Finkelstein et al. (2019). It is a generalization of the textbook vertical model in which products differ on quality (Q_j) and consumers differ on taste for quality (β_i), so that WTP equals: $W_{i,j} = \beta_i Q_j$ and utility equals $U_{i,j} = W_{i,j} - P_j = \beta_i Q_j - P_j$.


¹⁰See appendix B.2 for an alternative case where the outside option is preferred to H and L .

¹¹Even in settings without apparent vertical differentiation across plans within the market, our model can be useful in assessing counterfactual policies that might generate this type of differentiation. In particular, our examples below imply that a regulator encouraging better entrants may generate unintended cross-margin effects on the rates of uninsurance. Further, an apparent lack of vertical differentiation may itself be an equilibrium outcome in a vertical model, reflecting a situation where generous coverage has already unraveled.


¹²In the more general case where consumers receive subsidies for purchasing insurance or pay a penalty when choosing U , $W_L(s)$ and the (inverse) demand curve for insurance will diverge. Specifically, $D_L(s) = W_L(s) + S + M$. For simplicity, we ignore the subsidy and penalty terms here but fully incorporate consumer subsidies when we use the model to study the effects of common policies (section III) as well as in the empirical exercise (section V).

FIGURE 1.—ENROLLEE SORTING AND COST UNDER VERTICAL MODEL

(a) Demand and Consumer Sorting under Vertical Model



(b) Cost Curves under Vertical Model



Panel a shows demand and consumer sorting under the vertical model. $W_H(s)$ and $W_L(s)$ are willingness to pay for the H and L plans. $D_H(s; P_L)$ is the demand curve for H (as a function of P_H), which depends on the value of P_L . See the body text for additional description. Panel b shows the cost curves for H and L plans under the vertical model. $C_H(s)$ and $C_L(s)$ are the consumer type- s specific costs. $AC_H(s_{HL})$ and $AC_L(s_{LU}; s_{HL})$ are the average cost curves for H and L given that the intensive margin type is s_{HL} and the extensive margin type is s_{LU} . Adverse selection makes the price difference $P_H - P_L$ larger than the causal cost difference.

difference between the two.¹³ Rearranging equation (1) yields the (inverse) demand for H , given a fixed P_L :

$$D_H(s; P_L) \equiv W_H(s) - W_L(s) + P_L. \quad (2)$$

Figure 1a shows $D_H(s; P_L)$ with a dashed line. One can draw D_H by noting that it intersects the W_H curve at the cutoff type s_{LU} (since $W_L(s_{LU}) = P_L$).¹⁴ It then proceeds leftward

¹³Formally, the demand functions for the general case where $M \neq 0$ are defined by the following equations, where $\Delta P \equiv P_H - P_L$: $D_H(P) = s_{HL}(\Delta P)$; $D_L(P) = s_{LU}(P_L - M) - s_{HL}(\Delta P)$; $D_U(P) = 1 - s_{LU}(P_L - M)$.

¹⁴ D_H is not defined to the right of s_{LU} , since if P_H falls further than its level at this point, nobody buys L . As a result, the demand curve for H thereafter equals $W_H(s)$.

at a slope equal to that of ΔW_{HL} and its intersection with P_H determines s_{HL} . $D_H(s; P_L)$ is flatter than W_H because its slope equals that of $\Delta W_{HL}(s)$.

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

B. Costs

The model's cost primitives are expected insurer costs for consumers of type s in each plan j .¹⁵ These "type-specific costs" are defined as $C_j(s) = E[C_{ij} | s_i = s]$. $C_j(s)$ is analogous to "marginal cost" in the EFC model—so called because it refers to consumers on the margin of purchasing at a given price. However, to avoid confusion in our model where there are two purchasing margins, we refer to $C_j(s)$ as type-specific costs, or simply costs. In addition, we define $C_U(s)$ as the expected costs of uncompensated care of type- s consumers if they were uninsured. Along with adverse selection, external uncompensated care costs motivate subsidy and mandate policies.

Plan-specific average costs are defined as the average of $C_j(s)$ for all types who buy plan j at a given set of prices: $AC_j(P) = \frac{1}{D_j(P)} \int_{s \in D_j(P)} C_j(s) ds$, where (abusing notation slightly) $s \in D_j(P)$ refers to s -types who buy plan j at prices P .

We illustrate the construction of these cost curves in figure 1b. We show a case where cost curves C_H and C_L are downward sloping, indicating adverse selection. The gap between the two curves for a given s -type equals the difference in plan spending if the s -type consumer enrolls in H versus L . We refer to this as the "causal" plan effect, since it reflects the true difference in insurer spending for a given set of people.¹⁶

We start by deriving $AC_H(P)$, the average cost curve for the H plan. To avoid ambiguity later, it is helpful to redefine the argument of AC_H as the marginal type that buys

¹⁵A key insight of the EFC model is that—while costs may vary widely across consumers of a given WTP type—it is sufficient for welfare to consider the cost of the typical consumer of each type. The reason is that with community-rated pricing, consumers sort into plans based only on WTP. There is no way to segregate consumers more finely than WTP type, and since insurers are risk neutral, only the expected cost within type matters. We note, however, that this argument breaks down when leaving the world of community-rated prices, as pointed out by Bundorf et al. (2012), Geruso (2017), and Layton et al. (2017). Our model (like the model of EFC) thus cannot be used to assess the welfare consequences of policies that allow for consumer risk rating.

¹⁶As in EFC, the causal plan effect reflects both a difference in coverage (e.g., lower cost sharing) conditional on behavior, and any behavioral effect (or moral hazard) of the plans.

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

The average cost curve for L is more complicated because it is an average over a range of consumers, $s \in [s_{HL}, s_{LU}]$, with two endogenous margins. For each value of s_{LU} that defines sorting between U and L , there are many possible values of AC_L , depending on consumer sorting between H and L . This fact makes it impossible to plot a single fixed AC_L curve as we did with AC_H . Nonetheless, it is possible to plot $AC_L(s_{LU})$ conditional on $s_{HL}(P)$. We denote this curve $AC_L(s_{LU}; s_{HL})$ and illustrate it with a dashed line in figure 1b. There are many such iso- s_{HL} plots of AC_L (not pictured) that hold P_H fixed at various levels. The left-most point of the AC_L curve depends on the s_{HL} cutoff type determined by P_H . Higher values of s_{HL} imply that $AC_L(s_{LU}; s_{HL})$ starts from a higher point. Just as AC_H equals C_H at $s = 0$, AC_L equals C_L at $s = s_{HL}$. Moving rightward from $s = s_{HL}$, plan L adds more relatively healthy consumers, resulting in a downward-sloping average cost curve.

In summary, while AC_H is fixed and does not depend on the price of L , AC_L is an equilibrium object in that it changes as P_H , and therefore s_{HL} , changes. This implies that the average cost of L and thus the price of L in equilibrium depends on the price of H . Recognizing such dependencies is critical for analyzing policy interventions. For example, a subsidy targeted to H that results in a lower (net) P_H and a larger H enrollment (a rightward-shifted s_{HL}) would cause the left-most point on AC_L to shift down and rightward and would cause the curve to have a less-steep slope. In a competitive market, this would likely result in a lower P_L , causing additional consumers to enter the market.

C. Competitive Equilibrium

We consider competitive equilibria where plan prices, P , exactly equal their average costs:

$$P_H = AC_H(P) \quad \text{and} \quad P_L = AC_L(P), \quad (3)$$

In some settings, multiple price vectors will 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 et al. (2015) and limit attention to equilibria that meet the requirements of the Riley equilibrium (RE) notion. A policy satisfies the RE notion if there exists no “Riley deviation policy,” a competing policy that if offered, would earn a profit, render the old policy unprofitable, and for which there is no “safe response” that would render the Riley deviation unprofitable. A safe response is a policy offering that does not incur a loss when offered with the other existing policies in the market and renders the potential Riley deviation unprofitable. When we apply these requirements in our simulations, we find a unique equilibrium for all empirical settings that we simulate.¹⁷

Perfect competition is of course an approximation that will be imperfect in many relevant markets. We maintain this assumption, consistent with much prior work, to simplify the problem and provide a benchmark for thinking about cross-margin interactions.¹⁸

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

To provide intuition for equilibrium in our setting, we build up from the classic case in EFC, which includes only H and U as plan options.¹⁹ The EFC equilibrium can be seen in figure 2a if one ignores the W_L curve. It is defined by the intersection of W_H and AC_H , which determines the competitive equilibrium price. Absent an L plan, any s -type whose WTP for H exceeds the price of H will buy H , and all other s -types will opt to remain uninsured.

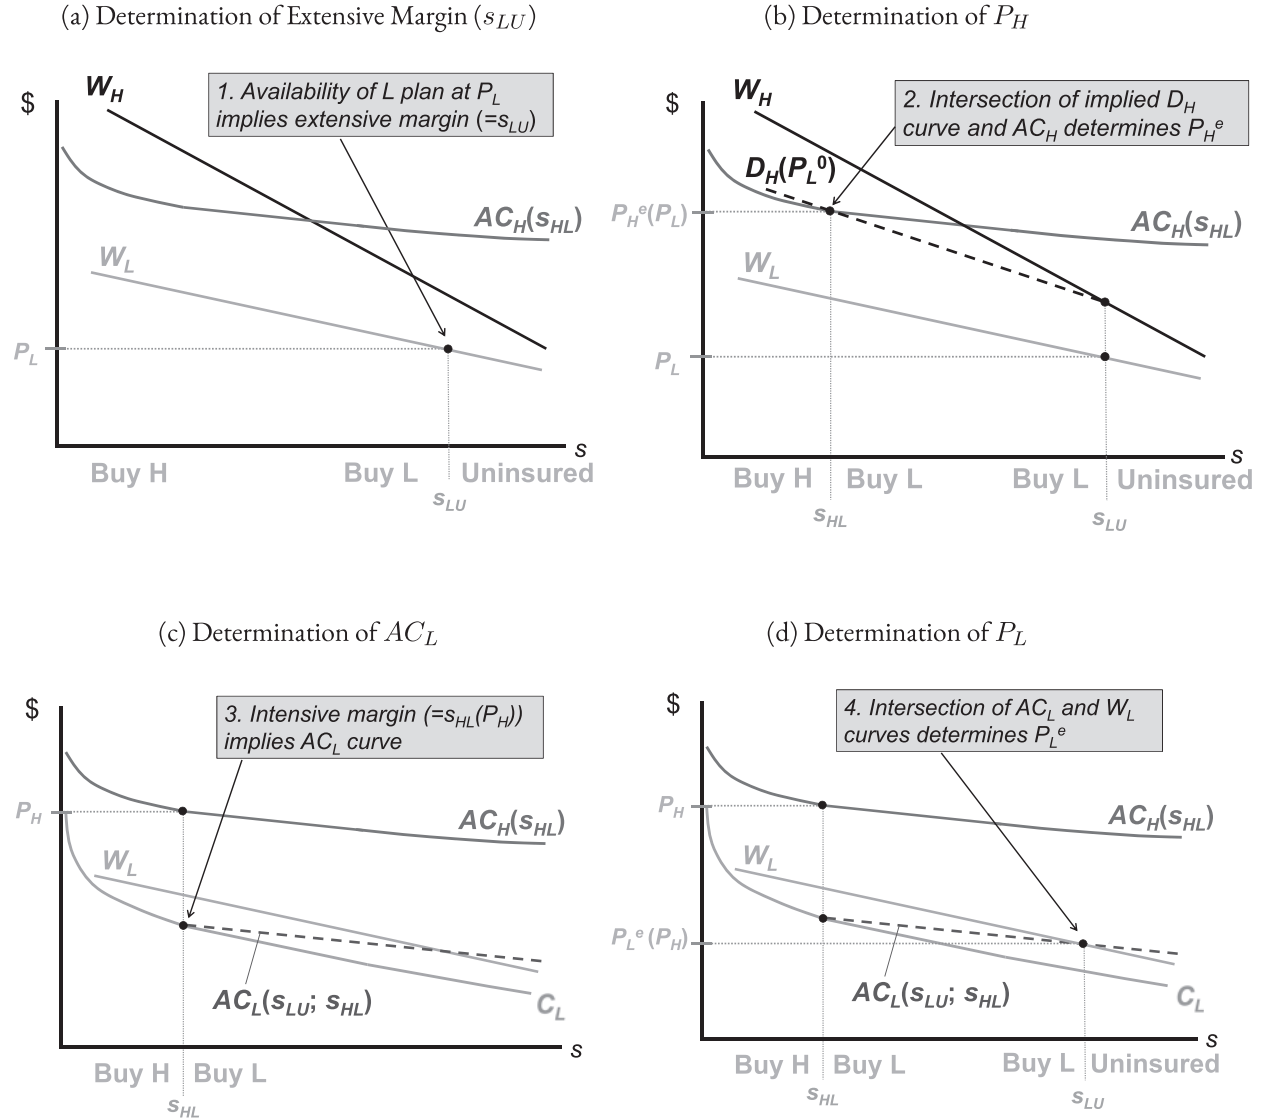
We next add L to the EFC choice set. To illustrate the equilibrium, we proceed in four steps, corresponding to the four panels in figure 2. Figures 2a and 2b show how P_H is determined, given a fixed price of L . Figure 2a shows that the fixed P_L implies a given extensive margin cutoff, s_{LU} . Figure 2b shows that this in turn implies an H plan demand curve, $D_H(P_L)$ (dashed). The intersection of $D_H(P_L)$ with H 's average cost curve determines P_H and the intensive margin cutoff

¹⁷A detailed discussion of these requirements and an algorithm for empirically identifying the RE are provided in appendixes C.3 and C.4, respectively.

¹⁸If there is free entry into both the H and the L contracts, prices will equal average costs in equilibrium, and there will be no cross-subsidization across the H and L contracts within a single firm. See the proofs in appendix A of Handel et al. (2015) and Azevedo and Gottlieb (2017). The intuition is that in such a setting, if one firm tried to cross-subsidize the adversely selected H contract with the L contract, another firm would enter the market and provide only the L contract at a lower price, with no need to cross-subsidize. This intuition would work less well in settings with a single fixed cost of firm entry, regardless of how many plans are offered.

¹⁹The correct analogy from EFC to our framework is a choice between H and U (rather than H and L) because the key feature of U is that its price is exogenously determined, like the lower coverage option in the EFC setting.

FIGURE 2.—DETERMINATION OF EQUILIBRIUM WITH H, L, AND OUTSIDE OPTION



Figures show how competitive equilibrium is determined in the vertical model with H and L plans and an outside option (uninsurance). Panels a and b show the determination of $P_H(P_L)$: a value of P_L implies the extensive margin (s_{LU}), which in turn implies the demand curve for H and the equilibrium P_H . Panels c and d show the determination of $P_L(P_H)$: a value of P_H implies the intensive margin (s_{HL}), which implies AC_L and the equilibrium value of P_L .

s_{HL} . This process determines the reaction function $P_H^e(P_L)$, the break-even price of H for a given price of L.

Figures 2c and 2d show how P_L is determined, given a fixed P_H . Figure 2c shows that the fixed P_H implies a given intensive margin cutoff (s_{HL}), which in turn fixes the AC_L curve. Figure 2d shows how the intersection of AC_L with W_L determines P_L and the extensive margin cutoff s_{LU} . This process determines the reaction function $P_L^e(P_H)$, which gives the break-even price of L for each price of H.

In equilibrium, the reaction functions must equal each other: $P_H = P_H^e(P_L)$ and $P_L = P_L^e(P_H)$. Figure 3 depicts the equilibrium, including the AC_L and D_H curves as dashed lines. These dashed lines are themselves equilibrium outcomes, even holding fixed consumer preferences and costs. In other words, there were many possible "iso- s_{HL} " AC_L curves and many possible "iso- P_L " D_H curves. The equilibrium vectors


of prices are the prices at which demand for L generates the equilibrium $D_H(P_L^e)$ and this demand for H simultaneously implies the equilibrium $AC_L(s_{HL})$ curve.

D. Social Welfare

We now show how our framework can be used to assess the welfare consequences of different policies. We define social welfare in the conventional way, as total social surplus (willingness-to-pay minus social resource cost). In order to make the figures simpler and more intuitive, we set C_U , the social cost of uninsurance, equal to 0. We nonetheless allow for a positive social cost of uninsurance in our empirical application below.

To build intuition, we start in figure 4a by illustrating the case where L is a pure cream skimmer. That is, L has low

FIGURE 3.—FINAL EQUILIBRIUM





The graph shows the final equilibrium under the vertical model with two plans (H and L) and an outside option (U). The dots mark the key intersections defining equilibrium prices and sorting. The intersection of AC_L and W_L determines P_L^e and the extensive margin type (s_{LU}^e). The D_H curve starts at this extensive margin (where it equals W_H), and its intersection with AC_H determines P_H^e and the intensive margin type (s_{HL}^e). This s_{HL}^e type marks the start of the AC_L curve (where it equals C_L).

average costs because it attracts low-cost individuals, but it has no causal effect on costs, so $C_L = C_H$ for any individual. For this case, given W_H , W_L , and $C_L = C_H$, we can find total social surplus for any allocation of consumers across plans described by the equilibrium cutoff values s_{HL}^e and s_{LU}^e .

Figure 4a shows that social surplus consists of two pieces. The first piece ($ABHG$) is the social surplus for consumers purchasing H , given by the area between W_H and $C_L = C_H$ for consumers with $s < s_{HL}$. The second piece ($EFIH$) is the social surplus for consumers purchasing L , given by the area between W_L and $C_L = C_H$ for consumers with $s \in [s_{HL}, s_{LU}]$. Figure 4a also illustrates forgone surplus for the allocation of consumers across plans. Here, the forgone surplus consists of three components. The first is the forgone surplus due to the fact that consumers with $s \in [s_{HL}, s_{LU}]$ purchased L when they would have generated more surplus by purchasing H , and it is described by the area between W_H and W_L for these consumers ($BCFE$). The second component is the forgone surplus due to the fact that consumers with $s > s_{LU}$ did not purchase insurance when they would have generated positive surplus by purchasing H , and it is described by the area between W_H and $\max\{W_L, C_L\}$ ($CDJF$). We refer to these two components as “intensive margin loss.” The third component is the forgone surplus due to the fact that consumers with $s \in [s_{LU}, s_{LU}^*]$ did not purchase insurance when they would have generated positive surplus by purchasing L , and it is described by the area between W_L and C_L for those consumers.

The figure thus shows how our graphical framework can be used to estimate welfare for any allocation of consumers across H , L , and U . Further, the framework makes it easy to determine the optimal allocation of consumers between insurance and uninsurance and between H and L . In the case of the particular demand and cost primitives drawn in figure 4a, the optimal allocation of consumers across plans is for all consumers to be in H . If H were not available, however,

FIGURE 4.—WELFARE

(a) Welfare when L Is a Pure Cream-Skimmer(b) Welfare when L Has a Cost Advantage

The graphs show welfare given equilibrium prices P^e and implied consumer sorting between H , L , and uninsured. Panel a shows the case where the L plan is a pure cream skimmer ($\Delta C_{HL} = C_H(s) - C_L(s) = 0$), while panel b shows the case where L has a causal cost advantage ($\Delta C_{HL} > 0$). The market surplus is shaded (light); the loss due to intensive margin misallocation (between H and L) is shaded (dark); and the loss due to extensive margin misallocation (between L and U) is shaded in thatched (darkest).

the optimal allocation of consumers across L and U would consist of all consumers with $s < s_{LU}^*$ purchasing L and all other consumers remaining uninsured.

In figure 4b, we apply our framework to the case where it is efficient for some consumers to be in L rather than in H and for others to remain uninsured. To do this, we change the assumption that L is a pure cream skimmer and instead assume that costs in H are higher than in L for each consumer and that the cost gap is constant across consumers: $\Delta C_{HL}(s) \equiv C_H(s) - C_L(s) = \delta > 0$. Intuitively, in this scenario, consumers prefer H because it provides more or better services—at a higher cost to the insurer. It is convenient to define a new curve $W_H^{Net}(s) = W_H(s) - \Delta C_{HL}(s)$, or WTP for H net of the incremental cost of H versus L . Under the

assumption that δ is constant, $W_H^{Net}(s)$ will be parallel to and below W_H . This is shown in figure 4b: as L 's cost advantage over H increases, W_H^{Net} shifts further down.²⁰

Given this new W_H^{Net} curve, social welfare is still fully characterized by the three curves, W_H^{Net} , W_L , and C_L , and social surplus and forgone surplus are defined in a similar manner to figure 4a. Social surplus still consists of two components. The first is the surplus generated by the consumers enrolled in H , and it is characterized by $ABHG$, the area between W_H^{Net} and C_L for consumers with $s < s_{HL}$.²¹ This component is smaller than it was in figure 4a due to the fact that now H has higher costs than L . In figure 4b, it is thus less socially advantageous for these consumers to be enrolled in H versus L . The second component is the surplus generated by the consumers enrolled in L , and it is characterized exactly as before by $EFIH$, the area between W_L and C_L for consumers with $s_{HL}^e < s < s_{LU}^e$. Forgone surplus is illustrated in figure 4b, similar to the illustration in figure 4a.²² In summary, figure 4 shows how our model can accommodate settings in which it is not socially efficient for all consumers to be enrolled in H or even in L , such as settings where there is moral hazard or administrative costs, for example.

Appendix B.3 derives a formal expression for welfare, allowing for cases where C_U is non-0, for example, if the outside option involves social costs like uncompensated care. This derivation formalizes what is shown graphically in figure 4.

III. Two-Margin Impacts of Risk Selection Policies

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

A. Benefit Regulation

We start by examining benefit regulation. In figure 5, we consider a rule that eliminates L plans from the market. This

thought experiment captures a variety of policies that set a binding floor on plan quality—for example, network adequacy rules, caps on out-of-pocket limits, and the ACA's "essential health benefits." These policies seek to address intensive margin adverse selection problems by eliminating low-quality, cream-skimming plans. But as we show, they can also have unintended extensive margin consequences.

Figure 5a shows the baseline equilibrium with both H and L plans, while figure 5b shows equilibrium with L plans eliminated, which reduces to the classic EFC equilibrium. Figure 5c shows the welfare impact of benefit regulation. This involves two competing effects: some consumers formerly in L shift to H (the intended consequence), and some consumers formerly in L become uninsured (the unintended consequence).

In the textbook cream-skimming case, where H is the socially efficient plan for everyone (though most consumers still generate more social surplus in L versus U), these two effects have opposing welfare consequences. The first (intended) effect increases social surplus by shifting people out of L —an inefficient plan that exists only by cream-skimming—and into H . The second (unintended) effect, however, lowers social surplus by shifting some L consumers into uninsurance. Thus, even in this textbook case where the L plan is an inefficient cream skimmer, banning it has ambiguous welfare consequences.²³

What explains this counterintuitive result? This can be thought of as an example of "theory of the second best"-style interactions that emerge with two margins of selection. Regulation that bans a pure cream-skimming L plan addresses an intensive margin selection problem. But it has the unintended side effect of worsening the extensive margin selection problem of too much uninsurance. Put differently, a pure cream-skimming L plan adds no social value within the market, but by segmenting the healthiest people into a low-price plan, it can improve welfare by bringing new consumers into the market.²⁴

B. Mandate Penalty on Uninsurance

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

The mandate penalty has both a direct effect and an indirect effect through equilibrium price adjustments. The direct

²⁰Heterogeneity in L 's cost advantage across s types could also be accommodated and would result in W_H^{Net} not being parallel to W_H .

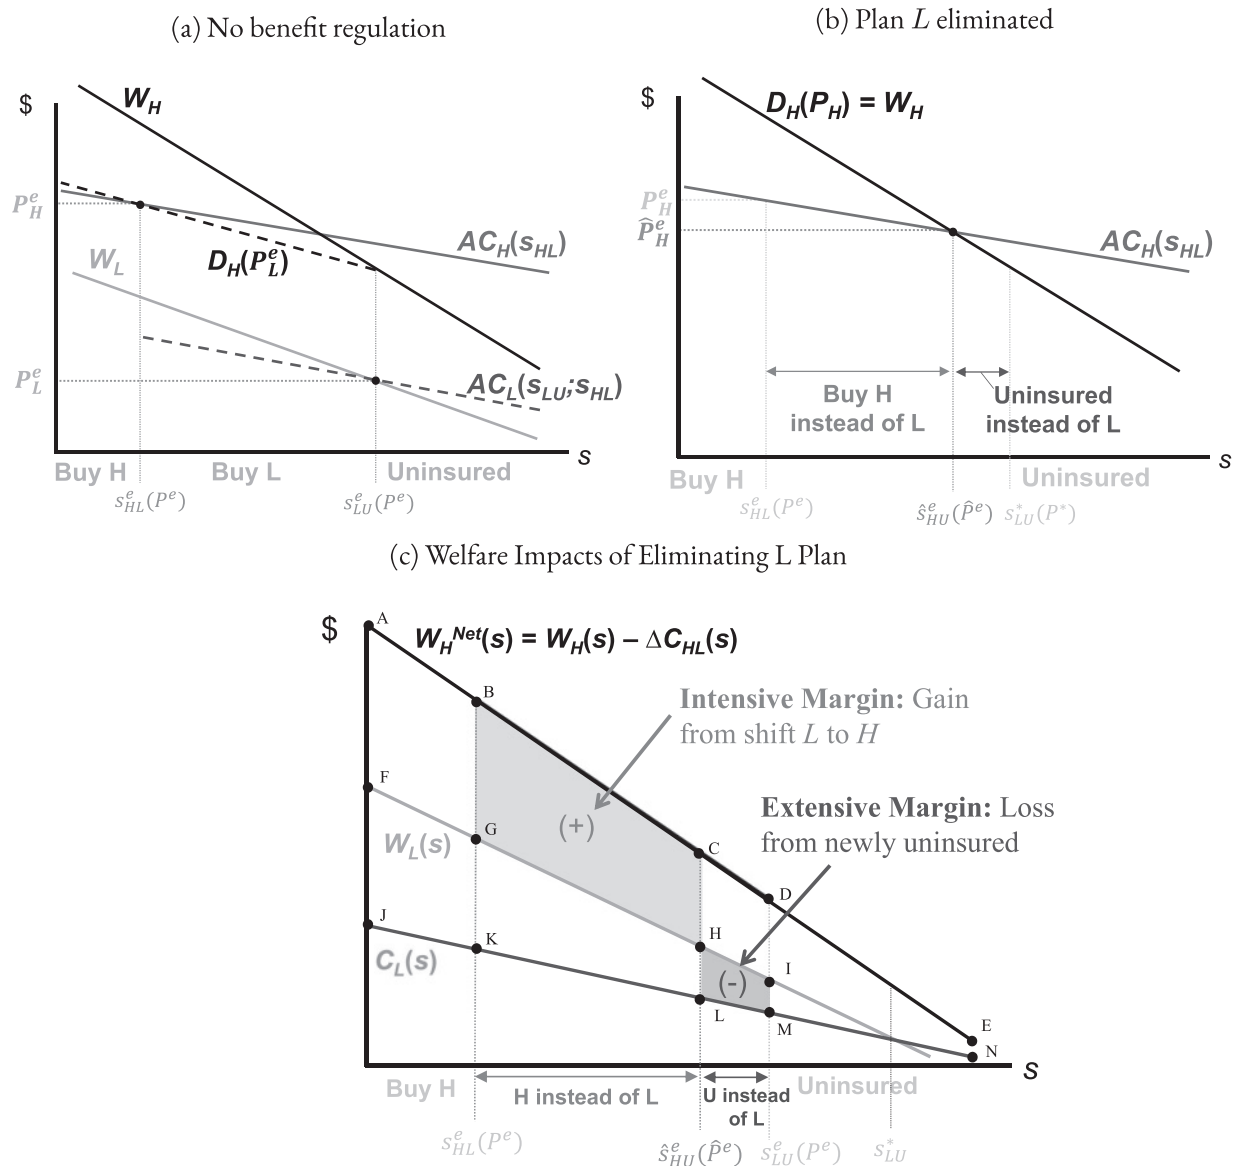
²¹To see this, note that this gap is equal to $W_H^{Net}(s) - C_L(s) = W_H(s) - (C_H(s) - C_L(s)) - C_L(s) = W_H(s) - C_H(s)$.

²²Here, forgone surplus again consists of two components. The first is the forgone intensive margin surplus due to the fact that consumers with $s \in [s_{HL}^e, s_{HL}^*]$ are enrolled in L but would generate more surplus if they were enrolled in H . It is characterized by the area between W_H^{Net} and W_L for these consumers (BKE). (Unlike in figure 4a, with H 's higher costs, it is now inefficient for any consumer with $s > s_{HL}^*$ to enroll in H .) The second component represents the extensive margin forgone surplus, and it is identical to the extensive margin forgone surplus in figure 4a.

²³The net welfare impact depends on the market primitives (W_H , W_L , C_H , C_L) and the social cost of uninsurance, C_U . Section II presents the framework for how these can be measured and the net welfare impact quantified.

²⁴Of course, this reasoning depends on the market stabilizing to a separating equilibrium where both H and L survive. If the market unravels to the L plan, insurance coverage will typically not be higher: the price of L will not be low (since it attracts all consumers), and because the quality of L is lower, uninsurance will typically be higher than in an H -only equilibrium where L is banned. Whether the market stabilizes to a separating equilibrium or unravels to L or upravel to H depends on the market primitives.

FIGURE 5.—IMPACT OF BENEFIT REGULATION



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

effect of a mandate penalty is to increase the demand for insurance. Figure 6a shows this via an upward shift in W_L and W_H by $\$M$, reflecting that both become cheaper relative to U (whose utility and price are normalized to 0). As a result of this shift, some people who were previously uninsured buy insurance in the L plan. This is the intended effect of the penalty.

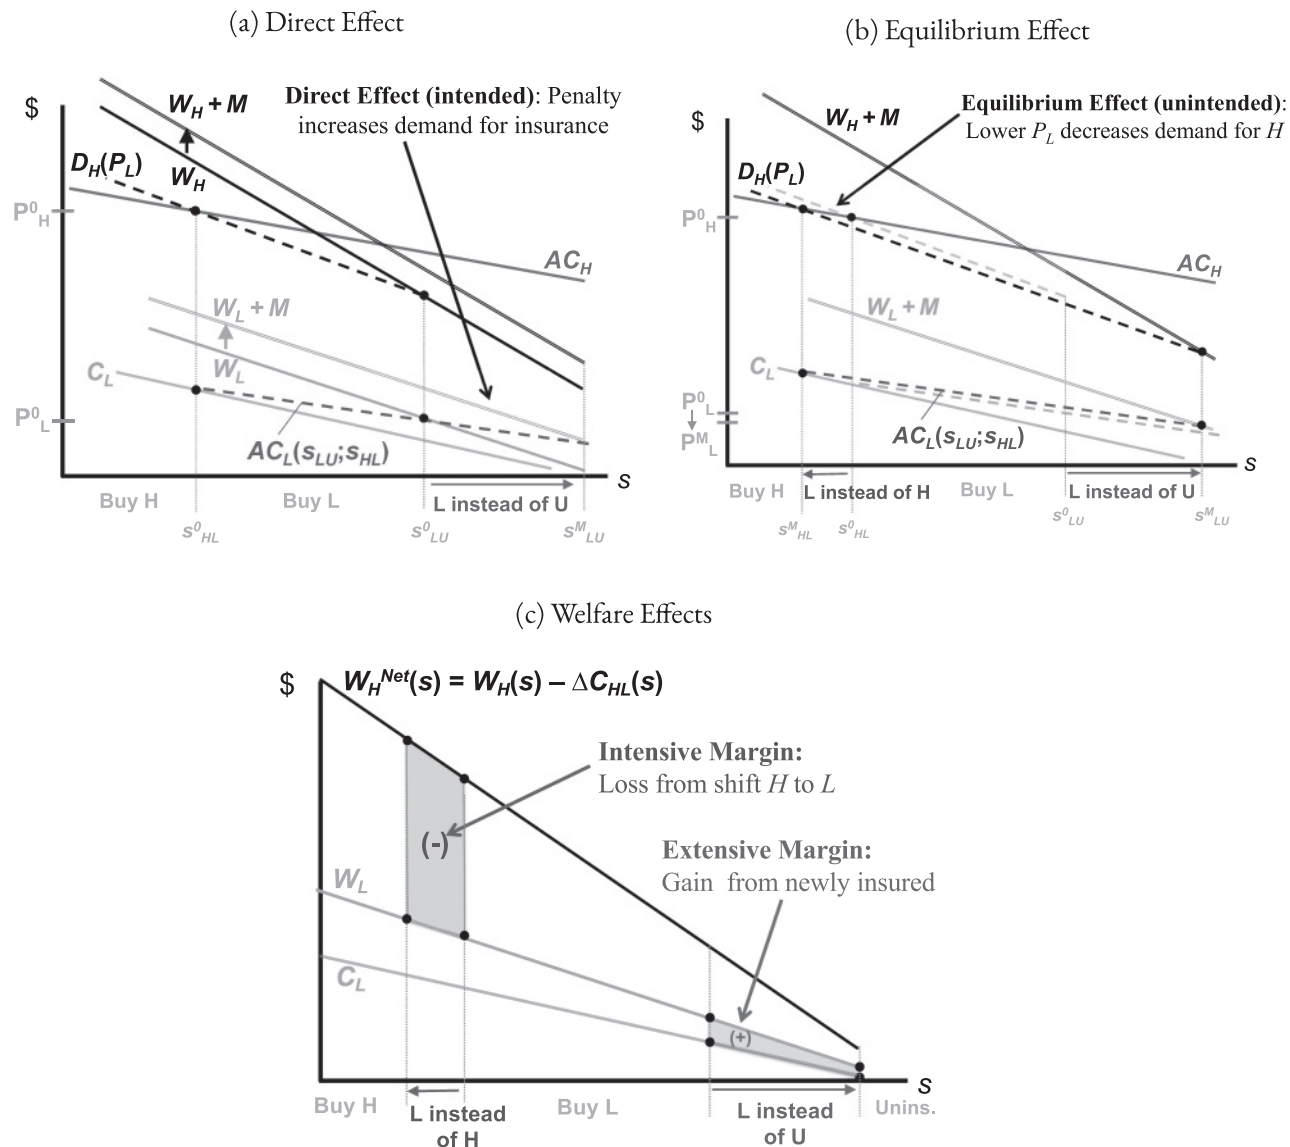
Figure 6b depicts the unintended, equilibrium effects of the penalty. By definition under extensive margin adverse selection, the newly insured individuals are relatively healthy. Because they buy the low-price L plan, they lower L 's average costs (i.e., a movement down the AC_L curve, not a shift in the AC_L curve) and therefore its price. The lower P_L leads some consumers to shift on the intensive margin from H to

L —as captured by the downward shift in H 's demand curve, $D_H(P_L)$. This is the main unintended effect of the penalty: although it is intended to reduce uninsurance, the penalty also shifts people toward lower-quality plans on the intensive margin.²⁵

There is a second equilibrium effect from this shift in consumers from H to L . The consumers who shift are high cost relative to L 's previous customers, pushing up L 's average costs. In figure 6b, this is depicted via an upward shift in the

²⁵We show in our simulations and in appendix A that this prediction is largely robust to relaxing the vertical model. It is driven by two properties: that the newly uninsured are relatively healthy (extensive margin adverse selection) and that the newly insured mostly choose the low-priced L plan.

FIGURE 6.—IMPACT OF MANDATE PENALTY ON UNINSURANCE



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

$AC_L(P_H)$ curve, which has to occur because of the higher P_H and the leftward shift in the marginal s_{HL} type. The higher average costs in L partly offset the fall in P_L due to the mandate and dampen the impact of the mandate on the price of L . Thus, our model shows how and why cross-margin effects may make a mandate less effective than one would predict from its direct effects alone: the penalty induces healthy people to enter the market but also induces relatively sick people to move from H to L . Nonetheless, as long as the original equilibrium is stable, one can show that on net, a larger penalty decreases P_L and uninsurance (see appendix A for a formal derivation).


Figure 6c shows the welfare effects in the textbook case where H is the efficient plan for all consumers. There are again competing effects: (intended) welfare gains from newly

insured consumers and (unintended) welfare losses from consumers moving from H to the lower-quality L plan. Thus, the interaction of the two margins of selection makes the welfare impact of a mandate ambiguous even in this textbook case. In the extreme, a penalty could even lead to a market where high-quality contracts are unavailable to consumers (i.e., market unraveling to L).

C. Risk Adjustment Transfers

Next we consider the impact of implementing risk adjustment, including the effects of strengthening or weakening risk adjustment transfers relative to the status quo. Of the three policies we consider, risk adjustment is the most difficult to illustrate graphically because the policy adds new

FIGURE 7.—WELFARE EFFECTS OF RISK ADJUSTMENT



The figure shows the welfare effects of a risk-adjustment policy that shifts consumers on the intensive margin from L to H (by lowering $P_H - P_L$) and on the extensive margin from L to U (by raising P_L). We show a case where H is globally more efficient than L , so the intensive margin shift is welfare improving, but where U is sometimes more efficient than L . Optimal sorting across the extensive margin occurs when $s_{LU}^e = \hat{s}_{LU}^e$.

risk-adjusted cost curves (for both L and H) that crowd the figure. (See figure A2 in the appendix.)

In the ACA marketplaces, the per-enrollee transfer to plan j is determined by a formula of the form²⁶

$$T_j(P) = \left(\frac{\bar{R}_j(P)}{\bar{R}(P)} - 1 \right) \cdot \bar{P}(P), \quad (4)$$

where $\bar{R}_j(P)$ is the average risk score of the consumers enrolling in plan j given price vector P , $\bar{R}(P)$ is the (share-weighted) average risk score among all consumers purchasing insurance, and $\bar{P}(P)$ is the (share-weighted) average price in the market. The transfer is positive as long as j 's average risk score is larger than $-j$'s average risk score. The sum of H 's and L 's transfers is always 0, making the transfer system budget neutral. Note that risk adjustment here is imperfect in the sense of not necessarily eliminating all variation in net enrollee costs.²⁷ This is consistent with our empirical findings below.

To understand the impact of risk adjustment on the two-margin problem, we tune its strength by introducing a parameter α . We define the transfer from L to H as $\alpha \times T(P)$.

²⁶The 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.

²⁷Perfect risk adjustment, where transfers exactly capture all variation in C_L across consumer types, is a useful thought experiment. But in practice markets include an imperfect form of risk adjustment, where transfers are based on individual risk scores computed from diagnoses appearing in health insurance claims. See Geruso and Layton (2020) for an overview. See the appendix for more discussion of the case of perfect risk adjustment.

With $\alpha = 0$, there is no risk adjustment. With $\alpha = 1$, there is ACA-level risk adjustment. Other values magnify or attenuate these transfers. For example, if a risk adjustment transfer were \$500 under $\alpha = 1$, it would be \$600 under $\alpha = 1.2$. Importantly, changes to α do not imply changes to the underlying risk scores (which are determined by enrollee diagnoses). Adjusting α corresponds to ongoing policy activity, as we discuss below.

In appendix A, we derive comparative statics describing the effect of an increase in α (i.e., a magnification of the imperfect transfers) on P_H and P_L . These comparative statics mimic the simulations we perform in the empirical section, where we simulate equilibria under no risk adjustment and with increasingly large risk-adjustment transfers (i.e., increasingly large values for α). Larger values of α unambiguously lower the price of H . The effect of an increase in α on the price of L , however, is ambiguous. In addition to risk adjustment's direct effect to push up L 's average costs by transferring money from L to H , there is a second, indirect effect. The consumers who shift from L to H tend to be L 's most expensive enrollees, even net of imperfect risk-adjustment transfers. This lowers L 's risk-adjusted average costs, pushing the price of L downward. This indirect effect will be larger when intensive margin adverse selection is severe (even after risk adjustment) and when consumers are highly price elastic on the intensive margin. Indeed, we find in some of our simulations that the indirect effect is large, and risk adjustment has minimal effects or even decreases P_L . We defer further discussion of the comparative statics to the results section.

Figure 7 depicts the welfare effects of a risk adjustment policy where the direct effect dominates such that the policy

shifts consumers from H to L and also has some effect on the extensive margin, shifting consumers from L to U . Again, we illustrate welfare for the textbook case where H is the efficient plan for all. As with benefit regulation and the mandate penalty, there are opposing effects: a welfare gain from the intensive margin shift from L to H and a welfare loss from the extensive margin shift from L to uninsurance. (There is also a welfare gain on the extensive margin due to the fact that some of the people induced to choose uninsurance instead of L generate negative social surplus when enrolled in L .) This suggests that, like the other policies, the welfare effects of risk adjustment are theoretically ambiguous.

D. Other Policies

The same price theory can be applied to other policies not explicitly discussed above, such as reinsurance. The key insight is that anything that affects selection on one margin has the potential to affect selection on the other margin, as firms adjust prices in equilibrium to compensate for the changing consumer risk pools.

Further, cross margin effects are relevant not only for policies that aim to address selection, but also for policies for which selection impacts are incidental or a nuisance. Handel (2013), for example, shows how addressing inertia through “nudging” can exacerbate intensive margin selection in an employer-sponsored plan setting. Our model implies that in other market settings, where uninsurance is a more empirically relevant concern, there is a further effect of nudging: worsening risk selection on the intensive margin (i.e., increasing the market segmentation of healthy enrollees into L and sick enrollees into H) through behavioral nudges may improve risk selection on the extensive margin by pushing down the equilibrium price of L . This may counterbalance the welfare harm documented in Handel (2013). Similar insights apply to any behavioral intervention that even incidentally affects the sorting of consumer risks (expected costs) across plans.²⁸ Similarly, behavioral interventions intended to increase take-up of insurance, such as information interventions or simplified enrollment pathways, may have important intensive margin consequences similar to the effects of a mandate.

IV. Simulations: Methods

To demonstrate how our model can be applied empirically, we draw on previously estimated model primitives from two separate Massachusetts pre-ACA individual health insur-

ance exchanges to simulate a hypothetical post-ACA market. Demand and cost curves from a low-income population are drawn from the subsidized health insurance exchange, known as Commonwealth Care (CommCare) as estimated by Finkelstein et al. (2019), which we abbreviate as FHS. A demand curve for higher-income individuals is drawn from the unsubsidized individual market “CommChoice” as estimated in Hackmann et al. (2015), which we abbreviate as HKK.²⁹ Our inclusion of both the low-income and high-income populations is motivated by the design of subsidies under the ACA. Low-income households receive subsidies that are linked to the price of insurance, a policy that limits cross-margin effects by fixing the extensive margin price of insurance. Higher-income households do not receive subsidies, meaning that cross-margin effects may be relevant. In order to capture these dynamics, we include both groups in our analysis. We apply the FHS cost curve to both populations. That is, people of a given s -type in either population would have the same expected cost conditional on plan.³⁰

We make two key modifications to the baseline FHS and HKK estimates. First, to allow for broader policy counterfactuals, we extrapolate the curves over the full range of s -types. Second, we combine the two sets of estimates to form one set of aggregated demand and cost curves, reflecting ACA markets that include subsidized (low-income) and unsubsidized (high-income) enrollees. Given these modifications, readers should consider these simulations illustrative of mechanisms rather than exact predictions for any specific market. The comingling of the subsidized and unsubsidized groups in the same market in our simulations is a choice aimed at illustrating the mechanisms we wish to highlight rather than as an accurate description of the Massachusetts market. Details on the construction of these demand and cost curves, as well as figures showing the final curves, are in appendix C.1.

Given these demand and cost curves, it is straightforward to estimate equilibrium prices and allocations of consumers across H , L , and U under a given set of policies. Our method for finding equilibrium is based on the approach described in figure 2a. We characterize equilibrium as a price vector P_H, P_L at which any plan that has nonzero enrollment breaks even. We then use a Riley equilibrium concept to choose which break-even price vector is the equilibrium price vector.³¹ This method results in a unique equilibrium for each policy environment we consider.

We then simulate market equilibrium under different specifications of two policies: a mandate penalty (ranging from

²⁸This is relevant not only as it relates to inertia (Polyakova, 2016) but also to misinformation (Handel & Kolstad, 2015), complexity (Ericson & Starc, 2016), and other behavioral concerns. It is also relevant for nonbehavioral policy changes in other markets, including Medicare. For example, Decarolis, Guglielmo, and Luscombe (2020) document that intensive margin risk selection was affected by a Medicare policy change that allowed midyear plan switching across Medicare Advantage plans. This could have extensive margin impacts on who chooses Medicare Advantage versus Traditional Medicare.

²⁹We import the HKK estimates to generate a demand curve for the high-income population, though in principle, simulating high income demand as an ad hoc shift or rotation to the estimated demand curve for the low-income population could have also served the purpose of illustrating the trade-offs in our model.

³⁰Both sets of demand and cost curves are well identified using exogenous variation in net consumer prices. FHS use a regression discontinuity design based on three household income cutoffs that generate discrete changes in consumer subsidies. HKK use a difference-in-differences design leveraging the introduction of an uninsurance penalty in Massachusetts.

³¹See appendix C.4 for additional details.

\$0 to \$60 per month) and risk-adjustment transfers (ranging from zero to three times the size of ACA transfers). We study the effects of these policies in a 2×2 matrix of market environments. The first dimension of the environment we vary is subsidy design, with two regimes: “ACA-like” subsidies that are linked to the price of the cheapest plan and “fixed” subsidies set at an exogenous dollar amount.³² In both cases, low-income consumers receive subsidies only if they purchase H or L , and the subsidy is identical for both plans. High-income consumers do not receive subsidies.

The second dimension we vary is whether L is a pure cream skimmer (i.e., $C_L(s) = C_H(s)$ for all s) or has a cost advantage. FHS find no evidence that L has lower costs than H in CommCare, motivating our cream-skimmer case. To illustrate another possibility, we simulate the case where L has a 15% cost advantage (i.e., $C_L(s) = 0.85C_H(s)$). Of particular interest is how the welfare consequences of risk adjustment and the uninsurance penalty vary across these two cases. We explore these in section VI.

V. Simulation Results: Prices and Enrollment

In this section, we present results on how prices and market shares change under stronger mandate penalties and stronger risk adjustment. In appendix D.2, we also present results on how prices and market shares change under benefit regulation, where we implement benefit regulation by eliminating L from the consumers’ choice set. In appendices D.4.1 and D.4.2, we explore the sensitivity of our results to relaxing the vertical model and modifying the primitives (specifically, consumers’ incremental WTP for H versus L), finding that the key results are quite robust. In presenting results, we vary consumer characteristics (demand and costs/selection), supply-side features (horizontal differentiation among plans), and policy interventions (mandates and subsidies, risk adjustment) to generate a catalog of findings that provide guidance on how these features interact to affect equilibrium prices and enrollment.

A. Mandate and Uninsurance Penalties

The first four panels of figure 8 present equilibrium market shares for each option, H , L , and U , under different levels of a mandate penalty for remaining uninsured ($P_U \equiv M$). We consider penalties in increments from \$0 to \$60, applied equally to both the subsidized and unsubsidized populations.³³ In all

cases we include ACA-style risk adjustment (described in detail in section VB). The top two panels of figure 8 contain the results for the case where L is a pure cream skimmer. The bottom two panels contain results for the case where L has a 15% cost advantage. The cases with ACA-like price-linked subsidies are shown in the left panels, and the cases with a fixed subsidy are in the right panels.³⁴ All results are also reported in appendix table A1.

For the two ACA-like subsidy cases (left), the patterns are qualitatively similar regardless of modeling L as a cream skimmer (top) or as having a cost advantage (bottom). When there is no mandate penalty, some consumers choose each of the three options, H , L , and U , though the share in H is extremely low in the cost advantage case. As the penalty increases, the uninsurance rate decreases, with no consumers remaining uninsured at a penalty of \$60 per month. However, there are also intensive margin consequences: as the penalty increases, there is a shift of consumers from H to L . In the case where L is a pure cream skimmer, H ’s market share decreases from 42% with no penalty to 23% with a penalty of \$60 per month. This represents a significant decline in H ’s market share and a significant deterioration of the average generosity of coverage among the insured. When L has a 15% cost advantage (bottom), the patterns are similar, though H ’s initial market share with no penalty is much lower (about 2%).

The two fixed subsidy cases are presented in figures 8b and 8d. When L is a pure cream skimmer (top), with zero penalty consumers are split across H , L , and U . As the penalty increases from zero consumers move from U to L , the intended effect of the policy. At a penalty of just under \$30 per month, the influx of inexpensive consumers into L causes P_L to get low enough that some consumers switch from H to L . As the penalty continues to increase, consumers move into L from both U and H until the mandate reaches just over \$40 per month and all consumers are insured. At this point, 23% of the market is enrolled in H , and 77% of the market is enrolled in L . This represents an intended decline in the uninsurance rate from 35% to 0% but also an unintended decline in H ’s market share from 42% to 23%.³⁵

In each of the cases in figures 8a to 8d, a larger mandate penalty has the intended consequence of decreasing uninsurance *and* the unintended consequence of shifting consumers from H to L .³⁶ This is consistent with implications of our

³²For stronger mandate penalties we follow the ACA rules by setting the subsidy such that the net-of-subsidy price of the index plan equals 4% of income for consumers at 150% of the federal poverty line (FPL) in 2011 (or \$55 per month), the year on which our estimated demand and cost curves are based. The ACA subsidy rules actually link the subsidy 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 .

³³We find that in all cases studied here, $P_U = 60$ is sufficient to drive the uninsurance rate to 0 in the presence of ACA risk adjustment transfers.

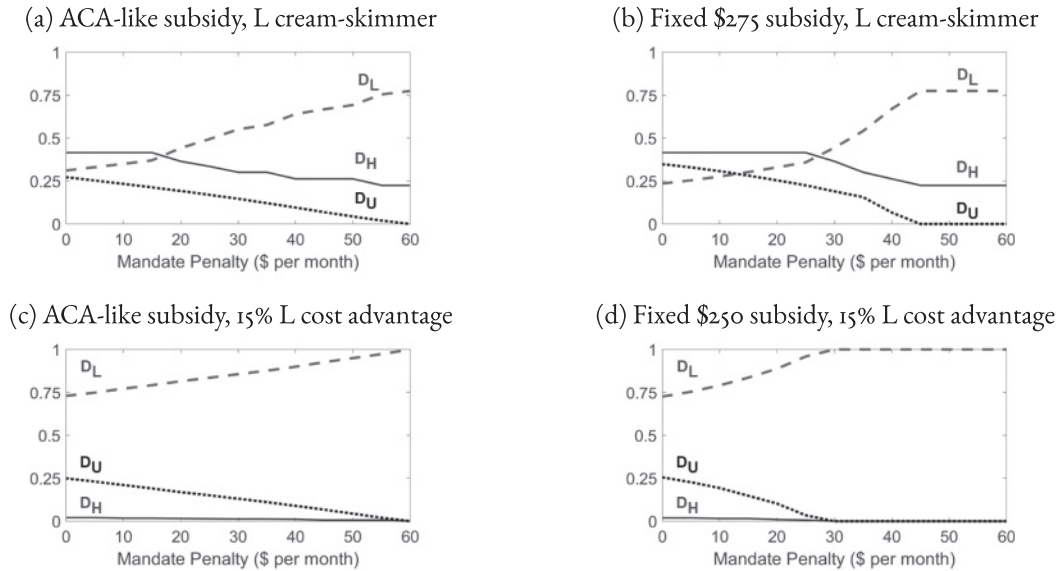
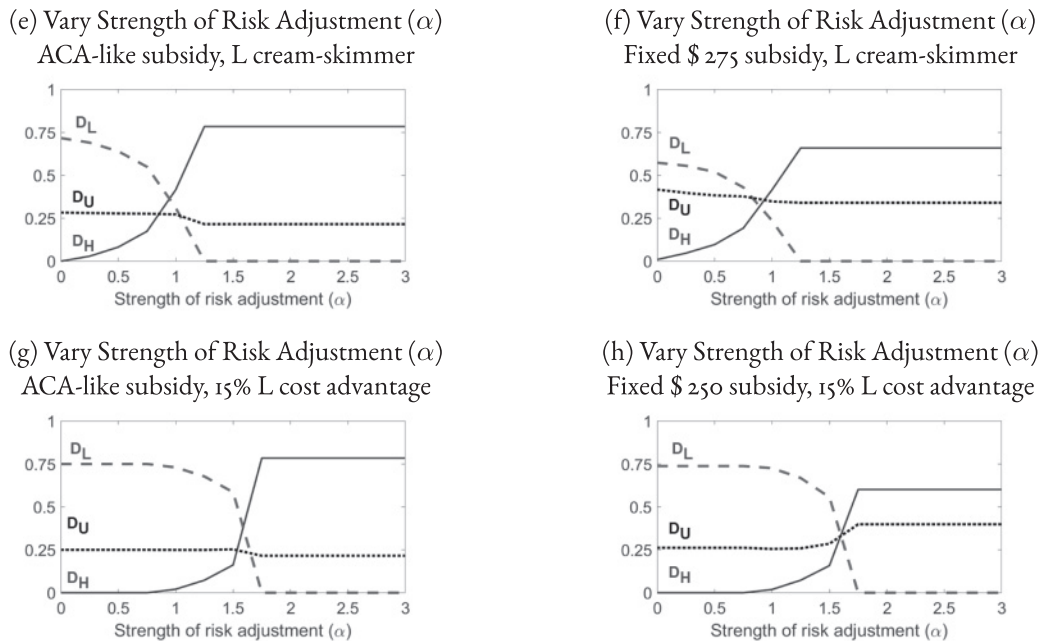
³⁴Fixed subsidies are equal to \$275 in the case where L is a pure cream skimmer and \$250 in the case where L has a 15% cost advantage. These values were chosen in order to ensure that risk adjustment and the uninsurance penalty have some effect on market shares. With subsidies that are “too large,” no consumers opt to be uninsured, and with subsidies that are “too small,” no consumers opt to purchase insurance, making the simulated policy modifications uninformative.

³⁵In the case where L has a 15% cost advantage, the penalty again decreases both the uninsurance rate (intended) and H ’s market share (unintended), but H ’s market share with a \$0 penalty is so low (around 3.5%) that the decline in H ’s market share (to 0) is relatively insignificant.

³⁶This finding also holds when we relax the vertical assumptions of the model, as we explore further in appendix D.4.1 and show in appendix figure

FIGURE 8.—MARKET SHARES VARYING SINGLE POLICY PARAMETERS

Mandate Penalty Simulations

Risk Adjustment (α) Simulations

Panels a to d show market shares for H , L , and uninsurance (U) from our simulations with varying sizes of the mandate penalty (x-axis, in \$ per month). Panels e to h show market shares for H , L , and uninsurance (U) from our simulations with varying strength of risk adjustment α (on the x-axis). As described in text, α is a multiplier on the risk adjustment transfer: $\alpha = 0$ implies no risk adjustment; $\alpha = 1$ is baseline risk adjustment using the ACA formula; and $\alpha > 1$ is overadjustment. The panels represent different subsidy designs and specifications for the L plan's causal cost advantage versus H (i.e., ΔC_{HL}). In panels a, b, e, and f, L is a pure cream skimmer ($\Delta C_{HL} = 0$), while in panels c, d, g, and h, L has a 15% cost advantage. Panels a and c have ACA-like subsidies linked to the price of L , while panels b and d have fixed subsidies of the indicated dollar amounts.

graphical model as well as the comparative statics we outline in sections II and III. The unintended intensive margin effect is starkest when L is a perfect cream skimmer, highlighting

how market primitives can amplify the cross-margin impacts of policy changes.³⁷

A10. In addition, in appendix D.4.2, we show that these results are robust to varying the incremental WTP for H versus L .

³⁷To see why the effect is larger for the cream-skimmer case, note that for fixed preferences, it is more difficult to achieve high enrollment in H when L has an actual cost advantage versus when L has similar costs to

B. Risk Adjustment

We now consider the effects of risk adjustment. We start with risk-adjustment transfers implied by the ACA risk adjustment transfer formula (see equation [4]). We first calculate risk scores for each individual using the HHS-HCC risk-adjustment model used in the ACA marketplaces. (This is a straightforward mechanical application of the regulator's algorithm to our individual-level claims data.) We then use those scores plus the FHS regression discontinuity design to estimate a risk score curve $RA(s)$ describing the average risk score across consumers of a given s -type. Because this curve is novel to this paper and not estimated by FHS, we describe the estimation of it in appendix C.2. We plot this curve alongside the cost curve in appendix figure A5. It is apparent that while risk scores explain part of the correlation between willingness to pay and costs, they do so only imperfectly. Specifically, we find that risk scores account for about one-third of the correlation between willingness to pay and costs, implying substantial selection on costs net of the ACA's imperfect risk adjustment policy. (Although incidental to our aims here, this is a novel finding.)

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

To vary the strength of risk-adjustment transfers, we maintain the original risk scores and structure of the transfer formula, but we multiply transfers by a scalar α (as in the discussion in section IIIC and comparative statics in appendix A) so that transfers from L to H are some multiple of the transfers implied by the ACA formula (i.e., $P_H = AC_H(p) - \alpha T(P)$ and $P_L = AC_L(P) + \alpha T(P)$). We allow α to vary from zero (no risk adjustment) to three (risk adjustment transfers 3 times the size of ACA transfers). The case of ACA transfers occurs where $\alpha = 1$. In these risk-adjustment simulations, we are not modifying the fit of risk adjustment or changing the scores in any way. Instead, we are enhancing the transfer implied by the same scores so that if a plan's risk-adjustment transfer was \$500 under $\alpha = 1$, it is \$600 under $\alpha = 1.2$. This approach to evaluating strengthening or weakening risk adjustment reflects real-world policy experimentation: In the early years of the ACA marketplaces, the federal government reduced α from 1 to 0.85 and gave states some flexibility to further reduce α with appropriate justification.³⁸ Our approach thus

H . This leads to lower enrollment in H even with a small penalty and less opportunity for a reduction in H 's market share.

³⁸The reduction of α from 1 to 0.85 occurred when the federal government decided to "remove administrative costs" from the benchmark premium that

maps to feasible policy interventions rather than assuming that the regulator can increase the predictive power of risk scores.

Equilibrium market shares for different levels of α in the cases without and with a cost advantage for L are found in the third and fourth rows of figure 8, respectively. Market shares under ACA-like subsidies are presented in the left panels, and market shares under fixed subsidies are found in the right panels. Results are also found in appendix table A2. With ACA-like subsidies, patterns are qualitatively similar when L is a pure cream skimmer and when L has a 15% cost advantage. In both cases, when there is no risk adjustment ($\alpha = 0$), the market unravels to L : no consumers choose H , and the market is split between L and uninsurance. As the strength of risk-adjustment transfers increases, consumers shift from L to H . This is the intended consequence of risk adjustment. When L is a pure cream skimmer, transfers about 1.25 times the size of ACA transfers are sufficient to cause the market to "upravel" to H . When L has a 15% cost advantage, transfers need to be 1.6 times the size of ACA transfers to generate the same outcome. In both cases, there is no extensive margin effect except at the level of α where the market initially upravel to H . At that point, there is a small reduction in the uninsurance rate. This reduction is due to the fact that there, the subsidy becomes linked to the (higher) price of H instead of the (lower) price of L due to the exit of L from the market. With the larger subsidy, more consumers purchase insurance.³⁹

Figures 8f and 8h present market shares under fixed subsidies with different levels of α . Here, we again see that stronger risk-adjustment transfers have the intended effect: Higher levels of α result in more consumers choosing H instead of L . In the case where L is a pure cream skimmer, we see only a small extensive margin effect, with a small decrease in the uninsurance rate as α increases. This is consistent with our comparative statics from section III: the direct effect of increasing the transfer from L to H is more than fully offset by the indirect effect of the costliest (net of imperfect risk adjustment) L enrollees leaving L and joining H , resulting in a decrease in P_L and a corresponding decrease in the uninsurance rate. (See section III and appendix A for a fuller discussion of this result.)

On the other hand, in the case where L has a 15% cost advantage, we see a different unintended extensive margin consequence of stronger risk-adjustment transfers: more consumers opt to remain uninsured. In this case, with no risk adjustment ($\alpha = 0$), all insured consumers opt for L , with no

multiplies insurer risk scores to determine transfers in the transfer formula described by equation (4).

³⁹This reduction seemingly goes against the intuition we present in section III where we showed that in many cases, risk adjustment may increase the uninsurance rate rather than decrease it as we see here. Note, however, that in the cases here, the subsidy is linked to the extensive margin price. This results in risk adjustment having no effect on the net-of-subsidy extensive margin price faced by the low-income consumers (except where L exits the market), limiting (and in this case eliminating) any unintended extensive margin consequence.

consumers choosing H and the market split between L and U . ACA risk-adjustment transfers ($\alpha = 1$) barely alter these market shares. As transfers are strengthened above ACA levels, consumers begin to opt for H instead of L . At the higher levels of α , extensive margin consequences also start to appear, with some consumers exiting the market and opting for uninsurance. When transfers are strengthened to two times the size of ACA transfers, the market upravel to H with all insured consumers opting for H instead of L . At $\alpha = 2$, the uninsurance rate reaches almost 50%, an increase of 15 percentage points (60%) compared to the case with no risk adjustment. This indicates that this shift of consumers to more generous coverage on the intensive margin had a substantial extensive margin impact. We show that the same result holds when we relax the vertical model assumptions in appendix figure A10.⁴⁰

These results provide important lessons for where the unintended extensive margin effects of risk adjustment will matter most. First, ACA-like price-linked subsidies protect against the unintended extensive margin effects of risk adjustment, even when those subsidies are only targeted to the low-income consumers making up 60% of the market (though there may be important effects on the size of the subsidies themselves, and thus government costs). Second, the unintended extensive margin effects are more likely to occur when L has a larger cost advantage. In cases where L and H have similar costs, extensive margin effects are likely to be small. But when L has a large cost advantage, stronger risk adjustment can have significant effects on the portion of consumers who opt to be uninsured.

VI. Simulation Results: Welfare

We next analyze the changes in social surplus associated with the policy simulations of section V. We characterize welfare at a baseline equilibrium, then trace the gains and losses associated with illustrative policy changes, and finally determine optimal policy. Importantly, we show that the optimal mandate penalty depends on the strength of risk adjustment and vice versa. One straightforward implication is that if mandate penalties were altered by legislative action or court outcomes, a constrained optimal response from a regulator would be to adjust risk-adjustment strength in concert. (Unlike mandate penalties, regulators typically have authority to tune risk adjustment without legal changes.)

We begin by noting the possibility that in many settings, social surplus may not be increased by policies that raise insurance take-up or move consumers from less generous to more generous coverage. This is because some consumers may not value insurance (or more generous coverage) more than its incremental cost. Further, policies may have opposing effects on the intensive and extensive margins, increasing

enrollment in more generous coverage while simultaneously decreasing overall insurance take-up, or vice versa. For these reasons, it is important to understand the effects of policies not just on market allocations (which section V presents) but also on welfare.

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

$$C_U(s) = \frac{(1-d)C_H(s)}{1+\phi} + \omega, \quad (5)$$

where d is the share of total uninsured health care costs that the uninsured pay out of pocket, ϕ is the assumed moral hazard from insurance, and ω is some fixed cost of uninsurance. For d and ϕ , we use the values as derived from Finkelstein et al. (2019) and assume that $d = 0.2$ and $\phi = 0.25$.⁴¹ We set the fixed cost $\omega = -\$97$ per month, which is the ω value consistent with 95% of the population being optimally insured when L has a 15% cost advantage.

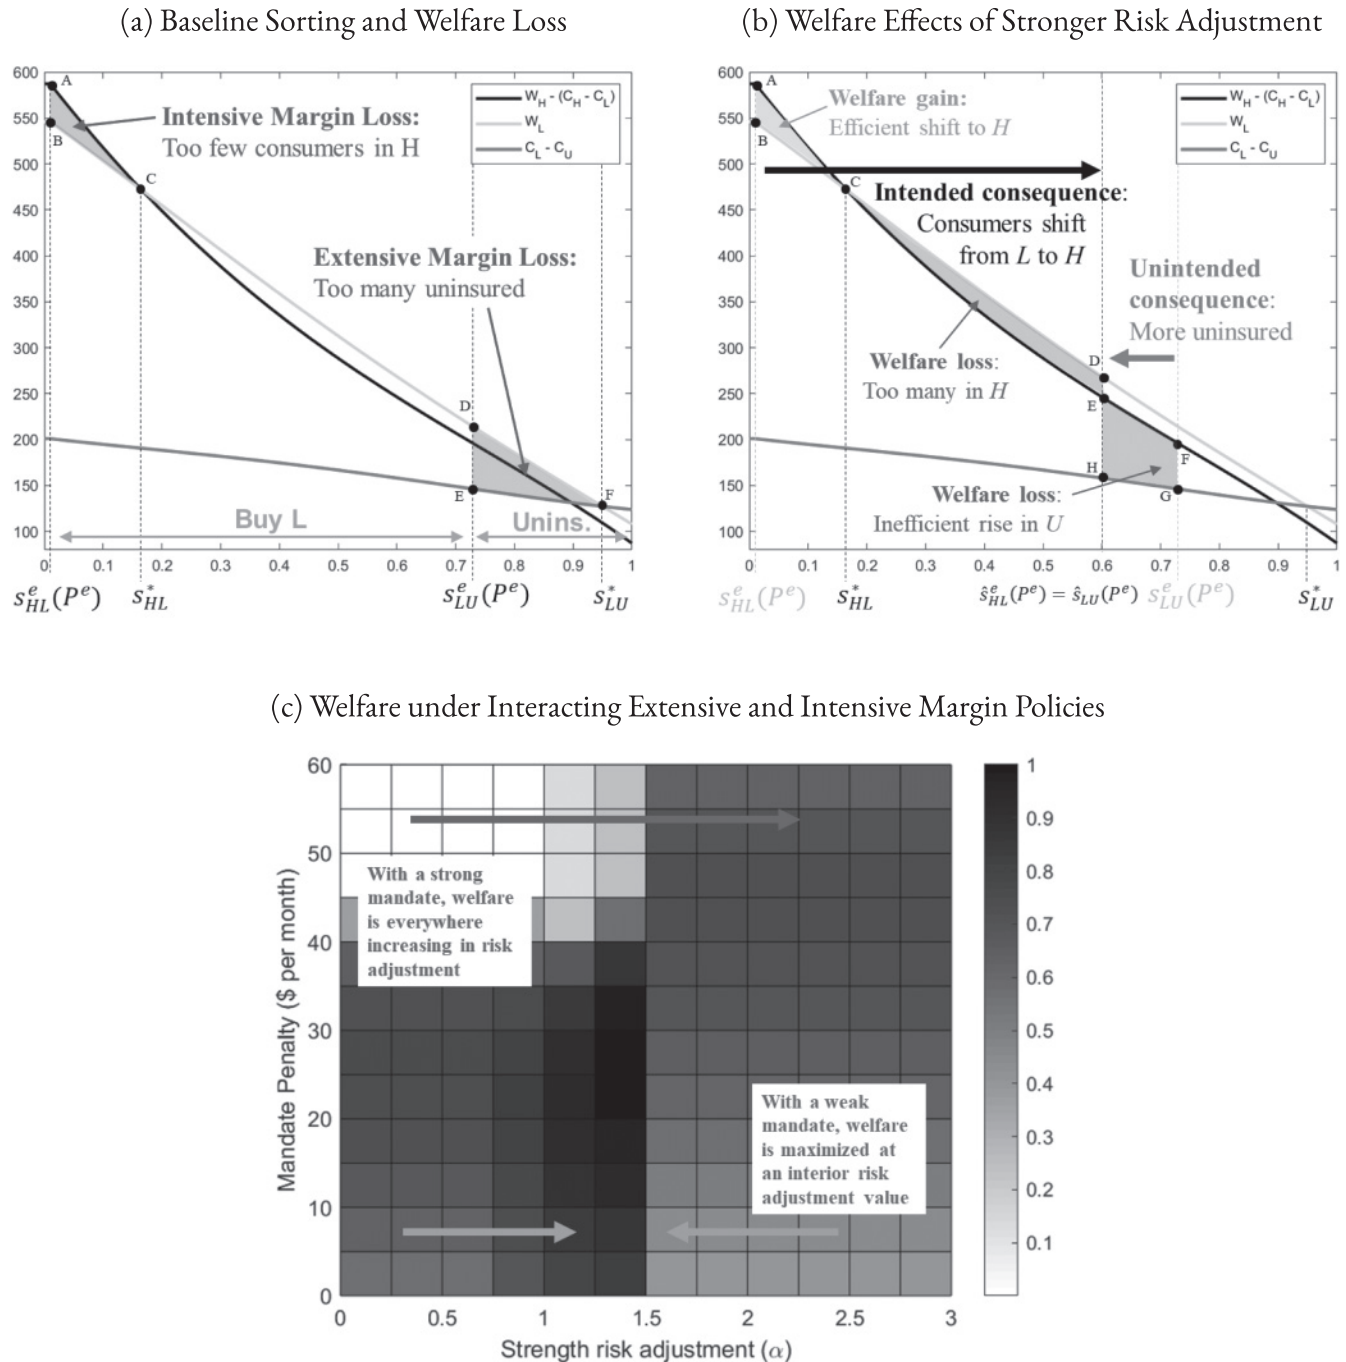
Before analyzing welfare, we provide an important caution: as is standard in the literature, welfare estimation depends on inferring consumer value from observed demand responses. In other words, our welfare estimates are accurate only to the extent that demand accurately reflects true valuations. Behavioral frictions might cause consumer demand to deviate from valuations (Handel, Kolstad, & Spinnewijn, 2019). Liquidity constraints could also cause valuation and demand to diverge. A separate issue is that our specification of C_U is ad hoc and may not reflect the actual social costs of uninsurance. Indeed, many of our welfare conclusions will necessarily be sensitive to assumptions about C_U . (See the results with alternative assumptions on C_U in appendix D.3.2.) We present this analysis to illustrate how to apply our framework but are cautious about drawing strong normative conclusions.⁴²

⁴¹We note that without this assumption (i.e., if we assume $C_U = 0$), it is inefficient for any consumer to purchase insurance, as no consumer values either H or L more than the cost of enrolling them in H or L . This fact plus a full discussion of the derivation of the assumed values of d and ϕ can be found in Finkelstein et al. (2019).

⁴²Importantly, considerations about choice frictions or about the difficulty of measuring C_U do not threaten the use of our model for the positive analysis of section V, which consists of predictions of prices and market shares

⁴⁰In appendix D.4.1, we explore the sensitivity of these results to the vertical model assumption, finding that the results are robust to modest relaxation of the assumption. See figure A10. Also, in appendix D.4.2, we show that these results are largely robust to varying the incremental WTP for H versus L .

FIGURE 9.—EMPIRICAL WELFARE EFFECTS FROM SIMULATIONS



In both panels a and b, we assume that there is a fixed subsidy equal to \$250 and L has a 15% cost advantage over H . Further, 60% of the population is low-income and 40% of the population is high-income, so WTP curves are weighted sums of both types. Panel a shows welfare losses in this setting under no mandate and $\alpha = 1$, relative to efficient sorting. Efficient cutoffs are indicated with a * while equilibrium outcomes are denoted with an e superscript. Panel b shows welfare changes under a risk-adjustment policy where $\alpha = 2$, relative to the baseline risk-adjustment policy where $\alpha = 1$. Panel c shows social welfare outcomes (darker = higher welfare) from the model simulations under different parameters for the strength of risk adjustment (α , x-axis) and for the size of the uninsurance mandate penalty (\$ per month, y-axis). The optimum for one policy depends on the other: with weak risk adjustment, a weaker mandate is optimal, while with strong risk adjustment, a strong mandate is optimal.

A. Welfare and Changes to Risk Adjustment

We now show how to estimate welfare with our graphical model. For parsimony, we focus on the case of strengthen-

under different counterfactual mandate penalties and risk adjustment. Such predictions do not rely on assumptions about C_U or about demand reflecting underlying consumer valuation.

ing risk adjustment transfers. In appendix D.3, we show the case of an uninsurance penalty. Figure 9 plots the empirical analogs to our welfare figures from section II. Figure 9a depicts forgone surplus relative to the social optimum under a baseline case with ACA risk adjustment ($\alpha = 1$), no mandate penalty, and a fixed subsidy equal to \$250. Figure 9b depicts the difference in social surplus between the baseline

case and a similar case where risk adjustment is strengthened ($\alpha = 2$), reflecting the simulation reported in figure 8h. Instead of plotting C_L , we plot $C_L^{Net} = C_L - C_U$, as in equation (18) to account for the fact that $C_U \neq 0$. We also plot $W_H^{Net} = W_H - (C_H - C_L)$ as in section II.

In figure 9a, we indicate the equilibrium s cutoffs for $\alpha = 1$. The intensive margin equilibrium cutoff is s_{HL}^e , and the extensive margin cutoff is s_{LU}^e . Thus, consumers with $s < s_{HL}^e$ enroll in H , consumers with $s_{HL}^e < s < s_{LU}^e$ enroll in L , and consumers with $s > s_{LU}^e$ remain uninsured.

Efficient sorting of consumers across options is indicated by s^* cutoff types. Consumers with $s < s_{HL}^*$ should be in H , consumers with $s_{HL}^* < s < s_{LU}^*$ should be in L , and the few consumers with $s > s_{LU}^*$ should be uninsured to maximize social surplus. In figure 9a, we depict the forgone surplus in the baseline ACA setting with shaded areas. Intensive margin forgone surplus (lost surplus due to consumers choosing L instead of H) is indicated by the welfare triangle ABC , representing a welfare loss of \$19.71.⁴³ Extensive margin forgone surplus is represented by the welfare triangle DEF . Welfare loss on this margin amounts to \$33.47. Combining these, the (average per consumer) forgone surplus in the baseline setting in figure 9a is thus \$53.18.

Figure 9b shows the welfare consequences of strengthening risk adjustment. To show the effects of strengthening risk adjustment, we increase α from 1 to 2, so that risk adjustment transfers are increased to two times the ACA transfers. We hold all other policy parameters fixed. Recall from the bottom-right panel of figure 8 that moving from $\alpha = 1$ to $\alpha = 2$ in this setting shifts nearly 60% of consumers in the market from L to H but also shifts 13% of consumers in the market from L to U . Overall, no consumers remain in L when $\alpha = 2$.

The first effect of increasing α is the intended consequence of risk adjustment, and here it implies both welfare gains and losses. Welfare gains occur when consumers whose incremental valuation for H versus L exceeds the incremental cost of H versus L (i.e., those with $W_H^{Net}(s) > W_L(s)$) enroll in H instead of L . These gains are represented by the welfare triangle ABC , and they amount to \$19.71. Welfare losses occur when consumers whose incremental valuation for H versus L is less than the incremental cost of H versus L (i.e., those with $W_H^{Net}(s) < W_L(s)$) enroll in H instead of L as L unravels. These offsetting welfare losses occur when “too many” consumers enroll in H , and they are represented by the welfare triangle CDE and amount to \$19.24. In other settings, where it is always more efficient for consumers to be enrolled in H instead of L (such as the pure cream-skimming case), there will only be welfare gains on this margin. In the case of figure 9b, the two effects nearly cancel each other out so that the net welfare gain due to the intended consequence of shifting consumers from L to H amounts to just \$0.47.

The second effect of increasing α is the unintended consequence of risk adjustment, and here it implies welfare losses. Because risk adjustment leads to a higher price of L , some consumers exit the market, increasing the uninsurance rate. In this case, all consumers who exit the market value insurance more than the (net) cost of insuring them, $C_L^{Net} = C_L - C_U$, causing the welfare consequences of this shift of consumers out of the market to be unambiguously negative. The size of the welfare loss is represented by the area of $EF GH$, which we estimate to be \$68.30. Combining the intended and unintended consequences of risk adjustment, we estimate that in this setting, doubling risk adjustment transfers by shifting from $\alpha = 1$ to $\alpha = 2$ would decrease welfare by \$67.83, on average per consumer.

Welfare results for all settings studied in figure 8, for the full range of levels of α , and under different assumptions about C_U are found in appendix D.3.2. These results indicate that under our baseline assumption of C_U with ACA-like subsidies, increasing the strength of risk adjustment transfers always improves welfare when L is a pure cream skimmer. In this case, there is no effect of risk adjustment on the extensive margin due to the linkage of the subsidy to the price, leaving only intensive margin consequences. The intensive margin effects of moving consumers from L to H are also unambiguously positive, as it is inefficient for any consumer to be enrolled in L versus H . When L has a cost advantage, increasing the strength of risk-adjustment transfers improves welfare given low initial levels of α but decreases welfare given higher initial levels of α , with the welfare-maximizing risk-adjustment policy having an α around 1.25, or 1.25 times the strength of ACA risk-adjustment transfers. This nonmonotonic result is due to the fact that increases in α from low initial levels of α induce only consumers who value H highest relative to L to enroll in H , with consumers whose incremental WTP does not exceed their incremental cost remaining enrolled in L .

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

⁴³These shapes are more triangle-ish than triangular.

B. Optimality under Interacting Policies

The findings above suggest the necessity of a second-best approach to policy: optimal extensive margin policy (penalties and subsidies) will often depend on the intensive margin policies (risk adjustment and benefit regulation) currently in use in a market. Here we show how our model can be used to assess optimal policy, allowing for these interactions.

We again consider uninsurance penalties and risk adjustment. We compute social welfare over a grid of uninsurance penalties and levels of α . We do this for the case in which L has a 10% cost advantage and low-income consumers (who comprise 60% of the market) receive a fixed subsidy equal to \$250 when purchasing insurance. The social cost of uninsurance is once again set to $C_U(s) = 0.25C_H(s) - 97$ as in the previous section. We cherry-pick this case because the two policies interact in interesting ways. For completeness, we perform similar analyses for all other settings studied in figure 8. Results are reported in appendix D.3.

Figure 9 presents the welfare estimates graphically as a heat map, where darker areas represent higher values of social surplus.⁴⁴ Under a 10% cost advantage, the socially efficient allocation is for 33% of the population to be in H , 60% of the population to be in L , and the remainder to be uninsured. We can examine how the optimal level of risk adjustment changes with different values of the mandate penalty. The figure shows that in this setting, when the mandate penalty is high, welfare is increasing in the strength of risk adjustment (i.e., higher α). At these high values of the mandate penalty, all consumers purchase insurance, eliminating any potential unintended extensive margin consequences. Under such high market enrollment, it is optimal to use strong risk adjustment to sort more people into H instead of L . With low levels of the mandate penalty, however, risk adjustment has important unintended extensive margin consequences. Thus, the benefits of shifting consumers from L to H must be traded off against the costs of shifting consumers out of the market and into U . The results in figure 9 indicate that with a small penalty, social surplus is maximized at $1.25 < \alpha < 1.5$, somewhat stronger than ACA risk adjustment but weaker than the optimal level of α under a strong penalty, which is > 1.5 .

We can also use figure 9 to consider the optimal mandate penalty for each level of α . With weak risk adjustment, starting from low levels of the mandate penalty, social surplus is increasing in the size of the penalty. However, starting from high levels of the penalty, the sign is opposite, with social surplus increasing rapidly as the penalty is reduced. This occurs because while a strong mandate penalty increases social surplus by inducing consumers to enroll in insurance, it also has the offsetting effect of shifting consumers from H to L . Ulti-

mately, an intermediate penalty level (around \$30) maximizes social surplus, though any level of the penalty below \$40 achieves much higher levels of social surplus than the level achieved by a penalty exceeding \$40. When risk adjustment is strong, social surplus is increasing in the mandate penalty. Here, strong risk adjustment causes the market to “upravel” to H , eliminating any potential unintended intensive margin consequences of increasing the level of the penalty. With strong risk adjustment, a stronger mandate thus only induces consumers to move from U to H , generating higher levels of social surplus.

In terms of optimal policy, figure 9 reveals that social surplus is highest for an intermediate level of both the uninsurance penalty and risk adjustment. Given such a combination of policies, consumers sort themselves to each of H , L , and U , the socially efficient outcome in this particular setting. Note that the lowest-surplus combinations are a strong mandate with weak risk adjustment or a weak mandate with strong risk adjustment.

In appendix D.3, we show that other settings have different optimal policies. In the case where L is a pure cream skimmer and subsidies are linked to prices (ACA-like subsidies), optimal policy is to have strong risk adjustment (high α) and a weak mandate. In the case where L has a cost advantage, a weak mandate with weak to moderate risk adjustment is the optimal policy. In all cases, it is clear that these two policies interact with each other, implying that evaluating one policy in isolation from the other can be misleading. Specifically, market designers should not only consider consumer preferences for high- versus low-quality coverage and consumer valuation of insurance but also the interaction between intensive and extensive margin selection when determining the optimal combination of policies.

VII. Conclusion

Adverse selection in insurance markets can occur on either the extensive (insurance versus uninsurance) or intensive (more versus less generous coverage) margin. While this possibility has long been recognized, most prior treatments of adverse selection focus on only one margin or the other, missing important cross-margin trade-offs inherent to many selection policies. In some cases, the unintended effects of policies are first-order with respect to welfare. This happens most often with a penalty for choosing to be uninsured. In particular, strengthening uninsurance penalties can increase insurance take-up while shifting some consumers from higher- to lower-quality coverage. Likewise, strengthening risk adjustment transfers can shift enrollment toward higher-quality coverage while also increasing uninsurance.

The simplicity of our approach is not without some costs. The assumption of perfect vertical ordering of demand is required to maintain simplicity in our graphs, though we show in both theory and empirics that our results are largely robust to relaxing this assumption. What matters is that the primary form of plan differentiation is vertical. Conclusions may

⁴⁴Consider a given α , mandate combination that generates a level of welfare $W(\alpha, \text{mandate})$. We scale/normalize the heat map shading as follows: $W^{\text{norm}}(\alpha, \text{mandate}) = \frac{W(\alpha, \text{mandate}) - \min(W)}{\max(W) - \min(W)}$, where the maximum and minimum are taken over all possible $(\alpha, \text{mandate})$ combinations for the setting.

differ in more complex cases, which are an important area for future research.

The issues we highlight are relevant for future reform of individual health insurance markets in the United States. Many have observed that the quality of coverage available in these settings is low, with most plans having tight provider networks, high deductibles, and strict utilization controls. Additionally, others have observed that take-up is far from complete, with many young and healthy consumers remaining uninsured (Domurat, Menashe, & Yin, 2021). These two observations are consistent with adverse selection on the intensive and extensive margins, respectively. Our framework highlights the unfortunate but important point that budget-neutral policies targeting one of these problems tend to exacerbate the other due to the trade-off between extensive and intensive margin selection. This point is often absent from reform discussions, and our intention is to correct this potentially costly omission.

REFERENCES

- Akerlof, George A., "The Market for 'Lemons': Quality Uncertainty and the Market Mechanism," *Quarterly Journal of Economics* 84:3 (1970), 488–500. 10.2307/1879431
- Azevedo, Eduardo, and Daniel Gottlieb, "Perfect Competition in Markets with Adverse Selection" *Econometrica* 85:1 (2017), 67–105. 10.3982/ECTA13434
- Bundorf, M. Kate, Jonathan Levin, and Neale Mahoney, "Pricing and Welfare in Health Plan Choice" *American Economic Review* 102:7 (2012), 3214–3248. 10.1257/aer.102.7.3214
- Carey, Colleen, "Technological Change and Risk Adjustment: Benefit Design Incentives in Medicare Part D," *American Economic Journal: Economic Policy* 9:1 (2017), 38–73. 10.1257/pol.20140171
- Cutler, David M., and Sarah J. Reber, "Paying for Health Insurance: The Trade-Off between Competition and Adverse Selection," *Quarterly Journal of Economics* 113:2 (1998), 433–466. 10.1162/003355398555649
- Decarolis, Francesco, Andrea Guglielmo, and Clavin Luscombe, "Open Enrollment Periods and Plan Choices," *Health Economics* 29:7 (2020), 733–747. 10.1002/hec.4014
- Domurat, Richard, Isaac Menashe, and Wesley Yin, "The Role of Behavioral Frictions in Health Insurance Marketplace Enrollment and Risk: Evidence from a Field Experiment," *American Economic Review* 111:5 (2021), 1549–1574.
- Einav, Liran, and Amy Finkelstein, "Selection in Insurance Markets: Theory and Empirics in Pictures," *Journal of Economic Perspectives* 25:1 (2011), 115–138. 10.1257/jep.25.1.115
- Einav, Liran, Amy Finkelstein, and Mark R. Cullen, "Estimating Welfare in Insurance Markets Using Variation in Prices," *Quarterly Journal of Economics* 125:3 (2010), 877–921. 10.1162/qjec.2010.125.3.877
- Ericson, Keith M. Marzilli, and Amanda Starc, "How Product Standardization Affects Choice: Evidence from the Massachusetts Health Insurance Exchange," *Journal of Health Economics* 50 (2016), 71–85. 10.1016/j.jhealeco.2016.09.005
- Finkelstein, Amy, Nathaniel Hendren, and Mark Shepard, "Subsidizing Health Insurance for Low-Income Adults: Evidence from Massachusetts," *American Economic Review* 109:3 (2019), 1530–1567. 10.1257/aer.20171455
- Geruso, Michael, "Demand Heterogeneity in Insurance Markets: Implications for Equity and Efficiency," *Quantitative Economics* 8:3 (2017), 929–975. 10.3982/QE794
- Geruso, Michael, and Timothy J. Layton, "Selection in Health Insurance Markets and Its Policy Remedies," *Journal of Economic Perspectives* 31:4 (2017), 23–50. 10.1257/jep.31.4.23
- , "Upcoding: Evidence from Medicare on Squishy Risk Adjustment," *Journal of Political Economy* 128:3 (2020).
- Geruso, Michael, Timothy Layton, and Daniel Prinz, "Screening in Contract Design: Evidence from the ACA Health Insurance Exchanges," *American Economic Journal: Economic Policy* 11:2 (2019), 64–107. 10.1257/pol.20170014
- Glazer, Jacob, and Thomas G. McGuire, "Optimal Risk Adjustment in Markets with Adverse Selection: An Application to Managed Care," *American Economic Review* 90:4 (2000), 1055–1071. 10.1257/aer.90.4.1055
- Gruber, Jonathan, "Delivering Public Health Insurance through Private Plan Choice in the United States," *Journal of Economic Perspectives* 31:4 (2017), 3–22. 10.1257/jep.31.4.3
- Hackmann, Martin B., Jonathan T. Kolstad, and Amanda E. Kowalski, "Adverse Selection and an Individual Mandate: When Theory Meets Practice," *American Economic Review* 105:3 (2015), 1030–1066. 10.1257/aer.20130758
- Handel, Benjamin R., "Adverse Selection and Inertia in Health Insurance Markets: When Nudging Hurts," *American Economic Review* 103:7 (2013), 2643–82. 10.1257/aer.103.7.2643
- Handel, Benjamin R., Igal Hendel, and Michael D. Whinston, "Equilibria in Health Exchanges: Adverse Selection versus Reclassification Risk," *Econometrica* 83:4 (2015), 1261–1313. 10.3982/ECTA12480
- Handel, Benjamin R., and Jonathan T. Kolstad, "Health Insurance for Humans: Information Frictions, Plan Choice, and Consumer Welfare," *American Economic Review* 105:8 (2015), 2449–2500. 10.1257/aer.20131126
- Handel, Benjamin R., Jonathan T. Kolstad, and Johannes Spinnewijn, "Information Frictions and Adverse Selection: Policy Interventions in Health Insurance Markets," this REVIEW 101:2 (2019), 326–340.
- Hendren, Nathaniel, "Private Information and Insurance Rejections," *Econometrica* 81:5 (2013), 1713–1762. 10.3982/ECTA10931
- Lavetti, Kurt, and Kosali Simon, "Strategic Formulary Design in Medicare Part D Plans," *American Economic Journal: Economic Policy* 10:3 (2018), 154–192. 10.1257/pol.20160248
- Layton, Timothy J., Randall P. Ellis, Thomas G. McGuire, and Richard Van Kleef, "Measuring Efficiency of Health Plan Payment Systems in Managed Competition Health Insurance Markets," *Journal of Health Economics* 56 (2017), 237–255. 10.1016/j.jhealeco.2017.05.004
- Mahoney, Neale, and E. Glen Weyl, "Imperfect Competition in Selection Markets," this REVIEW 99:4 (2017), 637–651.
- Polyakova, Maria, "Regulation of Insurance with Adverse Selection and Switching Costs: Evidence from Medicare Part D," *American Economic Journal: Applied Economics* 8:3 (2016), 165–95. 10.1257/app.20150004
- Rothschild, Michael, and Joseph Stiglitz, "Equilibrium in Competitive Insurance Markets: An Essay on the Economics of Imperfect Information," *Quarterly Journal of Economics* 90:4 (1976), 629–649. 10.2307/1885326
- Saltzman, Evan, "Managing Adverse Selection: Underinsurance versus Underenrollment," *The RAND Journal of Economics* 52:2 (2021), 359–381.
- Shepard, Mark, "Hospital Network Competition and Adverse Selection: Evidence from the Massachusetts Health Insurance Exchange," *American Economic Review* 112:2 (2022), 578–615.
- Tebaldi, Pietro, "Estimating Equilibrium in Health Insurance Exchanges: Price Competition and Subsidy Design under the ACA," Research/BFA working paper (2017).
- Veiga, André, and E. Glen Weyl, "Product Design in Selection Markets," *Quarterly Journal of Economics* 131:2 (2016), 1007–1056. 10.1093/qje/qjw007